The Racket Reference

Version 5.3

Matthew Flatt
and PLT

August 6, 2012

This manual defines the core Racket language and describes its most prominent libraries. The companion manual ***The Racket Guide*** provides a friendlier (though less precise and less complete) overview of the language.

```
#lang racket/base
#lang racket
```

Unless otherwise noted, the bindings defined in this manual are exported by the `racket/base` and `racket` languages.

The `racket/base` library is much smaller than the `racket` library and will typically load faster.
The `racket` library combines `racket/base`, `racket/bool`, `racket/bytes`, `racket/class`, `racket/cmdline`, `racket/contract`, `racket/dict`, `racket/file`, `racket/function`, `racket/future`, `racket/include`, `racket/list`, `racket/local`, `racket/match`, `racket/math`, `racket/path`, `racket/place`, `racket/port`, `racket/pretty`, `racket/promise`, `racket/sequence`, `racket/set`, `racket/shared`, `racket/stream`, `racket/string`, `racket/system`, `racket/tcp`, `racket/udp`, `racket/unit`, and `racket/vector`.

Contents

1 Language Model **18**

1.1 Evaluation Model 18

1.1.1 Sub-expression Evaluation and Continuations 18

1.1.2 Tail Position 18

1.1.3 Multiple Return Values 19

1.1.4 Top-Level Variables 19

1.1.5 Objects and Imperative Update 21

1.1.6 Object Identity and Comparisons 23

1.1.7 Garbage Collection 23

1.1.8 Procedure Applications and Local Variables 24

1.1.9 Variables and Locations 26

1.1.10 Modules and Module-Level Variables 26

1.1.11 Continuation Frames and Marks 28

1.1.12 Prompts, Delimited Continuations, and Barriers 28

1.1.13 Threads 29

1.1.14 Parameters 29

1.1.15 Exceptions 30

1.1.16 Custodians 30

1.2 Syntax Model 31

1.2.1 Identifiers and Binding 32

1.2.2 Syntax Objects 33

1.2.3 Expansion (Parsing) 34

1.2.4 Compilation 44

1.2.5 Namespaces 45

1.2.6 Inferred Value Names 47
1.3 The Reader 48
1.3.1 Delimiters and Dispatch 48
1.3.2 Reading Symbols 50
1.3.3 Reading Numbers 51
1.3.4 Reading Booleans 52
1.3.5 Reading Pairs and Lists 52
1.3.6 Reading Strings 54
1.3.7 Reading Quotes 55
1.3.8 Reading Comments 56
1.3.9 Reading Vectors 56
1.3.10 Reading Structures 57
1.3.11 Reading Hash Tables 58
1.3.12 Reading Boxes 58
1.3.13 Reading Characters 58
1.3.14 Reading Keywords 59
1.3.15 Reading Regular Expressions 59
1.3.16 Reading Graph Structure 60
1.3.17 Reading via an Extension 60
1.4 The Printer 62
1.4.1 Printing Symbols 63
1.4.2 Printing Numbers 63
1.4.3 Printing Booleans 64
1.4.4 Printing Pairs and Lists 64
1.4.5 Printing Strings 65

1.4.6 Printing Vectors 66
1.4.7 Printing Structures 66
1.4.8 Printing Hash Tables 67
1.4.9 Printing Boxes 68
1.4.10 Printing Characters 68
1.4.11 Printing Keywords 68
1.4.12 Printing Regular Expressions 69
1.4.13 Printing Paths 69
1.4.14 Printing Unreadable Values 69
1.4.15 Printing Compiled Code 69

2 Syntactic Forms **71**

2.1 Modules: `module`, `module*`, ... 71
2.2 Importing and Exporting: `require` and `provide` 75
2.2.1 Additional `require` Forms 96
2.2.2 Additional `provide` Forms 100
2.3 Literals: `quote` and `#%datum` 100
2.4 Expression Wrapper: `#%expression` 101
2.5 Variable References and `#%top` 101
2.6 Locations: `#%variable-reference` 102
2.7 Procedure Applications and `#%app` 103
2.8 Procedure Expressions: `lambda` and `case-lambda` 104
2.9 Local Binding: `let`, `let*`, `letrec`, ... 108
2.10 Local Definitions: `local` 112
2.11 Constructing Graphs: `shared` 112
2.12 Conditionals: `if`, `cond`, `and`, and `or` 115

2.13 Dispatch: `case` . . . 118

2.14 Definitions: `define`, `define-syntax`, ... 119

2.14.1 `require` Macros . . . 123

2.14.2 `provide` Macros . . . 123

2.15 Sequencing: `begin`, `begin0`, and `begin-for-syntax` . . . 124

2.16 Guarded Evaluation: `when` and `unless` . . . 125

2.17 Assignment: `set!` and `set!-values` . . . 126

2.18 Iterations and Comprehensions: `for`, `for/list`, ... 127

2.18.1 Iteration and Comprehension Forms . . . 128

2.18.2 Deriving New Iteration Forms . . . 133

2.18.3 Do Loops . . . 137

2.19 Continuation Marks: `with-continuation-mark` . . . 137

2.20 Quasiquoting: `quasiquote`, `unquote`, and `unquote-splicing` . . . 138

2.21 Syntax Quoting: `quote-syntax` . . . 139

2.22 Interaction Wrapper: `#%top-interaction` . . . 140

2.23 Limiting Scope: `define-package`, `open-package`, ... 140

2.24 Blocks: `block` . . . 143

2.25 Internal-Definition Limiting: `#%stratified-body` . . . 144

2.26 Performance Hints: `begin-encourage-inline` . . . 144

3 Datatypes **145**

3.1 Booleans and Equality . . . 145

3.1.1 Boolean Aliases . . . 150

3.2 Numbers . . . 152

3.2.1 Number Types . . . 153

3.2.2 Generic Numerics . . . 159

3.2.3 Flonums . . . 185
3.2.4 Fixnums . . . 190
3.3 Strings . . . 195
3.3.1 String Constructors, Selectors, and Mutators . . . 195
3.3.2 String Comparisons . . . 200
3.3.3 String Conversions . . . 203
3.3.4 Locale-Specific String Operations . . . 205
3.3.5 Additional String Functions . . . 206
3.4 Byte Strings . . . 209
3.4.1 Byte String Constructors, Selectors, and Mutators . . . 209
3.4.2 Byte String Comparisons . . . 214
3.4.3 Bytes to/from Characters, Decoding and Encoding . . . 215
3.4.4 Bytes to Bytes Encoding Conversion . . . 220
3.4.5 Additional Byte String Functions . . . 224
3.5 Characters . . . 225
3.5.1 Characters and Scalar Values . . . 225
3.5.2 Character Comparisons . . . 226
3.5.3 Classifications . . . 229
3.5.4 Character Conversions . . . 231
3.6 Symbols . . . 233
3.7 Regular Expressions . . . 235
3.7.1 Regexp Syntax . . . 236
3.7.2 Additional Syntactic Constraints . . . 241
3.7.3 Regexp Constructors . . . 243
3.7.4 Regexp Matching . . . 245

3.7.5 Regexp Splitting . . . 257
3.7.6 Regexp Substitution . . . 257
3.8 Keywords . . . 260
3.9 Pairs and Lists . . . 261
3.9.1 Pair Constructors and Selectors . . . 262
3.9.2 List Operations . . . 265
3.9.3 List Iteration . . . 266
3.9.4 List Filtering . . . 269
3.9.5 List Searching . . . 272
3.9.6 Pair Accessor Shorthands . . . 275
3.9.7 Additional List Functions and Synonyms . . . 282
3.9.8 Immutable Cyclic Data . . . 292
3.10 Mutable Pairs and Lists . . . 294
3.10.1 Mutable Pair Constructors and Selectors . . . 294
3.10.2 Mutable List Functions . . . 295
3.11 Vectors . . . 298
3.11.1 Additional Vector Functions . . . 300
3.12 Boxes . . . 306
3.13 Hash Tables . . . 307
3.14 Sequences and Streams . . . 315
3.14.1 Sequences . . . 316
3.14.2 Streams . . . 329
3.14.3 Generators . . . 333
3.15 Dictionaries . . . 337
3.16 Sets . . . 353

3.17 Procedures 358
3.17.1 Keywords and Arity 360
3.17.2 Reflecting on Primitives 369
3.17.3 Additional Higher-Order Functions 370
3.18 Void and Undefined 373

4 Structures **374**

4.1 Defining Structure Types: `struct` 375
4.2 Creating Structure Types 379
4.3 Structure Type Properties 384
4.4 Generic Interfaces 386
4.5 Copying and Updating Structures 390
4.6 Structure Utilities 391
4.7 Structure Type Transformer Binding 394

5 Classes and Objects **397**

5.1 Creating Interfaces 398
5.2 Creating Classes 399
5.2.1 Initialization Variables 414
5.2.2 Fields 416
5.2.3 Methods 417
5.3 Creating Objects 422
5.4 Field and Method Access 424
5.4.1 Methods 424
5.4.2 Fields 427
5.4.3 Generics 428

5.5 Mixins . . . 428
5.6 Traits . . . 429
5.7 Object and Class Contracts . . . 432
5.8 Object Equality and Hashing . . . 438
5.9 Object Serialization . . . 439
5.10 Object Printing . . . 441
5.11 Object, Class, and Interface Utilities . . . 441
5.12 Surrogates . . . 448

6 Units **450**

6.1 Creating Units . . . 450
6.2 Invoking Units . . . 454
6.3 Linking Units and Creating Compound Units . . . 455
6.4 Inferred Linking . . . 456
6.5 Generating A Unit from Context . . . 459
6.6 Structural Matching . . . 459
6.7 Extending the Syntax of Signatures . . . 461
6.8 Unit Utilities . . . 461
6.9 Unit Contracts . . . 462
6.10 Single-Unit Modules . . . 462
6.11 Single-Signature Modules . . . 463
6.12 Transformer Helpers . . . 463

7 Contracts **465**

7.1 Data-structure Contracts . . . 465
7.2 Function Contracts . . . 476

7.3 Parametric Contracts . . . 483
7.4 Lazy Data-structure Contracts . . . 485
7.5 Structure Type Property Contracts . . . 486
7.6 Attaching Contracts to Values . . . 489
7.6.1 Nested Contract Boundaries . . . 490
7.6.2 Low-level Contract Boundaries . . . 492
7.7 Building New Contract Combinators . . . 493
7.7.1 Blame Objects . . . 500
7.7.2 Contracts as structs . . . 505
7.7.3 Obligation Information in Check Syntax . . . 508
7.7.4 Utilities for Building New Combinators . . . 510
7.8 Contract Utilities . . . 510
7.9 racket/contract/base . . . 513
7.10 Legacy Contracts . . . 513
7.11 Random generation . . . 515

8 Pattern Matching **516**

8.1 Additional Matching Forms . . . 523
8.2 Extending match . . . 525
8.3 Library Extensions . . . 527

9 Control Flow **529**

9.1 Multiple Values . . . 529
9.2 Exceptions . . . 530
9.2.1 Error Message Conventions . . . 530
9.2.2 Raising Exceptions . . . 531

9.2.3 Handling Exceptions . . . 538
9.2.4 Configuring Default Handling . . . 540
9.2.5 Built-in Exception Types . . . 542
9.3 Delayed Evaluation . . . 547
9.3.1 Additional Promise Kinds . . . 548
9.4 Continuations . . . 550
9.4.1 Additional Control Operators . . . 556
9.5 Continuation Marks . . . 561
9.6 Breaks . . . 565
9.7 Exiting . . . 567

10 Concurrency and Parallelism **569**

10.1 Threads . . . 569
10.1.1 Creating Threads . . . 569
10.1.2 Suspending, Resuming, and Killing Threads . . . 570
10.1.3 Synchronizing Thread State . . . 572
10.1.4 Thread Mailboxes . . . 573
10.2 Synchronization . . . 574
10.2.1 Events . . . 574
10.2.2 Channels . . . 581
10.2.3 Semaphores . . . 582
10.2.4 Buffered Asynchronous Channels . . . 584
10.3 Thread-Local Storage . . . 587
10.3.1 Thread Cells . . . 587
10.3.2 Parameters . . . 589
10.4 Futures . . . 592

10.5 Futures Visualizer . . . 595
10.5.1 Execution Timeline . . . 596
10.5.2 Future Creation Tree . . . 597
10.6 Futures Tracing . . . 598
10.6.1 Future Performance Logging . . . 600
10.7 Places . . . 602
10.8 Distributed Places . . . 607
10.8.1 Define Remote Server . . . 628
10.8.2 Async Bidirectional Channels . . . 630
10.8.3 Distributed Places MPI . . . 631

11 Macros **635**

11.1 Pattern-Based Syntax Matching . . . 635
11.2 Syntax Object Content . . . 645
11.3 Syntax Object Bindings . . . 651
11.4 Syntax Transformers . . . 655
11.4.1 `require` Transformers . . . 670
11.4.2 `provide` Transformers . . . 673
11.4.3 Keyword-Argument Conversion Introspection . . . 675
11.5 Syntax Parameters . . . 676
11.5.1 Syntax Parameter Inspection . . . 677
11.6 Local Binding with Splicing Body . . . 677
11.7 Syntax Object Properties . . . 679
11.8 Syntax Taints . . . 681
11.9 Expanding Top-Level Forms . . . 684
11.9.1 Information on Expanded Modules . . . 686

11.10 File Inclusion 687
11.11 Syntax Utilities 688
11.11.1 Creating formatted identifiers 688
11.11.2 Pattern variables 689
11.11.3 Error reporting 690
11.11.4 Recording disappeared uses 691
11.11.5 Miscellaneous utilities 692

12 Input and Output **694**

12.1 Ports 694
12.1.1 Encodings and Locales 694
12.1.2 Managing Ports 696
12.1.3 Port Buffers and Positions 697
12.1.4 Counting Positions, Lines, and Columns 699
12.1.5 File Ports 701
12.1.6 String Ports 708
12.1.7 Pipes 710
12.1.8 Structures as Ports 710
12.1.9 Custom Ports 711
12.1.10 More Port Constructors, Procedures, and Events 731
12.2 Byte and String Input 747
12.3 Byte and String Output 756
12.4 Reading 759
12.5 Writing 765
12.6 Pretty Printing 772
12.6.1 Basic Pretty-Print Options 774

12.6.2 Per-Symbol Special Printing 775
12.6.3 Line-Output Hook 776
12.6.4 Value Output Hook 777
12.6.5 Additional Custom-Output Support 779
12.7 Reader Extension 780
12.7.1 Readtables 780
12.7.2 Reader-Extension Procedures 786
12.7.3 Special Comments 787
12.8 Printer Extension 787
12.9 Serialization 789
12.10 Fast-Load Serialization 797

13 Reflection and Security **799**

13.1 Namespaces 799
13.2 Evaluation and Compilation 807
13.3 The racket/load Language 816
13.4 Module Names and Loading 817
13.4.1 Resolving Module Names 817
13.4.2 Compiled Modules and References 820
13.4.3 Dynamic Module Access 824
13.5 Impersonators and Chaperones 826
13.5.1 Impersonator Constructors 829
13.5.2 Chaperone Constructors 835
13.5.3 Impersonator Properties 841
13.6 Security Guards 842
13.7 Custodians 844

13.8 Thread Groups . . . 846
13.9 Structure Inspectors . . . 847
13.10 Code Inspectors . . . 850
13.11 Sandboxed Evaluation . . . 851
13.11.1 Customizing Evaluators . . . 856
13.11.2 Interacting with Evaluators . . . 865
13.11.3 Miscellaneous . . . 869

14 Operating System **871**

14.1 Paths . . . 871
14.1.1 Manipulating Paths . . . 871
14.1.2 More Path Utilities . . . 879
14.1.3 Unix and Mac OS X Paths . . . 882
14.1.4 Windows Path Conventions . . . 882
14.2 Filesystem . . . 886
14.2.1 Locating Paths . . . 886
14.2.2 Files . . . 889
14.2.3 Directories . . . 893
14.2.4 Declaring Paths Needed at Run Time . . . 894
14.2.5 More File and Directory Utilities . . . 897
14.3 Networking . . . 907
14.3.1 TCP . . . 907
14.3.2 UDP . . . 912
14.4 Processes . . . 919
14.4.1 Simple Subprocesses . . . 924
14.5 Logging . . . 928

14.5.1 Creating Loggers . . . 928
14.5.2 Logging Events . . . 929
14.5.3 Receiving Logged Events . . . 930
14.6 Time . . . 931
14.6.1 Date Utilities . . . 933
14.7 Environment and Runtime Information . . . 935
14.8 Command-Line Parsing . . . 939

15 Memory Management **946**
15.1 Weak Boxes . . . 946
15.2 Ephemerons . . . 946
15.3 Wills and Executors . . . 947
15.4 Garbage Collection . . . 949

16 Unsafe Operations **951**
16.1 Unsafe Numeric Operations . . . 951
16.2 Unsafe Data Extraction . . . 955

17 Running Racket **959**
17.1 Running Racket or GRacket . . . 959
17.1.1 Initialization . . . 959
17.1.2 Exit Status . . . 960
17.1.3 Init Libraries . . . 960
17.1.4 Command Line . . . 961
17.1.5 Language Run-Time Configuration . . . 964
17.2 Libraries and Collections . . . 965
17.2.1 Collection Links . . . 968

17.3 Interactive Help 969
17.4 Interactive Module Loading 970
17.5 Debugging 971
17.5.1 Tracing 972

Index **977**

1 Language Model

1.1 Evaluation Model

Racket evaluation can be viewed as the simplification of expressions to obtain values. For example, just as an elementary-school student simplifies

```
1 + 1 = 2
```

Racket evaluation simplifies

```
(+ 1 1) → 2
```

The arrow → above replaces the more traditional = to emphasize that evaluation proceeds in a particular direction towards simpler expressions. In particular, a *value* is an expression that evaluation simplifies no further, such as the number `2`.

1.1.1 Sub-expression Evaluation and Continuations

Some simplifications require more than one step. For example:

```
(- 4 (+ 1 1)) → (- 4 2) → 2
```

An expression that is not a value can always be partitioned into two parts: a *redex*, which is the part that changed in a single-step simplification (highlighted), and the *continuation*, which is the surrounding expression context. In `(- 4 (+ 1 1))`, the redex is `(+ 1 1)`, and the continuation is `(- 4 [])`, where [] takes the place of the redex. That is, the continuation says how to "continue" after the redex is reduced to a value.

Before some things can be evaluated, some sub-expressions must be evaluated; for example, in the application `(- 4 (+ 1 1))`, the application of `-` cannot be reduced until the sub-expression `(+ 1 1)` is reduced.

Thus, the specification of each syntactic form specifies how (some of) its sub-expressions are evaluated, and then how the results are combined to reduce the form away.

The *dynamic extent* of an expression is the sequence of evaluation steps during which an expression contains the redex.

1.1.2 Tail Position

An expression *`expr1`* is in *tail position* with respect to an enclosing expression *`expr2`* if, whenever *`expr1`* becomes a redex, its continuation is the same as was the enclosing *`expr2`*'s

continuation.

For example, the `(+ 1 1)` expression is *not* in tail position with respect to `(- 4 (+ 1 1))`. To illustrate, we use the notation *C*[*expr*] to mean the expression that is produced by substituting *expr* in place of [] in the continuation *C*:

C[`(- 4 (+ 1 1))`] → *C*[`(- 4 2)`]

In this case, the continuation for reducing `(+ 1 1)` is *C*[`(- 4 [])`], not just *C*.

In contrast, `(+ 1 1)` is in tail position with respect to `(if (zero? 0) (+ 1 1) 3)`, because, for any continuation *C*,

C[`(if (zero? 0) (+ 1 1) 3)`] → *C*[`(if #t (+ 1 1) 3)`] → *C*[`(+ 1 1)`]

The steps in this reduction sequence are driven by the definition of `if`, and they do not depend on the continuation *C*. The "then" branch of an `if` form is always in tail position with respect to the `if` form. Due to a similar reduction rule for `if` and `#f`, the "else" branch of an `if` form is also in tail position.

Tail-position specifications provide a guarantee about the asymptotic space consumption of a computation. In general, the specification of tail positions goes with each syntactic form, like `if`.

1.1.3 Multiple Return Values

A Racket expression can evaluate to *multiple values*, in the same way that a procedure can accept multiple arguments.

Most continuations expect a particular number of result values. Indeed, most continuations, such as `(+ [] 1)` expect a single value. The continuation `(let-values ([(x y) []]) expr)` expects two result values; the first result replaces `x` in the body *expr*, and the second replaces `y` in *expr*. The continuation `(begin [] (+ 1 2))` accepts any number of result values, because it ignores the result(s).

In general, the specification of a syntactic form indicates the number of values that it produces and the number that it expects from each of its sub-expression. In addition, some procedures (notably `values`) produce multiple values, and some procedures (notably `call-with-values`) create continuations internally that accept a certain number of values.

1.1.4 Top-Level Variables

Given

```
x = 10
```

then an algebra student simplifies `x + 1` as follows:

```
x + 1 = 10 + 1 = 11
```

Racket works much the same way, in that a set of top-level variables are available for substitutions on demand during evaluation. For example, given

```
(define x 10)
```

then

```
(+ x 1) → (+ 10 1) → 11
```

In Racket, the way definitions appear is just as important as the way that they are used. Racket evaluation thus keeps track of both definitions and the current expression, and it extends the set of definitions in response to evaluating forms such as `define`.

Each evaluation step, then, takes the current set of definitions and program to a new set of definitions and program. Before a `define` can be moved into the set of definitions, its right-hand expression must be reduced to a value.

```
   defined:
   evaluate: (begin (define x (+ 9 1)) (+ x 1))
→ defined:
   evaluate: (begin (define x 10) (+ x 1))
→ defined: (define x 10)
   evaluate: (begin (void) (+ x 1))
→ defined: (define x 10)
   evaluate: (+ x 1)
→ defined: (define x 10)
   evaluate: (+ 10 1)
→ defined: (define x 10)
   evaluate: 11
```

Using `set!`, a program can change the value associated with an existing top-level variable:

```
   defined: (define x 10)
   evaluate: (begin (set! x 8) x)
→ defined: (define x 8)
   evaluate: (begin (void) x)
→ defined: (define x 8)
   evaluate: x
→ defined: (define x 8)
   evaluate: 8
```

1.1.5 Objects and Imperative Update

In addition to `set!` for imperative update of top-level variables, various procedures enable the modification of elements within a compound data structure. For example, `vector-set!` modifies the content of a vector.

To allow such modifications to data, we must distinguish between values, which are the results of expressions, and *objects*, which hold the data referenced by a value.

A few kinds of objects can serve directly as values, including booleans, `(void)`, and small exact integers. More generally, however, a value is a reference to an object. For example, a value can be a reference to a particular vector that currently holds the value `10` in its first slot. If an object is modified, then the modification is visible through all copies of the value that reference the same object.

In the evaluation model, a set of objects must be carried along with each step in evaluation, just like the definition set. Operations that create objects, such as `vector`, add to the set of objects:

```
   objects:
   defined:
   evaluate: (begin (define x (vector 10 20))
                    (define y x)
                    (vector-set! x 0 11)
                    (vector-ref y 0))
→  objects:  (define <o1> (vector 10 20))
   defined:
   evaluate: (begin (define x <o1>)
                    (define y x)
                    (vector-set! x 0 11)
                    (vector-ref y 0))
→  objects:  (define <o1> (vector 10 20))
   defined:  (define x <o1>)
   evaluate: (begin (void)
                    (define y x)
                    (vector-set! x 0 11)
                    (vector-ref y 0))
→  objects:  (define <o1> (vector 10 20))
   defined:  (define x <o1>)
   evaluate: (begin (define y x)
                    (vector-set! x 0 11)
                    (vector-ref y 0))
→  objects:  (define <o1> (vector 10 20))
   defined:  (define x <o1>)
```

```
  evaluate: (begin (define y <o1>)
                   (vector-set! x 0 11)
                   (vector-ref y 0))
→ objects: (define <o1> (vector 10 20))
  defined: (define x <o1>)
           (define y <o1>)
  evaluate: (begin (void)
                   (vector-set! x 0 11)
                   (vector-ref y 0))
→ objects: (define <o1> (vector 10 20))
  defined: (define x <o1>)
           (define y <o1>)
  evaluate: (begin (vector-set! x 0 11)
                   (vector-ref y 0))
→ objects: (define <o1> (vector 10 20))
  defined: (define x <o1>)
           (define y <o1>)
  evaluate: (begin (vector-set! <o1> 0 11)
                   (vector-ref y 0))
→ objects: (define <o1> (vector 11 20))
  defined: (define x <o1>)
           (define y <o1>)
  evaluate: (begin (void)
                   (vector-ref y 0))
→ objects: (define <o1> (vector 11 20))
  defined: (define x <o1>)
           (define y <o1>)
  evaluate: (vector-ref y 0)
→ objects: (define <o1> (vector 11 20))
  defined: (define x <o1>)
           (define y <o1>)
  evaluate: (vector-ref <o1> 0)
→ objects: (define <o1> (vector 11 20))
  defined: (define x <o1>)
           (define y <o1>)
  evaluate: 11
```

The distinction between a top-level variable and an object reference is crucial. A top-level variable is not a value; each time a variable expression is evaluated, the value is extracted from the current set of definitions. An object reference, in contrast is a value, and therefore needs no further evaluation. The model evaluation steps above use angle-bracketed `<o1>` for an object reference to distinguish it from a variable name.

A direct object reference can never appear in a text-based source program. A program rep-

resentation created with datum->syntax, however, can embed direct references to existing objects.

1.1.6 Object Identity and Comparisons

The eq? operator compares two values, returning #t when the values refer to the same object. This form of equality is suitable for comparing objects that support imperative update (e.g., to determine that the effect of modifying an object through one reference is visible through another reference). Also, an eq? test evaluates quickly, and eq?-based hashing is more lightweight than equal?-based hashing in hash tables.

In some cases, however, eq? is unsuitable as a comparison operator, because the generation of objects is not clearly defined. In particular, two applications of + to the same two exact integers may or may not produce results that are eq?, although the results are always equal?. Similarly, evaluation of a lambda form typically generates a new procedure object, but it may re-use a procedure object previously generated by the same source lambda form.

The behavior of a datatype with respect to eq? is generally specified with the datatype and its associated procedures.

1.1.7 Garbage Collection

See §15 “Memory Management” for functions related to garbage collection.

In the program state

```
objects: (define <o1> (vector 10 20))
         (define <o2> (vector 0))
defined: (define x <o1>)
evaluate: (+ 1 x)
```

evaluation cannot depend on <o2>, because it is not part of the program to evaluate, and it is not referenced by any definition that is accessible in the program. The object <o2> may therefore be removed from the evaluation by *garbage collection*.

A few special compound datatypes hold *weak references* to objects. Such weak references are treated specially by the garbage collector in determining which objects are reachable for the remainder of the computation. If an object is reachable only via a weak reference, then the object can be reclaimed, and the weak reference is replaced by a different value (typically #f).

As a special case, a fixnum is always considered reachable by the garbage collector. Many other values are always reachable due to the way they are implemented and used: A character in the Latin-1 range is always reachable, because equal? Latin-1 characters are always eq?, and all of the Latin-1 characters are referenced by an internal module. Similarly, null, #t, #f, eof, and #<void> and are always reachable. Values produced by quote remain

reachable when the `quote` expression itself is reachable.

1.1.8 Procedure Applications and Local Variables

Given

```
f(x) = x + 10
```

then an algebra student simplifies `f(7)` as follows:

```
f(7) = 7 + 10 = 17
```

The key step in this simplification is take the body of the defined function `f`, and then replace each `x` with the actual value `7`.

Racket procedure application works much the same way. A procedure is an object, so evaluating `(f 7)` starts with a variable lookup:

```
   objects:  (define <p1> (lambda (x) (+ x 10)))
   defined:  (define f <p1>)
   evaluate: (f 7)
→  objects:  (define <p1> (lambda (x) (+ x 10)))
   defined:  (define f <p1>)
   evaluate: (<p1> 7)
```

Unlike in algebra, however, the value associated with an argument can be changed in the body of a procedure by using `set!`, as in the example `(lambda (x) (begin (set! x 3) x))`. Since the value associated with `x` can be changed, an actual value cannot be substituted for `x` when the procedure is applied.

Instead, a new *location* is created for each variable on each application. The argument value is placed in the location, and each instance of the variable in the procedure body is replaced with the new location:

```
   objects:  (define <p1> (lambda (x) (+ x 10)))
   defined:  (define f <p1>)
   evaluate: (<p1> 7)
→  objects:  (define <p1> (lambda (x) (+ x 10)))
   defined:  (define f <p1>)
             (define xloc 7)
   evaluate: (+ xloc 10)
→  objects:  (define <p1> (lambda (x) (+ x 10)))
   defined:  (define f <p1>)
             (define xloc 7)
   evaluate: (+ 7 10)
```

```
→ objects:  (define <p1> (lambda (x) (+ x 10)))
  defined:  (define f <p1>)
            (define xloc 7)
  evaluate: 17
```

A location is the same as a top-level variable, but when a location is generated, it (conceptually) uses a name that has not been used before and that cannot be generated again or accessed directly.

Generating a location in this way means that `set!` evaluates for local variables in the same way as for top-level variables, because the local variable is always replaced with a location by the time the `set!` form is evaluated:

```
  objects:  (define <p1> (lambda (x) (begin (set! x 3) x)))
  defined:  (define f <p1>)
  evaluate: (f 7)
→ objects:  (define <p1> (lambda (x) (begin (set! x 3) x)))
  defined:  (define f <p1>)
  evaluate: (<p1> 7)
→ objects:  (define <p1> (lambda (x) (begin (set! x 3) x)))
  defined:  (define f <p1>)
            (define xloc 7)
  evaluate: (begin (set! xloc 3) xloc)
→ objects:  (define <p1> (lambda (x) (begin (set! x 3) x)))
  defined:  (define f <p1>)
            (define xloc 3)
  evaluate: (begin (void) xloc)
→ objects:  (define <p1> (lambda (x) (begin (set! x 3) x)))
  defined:  (define f <p1>)
            (define xloc 3)
  evaluate: xloc
→ objects:  (define <p1> (lambda (x) (begin (set! x 3) x)))
  defined:  (define f <p1>)
            (define xloc 3)
  evaluate: 3
```

The substitution and location-generation step of procedure application requires that the argument is a value. Therefore, in `((lambda (x) (+ x 10)) (+ 1 2))`, the `(+ 1 2)` subexpression must be simplified to the value `3`, and then `3` can be placed into a location for `x`. In other words, Racket is a *call-by-value* language.

Evaluation of a local-variable form, such as `(let ([x (+ 1 2)]) expr)`, is the same as for a procedure call. After `(+ 1 2)` produces a value, it is stored in a fresh location that replaces every instance of `x` in *expr*.

1.1.9 Variables and Locations

A *variable* is a placeholder for a value, and expressions in an initial program refer to variables. A *top-level variable* is both a variable and a location. Any other variable is always replaced by a location at run-time, so that evaluation of expressions involves only locations. A single *local variable* (i.e., a non-top-level, non-module-level variable), such as a procedure argument, can correspond to different locations through different instantiations.

For example, in the program

```
(define y (+ (let ([x 5]) x) 6))
```

both `y` and `x` are variables. The `y` variable is a top-level variable, and the `x` is a local variable. When this code is evaluated, a location is created for `x` to hold the value `5`, and a location is also created for `y` to hold the value `11`.

The replacement of a variable with a location during evaluation implements Racket's *lexical scoping*. For example, when a procedure-argument variable `x` is replaced by the location `xloc`, then it is replaced throughout the body of the procedure, including any nested `lambda` forms. As a result, future references of the variable always access the same location.

1.1.10 Modules and Module-Level Variables

See §2.1 "Modules: `module`, `module*`, ..." for the syntax of modules.

Most definitions in Racket are in modules. In terms of evaluation, a module is essentially a prefix on a defined name, so that different modules can define the name. That is, a *module-level variable* is like a top-level variable from the perspective of evaluation.

One difference between a module and a top-level definition is that a module can be declared without instantiating its module-level definitions. Evaluation of a `require` *instantiates* (i.e., triggers the *instantiation* of) a declared module, which creates variables that correspond to its module-level definitions.

For example, given the module declaration

```
(module m racket
  (define x 10))
```

the evaluation of `(require 'm)` creates the variable `x` and installs `10` as its value. This `x` is unrelated to any top-level definition of `x`.

Phases

A module can be instantiated in multiple *phases*. A phase is an integer that, again, is effectively a prefix on the names of module-level definitions. A top-level `require` instantiates a

module at phase 0, if the module is not already instantiated at phase 0. A top-level `(require (for-syntax ....))` instantiates a module at phase 1 (if it is not already instantiated at that level); `for-syntax` also has a different binding effect on further program parsing, as described in §1.2.3.4 "Introducing Bindings".

Within a module, some definitions are shifted by a phase already; the `begin-for-syntax` form is similar to `begin`, but it shifts expressions and definitions by a relative phase 1. Thus, if the module is instantiated at phase 1, the variables defined with `begin-for-syntax` are created at phase 2, and so on. Moreover, this relative phase acts as another layer of prefixing, so that a `define` of `x` and a `begin-for-syntax`-wrapped `define` of `x` can co-exist in a module without colliding. A `begin-for-syntax` form can be nested within a `begin-for-syntax` form, in which case definitions and expressions are in relative phase 2, and so on. Higher phases are mainly related to program parsing, instead of normal evaluation.

If a module instantiated at phase *n* `requires` another module, then the `required` module is first instantiated at phase *n*, and so on transitively. (Module `requires` cannot form cycles.) If a module instantiated at phase *n* `requires` `for-syntax` another module, the other module becomes *available* at phase *n*+1, and it may later be instantiated at phase *n*+1. If a module that is available at phase *n* for *n*>0 `requires` `for-template` another module, the other module becomes available at phase *n*-1, and so on. Instantiations of available modules above phase 0 are triggered on demand as described in §1.2.3.8 "Module Phases and Visits".

A final distinction among module instantiations is that multiple instantiations may exist at phase 1 and higher. These instantiations are created by the parsing of module forms (see §1.2.3.8 "Module Phases and Visits"), and are, again, conceptually distinguished by prefixes.

Top-level variables can exist in multiple phases in the same way as within modules. For example, `define` within `begin-for-syntax` creates a phase 1 variable. Furthermore, reflective operations like `make-base-namespace` and `eval` provide access to top-level variables in higher phases, while module instantiations (triggered by `require`) relative to such top-levels are in corresponding higher phases.

Module Re-declarations

When a module is declared using a name for which a module is already declared, the new declaration's definitions replace and extend the old declarations. If a variable in the old declaration has no counterpart in the new declaration, the old variable continues to exist, but its binding is not included in the lexical information for the module body. If a new variable definition has a counterpart in the old declaration, it effectively assigns to the old variable.

If a module is instantiated in any phases before it is re-declared, each re-declaration of the module is immediately instantiated in the same phases.

Submodules

A `module` or `module*` form within a top-level `module` form declares a *submodule*. A submodule is accessed relative to its enclosing module, usually with a `submod` path. Submodules can be nested to any depth.

Although a submodule is lexically nested within a module, it cannot necessarily access the bindings of its enclosing module directly. More specifically, a submodule declared with `module` cannot `require` from its enclosing module, but the enclosing module can `require` the submodule. In contrast, a submodule declared with `module*` conceptually follows its enclosing module, so can `require` from its enclosing module, but the enclosing module cannot `require` the submodule. Unless a submodule imports from its enclosing module or vice-versa, then visits or instantiations of the two modules are independent, and thier implementations may even be loaded from bytecode at different times.

When a submodule declaration has the form `(module* name #f ....)`, then all of the bindings of the enclosing module's bodies are visible in the submodule's body, and the submodule implicitly imports the enclosing module. The submodule can `provide` any bindings that it inherits from its enclosing module.

1.1.11 Continuation Frames and Marks

See §9.5 "Continuation Marks" for continuation-mark forms and functions.

Every continuation C can be partitioned into *continuation frames* C_1, C_2, ..., C_n such that $C = C_1[C_2[...[C_n]]]$, and no frame C_i can be itself partitioned into smaller continuations. Evaluation steps add and remove frames to the current continuation, typically one at a time.

Each frame is conceptually annotated with a set of *continuation marks*. A mark consists of a key and its value; the key is an arbitrary value, and each frame includes at most one mark for any key. Various operations set and extract marks from continuations, so that marks can be used to attach information to a dynamic extent. For example, marks can be used to record information for a "stack trace" to be used when an exception is raised, or to implement dynamic scope.

1.1.12 Prompts, Delimited Continuations, and Barriers

See §9.4 "Continuations" for continuation and prompt functions.

A *prompt* is a special kind of continuation frame that is annotated with a specific *prompt tag* (essentially a continuation mark). Various operations allow the capture of frames in the continuation from the redex position out to the nearest enclosing prompt with a particular prompt tag; such a continuation is sometimes called a *delimited continuation*. Other operations allow the current continuation to be extended with a captured continuation (specifically, a *composable continuation*). Yet other operations abort the computation to the nearest enclosing prompt with a particular tag, or replace the continuation to the nearest enclosing prompt with another one. When a delimited continuation is captured, the marks associated

with the relevant frames are also captured.

A *continuation barrier* is another kind of continuation frame that prohibits certain replacements of the current continuation with another. Specifically, a continuation can be replaced by another only when the replacement does not introduce any continuation barriers (but it may remove them). A continuation barrier thus prevents "downward jumps" into a continuation that is protected by a barrier. Certain operations install barriers automatically; in particular, when an exception handler is called, a continuation barrier prohibits the continuation of the handler from capturing the continuation past the exception point.

A *escape continuation* is essentially a derived concept. It combines a prompt for escape purposes with a continuation for mark-gathering purposes. As the name implies, escape continuations are used only to abort to the point of capture.

1.1.13 Threads

See §10 "Concurrency and Parallelism" for thread and synchronization functions.

Racket supports multiple *threads* of evaluation. Threads run concurrently, in the sense that one thread can preempt another without its cooperation, but threads currently all run on the same processor (i.e., the same underlying OS process and thread). See also §10.4 "Futures".

Threads are created explicitly by functions such as `thread`. In terms of the evaluation model, each step in evaluation actually consists of multiple concurrent expressions, up to one per thread, rather than a single expression. The expressions all share the same objects and top-level variables, so that they can communicate through shared state. Most evaluation steps involve a single step in a single expression, but certain synchronization primitives require multiple threads to progress together in one step.

In addition to the state that is shared among all threads, each thread has its own private state that is accessed through *thread cells*. A thread cell is similar to a normal mutable object, but a change to the value inside a thread cell is seen only when extracting a value from the cell from the same thread. A thread cell can be *preserved*; when a new thread is created, the creating thread's value for a preserved thread cell serves as the initial value for the cell in the created thread. For a non-preserved thread cell, a new thread sees the same initial value (specified when the thread cell is created) as all other threads.

1.1.14 Parameters

See §10.3.2 "Parameters" for parameter forms and functions.

Parameters are essentially a derived concept in Racket; they are defined in terms of continuation marks and thread cells. However, parameters are also built in, in the sense that some primitive procedures consult parameter values. For example, the default output stream for primitive output operations is determined by a parameter.

A parameter is a setting that is both thread-specific and continuation-specific. In the empty continuation, each parameter corresponds to a preserved thread cell; a corresponding *param-*

eter procedure accesses and sets the thread cell's value for the current thread.

In a non-empty continuation, a parameter's value is determined through a *parameterization* that is associated with the nearest enclosing continuation frame though a continuation mark (whose key is not directly accessible). A parameterization maps each parameter to a preserved thread cell, and the combination of thread cell and current thread yields the parameter's value. A parameter procedure sets or accesses the relevant thread cell for its parameter.

Various operations, such as `parameterize` or `call-with-parameterization`, install a parameterization into the current continuation's frame.

1.1.15 Exceptions

See §9.2 "Exceptions" for exception forms, functions, and types.

Exceptions are essentially a derived concept in Racket; they are defined in terms of continuations, prompts, and continuation marks. However, exceptions are also built in, in the sense that primitive forms and procedures may raise exceptions.

An *exception handler* to catch exceptions can be associated with a continuation frame though a continuation mark (whose key is not directly accessible). When an exception is raised, the current continuation's marks determine a chain of exception handler procedures that are consulted to handle the exception. A handler for uncaught exceptions is designated through a built-in parameter.

One potential action of an exception handler is to abort the current continuation up to an enclosing prompt with a particular prompt tag. The default handler for uncaught exceptions, in particular, aborts to a particular tag for which a prompt is always present, because the prompt is installed in the outermost frame of the continuation for any new thread.

1.1.16 Custodians

See §13.7 "Custodians" for custodian functions.

In GRacket, custodians also manage eventspaces.

A *custodian* manages a collection of threads, file-stream ports, TCP ports, TCP listeners, UDP sockets, and byte converters. Whenever a thread, etc., is created, it is placed under the management of the *current custodian* as determined by the `current-custodian` parameter.

Except for the root custodian, every custodian itself is managed by a custodian, so that custodians form a hierarchy. Every object managed by a subordinate custodian is also managed by the custodian's owner.

When a custodian is shut down via `custodian-shutdown-all`, it forcibly and immediately closes the ports, TCP connections, etc., that it manages, as well as terminating (or suspending) its threads. A custodian that has been shut down cannot manage new objects. If the current custodian is shut down before a procedure is called to create a managed resource (e.g., `open-input-port`, `thread`), the `exn:fail:contract` exception is raised.

A thread can have multiple managing custodians, and a suspended thread created with `thread/suspend-to-kill` can have zero custodians. Extra custodians become associated with a thread through `thread-resume` (see §10.1.2 "Suspending, Resuming, and Killing Threads"). When a thread has multiple custodians, it is not necessarily killed by a `custodian-shutdown-all`, but shut-down custodians are removed from the thread's managing set, and the thread is killed when its managing set becomes empty.

The values managed by a custodian are only weakly held by the custodian. As a result, a will can be executed for a value that is managed by a custodian. In addition, a custodian only weakly references its subordinate custodians; if a subordinate custodian is unreferenced but has its own subordinates, then the custodian may be collected, at which point its subordinates become immediately subordinate to the collected custodian's superordinate custodian.

In addition to the other entities managed by a custodian, a *custodian box* created with `make-custodian-box` strongly holds onto a value placed in the box until the box's custodian is shut down. The custodian only weakly retains the box itself, however (so the box and its content can be collected if there are no other references to them).

When Racket is compiled with support for per-custodian memory accounting (see `custodian-memory-accounting-available?`), the `current-memory-use` procedure can report a custodian-specific result. This result determines how much memory is occupied by objects that are reachable from the custodian's managed values, especially its threads, and including its sub-custodians' managed values. If an object is reachable from two custodians where neither is an ancestor of the other, an object is arbitrarily charged to one or the other, and the choice can change after each collection; objects reachable from both a custodian and its descendant, however, are reliably charged to the custodian and not to the descendants, unless the custodian can reach the objects only through a descendant custodian or a descendant's thread. Reachability for per-custodian accounting does not include weak references, references to threads managed by other custodians, references to other custodians, or references to custodian boxes for other custodians.

1.2 Syntax Model

The syntax of a Racket program is defined by

- a *read* pass that processes a character stream into a syntax object; and
- an *expand* pass that processes a syntax object to produce one that is fully parsed.

For details on the read pass, see §1.3 "The Reader". Source code is normally read in `read-syntax` mode, which produces a syntax object.

The expand pass recursively processes a syntax object to produce a complete parse of the program. Binding information in a syntax object drives the expansion process, and when the

expansion process encounters a binding form, it extends syntax objects for sub-expression with new binding information.

1.2.1 Identifiers and Binding

§4.2 "Identifiers and Binding" in ***The Racket Guide*** introduces binding.

An *identifier* is a source-program entity. Parsing (i.e., expanding) a Racket program reveals that some identifiers correspond to variables, some refer to syntactic forms, and some are quoted to produce a symbol or a syntax object.

An identifier *binds* another (i.e., it is a *binding*) when the former is parsed as a variable and the latter is parsed as a reference to the former; the latter is *bound*. The *scope* of a binding is the set of source forms to which it applies. The *environment* of a form is the set of bindings whose scope includes the form. A binding for a sub-expression *shadows* any bindings (i.e., it is *shadowing*) in its environment, so that uses of an identifier refer to the shadowing binding.

For example, as a bit of source, the text

```
(let ([x 5]) x)
```

includes two identifiers: `let` and `x` (which appears twice). When this source is parsed in a typical environment, `x` turns out to represent a variable (unlike `let`). In particular, the first `x` binds the second `x`.

A *top-level binding* is a binding from a definition at the top-level; a *module binding* is a binding from a definition in a module; all other bindings are *local bindings*. There is no difference between an *unbound* identifier and one with a top-level binding; within a module, references to top-level bindings are disallowed, and so such identifiers are called unbound in a module context.

Throughout the documentation, identifiers are typeset to suggest the way that they are parsed. A black, boldface identifier like `lambda` indicates a reference to a syntactic form. A plain blue identifier like `x` is a variable or a reference to an unspecified top-level variable. A hyperlinked identifier `cons` is a reference to a specific top-level variable.

Every binding has a *phase level* in which it can be referenced, where a phase level normally corresponds to an integer (but the special label phase level does not correspond to an integer). Phase level 0 corresponds to the run time of the enclosing module (or the run time of top-level expressions). Bindings in phase level 0 constitute the *base environment*. Phase level 1 corresponds to the time during which the enclosing module (or top-level expression) is expanded; bindings in phase level 1 constitute the *transformer environment*. Phase level -1 corresponds to the run time of a different module for which the enclosing module is imported for use at phase level 1 (relative to the importing module); bindings in phase level -1 constitute the *template environment*. The *label phase level* does not correspond to any execution time; it is used to track bindings (e.g., to identifiers within documentation) without implying an execution dependency.

If an identifier has a local binding, then it is the same for all phase levels, though the reference is allowed only at a particular phase level. Attempting to reference a local binding in a different phase level from the binding's context produces a syntax error. If an identifier has a top-level binding or module binding, then it can have different such bindings in different phase levels.

1.2.2 Syntax Objects

A *syntax object* combines a simpler Racket value, such as a symbol or pair, with *lexical information* about bindings, source-location information, syntax properties, and tamper status. In particular, an identifier is represented as a symbol object that combines a symbol with lexical and other information.

For example, a `car` identifier might have lexical information that designates it as the `car` from the `racket/base` language (i.e., the built-in `car`). Similarly, a `lambda` identifier's lexical information may indicate that it represents a procedure form. Some other identifier's lexical information may indicate that it references a top-level variable.

When a syntax object represents a more complex expression than an identifier or simple constant, its internal components can be extracted. Even for extracted identifiers, detailed information about binding is available mostly indirectly; two identifiers can be compared to determine whether they refer to the same binding (i.e., `free-identifier=?`), or whether each identifier would bind the other if one were in a binding position and the other in an expression position (i.e., `bound-identifier=?`).

For example, when the program written as

```
(let ([x 5]) (+ x 6))
```

is represented as a syntax object, then two syntax objects can be extracted for the two `x`s. Both the `free-identifier=?` and `bound-identifier=?` predicates will indicate that the `x`s are the same. In contrast, the `let` identifier is not `free-identifier=?` or `bound-identifier=?` to either `x`.

The lexical information in a syntax object is independent of the other half, and it can be copied to a new syntax object in combination with an arbitrary other Racket value. Thus, identifier-binding information in a syntax object is predicated on the symbolic name of the identifier as well as the identifier's lexical information; the same question with the same lexical information but different base value can produce a different answer.

For example, combining the lexical information from `let` in the program above to `'x` would not produce an identifier that is `free-identifier=?` to either `x`, since it does not appear in the scope of the `x` binding. Combining the lexical context of the `6` with `'x`, in contrast, would produce an identifier that is `bound-identifier=?` to both `x`s.

The `quote-syntax` form bridges the evaluation of a program and the representation of a program. Specifically, `(quote-syntax datum)` produces a syntax object that preserves all of the lexical information that *datum* had when it was parsed as part of the `quote-syntax` form.

1.2.3 Expansion (Parsing)

Expansion recursively processes a syntax object in a particular phase level, starting with phase level 0. Bindings from the syntax object's lexical information drive the expansion process, and cause new bindings to be introduced for the lexical information of sub-expressions. In some cases, a sub-expression is expanded in a deeper phase than the enclosing expression.

Fully Expanded Programs

A complete expansion produces a syntax object matching the following grammar:

Beware that the symbolic names of identifiers in a fully expanded program may not match the symbolic names in the grammar. Only the binding (according to `free-identifier=?`) matters.

```
top-level-form = general-top-level-form
               | (#%expression expr)
               | (module id module-path
                   (#%plain-module-begin
                    module-level-form ...))
               | (begin top-level-form ...)
               | (begin-for-syntax top-level-form ...)

module-level-form = general-top-level-form
                  | (#%provide raw-provide-spec ...)
                  | (begin-for-syntax module-level-form ...)
                  | (module id module-path
                      module-level-form ...)
                  | (module* id module-path
                      module-level-form ...)
                  | (module* id #f
                      module-level-form ...)

general-top-level-form = expr
                       | (define-values (id ...) expr)
                       | (define-syntaxes (id ...) expr)
                       | (#%require raw-require-spec ...)

expr = id
     | (#%plain-lambda formals expr ...+)
     | (case-lambda (formals expr ...+) ...)
     | (if expr expr expr)
```

```
                | (begin expr ...+)
                | (begin0 expr expr ...)
                | (let-values ([(id ...) expr] ...)
                    expr ...+)
                | (letrec-values ([(id ...) expr] ...)
                    expr ...+)
                | (set! id expr)
                | (quote datum)
                | (quote-syntax datum)
                | (with-continuation-mark expr expr expr)
                | (#%plain-app expr ...+)
                | (#%top . id)
                | (#%variable-reference id)
                | (#%variable-reference (#%top . id))
                | (#%variable-reference)

formals = (id ...)
        | (id ...+ . id)
        | id
```

A *fully-expanded* syntax object corresponds to a *parse* of a program (i.e., a *parsed* program), and lexical information on its identifiers indicates the parse.

More specifically, the typesetting of identifiers in the above grammar is significant. For example, the second case for *expr* is a syntax-object list whose first element is an identifier, where the identifier's lexical information specifies a binding to the `#%plain-lambda` of the `racket/base` language (i.e., the identifier is `free-identifier=?` to one whose binding is `#%plain-lambda`). In all cases, identifiers above typeset as syntactic-form names refer to the bindings defined in §2 "Syntactic Forms".

In a fully expanded program for a namespace whose base phase is 0, the relevant phase level for a binding in the program is *N* if the bindings has *N* surrounding `begin-for-syntax` and `define-syntaxes` forms—not counting any `begin-for-syntax` forms that wrap a `module` form for the body of the `module`. The *datum* in a `quote-syntax` form, however, always preserves its information for all phase levels.

In addition to the grammar above, `letrec-syntaxes+values` can appear in a fully local-expanded expression, as can `#%expression` in any expression position. For example, `letrec-syntaxes+values` and `#%expression` can appear in the result from `local-expand` when the stop list is empty.

Expansion Steps

In a recursive expansion, each single step in expanding a syntax object at a particular phase level depends on the immediate shape of the syntax object being expanded:

- If it is an identifier (i.e., a syntax-object symbol), then a binding is determined by the identifier's lexical information. If the identifier has a binding other than as a top-level variable, that binding is used to continue. If the identifier has no binding, a new syntax-object symbol `'#%top` is created using the lexical information of the identifier; if this `#%top` identifier has no binding (other than as a top-level variable), then parsing fails with an `exn:fail:syntax` exception. Otherwise, the new identifier is combined with the original identifier in a new syntax-object pair (also using the same lexical information as the original identifier), and the `#%top` binding is used to continue.
- If it is a syntax-object pair whose first element is an identifier, and if the identifier has a binding other than as a top-level variable, then the identifier's binding is used to continue.
- If it is a syntax-object pair of any other form, then a new syntax-object symbol `'#%app` is created using the lexical information of the pair. If the resulting `#%app` identifier has no binding, parsing fails with an `exn:fail:syntax` exception. Otherwise, the new identifier is combined with the original pair to form a new syntax-object pair (also using the same lexical information as the original pair), and the `#%app` binding is used to continue.
- If it is any other syntax object, then a new syntax-object symbol `'#%datum` is created using the lexical information of the original syntax object. If the resulting `#%datum` identifier has no binding, parsing fails with an `exn:fail:syntax` exception. Otherwise, the new identifier is combined with the original syntax object in a new syntax-object pair (using the same lexical information as the original pair), and the `#%datum` binding is used to continue.

Thus, the possibilities that do not fail lead to an identifier with a particular binding. This binding refers to one of three things:

- A transformer binding, such as introduced by `define-syntax` or `let-syntax`. If the associated value is a procedure of one argument, the procedure is called as a syntax transformer (described below), and parsing starts again with the syntax-object result. If the transformer binding is to any other kind of value, parsing fails with an `exn:fail:syntax` exception. The call to the syntax transformer is `parameterized` to set `current-namespace` to a namespace that shares bindings and variables with the namespace being used to expand, except that its base phase is one greater.
- A variable binding, such as introduced by a module-level `define` or by `let`. In this case, if the form being parsed is just an identifier, then it is parsed as a reference to the corresponding variable. If the form being parsed is a syntax-object pair, then an `#%app` is added to the front of the syntax-object pair in the same way as when the first item in the syntax-object pair is not an identifier (third case in the previous enumeration), and parsing continues.
- A core *syntactic form*, which is parsed as described for each form in §2 "Syntactic Forms". Parsing a core syntactic form typically involves recursive parsing of sub-forms, and may introduce bindings that determine the parsing of sub-forms.

Expansion Context

Each expansion step occurs in a particular *context*, and transformers and core syntactic forms may expand differently for different contexts. For example, a `module` form is allowed only in a top-level context, and it fails in other contexts. The possible contexts are as follows:

- *top-level context* : outside of any module, definition, or expression, except that sub-expressions of a top-level `begin` form are also expanded as top-level forms.
- *module-begin context* : inside the body of a module, as the only form within the module.
- *module context* : in the body of a module (inside the module-begin layer).
- *internal-definition context* : in a nested context that allows both definitions and expressions.
- *expression context* : in a context where only expressions are allowed.

Different core syntactic forms parse sub-forms using different contexts. For example, a `let` form always parses the right-hand expressions of a binding in an expression context, but it starts parsing the body in an internal-definition context.

Introducing Bindings

Bindings are introduced during expansion when certain core syntactic forms are encountered:

- When a `require` form is encountered at the top level or module level, all lexical information derived from the top level or the specific module's level is extended with bindings from the specified modules. If not otherwise indicated in the `require` form, bindings are introduced at the phase levels specified by the exporting modules: phase level 0 for each normal `provide`, phase level 1 for each `for-syntax provide`, and so on. The `for-meta provide` form allows exports at an arbitrary phase level (as long as a binding exists within the module at the phase level).

  A `for-syntax` sub-form within `require` imports similarly, but the resulting bindings have a phase level that is one more than the exported phase levels, when exports for the label phase level are still imported at the label phase level. More generally, a `for-meta` sub-form within `require` imports with the specified phase level shift; if the specified shift is `#f`, or if `for-label` is used to import, then all bindings are imported into the label phase level.

- When a `define`, `define-values`, `define-syntax`, or `define-syntaxes` form is encountered at the top level or module level, all lexical information derived from the

top level or the specific module's level is extended with bindings for the specified identifiers at phase level 0 (i.e., the base environment is extended).

- When a `begin-for-syntax` form is encountered at the top level or module level, bindings are introduced as for `define-values` and `define-syntaxes`, but at phase level 1 (i.e., the transformer environment is extended). More generally, `begin-for-syntax` forms can be nested, an each `begin-for-syntax` shifts its body definition by one phase level.
- When a `let-values` form is encountered, the body of the `let-values` form is extended (by creating new syntax objects) with bindings for the specified identifiers. The same bindings are added to the identifiers themselves, so that the identifiers in binding position are `bound-identifier=?` to uses in the fully expanded form, and so they are not `bound-identifier=?` to other identifiers. The bindings are available for use at the phase level at which the `let-values` form is expanded.
- When a `letrec-values` or `letrec-syntaxes+values` form is encountered, bindings are added as for `let-values`, except that the right-hand-side expressions are also extended with the bindings.
- Definitions in internal-definition contexts introduce bindings as described in §1.2.3.7 "Internal Definitions".

A new binding in lexical information maps to a new variable. The identifiers mapped to this variable are those that currently have the same binding (i.e., that are currently `bound-identifier=?`) to the identifier associated with the binding.

For example, in

```
(let-values ([(x) 10]) (+ x y))
```

the binding introduced for `x` applies to the `x` in the body, but not the `y` in the body, because (at the point in expansion where the `let-values` form is encountered) the binding `x` and the body `y` are not `bound-identifier=?`.

Transformer Bindings

In a top-level context or module context, when the expander encounters a `define-syntaxes` form, the binding that it introduces for the defined identifiers is a *transformer binding*. The value of the binding exists at expansion time, rather than run time (though the two times can overlap), though the binding itself is introduced with phase level 0 (i.e., in the base environment).

The value for the binding is obtained by evaluating the expression in the `define-syntaxes` form. This expression must be expanded (i.e., parsed) before it can be evaluated, and it is expanded at phase level 1 (i.e., in the transformer environment) instead of phase level 0.

If the resulting value is a procedure of one argument or the result of make-set!-transformer on a procedure, then it is used as a *syntax transformer* (a.k.a. *macro*). The procedure is expected to accept a syntax object and return a syntax object. A use of the binding (at phase level 0) triggers a call of the syntax transformer by the expander; see §1.2.3.2 "Expansion Steps".

Before the expander passes a syntax object to a transformer, the syntax object is extended with a *syntax mark* (that applies to all sub-syntax objects). The result of the transformer is similarly extended with the same syntax mark. When a syntax object's lexical information includes the same mark twice in a row, the marks effectively cancel. Otherwise, two identifiers are bound-identifier=? (that is, one can bind the other) only if they have the same binding and if they have the same marks—counting only marks that were added after the binding.

This marking process helps keep binding in an expanded program consistent with the lexical structure of the source program. For example, the expanded form of the program

```
(define x 12)
(define-syntax m
  (syntax-rules ()
    [(_ id) (let ([x 10]) id)]))
(m x)
```

is

```
(define x 12)
(define-syntax m
  (syntax-rules ()
    [(_ id) (let ([x 10]) id)]))
(let-values ([(x) 10]) x)
```

However, the result of the last expression is 12, not 10. The reason is that the transformer bound to m introduces the binding x, but the referencing x is present in the argument to the transformer. The introduced x is the one left with a mark, and the reference x has no mark, so the binding x is not bound-identifier=? to the body x.

The set! form works with the make-set!-transformer and prop:set!-transformer property to support *assignment transformers* that transform set! expressions. An assignment transformer contains a procedure that is applied by set! in the same way as a normal transformer by the expander.

The make-rename-transformer procedure or prop:rename-transformer property creates a value that is also handled specially by the expander and by set! as a transformer binding's value. When *id* is bound to a *rename transformer* produced by make-rename-transformer, it is replaced with the target identifier passed to make-rename-transformer. In addition, as long as the target identifier does not have a true value for

the 'not-free-identifier=? syntax property, the lexical information that contains the binding of *id* is also enriched so that *id* is free-identifier=? to the target identifier, identifier-binding returns the same results for both identifiers, and provide exports *id* as the target identifier. Finally, the binding is treated specially by syntax-local-value, and syntax-local-make-delta-introducer as used by syntax transformers.

In addition to using marks to track introduced identifiers, the expander tracks the expansion history of a form through syntax properties such as 'origin. See §11.7 "Syntax Object Properties" for more information.

Finally, the expander uses a tamper status to control the way that unexported and protected module bindings are used. See §11.8 "Syntax Taints" for more information on a tamper status.

The expander's handling of letrec-syntaxes+values is similar to its handling of define-syntaxes. A letrec-syntaxes+values can be expanded in an arbitrary phase level *n* (not just 0), in which case the expression for the transformer binding is expanded at phase level *n*+1.

The expressions in a begin-for-syntax form are expanded and evaluated in the same way as for define-syntaxes. However, any introduced bindings from definition within begin-for-syntax are at phase level 1 (not a transformer binding at phase level 0).

Partial Expansion

In certain contexts, such as an internal-definition context or module context, *partial expansion* is used to determine whether forms represent definitions, expressions, or other declaration forms. Partial expansion works by cutting off the normal recursion expansion when the relevant binding is for a primitive syntactic form.

As a special case, when expansion would otherwise add an #%app, #%datum, or #%top identifier to an expression, and when the binding turns out to be the primitive #%app, #%datum, or #%top form, then expansion stops without adding the identifier.

Internal Definitions

An internal-definition context supports local definitions mixed with expressions. Forms that allow internal definitions document such positions using the *body* meta-variable. Definitions in an internal-definition context are equivalent to local binding via letrec-syntaxes+values; macro expansion converts internal definitions to a letrec-syntaxes+values form.

Expansion of an internal-definition context relies on partial expansion of each *body* in an internal-definition sequence. Partial expansion of each *body* produces a form matching one

of the following cases:

- A `define-values` form: The lexical context of all syntax objects for the body sequence is immediately enriched with bindings for the `define-values` form. Further expansion of the definition is deferred, and partial expansion continues with the rest of the body.
- A `define-syntaxes` form: The right-hand side is expanded and evaluated (as for a `letrec-syntaxes+values` form), and a transformer binding is installed for the body sequence before partial expansion continues with the est of the body.
- A primitive expression form other than `begin`: Further expansion of the expression is deferred, and partial expansion continues with the rest of the body.
- A `begin` form: The sub-forms of the `begin` are spliced into the internal-definition sequence, and partial expansion continues with the first of the newly-spliced forms (or the next form, if the `begin` had no sub-forms).

After all body forms are partially expanded, if no definitions were encountered, then the expressions are collected into a `begin` form as he internal-definition context's expansion. Otherwise, at least one expression must appear after the last definition, and any *expr* that appears between definitions is converted to `(define-values () (begin` *expr* `(values)))`; the definitions are then converted to bindings in a `letrec-syntaxes+values` form, and all expressions after the last definition become the body of the `letrec-syntaxes+values` form.

Module Phases and Visits

A `require` form not only introduces bindings at expansion time, but also *visits* the referenced module when it is encountered by the expander. That is, the expander instantiates any variables defined in the module within `begin-for-syntax`, and it also evaluates all expressions for `define-syntaxes` transformer bindings.

Module visits propagate through `requires` in the same way as module instantiation. Moreover, when a module is visited at phase 0, any module that it `requires` `for-syntax` is instantiated at phase 1, while further `requires` `for-template` leading back to phase 0 causes the required module to be visited at phase 0 (i.e., not instantiated).

During compilation, the top-level of module context is itself implicitly visited. Thus, when the expander encounters `(require (for-syntax ....))`, it immediately instantiates the required module at phase 1, in addition to adding bindings at phase level 1 (i.e., the transformer environment). Similarly, the expander immediately evaluates any form that it encounters within `begin-for-syntax`.

Phases beyond 0 are visited on demand. For example, when the right-hand side of a phase-0 `let-syntax` is to be expanded, then modules that are available at phase 1 are visited. More

generally, initiating expansion at phase *n* visits modules at phase *n*, which in turn instantiates modules at phase *n*+1. These visits and instantiations apply to available modules in the enclosing namespace's module registry; a per-registry lock prevents multiple threads from concurrently instantiating and visiting available modules.

When the expander encounters `require` and `(require (for-syntax ....))` within a module context, the resulting visits and instantiations are specific to the expansion of the enclosing module, and are kept separate from visits and instantiations triggered from a top-level context or from the expansion of a different module. Along the same lines, when a module is attached to a namespace through `namespace-attach-module`, modules that it `require`s are transitively attached, but instances are attached only at phases at or below the namespace's base phase.

Macro-Introduced Bindings

When a top-level definition binds an identifier that originates from a macro expansion, the definition captures only uses of the identifier that are generated by the same expansion. This behavior is consistent with expansion in internal-definition contexts, where the defined identifier turns into a fresh lexical binding.

Examples:

```
> (define-syntax def-and-use-of-x
    (syntax-rules ()
      [(def-and-use-of-x val)
       ; x below originates from this macro:
       (begin (define x val) x)]))

> (define x 1)

> x
1
> (def-and-use-of-x 2)
2
> x
1
> (define-syntax def-and-use
    (syntax-rules ()
      [(def-and-use x val)
       ; "x" below was provided by the macro use:
       (begin (define x val) x)]))

> (def-and-use x 3)
3
> x
```

```
3
```

For a top-level definition (outside of a module), the order of evaluation affects the binding of a generated definition for a generated identifier use. If the use precedes the definition, then the use refers to a non-generated binding, just as if the generated definition were not present. (No such dependency on order occurs within a module, since a module binding covers the entire module body.) To support the declaration of an identifier before its use, the `define-syntaxes` form avoids binding an identifier if the body of the `define-syntaxes` declaration produces zero results.

Examples:

```
> (define bucket-1 0)

> (define bucket-2 0)

> (define-syntax def-and-set!-use-of-x
    (syntax-rules ()
      [(def-and-set!-use-of-x val)
       (begin (set! bucket-1 x) (define x val) (set! bucket-
2 x))]))

> (define x 1)

> (def-and-set!-use-of-x 2)

> x
1
> bucket-1
1
> bucket-2
2
> (define-syntax defs-and-uses/fail
    (syntax-rules ()
      [(def-and-use)
       (begin
         ; Initial reference to even precedes definition:
         (define (odd x) (if (zero? x) #f (even (sub1 x))))
         (define (even x) (if (zero? x) #t (odd (sub1 x))))
         (odd 17))]))

> (defs-and-uses/fail)
even: undefined;
 cannot reference undefined identifier
> (define-syntax defs-and-uses
    (syntax-rules ()
      [(def-and-use)
```

```
        (begin
          ; Declare before definition via no-values define-
syntaxes:
          (define-syntaxes (odd even) (values))
          (define (odd x) (if (zero? x) #f (even (sub1 x))))
          (define (even x) (if (zero? x) #t (odd (sub1 x))))
          (odd 17))]))

> (defs-and-uses)
#t
```

Macro-generated `"require"` and `"provide"` clauses also introduce and reference generation-specific bindings:

- In `require`, for a *`require-spec`* of the form `(rename-in [`*`orig-id bind-id`*`])` or `(only-in .... [`*`orig-id bind-id`*`])`, the *`bind-id`* is bound only for uses of the identifier generated by the same macro expansion as *`bind-id`*. In `require` for other *`require-spec`*s, the generator of the *`require-spec`* determines the scope of the bindings.

- In `provide`, for a *`provide-spec`* of the form *`id`*, the exported identifier is the one that binds *`id`* within the module in a generator-specific way, but the external name is the plain *`id`*. The exceptions for `all-except-out` are similarly determined in a generator-specific way, as is the *`orig-id`* binding of a `rename-out` form, but plain identifiers are used for the external names. For `all-defined-out`, only identifiers with definitions having the same generator as the `(all-defined-out)` form are exported; the external name is the plain identifier from the definition.

1.2.4 Compilation

Before expanded code is evaluated, it is first *compiled*. A compiled form has essentially the same information as the corresponding expanded form, though the internal representation naturally dispenses with identifiers for syntactic forms and local bindings. One significant difference is that a compiled form is almost entirely opaque, so the information that it contains cannot be accessed directly (which is why some identifiers can be dropped). At the same time, a compiled form can be marshaled to and from a byte string, so it is suitable for saving and re-loading code.

Although individual read, expand, compile, and evaluate operations are available, the operations are often combined automatically. For example, the `eval` procedure takes a syntax object and expands it, compiles it, and evaluates it.

1.2.5 Namespaces

See §13.1 "Namespaces" for functions that manipulate namespaces.

A *namespace* is a top-level mapping from symbols to binding information. It is the starting point for expanding an expression; a syntax object produced by `read-syntax` has no initial lexical context; the syntax object can be expanded after initializing it with the mappings of a particular namespace. A namespace is also the starting point evaluating expanded code, where the first step in evaluation is linking the code to specific module instances and top-level variables.

For expansion purposes, a namespace maps each symbol in each phase level to one of three possible bindings:

- a particular module binding from a particular module
- a top-level transformer binding named by the symbol
- a top-level variable named by the symbol

An "empty" namespace maps all symbols to top-level variables. Certain evaluations extend a namespace for future expansions; importing a module into the top-level adjusts the namespace bindings for all of the imported names, and evaluating a top-level `define` form updates the namespace's mapping to refer to a variable (in addition to installing a value into the variable).

A namespace also has a *module registry* that maps module names to module declarations (see §1.1.10 "Modules and Module-Level Variables"). This registry is shared by all phase levels.

For evaluation, each namespace encapsulates a distinct set of top-level variables at various phases, as well as a potentially distinct set of module instances in each phase. That is, even though module declarations are shared for all phase levels, module instances are distinct for each phase. Each namespace has a *base phase*, which corresponds to the phase used by reflective operations such as `eval` and `dynamic-require`. In particular, using `eval` on a `require` form instantiates a module in the namespace's base phase.

After a namespace is created, module instances from existing namespaces can be attached to the new namespace. In terms of the evaluation model, top-level variables from different namespaces essentially correspond to definitions with different prefixes, but attaching a module uses the same prefix for the module's definitions in namespaces where it is attached. The first step in evaluating any compiled expression is to link its top-level variable and module-level variable references to specific variables in the namespace.

At all times during evaluation, some namespace is designated as the *current namespace*. The current namespace has no particular relationship, however, with the namespace that was used to expand the code that is executing, or with the namespace that was used to link the compiled form of the currently evaluating code. In particular, changing the current namespace during

evaluation does not change the variables to which executing expressions refer. The current namespace only determines the behavior of reflective operations to expand code and to start evaluating expanded/compiled code.

Examples:

```
> (define x 'orig) ; define in the original namespace

; The following let expression is compiled in the original
; namespace, so direct references to x see 'orig.
> (let ([n (make-base-namespace)]) ; make new namespace
    (parameterize ([current-namespace n])
      (eval '(define x 'new)) ; evals in the new namespace
      (display x) ; displays 'orig
      (display (eval 'x)))) ; displays 'new
orignew
```

A namespace is purely a top-level entity, not to be confused with an environment. In particular, a namespace does not encapsulate the full environment of an expression inside local-binding forms.

If an identifier is bound to syntax or to an import, then defining the identifier as a variable shadows the syntax or import in future uses of the environment. Similarly, if an identifier is bound to a top-level variable, then binding the identifier to syntax or an import shadows the variable; the variable's value remains unchanged, however, and may be accessible through previously evaluated expressions.

Examples:

```
> (define x 5)

> (define (f) x)

> x
5
> (f)
5
> (define-syntax x (syntax-id-rules () [_ 10]))

> x
10
> (f)
5
> (define x 7)

> x
```

```
7
> (f)
7
> (module m racket (define x 8) (provide x))

> (require 'm)

> x
8
> (f)
7
```

1.2.6 Inferred Value Names

To improve error reporting, names are inferred at compile-time for certain kinds of values, such as procedures. For example, evaluating the following expression:

```
(let ([f (lambda () 0)]) (f 1 2 3))
```

produces an error message because too many arguments are provided to the procedure. The error message is able to report `f` as the name of the procedure. In this case, Racket decides, at compile-time, to name as `'f` all procedures created by the `let`-bound `lambda`.

Names are inferred whenever possible for procedures. Names closer to an expression take precedence. For example, in

```
(define my-f
  (let ([f (lambda () 0)]) f))
```

the procedure bound to `my-f` will have the inferred name `'f`.

When an `'inferred-name` property is attached to a syntax object for an expression (see §11.7 “Syntax Object Properties”), the property value is used for naming the expression, and it overrides any name that was inferred from the expression’s context. Normally, the property value should be a symbol. A `'inferred-name` property value of `#<void>` hides a name that would otherwise be inferred from context (perhaps because a binding identifier’s was automatically generated and should not be exposed).

When an inferred name is not available, but a source location is available, a name is constructed using the source location information. Inferred and property-assigned names are also available to syntax transformers, via `syntax-local-name`.

1.3 The Reader

Racket's reader is a recursive-descent parser that can be configured through a readtable and various other parameters. This section describes the reader's parsing when using the default readtable.

Reading from a stream produces one *datum*. If the result datum is a compound value, then reading the datum typically requires the reader to call itself recursively to read the component data.

The reader can be invoked in either of two modes: `read` mode, or `read-syntax` mode. In `read-syntax` mode, the result is always a syntax object that includes source-location and (initially empty) lexical information wrapped around the sort of datum that `read` mode would produce. In the case of pairs, vectors, and boxes, the content is also wrapped recursively as a syntax object. Unless specified otherwise, this section describes the reader's behavior in `read` mode, and `read-syntax` mode does the same modulo wrapping of the final result.

Reading is defined in terms of Unicode characters; see §12.1 "Ports" for information on how a byte stream is converted to a character stream.

Symbols, keywords, strings, byte strings, regexps, characters, and numbers produced by the reader in `read-syntax` mode are *interned*, which means that such values in the result of `read-syntax` are always `eq?` when they are `equal?` (whether from the same call or different calls to `read-syntax`). Symbols and keywords are interned in both `read` and `read-syntax` mode. Sending an interned value across a place channel does not necessarily produce an interned value at the receiving place. See also `datum-intern-literal` and `datum->syntax`.

1.3.1 Delimiters and Dispatch

Along with whitespace, the following characters are *delimiters*:

```
( ) [ ] { } " , ' ` ;
```

A delimited sequence that starts with any other character is typically parsed as either a symbol or number, but a few non-delimiter characters play special roles:

- `#` has a special meaning as an initial character in a delimited sequence; its meaning depends on the characters that follow; see below.
- `|` starts a subsequence of characters to be included verbatim in the delimited sequence (i.e., they are never treated as delimiters, and they are not case-folded when case-insensitivity is enabled); the subsequence is terminated by another `|`, and neither the initial nor terminating `|` is part of the subsequence.

- `\` outside of a `|` pair causes the following character to be included verbatim in a delimited sequence.

More precisely, after skipping whitespace, the reader dispatches based on the next character or characters in the input stream as follows:

`(` starts a pair or list; see §1.3.5 “Reading Pairs and Lists”
`[` starts a pair or list; see §1.3.5 “Reading Pairs and Lists”
`{` starts a pair or list; see §1.3.5 “Reading Pairs and Lists”
`)` matches `(` or raises `exn:fail:read`
`]` matches `[` or raises `exn:fail:read`
`}` matches `{` or raises `exn:fail:read`
`"` starts a string; see §1.3.6 “Reading Strings”
`'` starts a quote; see §1.3.7 “Reading Quotes”
`` ` `` starts a quasiquote; see §1.3.7 “Reading Quotes”
`,` starts a [splicing] unquote; see §1.3.7 “Reading Quotes”
`;` starts a line comment; see §1.3.8 “Reading Comments”
`#t` or `#T` true; see §1.3.4 “Reading Booleans”
`#f` or `#F` false; see §1.3.4 “Reading Booleans”
`#(` starts a vector; see §1.3.9 “Reading Vectors”
`#[` starts a vector; see §1.3.9 “Reading Vectors”
`#{` starts a vector; see §1.3.9 “Reading Vectors”
`#s(` starts a structure literal; see §1.3.10 “Reading Structures”
`#s[` starts a structure literal; see §1.3.10 “Reading Structures”
`#s{` starts a structure literal; see §1.3.10 “Reading Structures”
`#\` starts a character; see §1.3.13 “Reading Characters”
`#"` starts a byte string; see §1.3.6 “Reading Strings”
`#%` starts a symbol; see §1.3.2 “Reading Symbols”
`#:` starts a keyword; see §1.3.14 “Reading Keywords”
`#&` starts a box; see §1.3.12 “Reading Boxes”
`#|` starts a block comment; see §1.3.8 “Reading Comments”
`#;` starts an S-expression comment; see §1.3.8 “Reading Comments”
`#'` starts a syntax quote; see §1.3.7 “Reading Quotes”
`#!` starts a line comment; see §1.3.8 “Reading Comments”
`#!/` starts a line comment; see §1.3.8 “Reading Comments”
`#!` may start a reader extension; see §1.3.17 “Reading via an Extension”
`#`` starts a syntax quasiquote; see §1.3.7 “Reading Quotes”
`#,` starts a syntax [splicing] unquote; see §1.3.7 “Reading Quotes”
`#~` starts compiled code; see §1.4.15 “Printing Compiled Code”
`#i` or `#I` starts a number; see §1.3.3 “Reading Numbers”
`#e` or `#E` starts a number; see §1.3.3 “Reading Numbers”
`#x` or `#X` starts a number; see §1.3.3 “Reading Numbers”
`#o` or `#O` starts a number; see §1.3.3 “Reading Numbers”
`#d` or `#D` starts a number; see §1.3.3 “Reading Numbers”
`#b` or `#B` starts a number; see §1.3.3 “Reading Numbers”
`#<<` starts a string; see §1.3.6 “Reading Strings”

`#rx`	starts a regular expression; see §1.3.15 “Reading Regular Expressions”
`#px`	starts a regular expression; see §1.3.15 “Reading Regular Expressions”
`#ci`, `#cI`, `#Ci`, or `#CI`	switches case sensitivity; see §1.3.2 “Reading Symbols”
`#cs`, `#cS`, `#Cs`, or `#CS`	switches case sensitivity; see §1.3.2 “Reading Symbols”
`#hash`	starts a hash table; see §1.3.11 “Reading Hash Tables”
`#reader`	starts a reader extension use; see §1.3.17 “Reading via an Extension”
`#lang`	starts a reader extension use; see §1.3.17 “Reading via an Extension”
`#`⟨*digit*$_{10}$⟩$^{+}$`(`	starts a vector; see §1.3.9 “Reading Vectors”
`#`⟨*digit*$_{10}$⟩$^{+}$`[`	starts a vector; see §1.3.9 “Reading Vectors”
`#`⟨*digit*$_{10}$⟩$^{+}$`{`	starts a vector; see §1.3.9 “Reading Vectors”
`#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`=`	binds a graph tag; see §1.3.16 “Reading Graph Structure”
`#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`#`	uses a graph tag; see §1.3.16 “Reading Graph Structure”
otherwise	starts a symbol; see §1.3.2 “Reading Symbols”

1.3.2 Reading Symbols

§3.6 “Symbols” in ***The Racket Guide*** introduces the syntax of symbols.

A sequence that does not start with a delimiter or `#` is parsed as either a symbol or a number (see §1.3.3 “Reading Numbers”), except that `.` by itself is never parsed as a symbol or character (unless the `read-accept-dot` parameter is set to `#f`). A `#%` also starts a symbol. The resulting symbol is interned. A successful number parse takes precedence over a symbol parse.

When the `read-case-sensitive` parameter is set to `#f`, characters in the sequence that are not quoted by `|` or `\` are first case-normalized. If the reader encounters `#ci`, `#CI`, `#Ci`, or `#cI`, then it recursively reads the following datum in case-insensitive mode. If the reader encounters `#cs`, `#CS`, `#Cs`, or `#cS`, then it recursively reads the following datum in case-sensitive mode.

Examples:

`Apple`	reads equal to	`(string->symbol "Apple")`
`Ap#ple`	reads equal to	`(string->symbol "Ap#ple")`
`Ap ple`	reads equal to	`(string->symbol "Ap")`
`Ap| |ple`	reads equal to	`(string->symbol "Ap ple")`
`Ap\ ple`	reads equal to	`(string->symbol "Ap ple")`
`#ci Apple`	reads equal to	`(string->symbol "apple")`
`#ci |A|pple`	reads equal to	`(string->symbol "Apple")`
`#ci \Apple`	reads equal to	`(string->symbol "Apple")`
`#ci#cs Apple`	reads equal to	`(string->symbol "Apple")`
`#%Apple`	reads equal to	`(string->symbol "#%Apple")`

1.3.3 Reading Numbers

§3.2 “Numbers” in ***The Racket Guide*** introduces the syntax of numbers.

A sequence that does not start with a delimiter is parsed as a number when it matches the following grammar case-insensitively for ⟨*number*$_{10}$⟩ (decimal), where n is a meta-meta-variable in the grammar. The resulting number is interned in `read-syntax` mode.

A number is optionally prefixed by an exactness specifier, `#e` (exact) or `#i` (inexact), which specifies its parsing as an exact or inexact number; see §3.2 “Numbers” for information on number exactness. As the non-terminal names suggest, a number that has no exactness specifier and matches only ⟨*inexact-number*$_n$⟩ is normally parsed as an inexact number, otherwise it is parsed as an exact number. If the `read-decimal-as-inexact` parameter is set to `#f`, then all numbers without an exactness specifier are instead parsed as exact.

If the reader encounters `#b` (binary), `#o` (octal), `#d` (decimal), or `#x` (hexadecimal), it must be followed by a sequence that is terminated by a delimiter or end-of-file, and that matches the ⟨*general-number*$_2$⟩, ⟨*general-number*$_8$⟩, ⟨*general-number*$_{10}$⟩, or ⟨*general-number*$_{16}$⟩ grammar, respectively.

An ⟨*exponent-mark*$_n$⟩ in an inexact number serves both to specify an exponent and to specify a numerical precision. If single-precision IEEE floating point is supported (see §3.2 “Numbers”), the marks `f` and `s` specify single-precision. Otherwise, or with any other mark, double-precision IEEE floating point is used. In addition, single- and double-precision specials are distinct; specials with the `.0` suffix, like `+nan.0` are double-precision, whereas specials with the `.f` suffix are single-precision.

A `#` in an ⟨*inexact*$_n$⟩ number is the same as `0`, but `#` can be used to suggest that the digit's actual value is unknown.

⟨*number*$_n$⟩ ::= ⟨*exact*$_n$⟩ | ⟨*inexact*$_n$⟩
⟨*exact*$_n$⟩ ::= ⟨*exact-integer*$_n$⟩ | ⟨*exact-rational*$_n$⟩
| ⟨*exact-complex*$_n$⟩
⟨*exact-integer*$_n$⟩ ::= [⟨*sign*⟩] ⟨*unsigned-integer*$_n$⟩
⟨*unsigned-integer*$_n$⟩ ::= ⟨*digit*$_n$⟩$^+$
⟨*exact-rational*$_n$⟩ ::= ⟨*exact-integer*$_n$⟩ `/` ⟨*unsigned-integer*$_n$⟩
⟨*exact-complex*$_n$⟩ ::= ⟨*exact-rational*$_n$⟩ ⟨*sign*⟩ ⟨*exact-rational*$_n$⟩ `i`
⟨*inexact*$_n$⟩ ::= ⟨*inexact-real*$_n$⟩ | ⟨*inexact-complex*$_n$⟩
⟨*inexact-real*$_n$⟩ ::= [⟨*sign*⟩] ⟨*inexact-normal*$_n$⟩
| ⟨*sign*⟩ ⟨*inexact-special*$_n$⟩
⟨*inexact-unsigned*$_n$⟩ ::= ⟨*inexact-normal*$_n$⟩ | ⟨*inexact-special*$_n$⟩
⟨*inexact-normal*$_n$⟩ ::= ⟨*inexact-simple*$_n$⟩ [⟨*exp-mark*$_n$⟩ [⟨*sign*⟩] ⟨*digits#*$_n$⟩]
⟨*inexact-simple*$_n$⟩ ::= ⟨*digits#*$_n$⟩ [`.`] `#`*
| [⟨*exact-integer*$_n$⟩] `.` ⟨*digits#*$_n$⟩
| ⟨*digits#*$_n$⟩ `/` ⟨*digits#*$_n$⟩
⟨*inexact-special*$_n$⟩ ::= `inf.0` | `nan.0` | `inf.f` | `nan.f`
⟨*digits#*$_n$⟩ ::= ⟨*digit*$_n$⟩$^+$ `#`*
⟨*inexact-complex*$_n$⟩ ::= [⟨*inexact-real*$_n$⟩] ⟨*sign*⟩ ⟨*inexact-unsigned*$_n$⟩ `i`

| ⟨*inexact-real*$_n$⟩ `@` ⟨*inexact-real*$_n$⟩

⟨*sign*⟩ ::= `+` | `-`

⟨*digit*$_{16}$⟩ ::= ⟨*digit*$_{10}$⟩ | `a` | `b` | `c` | `d` | `e` | `f`

⟨*digit*$_{10}$⟩ ::= ⟨*digit*$_8$⟩ | `8` | `9`

⟨*digit*$_8$⟩ ::= ⟨*digit*$_2$⟩ | `2` | `3` | `4` | `5` | `6` | `7`

⟨*digit*$_2$⟩ ::= `0` | `1`

⟨*exp-mark*$_{16}$⟩ ::= `s` | `l`

⟨*exp-mark*$_{10}$⟩ ::= ⟨*exp-mark*$_{16}$⟩ | `d` | `e` | `f`

⟨*exp-mark*$_8$⟩ ::= ⟨*exp-mark*$_{10}$⟩

⟨*exp-mark*$_2$⟩ ::= ⟨*exp-mark*$_{10}$⟩

⟨*general-number*$_n$⟩ ::= [⟨*exactness*⟩] ⟨*number*$_n$⟩

⟨*exactness*⟩ ::= `#e` | `#i`

Examples:

```
-1          reads equal to -1
1/2         reads equal to (/ 1 2)
1.0         reads equal to (exact->inexact 1)
1+2i        reads equal to (make-complex 1 2)
1/2+3/4i    reads equal to (make-complex (/ 1 2) (/ 3 4))
1.0+3.0e7i  reads equal to (exact->inexact (make-complex 1 30000000))
2e5         reads equal to (exact->inexact 200000)
#i5         reads equal to (exact->inexact 5)
#e2e5       reads equal to 200000
#x2e5       reads equal to 741
#b101       reads equal to 5
```

1.3.4 Reading Booleans

A `#true`, `#t`, `#T` followed by a delimiter is the input syntax for the boolean constant “true,” and `#false`, `#f`, or `#F` followed by a delimiter is the complete input syntax for the boolean constant “false.”

1.3.5 Reading Pairs and Lists

When the reader encounters a `(`, `[`, or `{`, it starts parsing a pair or list; see §3.9 “Pairs and Lists” for information on pairs and lists.

To parse the pair or list, the reader recursively reads data until a matching `)`, `]`, or `}` (respectively) is found, and it specially handles a delimited `.`. Pairs `()`, `[]`, and `{}` are treated the same way, so the remainder of this section simply uses “parentheses” to mean any of these pair.

If the reader finds no delimited `.` among the elements between parentheses, then it produces

a list containing the results of the recursive reads.

If the reader finds two data between the matching parentheses that are separated by a delimited `.`, then it creates a pair. More generally, if it finds two or more data where the last datum is preceded by a delimited `.`, then it constructs nested pairs: the next-to-last element is paired with the last, then the third-to-last datum is paired with that pair, and so on.

If the reader finds three or more data between the matching parentheses, and if a pair of delimited `.`s surrounds any other than the first and last elements, the result is a list containing the element surrounded by `.`s as the first element, followed by the others in the read order. This convention supports a kind of infix notation at the reader level.

In `read-syntax` mode, the recursive reads for the pair/list elements are themselves in `read-syntax` mode, so that the result is a list or pair of syntax objects that is itself wrapped as a syntax object. If the reader constructs nested pairs because the input included a single delimited `.`, then only the innermost pair and outermost pair are wrapped as syntax objects. Whether wrapping a pair or list, if the pair or list was formed with `[` and `]`, then a `'paren-shape` property is attached to the result with the value `#\[`; if the list or pair was formed with `{` and `}`, then a `'paren-shape` property is attached to the result with the value `#\{`.

If a delimited `.` appears in any other configuration, then the `exn:fail:read` exception is raised. Similarly, if the reader encounters a `)`, `]`, or `}` that does not end a list being parsed, then the `exn:fail:read` exception is raised.

Examples:

`()`	reads equal to `(list)`
`(1 2 3)`	reads equal to `(list 1 2 3)`
`{1 2 3}`	reads equal to `(list 1 2 3)`
`[1 2 3]`	reads equal to `(list 1 2 3)`
`(1 (2) 3)`	reads equal to `(list 1 (list 2) 3)`
`(1 . 3)`	reads equal to `(cons 1 3)`
`(1 . (3))`	reads equal to `(list 1 3)`
`(1 . 2 . 3)`	reads equal to `(list 2 1 3)`

If the `read-square-bracket-as-paren` parameter is set to `#f`, then when the reader encounters `[` and `]`, the `"exn:fail:read"` exception is raised. Similarly, if the `read-curly-brace-as-paren` parameter is set to `#f`, then when the reader encounters `{` and `}`, the `"exn:fail:read"` exception is raised.

If the `read-accept-dot` parameter is set to `#f`, then a delimited `.` triggers an `exn:fail:read` exception. If the `read-accept-infix-dot` parameter is set to `#f`, then multiple delimited `.`s trigger an `exn:fail:read` exception, instead of the infix conversion.

1.3.6 Reading Strings

§3.4 “Strings (Unicode)” in ***The Racket Guide*** introduces the syntax of strings.

When the reader encounters `"`, it begins parsing characters to form a string. The string continues until it is terminated by another `"` (that is not escaped by `\`). The resulting string is interned in `read-syntax` mode.

Within a string sequence, the following escape sequences are recognized:

- `\a`: alarm (ASCII 7)
- `\b`: backspace (ASCII 8)
- `\t`: tab (ASCII 9)
- `\n`: linefeed (ASCII 10)
- `\v`: vertical tab (ASCII 11)
- `\f`: formfeed (ASCII 12)
- `\r`: return (ASCII 13)
- `\e`: escape (ASCII 27)
- `\"`: double-quotes (without terminating the string)
- `\'`: quote (i.e., the backslash has no effect)
- `\\`: backslash (i.e., the second is not an escaping backslash)
- `\`$\langle digit_8 \rangle^{\{1,3\}}$: Unicode for the octal number specified by $\text{digit}_8^{\{1,3\}}$ (i.e., 1 to 3 $\langle digit_8 \rangle$s), where each $\langle digit_8 \rangle$ is `0`, `1`, `2`, `3`, `4`, `5`, `6`, or `7`. A longer form takes precedence over a shorter form, and the resulting octal number must be between 0 and 255 decimal, otherwise the `exn:fail:read` exception is raised.
- `\x`$\langle digit_{16} \rangle^{\{1,2\}}$: Unicode for the hexadecimal number specified by $\langle digit_{16} \rangle^{\{1,2\}}$, where each $\langle digit_{16} \rangle$ is `0`, `1`, `2`, `3`, `4`, `5`, `6`, `7`, `8`, `9`, `a`, `b`, `c`, `d`, `e`, or `f` (case-insensitive). The longer form takes precedence over the shorter form.
- `\u`$\langle digit_{16} \rangle^{\{1,4\}}$: like `\x`, but with up to four hexadecimal digits (longer sequences take precedence). The resulting hexadecimal number must be a valid argument to `integer->char`, otherwise the `exn:fail:read` exception is raised—unless the encoding continues with another `\u` to form a surrogate-style encoding.
- `\u`$\langle digit_{16} \rangle^{\{4,4\}}$`\u`$\langle digit_{16} \rangle^{\{4,4\}}$: like `\u`, but for two hexadecimal numbers, where the first is in the range `#xD800` to `#xDBFF` and the second is in the range `#xDC00` to `#xDFFF`; the resulting character is the one represented by the numbers as a UTF-16 surrogate pair.

- `\U`⟨*digit*$_{16}$⟩$^{\{1,8\}}$: like `\x`, but with up to eight hexadecimal digits (longer sequences take precedence). The resulting hexadecimal number must be a valid argument to `integer->char`, otherwise the `exn:fail:read` exception is raised.

- `\`⟨*newline*⟩: elided, where ⟨*newline*⟩ is either a linefeed, carriage return, or carriage return–linefeed combination. This convention allows single-line strings to span multiple lines in the source.

If the reader encounters any other use of a backslash in a string constant, the `exn:fail:read` exception is raised.

§3.5 "Bytes and Byte Strings" in ***The Racket Guide*** introduces the syntax of byte strings.

A string constant preceded by `#` is parsed as a byte string. (That is, `#"` starts a byte-string literal.) See §3.4 "Byte Strings" for information on byte strings. The resulting byte string is interned in `read-syntax` mode. Byte-string constants support the same escape sequences as character strings, except `\u` and `\U`.

When the reader encounters `#<<`, it starts parsing a *here string*. The characters following `#<<` until a newline character define a terminator for the string. The content of the string includes all characters between the `#<<` line and a line whose only content is the specified terminator. More precisely, the content of the string starts after a newline following `#<<`, and it ends before a newline that is followed by the terminator, where the terminator is itself followed by either a newline or end-of-file. No escape sequences are recognized between the starting and terminating lines; all characters are included in the string (and terminator) literally. A return character is not treated as a line separator in this context. If no characters appear between `#<<` and a newline or end-of-file, or if an end-of-file is encountered before a terminating line, the `exn:fail:read` exception is raised.

Examples:

```
"Apple"      reads equal to "Apple"
"\x41pple"   reads equal to "Apple"
"\"Apple\""  reads equal to "\x22Apple\x22"
"\\"         reads equal to "\x5C"
#"Apple"     reads equal to (bytes 65 112 112 108 101)
```

1.3.7 Reading Quotes

When the reader encounters `'`, it recursively reads one datum and forms a new list containing the symbol `'quote` and the following datum. This convention is mainly useful for reading Racket code, where `'s` can be used as a shorthand for `(quote s)`.

Several other sequences are recognized and transformed in a similar way. Longer prefixes take precedence over short ones:

```
'   adds quote
`   adds quasiquote
```

`,` adds `unquote`
`,@` adds `unquote-splicing`
`#'` adds `syntax`
`#‘` adds `quasisyntax`
`#,` adds `unsyntax`
`#,@` adds `unsyntax-splicing`

Examples:
`'apple` reads equal to `(list 'quote 'apple)`
`‘(1 ,2)` reads equal to `(list 'quasiquote (list 1 (list 'unquote 2)))`

The `‘`, `,`, and `,@` forms are disabled when the `read-accept-quasiquote` parameter is set to `#f`, in which case the `exn:fail:read` exception is raised instead.

1.3.8 Reading Comments

A `;` starts a line comment. When the reader encounters `;`, it skips past all characters until the next linefeed (ASCII 10), carriage return (ASCII 13), next-line (Unicode 133), line-separator (Unicode 8232), or paragraph-separator (Unicode 8233) character.

A `#|` starts a nestable block comment. When the reader encounters `#|`, it skips past all characters until a closing `|#`. Pairs of matching `#|` and `|#` can be nested.

A `#;` starts an S-expression comment. When the reader encounters `#;`, it recursively reads one datum, and then discards it (continuing on to the next datum for the read result).

A `#!` (which is `#!` followed by a space) or `#!/` starts a line comment that can be continued to the next line by ending a line with `\`. This form of comment normally appears at the beginning of a Unix script file.

Examples:

`; comment`	reads equal to *nothing*
`#| a |# 1`	reads equal to `1`
`#| #| a |# 1 |# 2`	reads equal to `2`
`#;1 2`	reads equal to `2`
`#!/bin/sh`	reads equal to *nothing*
`#! /bin/sh`	reads equal to *nothing*

1.3.9 Reading Vectors

When the reader encounters a `#(`, `#[`, or `#{`, it starts parsing a vector; see §3.11 “Vectors” for information on vectors. The `#[` and `#{` forms can be disabled through the `read-square-bracket-as-paren` and `read-curly-brace-as-paren` parameters.

The elements of the vector are recursively read until a matching `)`, `]`, or `}` is found, just as for lists (see §1.3.5 “Reading Pairs and Lists”). A delimited `.` is not allowed among the vector elements.

An optional vector length can be specified between the `#` and `(`, `[`, or `{`. The size is specified using a sequence of decimal digits, and the number of elements provided for the vector must be no more than the specified size. If fewer elements are provided, the last provided element is used for the remaining vector slots; if no elements are provided, then `0` is used for all slots.

In `read-syntax` mode, each recursive read for the vector elements is also in `read-syntax` mode, so that the wrapped vector’s elements are also wrapped as syntax objects, and the vector is immutable.

Examples:

`#(1 apple 3)`	reads equal to	`(vector 1 'apple 3)`
`#3("apple" "banana")`	reads equal to	`(vector "apple" "banana" "banana")`
`#3()`	reads equal to	`(vector 0 0 0)`

1.3.10 Reading Structures

When the reader encounters a `#s(`, `#s[`, or `#s{`, it starts parsing an instance of a prefab structure type; see §4 “Structures” for information on structure types. The `#s[` and `#s{` forms can be disabled through the `read-square-bracket-as-paren` and `read-curly-brace-as-paren` parameters.

The elements of the structure are recursively read until a matching `)`, `]`, or `}` is found, just as for lists (see §1.3.5 “Reading Pairs and Lists”). A single delimited `.` is not allowed among the elements, but two `.`s can be used as in a list for an infix conversion.

The first element is used as the structure descriptor, and it must have the form (when quoted) of a possible argument to `make-prefab-struct`; in the simplest case, it can be a symbol. The remaining elements correspond to field values within the structure.

In `read-syntax` mode, the structure type must not have any mutable fields. The structure’s elements are read in `read-syntax` mode, so that the wrapped structure’s elements are also wrapped as syntax objects.

If the first structure element is not a valid prefab structure type key, or if the number of provided fields is inconsistent with the indicated prefab structure type, the `exn:fail:read` exception is raised.

1.3.11 Reading Hash Tables

A `#hash` starts an immutable hash-table constant with key matching based on `equal?`. The characters after `hash` must parse as a list of pairs (see §1.3.5 “Reading Pairs and Lists”) with a specific use of delimited `.`: it must appear between the elements of each pair in the list and nowhere in the sequence of list elements. The first element of each pair is used as the key for a table entry, and the second element of each pair is the associated value.

A `#hasheq` starts a hash table like `#hash`, except that it constructs a hash table based on `eq?` instead of `equal?`.

A `#hasheqv` starts a hash table like `#hash`, except that it constructs a hash table based on `eqv?` instead of `equal?`.

In all cases, the table is constructed by adding each mapping to the hash table from left to right, so later mappings can hide earlier mappings if the keys are equivalent.

Examples, where `make-...` stands for `make-immutable-hash`:

`#hash()`	reads equal to	`(make-... '())`
`#hasheq()`	reads equal to	`(make-...eq '())`
`#hash(("a" . 5))`	reads equal to	`(make-... '(("a" . 5)))`
`#hasheq((a . 5) (b . 7))`	reads equal to	`(make-...eq '((b . 7) (a . 5)))`
`#hasheq((a . 5) (a . 7))`	reads equal to	`(make-...eq '((a . 7)))`

1.3.12 Reading Boxes

When the reader encounters a `#&`, it starts parsing a box; see §3.12 “Boxes” for information on boxes. The content of the box is determined by recursively reading the next datum.

In `read-syntax` mode, the recursive read for the box content is also in `read-syntax` mode, so that the wrapped box’s content is also wrapped as a syntax object, and the box is immutable.

Examples:
`#&17` reads equal to `(box 17)`

1.3.13 Reading Characters

§3.3 “Characters” in ***The Racket Guide*** introduces the syntax of characters.

A `#\` starts a character constant, which has one of the following forms:

- `#\nul` or `#\null`: NUL (ASCII 0); the next character must not be alphabetic.
- `#\backspace`: backspace (ASCII 8); the next character must not be alphabetic.

- `#\tab`: tab (ASCII 9); the next character must not be alphabetic.
- `#\newline` or `#\linefeed`: linefeed (ASCII 10); the next character must not be alphabetic.
- `#\vtab`: vertical tab (ASCII 11); the next character must not be alphabetic.
- `#\page`: page break (ASCII 12); the next character must not be alphabetic.
- `#\return`: carriage return (ASCII 13); the next character must not be alphabetic.
- `#\space`: space (ASCII 32); the next character must not be alphabetic.
- `#\rubout`: delete (ASCII 127); the next character must not be alphabetic.
- `#\`$\langle digit_8 \rangle^{\{1,3\}}$: Unicode for the octal number specified by $\langle digit_8 \rangle^{\{1,3\}}$, as in string escapes (see §1.3.6 "Reading Strings").
- `#\u`$\langle digit_{16} \rangle^{\{1,4\}}$: like `#\x`, but with up to four hexadecimal digits.
- `#\U`$\langle digit_{16} \rangle^{\{1,6\}}$: like `#\x`, but with up to six hexadecimal digits.
- `#\`$\langle c \rangle$: the character $\langle c \rangle$, as long as `#\`$\langle c \rangle$ and the characters following it do not match any of the previous cases, and as long as the character after $\langle c \rangle$ is not alphabetic.

Examples:
```
#\newline  reads equal to (integer->char 10)
#\n        reads equal to (integer->char 110)
#\u3BB     reads equal to (integer->char 955)
#\λ        reads equal to (integer->char 955)
```

1.3.14 Reading Keywords

A `#:` starts a keyword. The parsing of a keyword after the `#:` is the same as for a symbol, including case-folding in case-insensitive mode, except that the part after `#:` is never parsed as a number. The resulting keyword is interned.

Examples:
```
#:Apple  reads equal to (string->keyword "Apple")
#:1      reads equal to (string->keyword "1")
```

1.3.15 Reading Regular Expressions

A `#rx` or `#px` starts a regular expression. The characters immediately after `#rx` or `#px` must parse as a string or byte string (see §1.3.6 "Reading Strings"). A `#rx` prefix starts a regular expression as would be constructed by `regexp`, `#px` as constructed by `pregexp`, `#rx#` as constructed by `byte-regexp`, and `#px#` as constructed by `byte-pregexp`. The resulting regular expression is interned in `read-syntax` mode.

Examples:

```
#rx".*"          reads equal to (regexp ".*")
#px"[\\s]*"      reads equal to (pregexp "[\\s]*")
#rx#".*"         reads equal to (byte-regexp #".*")
#px#"[\\s]*"     reads equal to (byte-pregexp #"[\\s]*")
```

1.3.16 Reading Graph Structure

A `#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`=` tags the following datum for reference via `#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`#`, which allows the reader to produce a datum that has graph structure.

For a specific ⟨*digit*$_{10}$⟩$^{\{1,8\}}$ in a single read result, each `#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`#` reference is replaced by the datum read for the corresponding `#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`=`; the definition `#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`=` also produces just the datum after it. A `#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`=` definition can appear at most once, and a `#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`=` definition must appear before a `#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`#` reference appears, otherwise the `exn:fail:read` exception is raised. If the `read-accept-graph` parameter is set to `#f`, then `#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`=` or `#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`#` triggers a `exn:fail:read` exception.

Although a comment parsed via `#;` discards the datum afterward, `#`⟨*digit*$_{10}$⟩$^{\{1,8\}}$`=` definitions in the discarded datum still can be referenced by other parts of the reader input, as long as both the comment and the reference are grouped together by some other form (i.e., some recursive read); a top-level `#;` comment neither defines nor uses graph tags for other top-level forms.

Examples:

```
(#1=100 #1# #1#)  reads equal to (list 100 100 100)
#0=(1 . #0#)      reads equal to (let* ([ph (make-placeholder #f)]
                                       [v (cons 1 ph)])
                                  (placeholder-set! ph v)
                                  (make-reader-graph v))
```

1.3.17 Reading via an Extension

§17.2 “Reader Extensions” in ***The Racket Guide*** introduces reader extension.

When the reader encounters `#reader`, it loads an external reader procedure and applies it to the current input stream.

The reader recursively reads the next datum after `#reader`, and passes it to the procedure that is the value of the `current-reader-guard` parameter; the result is used as a module path. The module path is passed to `dynamic-require` with either `'read` or `'read-syntax` (depending on whether the reader is in `read` or `read-syntax` mode).

The arity of the resulting procedure determines whether it accepts extra source-location information: a `read` procedure accepts either one argument (an input port) or five, and a `read-syntax` procedure accepts either two arguments (a name value and an input port) or

six. In either case, the four optional arguments are the reader's module path (as a syntax object in `read-syntax` mode) followed by the line (positive exact integer or `#f`), column (non-negative exact integer or `#f`), and position (positive exact integer or `#f`) of the start of the `#reader` form. The input port is the one whose stream contained `#reader`, where the stream position is immediately after the recursively read module path.

The procedure should produce a datum result. If the result is a syntax object in `read` mode, then it is converted to a datum using `syntax->datum`; if the result is not a syntax object in `read-syntax` mode, then it is converted to one using `datum->syntax`. See also §12.7.2 "Reader-Extension Procedures" for information on the procedure's results.

If the `read-accept-reader` parameter is set to `#f`, then if the reader encounters `#reader`, the `exn:fail:read` exception is raised.

§6.2.2 "The `#lang` Shorthand" in ***The Racket Guide*** introduces `#lang`.

The `#lang` reader form is similar to `#reader`, but more constrained: the `#lang` must be followed by a single space (ASCII 32), and then a non-empty sequence of alphanumeric ASCII, `+`, `-`, `_`, and/or `/` characters terminated by whitespace or an end-of-file. The sequence must not start or end with `/`. A sequence `#lang` ⟨*name*⟩ is equivalent to either `#reader (submod` ⟨*name*⟩ `reader)` or `#reader` ⟨*name*⟩`/lang/reader`, where the former is tried first guarded by a `module-declared?` check (but after filtering by `current-reader-guard`, so both are passed to the value of `current-reader-guard` if the latter is used). Note that the terminating whitespace (if any) is not consumed before the external reading procedure is called.

§17.3 "Defining new `#lang` Languages" in ***The Racket Guide*** introduces the creation languages for `#lang`.

Finally, `#!` is an alias for `#lang` followed by a space when `#!` is followed by alphanumeric ASCII, `+`, `-`, or `_`. Use of this alias is discouraged except as needed to construct programs that conform to certain grammars, such as that of R^{6}RS [Sperber07].

The `syntax/module-reader` library provides a domain-specific language for writing language readers.

By convention, `#lang` normally appears at the beginning of a file, possibly after comment forms, to specify the syntax of a module.

If the `read-accept-reader` or `read-accept-lang` parameter is set to `#f`, then if the reader encounters `#lang` or equivalent `#!`, the `exn:fail:read` exception is raised.

S-Expression Reader Language

```
#lang s-exp
```

§17.1.2 "Using `#lang s-exp`" in ***The Racket Guide*** introduces the `s-exp` meta-language.

The `s-exp` "language" is a kind of meta-language. It `read`s the S-expression that follows `#lang s-exp` and uses it as the language of a `module` form. It also reads all remaining S-expressions until an end-of-file, using them for the body of the generated `module`.

That is,

```
#lang s-exp module-path
form ...
```

is equivalent to

```
(module name-id module-path
  form ...)
```

where *name-id* is derived from the source input port's name: if the port name is a filename path, the filename without its directory path and extension is used for *name-id*, otherwise *name-id* is `anonymous-module`.

Chaining Reader Language

```
#lang reader
```

§17.3.2 "Using `#lang reader`" in ***The Racket Guide*** introduces the `reader` meta-language.

The `reader` "language" is a kind of meta-language. It `read`s the S-expression that follows `#lang reader` and uses it as a module path (relative to the module being read) that effectively takes the place of `reader`. In other words, the `reader` meta-language generalizes the syntax of the module specified after `#lang` to be a module path, and without the implicit addition of `/lang/reader` to the path.

1.4 The Printer

The Racket printer supports three modes:

- `write` mode prints core datatypes in such a way that using `read` on the output produces a value that is `equal?` to the printed value;
- `display` mode prints core datatypes in a more "end-user" style rather than "programmer" style; for example, a string `display`s as its content characters without surrounding `"`s or escapes;
- `print` mode by default—when `print-as-expression` is `#t`—prints most datatypes in such a way that evaluating the output as an expression produces a value that is `equal?` to the printed value; when `print-as-expression` is set to `#f`, then `print` mode is like `write` mode.

In `print` mode when `print-as-expression` is `#t` (as is the default), a value prints at a *quoting depth* of either `0` (unquoted) or `1` (quoted). The initial quoting depth is accepted as an optional argument by `print`, and printing of some compound datatypes adjusts the print depth for component values. For example, when a list is printed at quoting depth `0` and all of its elements are *quotable*, the list is printed with a `'` prefix, and the list's elements are printed at quoting depth `1`.

When the print-graph parameter is set to #t, then the printer first scans an object to detect cycles. The scan traverses the components of pairs, mutable pairs, vectors, boxes (when print-box is #t), hash tables (when print-hash-table is #t), fields of structures exposed by struct->vector (when print-struct is #t), and fields of structures exposed by printing when the structure's type has the prop:custom-write property. If print-graph is #t, then this information is used to print sharing through graph definitions and references (see §1.3.16 "Reading Graph Structure"). If a cycle is detected in the initial scan, then print-graph is effectively set to #t automatically.

With the exception of displaying byte strings, printing is defined in terms of Unicode characters; see §12.1 "Ports" for information on how a character stream is written to a port's underlying byte stream.

1.4.1 Printing Symbols

Symbols containing spaces or special characters write using escaping \ and quoting |s. When the read-case-sensitive parameter is set to #f, then symbols containing uppercase characters also use escaping \ and quoting |s. In addition, symbols are quoted with |s or leading \ when they would otherwise print the same as a numerical constant or as a delimited . (when read-accept-dot is #t).

When read-accept-bar-quote is #t, |s are used in printing when one | at the beginning and one | at the end suffice to correctly print the symbol. Otherwise, \s are always used to escape special characters, instead of quoting them with |s.

When read-accept-bar-quote is #f, then | is not treated as a special character. The following are always special characters:

```
( ) [ ] { } " , ' ` ; \
```

In addition, # is a special character when it appears at the beginning of the symbol, and when it is not followed by %.

Symbols display without escaping or quoting special characters. That is, the display form of a symbol is the same as the display form of symbol->string applied to the symbol.

Symbols print the same as they write, unless print-as-expression is set to #t (as is the default) and the current quoting depth is 0. In that case, the symbol's printed form is prefixed with '. For the purposes of printing enclosing datatypes, a symbol is quotable.

1.4.2 Printing Numbers

A number prints the same way in write, display, and print modes. For the purposes of printing enclosing datatypes, a symbol is quotable.

A complex number that is not a real number always prints as ⟨*m*⟩+⟨*n*⟩i, where ⟨*m*⟩ and ⟨*n*⟩ are the printed forms of its real and imaginary parts, respectively.

An inexact real number prints with either a . decimal point, an e exponent marker, or both. The form is selected so that the output is as short as possible, with the constraint that reading the printed form back in produces an equal? number.

An exact 0 prints as 0.

A positive, exact integer prints as a sequence of decimal digits that does not start with 0.

A positive, exact, real, non-integer number prints as ⟨*m*⟩/⟨*n*⟩, where ⟨*m*⟩ and ⟨*n*⟩ are the printed forms of the number's numerator and denominator (as determined by numerator and denominator).

A negative exact number prints with a - prefix on the printed form of the number's exact negation.

1.4.3 Printing Booleans

The boolean constant #t prints as #true or #t in all modes (display, write, and print), depending on the value of print-boolean-long-form, and the constant #f prints as #false or #f. For the purposes of printing enclosing datatypes, a symbol is quotable.

1.4.4 Printing Pairs and Lists

In write and display modes, an empty list prints as (). A pair normally prints starting with (followed by the printed form of its car. The rest of the printed form depends on the cdr:

- If the cdr is a pair or the empty list, then the printed form of the pair completes with the printed form of the cdr, except that the leading (in the cdr's printed form is omitted.
- Otherwise, the printed for of the pair continues with a space, ., another space, the printed form of the cdr, and a).

If print-reader-abbreviations is set to #t, then pair printing in write mode is adjusted in the case of a pair that starts a two-element list whose first element is 'quote, 'quasiquote, 'unquote, 'unquote-splicing, 'syntax, 'quasisyntax, 'unsyntax, or 'unsyntax-splicing. In that case, the pair is printed with the corresponding reader syntax: ', `, ,, ,@, #', #`, #,, or #,@, respectively. After the reader syntax, the second element of the list is printed. When the list is a tail of an enclosing list, the tail is printed after

a . in the enclosing list (after which the reader abbreviations work), instead of including the tail as two elements of the enclosing list.

The printed form of a pair is the same in both write and display modes, except as the printed form of the pair's car and cdr vary with the mode. The print form is also the same if print-as-expression is #f or the quoting depth is 1.

For print mode when print-as-expression is #t and the quoting depth is 0, then the empty list prints as '(). For a pair whose car and cdr are quotable, the pair prints in write mode but with a ' prefix; the pair's content is printed with quoting depth 1. Otherwise, when the car or cdr is not quotable, then pair prints with either cons (when the cdr is not a pair), list (when the pair is a list), or list* (otherwise) after the opening (, any . that would otherwise be printed is suppressed, and the pair content is printed at quoting depth 0. In all cases, when print-as-expression is #t for print mode, then the value of print-reader-abbreviations is ignored and reader abbreviations are always used for lists printed at quoting depth 1.

By default, mutable pairs (as created with mcons) print the same as pairs for write and display, except that { and } are used instead of (and). Note that the reader treats {...} and (...) equivalently on input, creating immutable pairs in both cases. Mutable pairs in print mode with print-as-expression as #f or a quoting depth of 1 also use { and }. In print mode with print-as-expression as #t and a quoting depth of 0, a mutable pair prints as (mcons , the mcar and mcdr printed at quoting depth 0 and separated by a space, and a closing).

If the print-pair-curly-braces parameter is set to #t, then pairs print using { and } when not using print mode with print-as-expression as #t and a quoting depth of 0. If the print-mpair-curly-braces parameter is set to #f, then mutable pairs print using (and) in that mode.

For the purposes of printing enclosing datatypes, an empty list is always quotable, a pair is quotable when its car and cdr are quotable, and a mutable list is never quotable.

1.4.5 Printing Strings

All strings display as their literal character sequences.

The write or print form of a string starts with " and ends with another ". Between the "s, each character is represented. Each graphic or blank character is represented as itself, with two exceptions: " is printed as \", and \ is printed as \\. Each non-graphic, non-blank character (according to char-graphic? and char-blank?) is printed using the escape sequences described in §1.3.6 "Reading Strings", using \a, \b, \t, \n, \v, \f, \r, or \e if possible, otherwise using \u with four hexadecimal digits or \U with eight hexadecimal digits (using the latter only if the character value does not fit into four digits).

All byte strings `display` as their literal byte sequence; this byte sequence may not be a valid UTF-8 encoding, so it may not correspond to a sequence of characters.

The `write` or `print` form of a byte string starts with `#"` and ends with a `"`. Between the `"`s, each byte is written using the corresponding ASCII decoding if the byte is between 0 and 127 and the character is graphic or blank (according to `char-graphic?` and `char-blank?`). Otherwise, the byte is written using `\a`, `\b`, `\t`, `\n`, `\v`, `\f`, `\r`, or `\e` if possible, otherwise using `\` followed by one to three octal digits (only as many as necessary).

For the purposes of printing enclosing datatypes, a string or a byte string is quotable.

1.4.6 Printing Vectors

In `display` mode, the printed form of a vector is `#` followed by the printed form of `vector->list` applied to the vector. In `write` mode, the printed form is the same, except that when the `print-vector-length` parameter is `#t`, a decimal integer is printed after the `#`, and a repeated last element is printed only once.

Vectors `print` the same as they `write`, unless `print-as-expression` is set to `#t` and the current quoting depth is `0`. In that case, if all of the vector's elements are quotable, then the vector's `print`ed form is prefixed with `'` and its elements printed with quoting depth `1`. If its elements are not all quotable, then the vector `print`s as `(vector` , the elements at quoting depth `0`, and a closing `)`. A vector is quotable when all of its elements are quotable.

1.4.7 Printing Structures

When the `print-struct` parameter is set to `#t`, then the way that structures print depends on details of the structure type for which the structure is an instance:

- If the structure type is a prefab structure type, then it prints in `write` or `display` mode using `#s(` followed by the prefab structure type key, then the printed form of each field in the structure, and then `)`.

  In `print` mode when `print-as-expression` is set to `#t` and the current quoting depth is `0`, if the structure's content is all quotable, then the structure's `print`ed form is prefixed with `'` and its content is printed with quoting depth `1`. If any of its content is not quotable, then the structure type prints the same as a non-prefab structure type.

  An instance of a prefab structure type is quotable when all of its content is quotable.

- If the structure has a `prop:custom-write` property value, then the associated procedure is used to print the structure, unless the `print-unreadable` parameter is set to `#f`.

  For `print` mode, an instance of a structure type with a `prop:custom-write` property is treated as quotable if it has the `prop:custom-print-quotable` property with a

value of 'always. If it has 'maybe as the property value, then the structure is treated as quotable if its content is quotable, where the content is determined by the values recursively printed by the structure's prop:custom-write procedure. Finally, if the structure has 'self as the property value, then it is treated as quotable.

In print mode when print-as-expression is #t, the structure's prop:custom-write procedure is called with either 0 or 1 as the quoting depth, normally depending on the structure's prop:custom-print-quotable property value. If the property value is 'always, the quoting depth is normally 1. If the property value is 'maybe, then the quoting depth is 1 if the structure is quotable, or normally 0 otherwise. If the property value is 'self, then the quoting depth may be 0 or 1; it is normally 0 if the structure is not printed as a part of an enclosing quotable value, even though the structure is treated as quotable. Finally, if the property value is 'never, then the quoting depth is normally 0. The quoting depth can vary from its normal value if the structure is printed with an explicit quoting depth of 1.

- If the structure's type is transparent or if any ancestor is transparent (i.e., struct? on the instance produces #t), then the structure prints as the vector produced by struct->vector in display mode, in write mode, or in print mode when print-as-expression is set to #f or when the quoting depth is 0.

  In print mode with print-as-expression as #t and a quoting depth of 0, the structure content is printed with a (followed by the structure's type name (as determined by object-name) in write mode; the remaining elements are printed at quoting depth 0 and separated by a space, and finally a closing).

  A transparent structure type that is not a prefab structure type is never quotable.

- For any other structure type, the structure prints as an unreadable value; see §1.4.14 "Printing Unreadable Values" for more information.

If the print-struct parameter is set to #f, then all structures without a prop:custom-write property print as unreadable values (see §1.4.14 "Printing Unreadable Values") and count as quotable.

1.4.8 Printing Hash Tables

When the print-hash-table parameter is set to #t, in write and display modes, a hash table prints starting with #hash(, #hasheqv(, or #hasheq(for a table using equal?, eqv?, or eq? key comparisons, respectively. After the prefix, each key–value mapping is shown as (, the printed form of a key, a space, ., a space, the printed form the corresponding value, and), with an additional space if the key–value pair is not the last to be printed. After all key–value pairs, the printed form completes with).

In print mode when print-as-expression is #f or the quoting depth is 1, the printed form is the same as for write. Otherwise, if the hash table's keys and values are all quotable, the table prints with a ' prefix, and the table's key and values are printed at quoting depth

1. If some key or value is not quotable, the hash table prints as `(hash` , `(hasheqv` , or `(hasheq` followed by alternating keys and values printed at quoting depth 1 and separated by spaces, and finally a closing `)`. A hash table is quotable when all of its keys and values are quotable.

When the `print-hash-table` parameter is set to `#f`, a hash table prints as `#<hash>` and counts as quotable.

1.4.9 Printing Boxes

When the `print-box` parameter is set to `#t`, a box prints as `#&` followed by the printed form of its content in `write`, `display`, or `print` mode when `print-as-expression` is `#f` or the quoting depth is 1.

In `print` mode when `print-as-expression` is `#t` and the quoting depth is 0, a box prints with a `'` prefix and its value is printed at quoting depth 1 when its content is quotable, otherwise the box prints a `(box` followed by the content at quoting depth 0 and a closing `)`. A box is quotable when its content is quotable.

When the `print-box` parameter is set to `#f`, a box prints as `#<box>` and counts as quotable.

1.4.10 Printing Characters

Characters with the special names described in §1.3.13 “Reading Characters” `write` and `print` using the same name. (Some characters have multiple names; the `#\newline` and `#\nul` names are used instead of `#\linefeed` and `#\null`.) Other graphic characters (according to `char-graphic?`) `write` as `#\` followed by the single character, and all others characters are written in `#\u` notation with four digits or `#\U` notation with eight digits (using the latter only if the character value does not fit in four digits).

All characters `display` directly as themselves (i.e., a single character).

For the purposes of printing enclosing datatypes, a character is quotable.

1.4.11 Printing Keywords

Keywords `write`, `print`, and `display` the same as symbols (see §1.4.1 “Printing Symbols”) except with a leading `#:` (after any `'` prefix added in `print` mode), and without special handling for an initial `#` or when the printed form would match a number or a delimited `.` (since `#:` distinguishes the keyword).

For the purposes of printing enclosing datatypes, a keyword is quotable.

1.4.12 Printing Regular Expressions

Regexp values `write`, `display`, and `print` starting with `#px` (for `pregexp`-based regexps) or `#rx` (for `regexp`-based regexps) followed by the `write` form of the regexp's source string or byte string.

For the purposes of printing enclosing datatypes, a regexp value is quotable.

1.4.13 Printing Paths

Paths `write` and `print` as `#<path:....>`. A path `display`s the same as the string produced by `path->string`. For the purposes of printing enclosing datatypes, a path counts as quotable.

Although a path can be converted to a string with `path->string` or to a byte string with `path->bytes`, neither is clearly the right choice for printing a path and reading it back. If the path value is meant to be moved among platforms, then a string is probably the right choice, despite the potential for losing information when converting a path to a string. For a path that is intended to be re-read on the same platform, a byte string is probably the right choice, since it preserves information in an unportable way. Paths do not print in a readable way so that programmers are not misled into thinking that either choice is always appropriate.

1.4.14 Printing Unreadable Values

For any value with no other printing specification, assuming that the `print-unreadable` parameter is set to `#t`, the output form is `#<`⟨*something*⟩`>`, where ⟨*something*⟩ is specific to the type of the value and sometimes to the value itself. If `print-unreadable` is set to `#f`, then attempting to print an unreadable value raises `exn:fail`.

For the purposes of printing enclosing datatypes, a value that prints unreadably nevertheless counts as quotable.

1.4.15 Printing Compiled Code

Compiled code as produced by `compile` prints using `#~`. Compiled code printed with `#~` is essentially assembly code for Racket, and reading such a form produces a compiled form when the `read-accept-compiled` parameter is set to `#t`.

When a compiled form contains syntax object constants, they must not be tainted or armed; the `#~`-marshaled form drops source-location information and properties (see §11.7 "Syntax Object Properties") for the syntax objects.

Compiled code parsed from `#~` may contain references to unexported or protected bindings from a module. At read time, such references are associated with the current code inspector (see `current-code-inspector`), and the code will only execute if that inspector controls the relevant module invocation (see §13.10 “Code Inspectors”).

A compiled-form object may contain uninterned symbols (see §3.6 “Symbols”) that were created by `gensym` or `string->uninterned-symbol`. When the compiled object is read via `#~`, each uninterned symbol in the original form is mapped to a new uninterned symbol, where multiple instances of a single symbol are consistently mapped to the same new symbol. The original and new symbols have the same printed representation. Unreadable symbols, which are typically generated indirectly during expansion and compilation, are saved and restored consistently through `#~`.

The dynamic nature of uninterned symbols and their localization within `#~` can cause problems when `gensym` or `string->uninterned-symbol` is used to construct an identifier for a top-level or module binding (depending on how the identifier and its references are compiled). To avoid problems, generate distinct identifiers either with `generate-temporaries` or by applying the result of `make-syntax-introducer` to an existing identifier; those functions lead to top-level and module variables with unreadable symbolic names, and the names are deterministic as long as expansion is otherwise deterministic.

Despite the problems inherent with uninterned symbols as variable names, they are partially supported even across multiple `#~`s: When compiled code contains a reference to a module-defined variable whose name is an uninterned symbol, the relative position of the variable among the module’s definitions is recorded, and the reference can be linked back to the definition based on its position and the characters in its name. This accommodation works only for variable references in compiled code; it does not work for `syntax`-quoted identifiers, for example.

Finally, a compiled form may contain path literals. Although paths are not normally printed in a way that can be read back in, path literals can be written and read as part of compiled code. The `current-write-relative-directory` parameter is used to convert the path to a relative path as is it written, and then `current-load-relative-directory` parameter is used to convert any relative path back as it is read. The relative-path conversion applies on reading whether the path was originally relative or not.

2 Syntactic Forms

This section describes the core syntax forms that appear in a fully expanded expression, plus many closely related non-core forms. See §1.2.3.1 “Fully Expanded Programs” for the core grammar.

Notation

Each syntactic form is described by a BNF-like notation that describes a combination of (syntax-wrapped) pairs, symbols, and other data (not a sequence of characters). These grammatical specifications are shown as in the following specification of a `something` form:

```
(something id thing-expr ...)

  thing-expr : number?
```

Within such specifications,

- `...` indicates zero or more repetitions of the preceding datum; more generally, *N* consecutive `...`s a row indicate a consecutive repetition of the preceding *N* datums.
- `...+` indicates one or more repetitions of the preceding datum.
- Italic meta-identifiers play the role of non-terminals. Some meta-identifier names imply syntactic constraints:
  - A meta-identifier that ends in *id* stands for an identifier.
  - A meta-identifier that ends in *keyword* stands for a keyword.
  - A meta-identifier that ends with *expr* (such as *thing-expr*) stands for a sub-form that is expanded as an expression.
  - A meta-identifier that ends with *body* stands for a sub-form that is expanded in an internal-definition context (see §1.2.3.7 “Internal Definitions”).
- Contracts indicate constraints on sub-expression results. For example, *thing-expr* : `number?` indicates that the expression *thing-expr* must produce a number.

2.1 Modules: `module`, `module*`, ...

```
(module id module-path form ...)
```

§6.2.1 “The `module` Form” in ***The Racket Guide*** introduces `module`.

Declares a top-level module or a submodule. For a top-level module, if the `current-module-declare-name` parameter is set, the parameter value is used for the module name

and *id* is ignored, otherwise (quote *id*) is the name of the declared module. For a submodule, *id* is the name of the submodule to be used as an element within a submod module path.

For a module-like form that works in definitions context other than the top level or a module body, see define-package.

The *module-path* form must be as for require, and it supplies the initial bindings for the body *form*s. That is, it is treated like a (require *module-path*) prefix before the *form*s, except that the bindings introduced by *module-path* can be shadowed by definitions and requires in the module body *form*s.

If a single *form* is provided, then it is partially expanded in a module-begin context. If the expansion leads to #%plain-module-begin, then the body of the #%plain-module-begin is the body of the module. If partial expansion leads to any other primitive form, then the form is wrapped with #%module-begin using the lexical context of the module body; this identifier must be bound by the initial *module-path* import, and its expansion must produce a #%plain-module-begin to supply the module body. Finally, if multiple *form*s are provided, they are wrapped with #%module-begin, as in the case where a single *form* does not expand to #%plain-module-begin.

After such wrapping, if any, and before any expansion, an 'enclosing-module-name property is attached to the #%module-begin syntax object (see §11.7 "Syntax Object Properties"); the property's value is a symbol corresponding to *id*.

Each *form* is partially expanded (see §1.2.3.6 "Partial Expansion") in a module context. Further action depends on the shape of the form:

- If it is a begin form, the sub-forms are flattened out into the module's body and immediately processed in place of the begin.
- If it is a define-syntaxes form, then the right-hand side is evaluated (in phase 1), and the binding is immediately installed for further partial expansion within the module. Evaluation of the right-hand side is parameterized to set current-namespace as in let-syntax.
- If it is a begin-for-syntax form, then the body is expanded (in phase 1) and evaluated. Expansion within a begin-for-syntax form proceeds with the same partial-expansion process as for a module body, but in a higher phase, and saving all #%provide forms for all phases until the end of the module's expansion. Evaluation of the body is parameterized to set current-namespace as in let-syntax.
- If the form is a #%require form, bindings are introduced immediately, and the imported modules are instantiated or visited as appropriate.
- If the form is a #%provide form, then it is recorded for processing after the rest of the body.
- If the form is a define-values form, then the binding is installed immediately, but the right-hand expression is not expanded further.

- If the form is a `module` form, then it is immediately expanded and declared for the extent of the current top-level enclosing module's expansion.
- If the form is a `module*` form, then it is not expanded further.
- Similarly, if the form is an expression, it is not expanded further.

After all *form*s have been partially expanded this way, then the remaining expression forms (including those on the right-hand side of a definition) are expanded in an expression context. After all expression forms, `#%provide` forms are processed in the order in which they appear (independent of phase) in the expanded module. Finally, all `module*` forms are expanded in order, so that each becomes available for use by subsequent `module*` forms; the enclosing module itself is also available for use by `module*` submodules.

The scope of all imported identifiers covers the entire module body, except for nested `module` and `module*` forms (assuming a non-`#f` *module-path* in the latter case). The scope of any identifier defined within the module body similarly covers the entire module body except for such nested `module` and `module*` forms. The ordering of syntax definitions does not affect the scope of the syntax names; a transformer for `A` can produce expressions containing `B`, while the transformer for `B` produces expressions containing `A`, regardless of the order of declarations for `A` and `B`. However, a syntactic form that produces syntax definitions must be defined before it is used.

No identifier can be imported or defined more than once at any phase level within a single module. Every exported identifier must be imported or defined. No expression can refer to a top-level variable. A `module*` form in which the enclosing module's bindings are visible (i.e., a nested `module*` with `#f` instead of a *module-path*) can define or import bindings that shadow the enclosing module's bindings.

The evaluation of a `module` form does not evaluate the expressions in the body of the module. Evaluation merely declares a module, whose full name depends both on *id* or `(current-module-declare-name)`.

A module body is executed only when the module is explicitly instantiated via `require` or `dynamic-require`. On invocation, imported modules are instantiated in the order in which they are `required` into the module (although earlier instantiations or transitive `requires` can trigger the instantiation of a module before its order within a given module). Then, expressions and definitions are evaluated in order as they appear within the module. Each evaluation of an expression or definition is wrapped with a continuation prompt (see `call-with-continuation-prompt`) for the default continuation and using the default prompt handler.

Accessing a module-level variable before it is defined signals a run-time error, just like accessing an undefined global variable. If a module (in its fully expanded form) does not contain a `set!` for an identifier that defined within the module, then the identifier is a *constant* after it is defined; its value cannot be changed afterward, not even through reflective mechanisms. The `compile-enforce-module-constants` parameter, however, can be used to

disable enforcement of constants.

When a syntax object representing a `module` form has a `'module-language` syntax property attached, and when the property value is a vector of three elements where the first is a module path (in the sense of `module-path?`) and the second is a symbol, then the property value is preserved in the corresponding compiled and/or declared module. The third component of the vector should be printable and `read`able, so that it can be preserved in marshaled bytecode. The `racket/base` and `racket` languages attach `'#(racket/language-info get-info #f)` to a `module` form. See also `module-compiled-language-info`, `module->language-info`, and `racket/language-info`.

If a `module` form has a single body *`form`* and if the form is a `#%plain-module-begin` form, then the body *`form`* is traversed to find `module` and `module*` forms that are either immediate, under `begin`, or under `begin-for-syntax`. (That is, the body is searched before adding any lexical context due to the module's initial *`module-path`* import.) Each such module form is given a `'submodule` syntax property that whose value is the initial module form. Then, when `module` or `module*` is expanded in a submodule position, if the form has a `'submodule` syntax property, the property value is used as the form to expand. This protocol avoids the contamination of submodule lexical scope when re-expanding `module` forms that contain submodules.

See also §1.1.10 "Modules and Module-Level Variables" and §1.2.3.8 "Module Phases and Visits".

Example:

```
> (module duck racket/base
    (provide num-eggs quack)
    (define num-eggs 2)
    (define (quack n)
      (unless (zero? n)
        (printf "quack\n")
        (quack (sub1 n)))))
```

```
(module* id module-path form ...)
(module* id #f form ...)
```

§6.2.3 "Submodules" in ***The Racket Guide*** introduces `module*`.

Like `module`, but only for declaring a submodule within a module, and for submodules that may `require` the enclosing module.

Instead of a *`module-path`* after *`id`*, `#f` indicates that all bindings from the enclosing module are visible in the submodule; `begin-for-syntax` forms that wrap the `module*` form shift the phase level of the enclosing module's bindings relative to the submodule. When a `module*` form has a *`module-path`*, the submodule starts with an empty lexical context in the same way as a top-level `module` form, and enclosing `begin-for-syntax` forms have no effect on the submodule.

```
(module+ id form ...)
```

§6.2.4 “Main and Test Submodules” in ***The Racket Guide*** introduces `module+`.

Declares and/or adds to a submodule named *id*.

Each addition for *id* is combined in order to form the entire submodule using `(module* id #f ....)` at the end of the enclosing module. If there is only one `module+` for a given *id*, then `(module+ id form ...)` is equivalent to `(module* id #f form ...)`.

A submodule must not be defined using `module+` *and* `module` or `module*`. That is, if a submodule is made of `module+` pieces, then it must be made *only* of `module+` pieces.

```
(#%module-begin form ...)
```

Legal only in a module begin context, and handled by the `module` and `module*` forms.

The `#%module-begin` form of `racket/base` wraps every top-level expression to print non-`#<void>` results using `current-print`.

```
(#%plain-module-begin form ...)
```

Legal only in a module begin context, and handled by the `module` and `module*` forms.

2.2 Importing and Exporting: `require` and `provide`

§6.4 “Imports: `require`” in ***The Racket Guide*** introduces `require`.

```
(require require-spec ...)
```

```
      require-spec = module-path
                   | (only-in require-spec id-maybe-renamed ...)
                   | (except-in require-spec id ...)
                   | (prefix-in prefix-id require-spec)
                   | (rename-in require-spec [orig-id bind-id] ...)
                   | (combine-in require-spec ...)
                   | (relative-in module-path require-spec ...)
                   | (only-meta-in phase-level require-spec ...)
                   | (for-syntax require-spec ...)
                   | (for-template require-spec ...)
                   | (for-label require-spec ...)
                   | (for-meta phase-level require-spec ...)
                   | derived-require-spec

       module-path = root-module-path
                   | (submod root-module-path submod-path-element ...)
                   | (submod "." submod-path-element ...)
                   | (submod ".." submod-path-element ...)

  root-module-path = (quote id)
                   | rel-string
                   | (lib rel-string ...+)
                   | id
                   | (file string)
                   | (planet id)
                   | (planet string)
                   | (planet rel-string
                             (user-string pkg-string vers)
                             rel-string ...)

submod-path-element = id
                    | ".."

  id-maybe-renamed = id
                   | [orig-id bind-id]

       phase-level = exact-integer
                   | #f

              vers =
                   | nat
                   | nat minor-vers

        minor-vers = nat
                   | (nat nat)
                   | (= nat)
                   | (+ nat)
                   | (- nat)
```

In a top-level context, `require` instantiates modules (see §1.1.10 “Modules and Module-Level Variables”). In a top-level context or module context, expansion of `require` visits modules (see §1.2.3.8 “Module Phases and Visits”). In both contexts and both evaluation and expansion, `require` introduces bindings into a namespace or a module (see §1.2.3.4 “Introducing Bindings”). A `require` form in a expression context or internal-definition context is a syntax error.

A *require-spec* designates a particular set of identifiers to be bound in the importing context. Each identifier is mapped to a particular export of a particular module; the identifier to bind may be different from the symbolic name of the originally exported identifier. Each identifier also binds at a particular phase level.

No identifier can be bound multiple times in a given phase level by an import, unless all of the bindings refer to the same original definition in the same module. In a module context, an identifier can be either imported or defined for a given phase level, but not both.

The syntax of *require-spec* can be extended via `define-require-syntax`, and when multiple *require-spec*s are specified in a `require`, the bindings of each *require-spec* are visible for expanding later *require-spec*s. The pre-defined forms (as exported by `racket/base`) are as follows:

module-path

Imports all exported bindings from the named module, using the export identifiers as the local identifiers. (See below for information on *module-path*.) The lexical context of the *module-path* form determines the context of the introduced identifiers.

(only-in *require-spec id-maybe-renamed* ...)

Like *require-spec*, but constrained to those exports for which the identifiers to bind match *id-maybe-renamed*: as *id* or as *orig-id* in [*orig-id bind-id*]. If the *id* or *orig-id* of any *id-maybe-renamed* is not in the set that *require-spec* describes, a syntax error is reported.

Examples:

```
> (require (only-in racket/tcp
                    tcp-listen
                    [tcp-accept my-accept]))

> tcp-listen
#<procedure:tcp-listen>
> my-accept
#<procedure:tcp-accept>
```

```
> tcp-accept
tcp-accept: undefined;
 cannot reference undefined identifier
```

(except-in *require-spec id ...*)

Like *require-spec*, but omitting those imports for which *id*s are the identifiers to bind; if any *id* is not in the set that *require-spec* describes, a syntax error is reported.

Examples:

```
> (require (except-in racket/tcp
                      tcp-listen))

> tcp-accept
#<procedure:tcp-accept>
> tcp-listen
tcp-listen: undefined;
 cannot reference undefined identifier
```

(prefix-in *prefix-id require-spec*)

Like *require-spec*, but adjusting each identifier to be bound by prefixing it with *prefix-id*. The lexical context of the *prefix-id* is ignored, and instead preserved from the identifiers before prefixing.

Examples:

```
> (require (prefix-in tcp: racket/tcp))

> tcp:tcp-accept
#<procedure:tcp-accept>
> tcp:tcp-listen
#<procedure:tcp-listen>
```

(rename-in *require-spec* [*orig-id bind-id*] ...)

Like *require-spec*, but replacing the identifier to bind *orig-id* with *bind-id*; if any *orig-id* is not in the set that *require-spec* describes, a syntax error is reported.

Examples:

```
> (require (rename-in racket/tcp
                      (tcp-accept accept)
                      (tcp-listen listen)))

> accept
#<procedure:tcp-accept>
> listen
#<procedure:tcp-listen>
```

(combine-in *require-spec* ...)

The union of the *require-spec*s.

Examples:

```
> (require (combine-in (only-in racket/tcp tcp-accept)
                       (only-in racket/tcp tcp-listen)))

> tcp-accept
#<procedure:tcp-accept>
> tcp-listen
#<procedure:tcp-listen>
```

(relative-in *module-path require-spec* ...)

Like the union of the *require-spec*s, but each relative module path in a *require-spec* is treated as relative to *module-path* instead of the enclosing context.

The require transformer that implements relative-in sets current-require-module-path to adjust module paths in the *require-spec*s.

(only-meta-in *phase-level require-spec* ...)

Like the combination of *require-spec*s, but removing any binding that is not for *phase-level*, where #f for *phase-level* corresponds to the label phase level.

The following example imports bindings only at phase level 1, the transform phase:

```
> (module nest racket
    (provide (for-syntax meta-eggs)
             (for-meta 1 meta-chicks)
             num-eggs)
    (define-for-syntax meta-eggs 2)
    (define-for-syntax meta-chicks 3)
    (define num-eggs 2))
```

```
> (require (only-meta-in 1 'nest))

> (define-syntax (desc stx)
    (printf "~s ~s\n" meta-eggs meta-chicks)
    #'(void))

> (desc)
2 3

> num-eggs
num-eggs: undefined;
 cannot reference undefined identifier
```

The following example imports only bindings at phase level 0, the normal phase.

```
> (require (only-meta-in 0 'nest))

> num-eggs
2
```

(for-meta *phase-level require-spec ...*)

Like the combination of *require-spec*s, but the binding specified by each *require-spec* is shifted by *phase-level*. The label phase level corresponds to `#f`, and a shifting combination that involves `#f` produces `#f`.

Examples:

```
> (module nest racket
    (provide num-eggs)
    (define num-eggs 2))

> (require (for-meta 0 'nest))

> num-eggs
2
> (require (for-meta 1 'nest))

> (define-syntax (roost stx)
    (datum->syntax stx num-eggs))

> (roost)
2
```

(for-syntax *require-spec* ...)

Same as (for-meta 1 *require-spec* ...).

(for-template *require-spec* ...)

Same as (for-meta -1 *require-spec* ...).

(for-label *require-spec* ...)

Same as (for-meta #f *require-spec* ...).

derived-require-spec

See define-require-syntax for information on expanding the set of *require-spec* forms.

§6.3 "Module Paths" in ***The Racket Guide*** introduces module paths.

A *module-path* identifies a module, either a root module or a submodule that is declared lexically within another module. A root module is identified either through a concrete name in the form of an identifier, or through an indirect name that can trigger automatic loading of the module declaration. Except for the (quote *id*) case below, the actual resolution of a root module path is up to the current module name resolver (see current-module-name-resolver), and the description below corresponds to the default module name resolver.

(quote *id*)

Refers to a submodule previously declared with the name *id* or a module previously declared interactively with the name *id*. When *id* refers to a submodule, (quote *id*) is equivalent to (submod "." *id*).

Examples:

```
; a module declared interactively as test:
> (require 'test)
```

rel-string

A path relative to the containing source (as determined by current-load-relative-directory or current-directory). Regardless of the current

platform, *rel-string* is always parsed as a Unix-format relative path: / is the path delimiter (multiple adjacent /s are treated as a single delimiter), .. accesses the parent directory, and . accesses the current directory. The path cannot be empty or contain a leading or trailing slash, path elements before than the last one cannot include a file suffix (i.e., a . in an element other than . or ..), and the only allowed characters are ASCII letters, ASCII digits, -, +, _, ., /, and %. Furthermore, a % is allowed only when followed by two lowercase hexadecimal digits, and the digits must form a number that is not the ASCII value of a letter, digit, -, +, or _.

The % provision is intended to support a one-to-one encoding of arbitrary strings as path elements (after UTF-8 encoding). Such encodings are not decoded to arrive at a filename, but instead preserved in the file access.

If *rel-string* ends with a ".ss" suffix, it is converted to a ".rkt" suffix. The compiled-load handler may reverse that conversion if a ".rkt" file does not exist and a ".ss" exists.

Examples:

```
; a module named "x.rkt" in the same
; directory as the enclosing module's file:
> (require "x.rkt")

; a module named "x.rkt" in the parent directory
; of the enclosing module file's directory:
> (require "../x.rkt")
```

(lib *rel-string* ...+)

A path to a module installed into a collection (see §17.2 "Libraries and Collections"). The *rel-string*s in lib are constrained similar to the plain *rel-string* case, with the additional constraint that a *rel-string* cannot contain . or .. directory indicators.

The specific interpretation of the path depends on the number and shape of the *rel-string*s:

- If a single *rel-string* is provided, and if it consists of a single element (i.e., no /) with no file suffix (i.e., no .), then *rel-string* names a collection, and "main.rkt" is the library file name.

  Examples:

  ```
  ; the main swindle library:
  > (require (lib "swindle"))

  ; the same:
  > (require (lib "swindle/main.rkt"))
  ```

- If a single *rel-string* is provided, and if it consists of multiple /-separated elements, then each element up to the last names a collection, subcollection, etc., and the last element names a file. If the last element has no file suffix, ".rkt" is added, while a ".ss" suffix is converted to ".rkt".

  Examples:

  ```
  ; "turbo.rkt" from the "swindle" collection:
  > (require (lib "swindle/turbo"))

  ; the same:
  > (require (lib "swindle/turbo.rkt"))

  ; the same:
  > (require (lib "swindle/turbo.ss"))
  ```

- If a single *rel-string* is provided, and if it consists of a single element *with* a file suffix (i.e, with a .), then *rel-string* names a file within the "mzlib" collection. A ".ss" suffix is converted to ".rkt". (This convention is for compatibility with older version of Racket.)

  Examples:

  ```
  ; "tar.rkt" module from the "mzlib" collection:
  > (require (lib "tar.ss"))
  ```

- Otherwise, when multiple *rel-string*s are provided, the first *rel-string* is effectively moved after the others, and all *rel-string*s are appended with / separators. The resulting path names a collection, then subcollection, etc., ending with a file name. No suffix is added automatically, but a ".ss" suffix is converted to ".rkt". (This convention is for compatibility with older version of Racket.)

  Examples:

  ```
  ; "tar.rkt" module from the "mzlib" collection:
  > (require (lib "tar.ss" "mzlib"))
  ```

id

A shorthand for a lib form with a single *rel-string* whose characters are the same as in the symbolic form of *id*. In addition to the constraints of a lib *rel-string*, *id* must not contain ..

Example:

```
> (require racket/tcp)
```

(file *string*)

Similar to the plain *rel-string* case, but *string* is a path—possibly absolute—using the current platform's path conventions and expand-user-path. A ".ss" suffix is converted to ".rkt".

Example:

```
> (require (file "~/tmp/x.rkt"))
```

(planet *id*)
(planet *string*)
(planet *rel-string* (*user-string pkg-string vers*)
rel-string ...)

Specifies a library available via the PLaneT server.

The first form is a shorthand for the last one, where the *id*'s character sequence must match the following ⟨*spec*⟩ grammar:

⟨*spec*⟩ ::= ⟨*owner*⟩ / ⟨*pkg*⟩ ⟨*lib*⟩
⟨*owner*⟩ ::= ⟨*elem*⟩
⟨*pkg*⟩ ::= ⟨*elem*⟩ | ⟨*elem*⟩ : ⟨*version*⟩
⟨*version*⟩ ::= ⟨*int*⟩ | ⟨*int*⟩ : ⟨*minor*⟩
⟨*minor*⟩ ::= ⟨*int*⟩ | <= ⟨*int*⟩ | >= ⟨*int*⟩ | = ⟨*int*⟩
| ⟨*int*⟩ - ⟨*int*⟩
⟨*lib*⟩ ::= ⟨*empty*⟩ | / ⟨*path*⟩
⟨*path*⟩ ::= ⟨*elem*⟩ | ⟨*elem*⟩ / ⟨*path*⟩

and where an ⟨*elem*⟩ is a non-empty sequence of characters that are ASCII letters, ASCII digits, -, +, _, or % followed by lowercase hexadecimal digits (that do not encode one of the other allowed characters), and an ⟨*int*⟩ is a non-empty sequence of ASCII digits. As this shorthand is expended, a ".plt" extension is added to ⟨*pkg*⟩, and a ".rkt" extension is added to ⟨*path*⟩; if no ⟨*path*⟩ is included, "main.rkt" is used in the expansion.

A (planet *string*) form is like a (planet *id*) form with the identifier converted to a string, except that the *string* can optionally end with a file extension (i.e., a .) for a ⟨*path*⟩. A ".ss" file extension is converted to ".rkt".

In the more general last form of a planet module path, the *rel-string*s are similar to the lib form, except that the (*user-string pkg-string vers*) names a PLaneT-based package instead of a collection. A version specification can include an optional major and minor version, where the minor version can

be a specific number or a constraint: `(nat nat)` specifies an inclusive range, `(= nat)` specifies an exact match, `(+ nat)` specifies a minimum version and is equivalent to just *nat*, and `(- nat)` specifies a maximum version. The `=`, `+`, and `-` identifiers in a minor-version constraint are recognized symbolically.

Examples:

```
; "main.rkt" in package "farm" by "mcdonald":
> (require (planet mcdonald/farm))

; "main.rkt" in version >= 2.0 of "farm" by "mcdonald":
> (require (planet mcdonald/farm:2))

; "main.rkt" in version >= 2.5 of "farm" by "mcdonald":
> (require (planet mcdonald/farm:2:5))

; "duck.rkt" in version >= 2.5 of "farm" by "mcdonald":
> (require (planet mcdonald/farm:2:5/duck))
```

```
(submod root-module submod-path-element ...)
(submod "." submod-path-element ...)
(submod ".." submod-path-element ...)
```

Identifies a submodule within the module specified by *root-module* or relative to the current module in the case of `(submod "." ....)`, where `(submod ".." submod-path-element ...)` is equivalent to `(submod "." ".." submod-path-element ...)`. Submodules have symbolic names, and a sequence of identifiers as *submod-path-element*s determine a path of successively nested submodules with the given names. A `".."` as a *submod-path-element* names the enclosing module of a submodule, and it's intended for use in `(submod "." ....)` and `(submod ".." ....)` forms.

```
(local-require require-spec ...)
```

Like `require`, but for use in a internal-definition context to import just into the local context. Only bindings from phase level 0 are imported.

```
(provide provide-spec ...)
```

§6.5 "Exports: `provide`" in ***The Racket Guide*** introduces `provide`.

```
provide-spec = id
             | (all-defined-out)
             | (all-from-out module-path ...)
             | (rename-out [orig-id export-id] ...)
             | (except-out provide-spec provide-spec ...)
             | (prefix-out prefix-id provide-spec)
             | (struct-out id)
             | (combine-out provide-spec ...)
             | (protect-out provide-spec ...)
             | (for-meta phase-level provide-spec ...)
             | (for-syntax provide-spec ...)
             | (for-template provide-spec ...)
             | (for-label provide-spec ...)
             | derived-provide-spec

 phase-level = exact-integer
             | #f
```

Declares exports from a module. A `provide` form must appear in a module context or a module-begin context.

A *provide-spec* indicates one or more bindings to provide. For each exported binding, the external name is a symbol that can be different from the symbolic form of the identifier that is bound within the module. Also, each export is drawn from a particular phase level and exported at the same phase level; by default, the relevant phase level is the number of `begin-for-syntax` forms that enclose the `provide` form.

The syntax of *provide-spec* can be extended by bindings to provide transformers or provide pre-transformers, such as via `define-provide-syntax`, but the pre-defined forms are as follows.

id

Exports *id*, which must be bound within the module (i.e., either defined or imported) at the relevant phase level. The symbolic form of *id* is used as the external name, and the symbolic form of the defined or imported identifier must match (otherwise, the external name could be ambiguous).

Examples:

```
> (module nest racket
    (provide num-eggs)
    (define num-eggs 2))

> (require 'nest)
```

```
> num-eggs
2
```

If *id* has a transformer binding to a rename transformer, then the transformer affects the exported binding. See make-rename-transformer for more information.

(all-defined-out)

Exports all identifiers that are defined at the relevant phase level within the exporting module, and that have the same lexical context as the (all-defined-out) form, excluding bindings to rename transformers where the target identifier has the 'not-provide-all-defined syntax property. The external name for each identifier is the symbolic form of the identifier. Only identifiers accessible from the lexical context of the (all-defined-out) form are included; that is, macro-introduced imports are not re-exported, unless the (all-defined-out) form was introduced at the same time.

Examples:

```
> (module nest racket
    (provide (all-defined-out))
    (define num-eggs 2))

> (require 'nest)

> num-eggs
2
```

(all-from-out *module-path* ...)

Exports all identifiers that are imported into the exporting module using a require-spec built on each *module-path* (see §2.2 “Importing and Exporting: require and provide”) with no phase-level shift. The symbolic name for export is derived from the name that is bound within the module, as opposed to the symbolic name of the export from each *module-path*. Only identifiers accessible from the lexical context of the *module-path* are included; that is, macro-introduced imports are not re-exported, unless the *module-path* was introduced at the same time.

Examples:

```
> (module nest racket
    (provide num-eggs)
    (define num-eggs 2))
```

```
> (module hen-house racket
    (require 'nest)
    (provide (all-from-out 'nest)))

> (require 'hen-house)

> num-eggs
2
```

(rename-out [*orig-id* *export-id*] ...)

Exports each *orig-id*, which must be bound within the module at the relevant phase level. The symbolic name for each export is *export-id* instead orig-d.

Examples:

```
> (module nest racket
    (provide (rename-out [count num-eggs]))
    (define count 2))

> (require 'nest)

> num-eggs
2
> count
count: undefined;
 cannot reference undefined identifier
```

(except-out *provide-spec provide-spec* ...)

Like the first *provide-spec*, but omitting the bindings listed in each subsequent *provide-spec*. If one of the latter bindings is not included in the initial *provide-spec*, a syntax error is reported. The symbolic export name information in the latter *provide-spec*s is ignored; only the bindings are used.

Examples:

```
> (module nest racket
    (provide (except-out (all-defined-out)
                         num-chicks))
    (define num-eggs 2)
    (define num-chicks 3))

> (require 'nest)
```

```
> num-eggs
2
> num-chicks
num-chicks: undefined;
 cannot reference undefined identifier
```

(prefix-out *prefix-id provide-spec*)

Like *provide-spec*, but with each symbolic export name from *provide-spec* prefixed with *prefix-id*.

Examples:

```
> (module nest racket
    (provide (prefix-out chicken: num-eggs))
    (define num-eggs 2))

> (require 'nest)

> chicken:num-eggs
2
```

(struct-out *id*)

Exports the bindings associated with a structure type *id*. Typically, *id* is bound with `(struct id ....)`; more generally, *id* must have a transformer binding of structure-type information at the relevant phase level; see §4.7 "Structure Type Transformer Binding". Furthermore, for each identifier mentioned in the structure-type information, the enclosing module must define or import one identifier that is `free-identifier=?`. If the structure-type information includes a super-type identifier, and if the identifier has a transformer binding of structure-type information, the accessor and mutator bindings of the super-type are *not* included by `struct-out` for export.

Examples:

```
> (module nest racket
    (provide (struct-out egg))
    (struct egg (color wt)))

> (require 'nest)

> (egg-color (egg 'blue 10))
'blue
```

(combine-out *provide-spec* ...)

The union of the *provide-spec*s.

Examples:

```
> (module nest racket
    (provide (combine-out num-eggs num-chicks))
    (define num-eggs 2)
    (define num-chicks 1))

> (require 'nest)

> num-eggs
2
> num-chicks
1
```

(protect-out *provide-spec* ...)

Like the union of the *provide-spec*s, except that the exports are protected; requiring modules may refer to these bindings, but may not extract these bindings from macro expansions or access them via `eval` without access privileges. For more details, see §13.10 “Code Inspectors”. The *provide-spec* must specify only bindings that are defined within the exporting module.

Examples:

```
> (module nest racket
    (provide num-eggs (protect-out num-chicks))
    (define num-eggs 2)
    (define num-chicks 3))

> (define weak-inspector (make-inspector (current-code-
inspector)))

> (define (weak-eval x)
    (parameterize ([current-code-inspector weak-
inspector])
      (define weak-ns (make-base-namespace))
      (namespace-attach-module (current-namespace)
                               ''nest
                               weak-ns)
      (parameterize ([current-namespace weak-ns])
        (namespace-require ''nest)
        (eval x))))
```

```
> (require 'nest)

> (list num-eggs num-chicks)
'(2 3)
> (weak-eval 'num-eggs)
2
> (weak-eval 'num-chicks)
num-chicks: access disallowed by code inspector to protected
variable from module: 'nest
  in: num-chicks
```

(for-meta *phase-level* *provide-spec* ...)

Like the union of the *provide-spec*s, but adjusted to apply to the phase level specified by *phase-level* relative to the current phase level (where `#f` corresponds to the label phase level). In particular, an *id* or `rename-out` form as a *provide-spec* refers to a binding at *phase-level* relative to the current level, an `all-defined-out` exports only definitions at *phase-level* relative to the current phase level, and an `all-from-out` exports bindings imported with a shift by *phase-level*.

Examples:

```
> (module nest racket
    (begin-for-syntax
     (define eggs 2))
    (define chickens 3)
    (provide (for-syntax eggs)
             chickens))

> (require 'nest)

> (define-syntax (test-eggs stx)
    (printf "Eggs are ~a\n" eggs)
    #'0)

> (test-eggs)
Eggs are 2
0
> chickens
3
> (module broken-nest racket
    (define eggs 2)
    (define chickens 3)
    (provide (for-syntax eggs)
             chickens))
```

```
eval:7:0: module: provided identifier not defined or
imported for phase 1
  at: eggs
  in: (#%module-begin (printing-module-begin (define eggs
2)) (printing-module-begin (define chickens 3))
(printing-module-begin (provide (for-syntax eggs)
chickens)))
> (module nest2 racket
    (begin-for-syntax
     (define eggs 2))
    (provide (for-syntax eggs)))

> (require (for-meta 2 racket/base)
           (for-syntax 'nest2))

> (define-syntax (test stx)
    (define-syntax (show-eggs stx)
      (printf "Eggs are ~a\n" eggs)
      #'0)
    (begin
      (show-eggs)
      #'0))
Eggs are 2

> (test)
0
```

(for-syntax *provide-spec* ...)

Same as (for-meta 1 *provide-spec* ...).

(for-template *provide-spec* ...)

Same as (for-meta -1 *provide-spec* ...).

(for-label *provide-spec* ...)

Same as (for-meta #f *provide-spec* ...).

derived-provide-spec

See define-provide-syntax for information on expanding the set of *provide-spec* forms.

Each export specified within a module must have a distinct symbolic export name, though the same binding can be specified with the multiple symbolic names.

```
(for-meta phase-level require-spec ...)
```

See `require` and `provide`.

```
(for-syntax require-spec ...)
```

See `require` and `provide`.

```
(for-template require-spec ...)
```

See `require` and `provide`.

```
(for-label require-spec ...)
```

See `require` and `provide`.

```
(#%require raw-require-spec ...)
```

```
    raw-require-spec = phaseless-spec
                     | (for-meta phase-level phaseless-spec ...)
                     | (for-syntax phaseless-spec ...)
                     | (for-template phaseless-spec ...)
                     | (for-label phaseless-spec ...)
                     | (just-meta phase-level raw-require-spec ...)

        phase-level = exact-integer
                     | #f

      phaseless-spec = raw-module-path
                     | (only raw-module-path id ...)
                     | (prefix prefix-id raw-module-path)
                     | (all-except raw-module-path id ...)
                     | (prefix-all-except prefix-id
                                          raw-module-path id ...)
                     | (rename raw-module-path local-id exported-id)

     raw-module-path = (submod raw-root-module-path id ...+)
                     | (submod "." id ...+)

raw-root-module-path = (quote id)
                     | rel-string
                     | (lib rel-string ...)
                     | id
                     | (file string)
                     | (planet rel-string
                               (user-string pkg-string vers ...))
                     | literal-path
```

The primitive import form, to which `require` expands. A *raw-require-spec* is similar to a *require-spec* in a `require` form, except that the syntax is more constrained, not composable, and not extensible. Also, sub-form names like `for-syntax` and `lib` are recognized symbolically, instead of via bindings. Although not formalized in the grammar above, a `just-meta` form cannot appear within a `just-meta` form.

Each *raw-require-spec* corresponds to the obvious *require-spec*, but the `rename` sub-form has the identifiers in reverse order compared to `rename-in`.

For most *raw-require-spec*s, the lexical context of the *raw-require-spec* determines the context of introduced identifiers. The exception is the `rename` sub-form, where the lexical context of the *local-id* is preserved.

A *literal-path* as a *raw-root-module-path* corresponds to a path in the sense of `path?`. Since path values are never produced by `read-syntax`, they appear only in programmatically constructed expressions. They also appear naturally as arguments to func-

tions such as namespace-require, with otherwise take a quoted raw-module-spec.

```
(#%provide raw-provide-spec ...)

raw-provide-spec = phaseless-spec
                 | (for-meta phase-level phaseless-spec)
                 | (for-syntax phaseless-spec)
                 | (for-label phaseless-spec)
                 | (protect raw-provide-spec)

    phase-level = exact-integer
                 | #f

  phaseless-spec = id
                 | (rename local-id export-id)
                 | (struct struct-id (field-id ...))
                 | (all-from raw-module-path)
                 | (all-from-except raw-module-path id ...)
                 | (all-defined)
                 | (all-defined-except id ...)
                 | (prefix-all-defined prefix-id)
                 | (prefix-all-defined-except prefix-id id ...)
                 | (protect phaseless-spec ...)
                 | (expand (id . datum))
```

The primitive export form, to which provide expands. A *raw-module-path* is as for #%require. A protect sub-form cannot appear within a protect sub-form.

Like #%require, the sub-form keywords for #%provide are recognized symbolically, and nearly every *raw-provide-spec* has an obvious equivalent *provide-spec* via provide, with the exception of the struct and expand sub-forms.

A (struct *struct-id* (*field-id* ...)) sub-form expands to *struct-id*, make-*struct-id*, struct:*struct-id*, *struct-id*?, *struct-id-field-id* for each *field-id*, and set-*struct-id-field-id*! for each *field-id*. The lexical context of the *struct-id* is used for all generated identifiers.

Unlike #%require, the #%provide form is macro-extensible via an explicit expand sub-form; the (*id* . *datum*) part is locally expanded as an expression (even though it is not actually an expression), stopping when a begin form is produced; if the expansion result is (begin *raw-provide-spec* ...), it is spliced in place of the expand form, otherwise a syntax error is reported. The expand sub-form is not normally used directly; it provides a hook for implementing provide and provide transformers.

The all-from and all-from-except forms re-export only identifiers that are accessible in lexical context of the all-from or all-from-except form itself. That is, macro-introduced imports are not re-exported, unless the all-from or all-from-except form

was introduced at the same time. Similarly, `all-defined` and its variants export only definitions accessible from the lexical context of the *phaseless-spec* form.

2.2.1 Additional `require` Forms

```
(require racket/require)
```

The bindings documented in this section are provided by the `racket/require` library, not `racket/base` or `racket`.

The following forms support more complex selection and manipulation of sets of imported identifiers.

```
(matching-identifiers-in regexp require-spec)
```

Like *require-spec*, but including only imports whose names match *regexp*. The *regexp* must be a literal regular expression (see §3.7 “Regular Expressions”).

Examples:

```
> (module zoo racket/base
    (provide tunafish swordfish blowfish
             monkey lizard ant)
    (define tunafish 1)
    (define swordfish 2)
    (define blowfish 3)
    (define monkey 4)
    (define lizard 5)
    (define ant 6))

> (require racket/require)

> (require (matching-identifiers-in #rx"\\w*fish" 'zoo))

> tunafish
1
> swordfish
2
> blowfish
3
> monkey
monkey: undefined;
 cannot reference undefined identifier
```

```
(subtract-in require-spec subtracted-spec ...)
```

Like *require-spec*, but omitting those imports that would be imported by one of the *subtracted-spec*s.

Examples:

```
> (module earth racket
    (provide land sea air)
    (define land 1)
    (define sea 2)
    (define air 3))

> (module mars racket
    (provide aliens)
    (define aliens 4))

> (module solar-system racket
    (require 'earth 'mars)
    (provide (all-from-out 'earth)
             (all-from-out 'mars)))

> (require racket/require)

> (require (subtract-in 'solar-system 'earth))

> land
land: undefined;
 cannot reference undefined identifier
> aliens
4
```

(filtered-in *proc-expr require-spec*)

Applies an arbitrary transformation on the import names (as strings) of *require-spec*. The *proc-expr* must evaluate at expansion time to a single-argument procedure, which is applied on each of the names from *require-spec*. For each name, the procedure must return either a string for the import's new name or `#f` to exclude the import.

For example,

```
(require (filtered-in
          (lambda (name)
            (and (regexp-match? #rx"^[a-z-]+$" name)
                 (regexp-replace #rx"-" (string-
titlecase name) "")))
          racket/base))
```

imports only bindings from racket/base that match the pattern #rx"^[a-z-]+$", and it converts the names to “camel case.”

(path-up rel-string ...)

Specifies paths to modules named by the rel-strings similar to using the rel-strings directly, except that if a required module file is not found relative to the enclosing source, it is searched for in the parent directory, and then in the grand-parent directory, etc., all the way to the root directory. The discovered path relative to the enclosing source becomes part of the expanded form.

This form is useful in setting up a “project environment.” For example, using the following "config.rkt" file in the root directory of your project:

```
#lang racket/base
(require racket/require-syntax
         (for-syntax "utils/in-here.rkt"))

(provide utils-in)
(define-require-syntax utils-in in-here-transformer)
```

and using "utils/in-here.rkt" under the same root directory:

```
#lang racket/base
(require racket/runtime-path)
(provide in-here-transformer)
(define-runtime-path here ".")
(define (in-here-transformer stx)
  (syntax-case stx ()
    [(_ sym)
     (identifier? #'sym)
     (let ([path (build-path here (format "~a.rkt" (syntax-
e #'sym)))])
       (datum->syntax stx `(file ,(path->string path)) stx))]))
```

then path-up works for any other module under the project directory to find "config.rkt":

```
(require racket/require
         (path-up "config.rkt")
         (utils-in foo))
```

Note that the order of requires in the example is important, as each of the first two bind the identifier used in the following.

An alternative in this scenario is to use `path-up` directly to find the utility module:

```
(require racket/require
         (path-up "utils/foo.rkt"))
```

but then sub-directories that are called `"utils"` override the one in the project's root. In other words, the previous method requires only a single unique name.

```
(multi-in subs ...+)

    subs = sub-path
         | (sub-path ...)

sub-path = rel-string
         | id
```

Specifies multiple files to be required from a hierarchy of directories or collections. The set of required module paths is computed as the Cartesian product of the *subs* groups, where each *sub-path* is combined with other *sub-path*s in order using a `/` separator. A *sub-path* as a *subs* is equivalent to `(`*sub-path*`)`. All *sub-path*s in a given `multi-in` form must be either strings or identifiers.

Examples:

```
(require (multi-in racket (dict list)))
```

is equivalent to `(require racket/dict racket/list)`

```
(require (multi-in "math" "matrix" "utils.rkt"))
```

is equivalent to `(require "math/matrix/utils.rkt")`

```
(require (multi-in "utils" ("math.rkt" "matrix.rkt")))
```

is equivalent to `(require "utils/math.rkt" "utils/matrix.rkt")`

```
(require (multi-in ("math" "matrix") "utils.rkt"))
```

is equivalent to `(require "math/utils.rkt" "matrix/utils.rkt")`

```
(require (multi-in ("math" "matrix") ("utils.rkt" "helpers.rkt")))
```

is equivalent to `(require "math/utils.rkt" "math/helpers.rkt"`
`"matrix/utils.rkt" "matrix/helpers.rkt")`

2.2.2 Additional `provide` Forms

```
(require racket/provide)
```

The bindings documented in this section are provided by the `racket/provide` library, not `racket/base` or `racket`.

```
(matching-identifiers-out regexp provide-spec)
```

Like *provide-spec*, but including only exports of bindings with an external name that matches *regexp*. The *regexp* must be a literal regular expression (see §3.7 "Regular Expressions").

```
(filtered-out proc-expr provide-spec)
```

Analogous to `filtered-in`, but for filtering and renaming exports.

For example,

```
(provide (filtered-out
          (lambda (name)
            (and (regexp-match? #rx"^[a-z-]+$" name)
                 (regexp-replace
                  #rx"-" (string-titlecase name) "")))
          (all-defined-out)))
```

exports only bindings that match the pattern `#rx"^[a-z-]+$"`, and it converts the names to "camel case."

2.3 Literals: `quote` and `#%datum`

Many forms are implicitly quoted (via `#%datum`) as literals. See §1.2.3.2 "Expansion Steps" for more information.

§4.10 "Quoting: `quote` and `'`" in ***The Racket Guide*** introduces `quote`.

```
(quote datum)
```

Produces a constant value corresponding to *datum* (i.e., the representation of the program fragment) without its lexical information, source location, etc. Quoted pairs, vectors, and boxes are immutable.

Examples:

```
> (quote x)
'x
> (quote (+ 1 2))
'(+ 1 2)
> (+ 1 2)
3
```

```
(#%datum . datum)
```

Expands to `(quote datum)`, as long as *datum* is not a keyword. If *datum* is a keyword, a syntax error is reported.

See also §1.2.3.2 “Expansion Steps” for information on how the expander introduces `#%datum` identifiers.

Examples:

```
> (#%datum . 10)
10
> (#%datum . x)
'x
> (#%datum . #:x)
eval:6:0: #%datum: keyword used as an expression
  in: #:x
```

2.4 Expression Wrapper: `#%expression`

```
(#%expression expr)
```

Produces the same result as *expr*. The only use of `#%expression` is to force the parsing of a form as an expression.

Examples:

```
> (#%expression (+ 1 2))
3
> (#%expression (define x 10))
eval:8:0: define: not allowed in an expression context
  in: (define x 10)
```

2.5 Variable References and `#%top`

```
id
```

Refers to a module-level or local binding, when *id* is not bound as a transformer (see §1.2.3 “Expansion (Parsing)”). At run-time, the reference evaluates to the value in the location associated with the binding.

When the expander encounters an *id* that is not bound by a module-level or local binding, it converts the expression to (#%top . *id*) giving #%top the lexical context of the *id*; typically, that context refers to #%top. See also §1.2.3.2 “Expansion Steps”.

Examples:

```
> (define x 10)

> x
10
> (let ([x 5]) x)
5
> ((lambda (x) x) 2)
2
```

```
(#%top . id)
```

Refers to a module-level or top-level definition that could bind *id*, even if *id* has a local binding in its context.

Within a module form, (#%top . *id*) expands to just *id*—with the obligation that *id* is defined within the module. At phase level 0, (#%top . *id*) is an immediate syntax error if *id* is not bound. At phase level 1 and higher, a syntax error is reported if *id* is not defined at the corresponding phase by the end of module-body partial expansion.

See also §1.2.3.2 “Expansion Steps” for information on how the expander introduces #%top identifiers.

Examples:

```
> (define x 12)

> (let ([x 5]) (#%top . x))
12
```

2.6 Locations: #%variable-reference

```
(#%variable-reference id)
(#%variable-reference (#%top . id))
(#%variable-reference)
```

Produces an opaque *variable reference* value representing the location of *id*, which must be bound as a variable. If no *id* is supplied, the resulting value refers to an “anonymous”

variable defined within the enclosing context (i.e., within the enclosing module, or at the top level if the form is not inside a module).

A variable reference can be used with variable-reference->empty-namespace, variable-reference->resolved-module-path, and variable-reference->namespace, but facilities like define-namespace-anchor and namespace-anchor->namespace wrap those to provide a clearer interface. A variable reference is also useful to low-level extensions; see *Inside: Racket C API*.

2.7 Procedure Applications and #%app

§4.3 "Function Calls (Procedure Applications)" in ***The Racket Guide*** introduces procedure applications.

(*proc-expr arg ...*)

Applies a procedure, when *proc-expr* is not an identifier that has a transformer binding (see §1.2.3 "Expansion (Parsing)").

More precisely, the expander converts this form to (#%app *proc-expr arg ...*), giving #%app the lexical context that is associated with the original form (i.e., the pair that combines *proc-expr* and its arguments). Typically, the lexical context of the pair indicates the procedure-application #%app that is described next. See also §1.2.3.2 "Expansion Steps".

Examples:

```
> (+ 1 2)
3
> ((lambda (x #:arg y) (list y x)) #:arg 2 1)
'(2 1)
```

(#%app *proc-expr arg ...*)

Applies a procedure. Each *arg* is one of the following:

arg-expr

The resulting value is a non-keyword argument.

keyword arg-expr

The resulting value is a keyword argument using *keyword*. Each *keyword* in the application must be distinct.

The *proc-expr* and *arg-expr*s are evaluated in order, left to right. If the result of *proc-expr* is a procedure that accepts as many arguments as non-*keyword arg-expr*s, if it accepts arguments for all of the *keyword*s in the application, and if all required keyword-based arguments are represented among the *keyword*s in the application, then the procedure is called with the values of the `arg-expr`s. Otherwise, the `exn:fail:contract` exception is raised.

The continuation of the procedure call is the same as the continuation of the application expression, so the results of the procedure are the results of the application expression.

The relative order of *keyword*-based arguments matters only for the order of *arg-expr* evaluations; the arguments are associated with argument variables in the applied procedure based on the *keyword*s, and not their positions. The other *arg-expr* values, in contrast, are associated with variables according to their order in the application form.

See also §1.2.3.2 “Expansion Steps” for information on how the expander introduces `#%app` identifiers.

Examples:

```
> (#%app + 1 2)
3
> (#%app (lambda (x #:arg y) (list y x)) #:arg 2 1)
'(2 1)
> (#%app cons)
cons: arity mismatch;
 the expected number of arguments does not match the given
number
  expected: 2
  given: 0
```

```
(#%plain-app proc-expr arg-expr ...)
(#%plain-app)
```

Like `#%app`, but without support for keyword arguments. As a special case, `(#%plain-app)` produces `'()`.

2.8 Procedure Expressions: `lambda` and `case-lambda`

```
(lambda kw-formals body ...+)
```

§4.4 “Functions (Procedures): `lambda`” in ***The Racket Guide*** introduces procedure expressions.

```
(λ kw-formals body ...+)

kw-formals = (arg ...)
           | (arg ...+ . rest-id)
           | rest-id

       arg = id
           | [id default-expr]
           | keyword id
           | keyword [id default-expr]
```

Produces a procedure. The *kw-formals* determines the number of arguments and which keyword arguments that the procedure accepts.

Considering only the first *arg* case, a simple *kw-formals* has one of the following three forms:

```
(id ...)
```

The procedure accepts as many non-keyword argument values as the number of *id*s. Each *id* is associated with an argument value by position.

```
(id ...+ . rest-id)
```

The procedure accepts any number of non-keyword arguments greater or equal to the number of *id*s. When the procedure is applied, the *id*s are associated with argument values by position, and all leftover arguments are placed into a list that is associated to *rest-id*.

```
rest-id
```

The procedure accepts any number of non-keyword arguments. All arguments are placed into a list that is associated with *rest-id*.

More generally, an *arg* can include a keyword and/or default value. Thus, the first two cases above are more completely specified as follows:

```
(arg ...)
```

Each *arg* has the following four forms:

id

Adds one to both the minimum and maximum number of non-keyword arguments accepted by the procedure. The *id* is associated with an actual argument by position.

[*id default-expr*]

Adds one to the maximum number of non-keyword arguments accepted by the procedure. The *id* is associated with an actual argument by position, and if no such argument is provided, the *default-expr* is evaluated to produce a value associated with *id*. No *arg* with a *default-expr* can appear before an *id* without a *default-expr* and without a *keyword*.

keyword id

The procedure requires a keyword-based argument using *keyword*. The *id* is associated with a keyword-based actual argument using *keyword*.

keyword [*id default-expr*]

The procedure accepts a keyword-based using *keyword*. The *id* is associated with a keyword-based actual argument using *keyword*, if supplied in an application; otherwise, the *default-expr* is evaluated to obtain a value to associate with *id*.

The position of a *keyword arg* in *kw-formals* does not matter, but each specified *keyword* must be distinct.

(*arg ...+ . rest-id*)

Like the previous case, but the procedure accepts any number of non-keyword arguments beyond its minimum number of arguments. When more arguments are provided than non-*keyword* arguments among the *arg*s, the extra arguments are placed into a list that is associated to *rest-id*.

The *kw-formals* identifiers are bound in the *body*s. When the procedure is applied, a new location is created for each identifier, and the location is filled with the associated argument value. The locations are created and filled in order, with *default-expr*s evaluated as needed to fill locations.

In other words, argument bindings with default-value expressions are evaluated analogous to `let*`.

If any identifier appears in the *body*s that is not one of the identifiers in *kw-formals*, then it refers to the same location that it would if it appeared in place of the `lambda` expression. (In other words, variable reference is lexically scoped.)

When multiple identifiers appear in a *kw-formals*, they must be distinct according to `bound-identifier=?`.

If the procedure produced by `lambda` is applied to fewer or more by-position or by-keyword arguments than it accepts, to by-keyword arguments that it does not accept, or without required by-keyword arguments, then the `exn:fail:contract` exception is raised.

The last *body* expression is in tail position with respect to the procedure body.

Examples:

```
> ((lambda (x) x) 10)
10
> ((lambda (x y) (list y x)) 1 2)
'(2 1)
> ((lambda (x [y 5]) (list y x)) 1 2)
'(2 1)
> (let ([f (lambda (x #:arg y) (list y x))])
   (list (f 1 #:arg 2)
         (f #:arg 2 1)))
'((2 1) (2 1))
```

When compiling a `lambda` or `case-lambda` expression, Racket looks for a `'method-arity-error` property attached to the expression (see §11.7 “Syntax Object Properties”). If it is present with a true value, and if no case of the procedure accepts zero arguments, then the procedure is marked so that an `exn:fail:contract:arity` exception involving the procedure will hide the first argument, if one was provided. (Hiding the first argument is useful when the procedure implements a method, where the first argument is implicit in the original source). The property affects only the format of `exn:fail:contract:arity` exceptions, not the result of `procedure-arity`.

When a keyword-accepting procedure is bound to an identifier in certain ways, and when the identifier is used in the function position of an application form, then the application form may be expanded in such a way that the original binding is obscured as the target of the application. To help expose the connection between the function application and function declaration, an identifier in the expansion of the function application is tagged with a syntax property accessible via `syntax-procedure-alias-property` if it is effectively an alias for the original identifier. An identifier in the expansion is tagged with a syntax property accessible via `syntax-procedure-converted-arguments-property` if it is like the original identifier except that the arguments are converted to a flattened form: keyword arguments, required by-position arguments, by-position optional arguments, and rest arguments—all as required, by-position arguments; the keyword arguments are sorted by keyword name, each optional keyword argument is followed by a boolean to indicate

whether a value is provided, and `#f` is used for an optional keyword argument whose value is not provided; optional by-position arguments include `#f` for each non-provided argument, and then the sequence of optional-argument values is followed by a parallel sequence of booleans to indicate whether each optional-argument value was provided.

```
(case-lambda [formals body ...+] ...)

formals = (id ...)
        | (id ...+ . rest-id)
        | rest-id
```

Produces a procedure. Each `[formals body ...+]` clause is analogous to a single `lambda` procedure; applying the `case-lambda`-generated procedure is the same as applying a procedure that corresponds to one of the clauses—the first procedure that accepts the given number of arguments. If no corresponding procedure accepts the given number of arguments, the `exn:fail:contract` exception is raised.

Note that a `case-lambda` clause supports only *formals*, not the more general *kw-formals* of `lambda`. That is, `case-lambda` does not directly support keyword and optional arguments.

Example:

```
> (let ([f (case-lambda
            [() 10]
            [(x) x]
            [(x y) (list y x)]
            [r r])])
    (list (f)
          (f 1)
          (f 1 2)
          (f 1 2 3)))
'(10 1 (2 1) (1 2 3))
```

```
(#%plain-lambda formals body ...+)
```

Like `lambda`, but without support for keyword or optional arguments.

2.9 Local Binding: `let`, `let*`, `letrec`, ...

§4.6 “Local Binding” in ***The Racket Guide*** introduces local binding.

```
(let ([id val-expr] ...) body ...+)
(let proc-id ([id init-expr] ...) body ...+)
```

The first form evaluates the *val-expr*s left-to-right, creates a new location for each *id*, and places the values into the locations. It then evaluates the *body*s, in which the *id*s are bound.

The last *body* expression is in tail position with respect to the `let` form. The *id*s must be distinct according to `bound-identifier=?`.

Examples:

```
> (let ([x 5]) x)
5
> (let ([x 5])
    (let ([x 2]
          [y x])
      (list y x)))
'(5 2)
```

The second form evaluates the *init-expr*s; the resulting values become arguments in an application of a procedure `(lambda (`*id* `...) `*body* `...+)`, where *proc-id* is bound within the *body*s to the procedure itself.

Example:

```
> (let fac ([n 10])
    (if (zero? n)
        1
        (* n (fac (sub1 n)))))
3628800
```

`(let* ([`*id val-expr*`] ...) `*body* `...+)`

Like `let`, but evaluates the *val-expr*s one by one, creating a location for each *id* as soon as the value is available. The *id*s are bound in the remaining *val-expr*s as well as the *body*s, and the *id*s need not be distinct; later bindings shadow earlier bindings.

Example:

```
> (let* ([x 1]
         [y (+ x 1)])
    (list y x))
'(2 1)
```

`(letrec ([`*id val-expr*`] ...) `*body* `...+)`

Like `let`, including left-to-right evaluation of the *val-expr*s, but the locations for all *id*s are created first and filled with `#<undefined>`, all *id*s are bound in all *val-expr*s as well as the *body*s, and each *id* is set immediately after the corresponding *val-expr* is evaluated. The *id*s must be distinct according to `bound-identifier=?`.

Example:

```
> (letrec ([is-even? (lambda (n)
                       (or (zero? n)
                           (is-odd? (sub1 n))))]
           [is-odd? (lambda (n)
                      (and (not (zero? n))
                           (is-even? (sub1 n))))])
    (is-odd? 11))
#t
```

```
(let-values ([(id ...) val-expr] ...) body ...+)
```

Like `let`, except that each *val-expr* must produce as many values as corresponding *id*s, otherwise the `exn:fail:contract` exception is raised. A separate location is created for each *id*, all of which are bound in the *body*s.

Example:

```
> (let-values ([(x y) (quotient/remainder 10 3)])
    (list y x))
'(1 3)
```

```
(let*-values ([(id ...) val-expr] ...) body ...+)
```

Like `let*`, except that each *val-expr* must produce as many values as corresponding *id*s. A separate location is created for each *id*, all of which are bound in the later *val-expr*s and in the *body*s.

Example:

```
> (let*-values ([(x y) (quotient/remainder 10 3)]
                [(z) (list y x)])
    z)
'(1 3)
```

```
(letrec-values ([(id ...) val-expr] ...) body ...+)
```

Like `letrec`, except that each *val-expr* must produce as many values as corresponding *id*s. A separate location is created for each *id*, all of which are initialized to `#<undefined>` and bound in all *val-expr*s and in the *body*s.

Example:

```
> (letrec-values ([(is-even? is-odd?)
                   (values
                     (lambda (n)
                       (or (zero? n)
                           (is-odd? (sub1 n))))
```

```
                      (lambda (n)
                        (or (= n 1)
                            (is-even? (sub1 n))))))])
     (is-odd? 11))
#t
```

```
(let-syntax ([id trans-expr] ...) body ...+)
```

See also `splicing-let-syntax`.

Creates a transformer binding (see §1.2.3.5 “Transformer Bindings”) of each *id* with the value of *trans-expr*, which is an expression at phase level 1 relative to the surrounding context. (See §1.2.1 “Identifiers and Binding” for information on phase levels.)

The evaluation of each *trans-expr* is `parameterized` to set `current-namespace` to a namespace that shares bindings and variables with the namespace being used to expand the `let-syntax` form, except that its base phase is one greater.

Each *id* is bound in the *body*s, and not in other *trans-expr*s.

```
(letrec-syntax ([id trans-expr] ...) body ...+)
```

See also `splicing-letrec-syntax`.

Like `let-syntax`, except that each *id* is also bound within all *trans-expr*s.

```
(let-syntaxes ([(id ...) trans-expr] ...) body ...+)
```

See also `splicing-let-syntaxes`.

Like `let-syntax`, but each *trans-expr* must produce as many values as corresponding *id*s, each of which is bound to the corresponding value.

```
(letrec-syntaxes ([(id ...) trans-expr] ...) body ...+)
```

See also `splicing-letrec-syntaxes`.

Like `let-syntax`, except that each *id* is also bound within all *trans-expr*s.

```
(letrec-syntaxes+values ([(trans-id ...) trans-expr] ...)
                        ([(val-id ...) val-expr] ...)
  body ...+)
```

Combines `letrec-syntaxes` with a variant of `letrec-values`: each *trans-id* and *val-id* is bound in all *trans-expr*s and *val-expr*s.

The `letrec-syntaxes+values` form is the core form for local compile-time bindings, since forms like `letrec-syntax` and internal-definition contexts expand to it. In a fully expanded expression (see §1.2.3.1 “Fully Expanded Programs”), the *trans-id* bindings are discarded and the form reduces to a combination of `letrec-values` or `let-values`, but `letrec-syntaxes+values` can appear in the result of `local-expand` with an empty stop list.

For variables bound by `letrec-syntaxes+values`, the location-creation rules differ slightly from `letrec-values`. The `[(val-id ...) val-expr]` binding clauses are partitioned into minimal sets of clauses that satisfy the following rule: if a clause has a *val-id* binding that is referenced (in a full expansion) by the *val-expr* of an earlier clause, the two clauses and all in between are in the same set. If a set consists of a single clause whose *val-expr* does not refer to any of the clause's *val-id*s, then locations for the *val-id*s are created *after* the *val-expr* is evaluated. Otherwise, locations for all *val-id*s in a set are created just before the first *val-expr* in the set is evaluated.

The end result of the location-creation rules is that scoping and evaluation order are the same as for `letrec-values`, but the compiler has more freedom to optimize away location creation. The rules also correspond to a nesting of `let-values` and `letrec-values`, which is how `letrec-syntaxes+values` for a fully-expanded expression.

See also `local`, which supports local bindings with `define`, `define-syntax`, and more.

2.10 Local Definitions: `local`

```
(require racket/local)
```

The bindings documented in this section are provided by the `racket/local` and `racket` libraries, but not `racket/base`.

```
(local [definition ...] body ...+)
```

Like `letrec-syntaxes+values`, except that the bindings are expressed in the same way as in the top-level or in a module body: using `define`, `define-values`, `define-syntax`, `struct`, etc. Definitions are distinguished from non-definitions by partially expanding *definition* forms (see §1.2.3.6 "Partial Expansion"). As in the top-level or in a module body, a `begin`-wrapped sequence is spliced into the sequence of *definition*s.

2.11 Constructing Graphs: `shared`

```
(require racket/shared)
```

The bindings documented in this section are provided by the `racket/shared` and `racket` libraries, but not `racket/base`.

```
(shared ([id expr] ...) body ...+)
```

Binds *id*s with shared structure according to `exprs` and then evaluates the `body-exprs`, returning the result of the last expression.

The `shared` form is similar to `letrec`, except that special forms of *`expr`* are recognized (after partial macro expansion) to construct graph-structured data, where the corresponding `letrec` would instead produce `#<undefined>`s.

Each *`expr`* (after partial expansion) is matched against the following *`shared-expr`* grammar, where earlier variants in a production take precedence over later variants:

```
        shared-expr = shell-expr
                    | plain-expr

         shell-expr = (cons in-immutable-expr in-immutable-expr)
                    | (list in-immutable-expr ...)
                    | (list* in-immutable-expr ...)
                    | (append early-expr ... in-immutable-expr)
                    | (vector-immutable in-immutable-expr ...)
                    | (box-immutable in-immutable-expr)
                    | (mcons patchable-expr patchable-expr)
                    | (vector patchable-expr ...)
                    | (box patchable-expr ...)
                    | (prefix:make-id patchable-expr ...)

  in-immutable-expr = shell-id
                    | shell-expr
                    | early-expr

           shell-id = id

     patchable-expr = expr

         early-expr = expr

         plain-expr = expr
```

The *`prefix:make-id`* identifier above matches three kinds of references. The first kind is any binding whose name has `make-` in the middle, and where *`prefix:id`* has a transformer binding to structure information with a full set of mutator bindings; see §4.7 “Structure Type Transformer Binding”. The second kind is an identifier that itself has a transformer binding to structure information. The third kind is an identifier that has a `'constructor-for` syntax property whose value is an identifier with a transformer binding to structure information. A *`shell-id`*, meanwhile, must be one of the *`id`*s bound by the `shared` form to a *`shell-expr`*.

When the *`expr`*s of the `shared` form are parsed as *`shared-expr`* (taking into account the order of the variants for parsing precedence), the sub-expressions that were parsed via *`early-expr`* will be evaluated first when the `shared` form is evaluated. Among such expressions, they are evaluated in the order as they appear within the `shared` form. However, any ref-

erence to an *id* bound by `shared` produces `#<undefined>`, even if the binding for the *id* appears before the corresponding *early-expr* within the `shared` form.

The *shell-ids* and *shell-exprs* (not counting *patchable-expr* and *early-expr* sub-expressions) are effectively evaluated next. A *shell-id* reference produces the same value as the corresponding *id* will produce within the *body*s, assuming that *id* is never mutated with `set!`. This special handling of a *shell-id* reference is one way in which `shared` supports the creation of cyclic data, including immutable cyclic data.

Next, the *plain-expr*s are evaluated as for `letrec`, where a reference to an *id* produces `#<undefined>` if it is evaluated before the right-hand side of the *id* binding.

Finally, the *patchable-expr*s are evaluated. At this point, all *id*s are bound, so *patchable-expr*s also creates data cycles (but only with cycles that can be created via mutation).

Examples:

```
> (shared ([a (cons 1 a)])
    a)
#0='(1 . #0#)
> (shared ([a (cons 1 b)]
           [b (cons 2 a)])
    a)
#0='(1 2 . #0#)
> (shared ([a (cons 1 b)]
           [b 7])
    a)
'(1 . 7)
> (shared ([a a]) ; no indirection...
    a)
#<undefined>
> (shared ([a (cons 1 b)] ; b is early...
           [b a])
    a)
'(1 . #<undefined>)
> (shared ([a (mcons 1 b)] ; b is patchable...
           [b a])
    a)
#0=(mcons 1 #0#)
> (shared ([a (vector b b b)]
           [b (box 1)])
    (set-box! b 5)
    a)
'#(#&5 #&5 #&5)
> (shared ([a (box b)]
           [b (vector (unbox a)   ; unbox after a is patched
```

```
                    (unbox c))] ; unbox before c is patched
          [c (box b)])
    b)
#0='#(#0# #<undefined>)
```

2.12 Conditionals: `if`, `cond`, `and`, and `or`

§4.7 "Conditionals" in ***The Racket Guide*** introduces conditionals.

```
(if test-expr then-expr else-expr)
```

Evaluates *test-expr*. If it produces any value other than `#f`, then *then-expr* is evaluated, and its results are the result for the `if` form. Otherwise, *else-expr* is evaluated, and its results are the result for the `if` form. The *then-expr* and *else-expr* are in tail position with respect to the `if` form.

Examples:

```
> (if (positive? -5) (error "doesn't get here") 2)
2
> (if (positive? 5) 1 (error "doesn't get here"))
1
> (if 'we-have-no-bananas "yes" "no")
"yes"
```

```
(cond cond-clause ...)

cond-clause = [test-expr then-body ...+]
            | [else then-body ...+]
            | [test-expr => proc-expr]
            | [test-expr]
```

§4.7.3 "Chaining Tests: `cond`" in ***The Racket Guide*** introduces `cond`.

A *cond-clause* that starts with `else` must be the last *cond-clause*.

If no *cond-clause*s are present, the result is `#<void>`.

If only a `[else then-body ...+]` is present, then the *then-body*s are evaluated. The results from all but the last *then-body* are ignored. The results of the last *then-body*, which is in tail position with respect to the `cond` form, are the results for the whole `cond` form.

Otherwise, the first *test-expr* is evaluated. If it produces `#f`, then the result is the same as a `cond` form with the remaining *cond-clause*s, in tail position with respect to the original `cond` form. Otherwise, evaluation depends on the form of the *cond-clause*:

[*test-expr* *then-body* ...+]

The *then-body*s are evaluated in order, and the results from all but the last *then-body* are ignored. The results of the last *then-body*, which is in tail position with respect to the `cond` form, provides the result for the whole `cond` form.

[*test-expr* => *proc-expr*]

The *proc-expr* is evaluated, and it must produce a procedure that accepts one argument, otherwise the `exn:fail:contract` exception is raised. The procedure is applied to the result of *test-expr* in tail position with respect to the `cond` expression.

[*test-expr*]

The result of the *test-expr* is returned as the result of the `cond` form. The *test-expr* is not in tail position.

Examples:

```
> (cond)

> (cond
    [else 5])
5
> (cond
   [(positive? -5) (error "doesn't get here")]
   [(zero? -5) (error "doesn't get here, either")]
   [(positive? 5) 'here])
'here
> (cond
   [(member 2 '(1 2 3)) => (lambda (l) (map - l))])
'(-2 -3)
> (cond
   [(member 2 '(1 2 3))])
'(2 3)
```

`else`

Recognized specially within forms like `cond`. An `else` form as an expression is a syntax error.

`=>`

Recognized specially within forms like `cond`. A `=>` form as an expression is a syntax error.

```
(and expr ...)
```

§4.7.2 “Combining Tests: `and` and `or`” in ***The Racket Guide*** introduces `and`.

If no *expr*s are provided, then result is `#t`.

If a single *expr* is provided, then it is in tail position, so the results of the `and` expression are the results of the *expr*.

Otherwise, the first *expr* is evaluated. If it produces `#f`, the result of the `and` expression is `#f`. Otherwise, the result is the same as an `and` expression with the remaining *expr*s in tail position with respect to the original `and` form.

Examples:

```
> (and)
#t
> (and 1)
1
> (and (values 1 2))
1
2
> (and #f (error "doesn't get here"))
#f
> (and #t 5)
5
```

```
(or expr ...)
```

§4.7.2 “Combining Tests: `and` and `or`” in ***The Racket Guide*** introduces `or`.

If no *expr*s are provided, then result is `#f`.

If a single *expr* is provided, then it is in tail position, so the results of the `or` expression are the results of the *expr*.

Otherwise, the first *expr* is evaluated. If it produces a value other than `#f`, that result is the result of the `or` expression. Otherwise, the result is the same as an `or` expression with the remaining *expr*s in tail position with respect to the original `or` form.

Examples:

```
> (or)
#f
> (or 1)
1
> (or (values 1 2))
1
2
```

```
> (or 5 (error "doesn't get here"))
5
> (or #f 5)
5
```

2.13 Dispatch: `case`

```
(case val-expr case-clause ...)

case-clause = [(datum ...) then-body ...+]
            | [else then-body ...+]
```

Evaluates *val-expr* and uses the result to select a *case-clause*. The selected clause is the first one with a *datum* whose quoted form is `eqv?` to the result of *val-expr*. If no such *datum* is present, the `else` *case-clause* is selected; if no `else` *case-clause* is present, either, then the result of the `case` form is `#<void>`.

For the selected *case-clause*, the results of the last *then-body*, which is in tail position with respect to the `case` form, are the results for the whole `case` form.

A *case-clause* that starts with `else` must be the last *case-clause*.

Examples:

```
> (case (+ 7 5)
    [(1 2 3) 'small]
    [(10 11 12) 'big])
'big
> (case (- 7 5)
    [(1 2 3) 'small]
    [(10 11 12) 'big])
'small

(define (classify c)
  (case (char-general-category c)
    [(ll lu lt ln lo) "letter"]
    [(nd nl no) "number"]
    [else "other"]))

> (classify #\A)
"letter"
> (classify #\1)
"number"
> (classify #\!)
"other"
```

2.14 Definitions: `define`, `define-syntax`, ...

§4.5 “Definitions: `define`” in ***The Racket Guide*** introduces definitions.

```
(define id expr)
(define (head args) body ...+)

head = id
     | (head args)

args = arg ...
     | arg ... . rest-id

 arg = arg-id
     | [arg-id default-expr]
     | keyword arg-id
     | keyword [arg-id default-expr]
```

The first form binds *id* to the result of *expr*, and the second form binds *id* to a procedure. In the second case, the generated procedure is `(CVT (head args) body ...+)`, using the `CVT` meta-function defined as follows:

```
(CVT (id . kw-formals) . datum)   = (lambda kw-formals . datum)
(CVT (head . kw-formals) . datum) = (lambda kw-formals expr)
                                    if (CVT head . datum) = expr
```

In an internal-definition context (see §1.2.3.7 “Internal Definitions”), a `define` form introduces a local binding. At the top level, the top-level binding for *id* is created after evaluating *expr*, if it does not exist already, and the top-level mapping of *id* (in the namespace linked with the compiled definition) is set to the binding at the same time.

In a context that allows liberal expansion of `define`, *id* is bound as syntax if *expr* is an immediate `lambda` form with keyword arguments or *args* include keyword arguments.

Examples:

```
(define x 10)

> x
10


(define (f x)
  (+ x 1))

> (f 10)
11
```

```
(define ((f x) [y 20])
  (+ x y))

> ((f 10) 30)
40
> ((f 10))
30
```

```
(define-values (id ...) expr)
```

Evaluates the *expr*, and binds the results to the *id*s, in order, if the number of results matches the number of *id*s; if *expr* produces a different number of results, the exn:fail:contract exception is raised.

In an internal-definition context (see §1.2.3.7 “Internal Definitions”), a define-values form introduces local bindings. At the top level, the top-level binding for each *id* is created after evaluating *expr*, if it does not exist already, and the top-level mapping of each *id* (in the namespace linked with the compiled definition) is set to the binding at the same time.

Examples:

```
(define-values () (values))

(define-values (x y z) (values 1 2 3))

> z
3
```

If a define-values form for a function definition in a module body has a 'compiler-hint:cross-module-inline syntax property with a true value, then the Racket treats the property as a performance hint. See §18.4 “Function-Call Optimizations” in ***The Racket Guide*** for more information, and see also begin-encourage-inline.

```
(define-syntax id expr)
(define-syntax (head args) body ...+)
```

The first form creates a transformer binding (see §1.2.3.5 “Transformer Bindings”) of *id* with the value of *expr*, which is an expression at phase level 1 relative to the surrounding context. (See §1.2.1 “Identifiers and Binding” for information on phase levels.) Evaluation of *expr* side is parameterized to set current-namespace as in let-syntax.

The second form is a shorthand the same as for define; it expands to a definition of the first form where the *expr* is a lambda form.

In an internal-definition context (see §1.2.3.7 “Internal Definitions”), a define-syntax form introduces a local binding.

Examples:

```
> (define-syntax foo
    (syntax-rules ()
      ((_ a ...)
       (printf "~a\n" (list a ...)))))

> (foo 1 2 3 4)
(1 2 3 4)

> (define-syntax (bar syntax-object)
    (syntax-case syntax-object ()
      ((_ a ...)
       #'(printf "~a\n" (list a ...)))))

> (bar 1 2 3 4)
(1 2 3 4)
```

```
(define-syntaxes (id ...) expr)
```

Like `define-syntax`, but creates a transformer binding for each *id*. The *expr* should produce as many values as *id*s, and each value is bound to the corresponding *id*.

When *expr* produces zero values for a top-level `define-syntaxes` (i.e., not in a module or internal-definition position), then the *id*s are effectively declared without binding; see §1.2.3.9 “Macro-Introduced Bindings”.

In an internal-definition context (see §1.2.3.7 “Internal Definitions”), a `define-syntaxes` form introduces local bindings.

Examples:

```
> (define-syntaxes (foo1 foo2 foo3)
    (let ([transformer1 (lambda (syntax-object)
                          (syntax-case syntax-object ()
                            [(_) #'1]))]
          [transformer2 (lambda (syntax-object)
                          (syntax-case syntax-object ()
                            [(_) #'2]))]
          [transformer3 (lambda (syntax-object)
                          (syntax-case syntax-object ()
                            [(_) #'3]))])
      (values transformer1
              transformer2
              transformer3)))
```

```
> (foo1)
1
> (foo2)
2
> (foo3)
3
```

```
(define-for-syntax id expr)
(define-for-syntax (head args) body ...+)
```

Like `define`, except that the binding is at phase level 1 instead of phase level 0 relative to its context. The expression for the binding is also at phase level 1. (See §1.2.1 “Identifiers and Binding” for information on phase levels.) The form is a shorthand for `(begin-for-syntax (define id expr))` or `(begin-for-syntax (define (head args) body ...+))`.

Within a module, bindings introduced by `define-for-syntax` must appear before their uses or in the same `define-for-syntax` form (i.e., the `define-for-syntax` form must be expanded before the use is expanded). In particular, mutually recursive functions bound by `define-for-syntax` must be defined by the same `define-for-syntax` form.

Examples:

```
> (define-for-syntax helper 2)

> (define-syntax (make-two syntax-object)
    (printf "helper is ~a\n" helper)
    #'2)

> (make-two)
helper is 2
2
; 'helper' is not bound in the runtime phase
> helper
helper: undefined;
 cannot reference undefined identifier
> (define-for-syntax (filter-ids ids)
    (filter identifier? ids))

> (define-syntax (show-variables syntax-object)
    (syntax-case syntax-object ()
      [(_ expr ...)
       (with-syntax ([(only-ids ...)
                      (filter-ids (syntax->list #'(expr ...)))])
         #'(list only-ids ...))]))
```

```
> (let ([a 1] [b 2] [c 3])
    (show-variables a 5 2 b c))
'(1 2 3)
```

```
(define-values-for-syntax (id ...) expr)
```

Like `define-for-syntax`, but *expr* must produce as many values as supplied *id*s, and all of the *id*s are bound (at phase level 1).

Examples:

```
> (define-values-for-syntax (foo1 foo2) (values 1 2))

> (define-syntax (bar syntax-object)
    (printf "foo1 is ~a foo2 is ~a\n" foo1 foo2)
    #'2)

> (bar)
foo1 is 1 foo2 is 2
2
```

2.14.1 `require` Macros

```
(require racket/require-syntax)
```

The bindings documented in this section are provided by the `racket/require-syntax` library, not `racket/base` or `racket`.

```
(define-require-syntax id expr)
(define-require-syntax (id args ...) body ...+)
```

The first form is like `define-syntax`, but for a `require` sub-form. The `proc-expr` must produce a procedure that accepts and returns a syntax object representing a `require` sub-form.

This form expands to `define-syntax` with a use of `make-require-transformer`; see §11.4.1 “`require` Transformers” for more information.

The second form is a shorthand the same as for `define-syntax`; it expands to a definition of the first form where the *expr* is a `lambda` form.

2.14.2 `provide` Macros

```
(require racket/provide-syntax)
```

The bindings documented in this section are provided by the racket/provide-syntax library, not racket/base or racket.

```
(define-provide-syntax id expr)
(define-provide-syntax (id args ...) body ...+)
```

The first form is like define-syntax, but for a provide sub-form. The proc-expr must produce a procedure that accepts and returns a syntax object representing a provide sub-form.

This form expands to define-syntax with a use of make-provide-transformer; see §11.4.2 "provide Transformers" for more information.

The second form is a shorthand the same as for define-syntax; it expands to a definition of the first form where the *expr* is a lambda form.

2.15 Sequencing: begin, begin0, and begin-for-syntax

§4.8 "Sequencing" in ***The Racket Guide*** introduces begin and begin0.

```
(begin form ...)
(begin expr ...+)
```

The first form applies when begin appears at the top level, at module level, or in an internal-definition position (before any expression in the internal-definition sequence). In that case, the begin form is equivalent to splicing the *form*s into the enclosing context.

The second form applies for begin in an expression position. In that case, the *expr*s are evaluated in order, and the results are ignored for all but the last *expr*. The last *expr* is in tail position with respect to the begin form.

Examples:

```
> (begin
    (define x 10)
    x)
10
> (+ 1 (begin
         (printf "hi\n")
         2))
hi
3
> (let-values ([(x y) (begin
                        (values 1 2 3)
                        (values 1 2))])
   (list x y))
```

```
'(1 2)
```

```
(begin0 expr body ...+)
```

Evaluates the *expr*, then evaluates the *body*s, ignoring the *body* results. The results of the *expr* are the results of the `begin0` form, but the *expr* is in tail position only if no *body*s are present.

Example:

```
> (begin0
    (values 1 2)
    (printf "hi\n"))
hi
1
2
```

```
(begin-for-syntax form ...)
```

Allowed only in a top-level context or module context, shifts the phase level of each *form* by one:

- expressions reference bindings at a phase level one greater than in the context of the `begin-for-syntax` form;
- `define`, `define-values`, `define-syntax`, and `define-syntaxes` forms bind at a phase level one greater than in the context of the `begin-for-syntax` form;
- in `require` and `provide` forms, the default phase level is greater, which is roughly like wrapping the content of the `require` form with `for-syntax`;
- expression form *expr*: converted to `(define-values-for-syntax () (begin` *expr* `(values)))`, which effectively evaluates the expression at expansion time and, in the case of a module context, preserves the expression for future visits of the module.

See also `module` for information about expansion order and partial expansion for `begin-for-syntax` within a module context. Evaluation of an `expr` within `begin-for-syntax` is `parameterized` to set `current-namespace` as in `let-syntax`.

2.16 Guarded Evaluation: `when` and `unless`

§4.8.3 “Effects If...: `when` and `unless`” in ***The Racket Guide*** introduces `when` and `unless`.

```
(when test-expr body ...+)
```

Evaluates *test-expr*. If the result is `#f`, then the result of the `when` expression is `#<void>`. Otherwise, the *body*s are evaluated, and the last *body* is in tail position with respect to the `when` form.

Examples:

```
> (when (positive? -5)
    (display "hi"))

> (when (positive? 5)
    (display "hi")
    (display " there"))
hi there
```

```
(unless test-expr body ...+)
```

Equivalent to `(when (not test-expr) body ...+)`.

Examples:

```
> (unless (positive? 5)
    (display "hi"))

> (unless (positive? -5)
    (display "hi")
    (display " there"))
hi there
```

2.17 Assignment: `set!` and `set!-values`

§4.9 "Assignment: `set!`" in ***The Racket Guide*** introduces `set!`.

```
(set! id expr)
```

If *id* has a transformer binding to an assignment transformer, as produced by `make-set!-transformer` or as an instance of a structure type with the `prop:set!-transformer` property, then this form is expanded by calling the assignment transformer with the full expressions. If *id* has a transformer binding to a rename transformer as produced by `make-rename-transformer` or as an instance of a structure type with the `prop:rename-transformer` property, then this form is expanded by replacing *id* with the target identifier (e.g., the one provided to `make-rename-transformer`). If a transformer binding has both `prop:set!-transformer` and `prop:rename-transformer` properties, the latter takes precedence.

Otherwise, evaluates *expr* and installs the result into the location for *id*, which must be bound as a local variable or defined as a top-level variable or module-level variable. If *id*

refers to an imported binding, a syntax error is reported. If *id* refers to a top-level variable that has not been defined, the `exn:fail:contract` exception is raised.

See also `compile-allow-set!-undefined`.

Examples:

```
(define x 12)

> (set! x (add1 x))

> x
13
> (let ([x 5])
    (set! x (add1 x))
    x)
6
> (set! i-am-not-defined 10)
set!: assignment disallowed;
 cannot set undefined
  variable: i-am-not-defined
```

(set!-values (*id* ...) *expr*)

Assuming that all *id*s refer to variables, this form evaluates *expr*, which must produce as many values as supplied *id*s. The location of each *id* is filled with the corresponding value from *expr* in the same way as for `set!`.

Example:

```
> (let ([a 1]
        [b 2])
    (set!-values (a b) (values b a))
    (list a b))
'(2 1)
```

More generally, the `set!-values` form is expanded to

```
(let-values ([(tmp-id ...) expr])
  (set! id tmp-id) ...)
```

which triggers further expansion if any *id* has a transformer binding to an assignment transformer.

2.18 Iterations and Comprehensions: `for`, `for/list`, ...

§11 "Iterations and Comprehensions" in ***The Racket Guide*** introduces iterations and comprehensions.

The `for` iteration forms are based on SRFI-42 [SRFI-42].

2.18.1 Iteration and Comprehension Forms

```
(for (for-clause ...) body ...+)

for-clause = [id seq-expr]
           | [(id ...) seq-expr]
           | #:when guard-expr
           | #:unless guard-expr

  seq-expr : sequence?
```

Iteratively evaluates *body*. The *for-clause*s introduce bindings whose scope includes *body* and that determine the number of times that *body* is evaluated.

In the simple case, each *for-clause* has one of its first two forms, where `[id seq-expr]` is a shorthand for `[(id) seq-expr]`. In this simple case, the *seq-expr*s are evaluated left-to-right, and each must produce a sequence value (see §3.14.1 “Sequences”).

The `for` form iterates by drawing an element from each sequence; if any sequence is empty, then the iteration stops, and `#<void>` is the result of the `for` expression. Otherwise a location is created for each *id* to hold the values of each element; the sequence produced by a *seq-expr* must return as many values for each iteration as corresponding *id*s.

The *id*s are then bound in the *body*, which is evaluated, and whose results are ignored. Iteration continues with the next element in each sequence and with fresh locations for each *id*.

A `for` form with zero *for-clause*s is equivalent to a single *for-clause* that binds an unreferenced *id* to a sequence containing a single element. All of the *id*s must be distinct according to `bound-identifier=?`.

If any *for-clause* has the form `#:when` *guard-expr*, then only the preceding clauses (containing no `#:when` or `#:unless`) determine iteration as above, and the *body* is effectively wrapped as

```
(when guard-expr
  (for (for-clause ...) body ...+))
```

using the remaining `for-clauses`. A *for-clause* of the form `#:unless` *guard-expr* corresponds to the same transformation with `unless` in place of `when`.

Examples:

```
> (for ([i '(1 2 3)]
        [j "abc"]
        #:when (odd? i)
```

```
        [k #(#t #f)])
    (display (list i j k)))
(1 a #t)(1 a #f)(3 c #t)(3 c #f)

> (for ([(i j) #hash(("a" . 1) ("b" . 20))])
    (display (list i j)))
(b 20)(a 1)

> (for ()
    (display "here"))
here

> (for ([i '()])
    (error "doesn't get here"))
```

(for/list (*for-clause* ...) *body* ...+)

Iterates like `for`, but that the last expression in the *body*s must produce a single value, and the result of the `for/list` expression is a list of the results in order. When evaluation of a *body* is skipped due to a `#:when` or `#:unless` clause, the result list includes no corresponding element.

Examples:

```
> (for/list ([i '(1 2 3)]
             [j "abc"]
             #:when (odd? i)
             [k #(#t #f)])
    (list i j k))
'((1 #\a #t) (1 #\a #f) (3 #\c #t) (3 #\c #f))
> (for/list () 'any)
'(any)
> (for/list ([i '()])
    (error "doesn't get here"))
'()
```

(for/vector (*for-clause* ...) *body* ...+)
(for/vector #:length *length-expr* (*for-clause* ...) *body* ...+)

Iterates like `for/list`, but the result are accumulated into a vector instead of a list. If the optional `#:length` form is used, then *length-expr* must evaluate to an `exact-nonnegative-integer?`, and the result vector is constructed with this length. In this case, the iteration can be performed more efficiently, and terminates when the vector is full or the requested number of iterations have been performed, whichever comes first. If the provided *length-expr* evaluates to a length longer than the number of iterations then the remaining slots of the vector are initialized to the default argument of `make-vector`.

```
(for/hash (for-clause ...) body ...+)
(for/hasheq (for-clause ...) body ...+)
(for/hasheqv (for-clause ...) body ...+)
```

Like `for/list`, but the result is an immutable hash table; `for/hash` creates a table using `equal?` to distinguish keys, `for/hasheq` produces a table using `eq?`, and `for/hasheqv` produces a table using `eqv?`. The last expression in the *body*s must return two values: a key and a value to extend the hash table accumulated by the iteration.

Example:

```
> (for/hash ([i '(1 2 3)])
    (values i (number->string i)))
'#hash((1 . "1") (2 . "2") (3 . "3"))
```

```
(for/and (for-clause ...) body ...+)
```

Iterates like `for`, but when last expression of *body* produces `#f`, then iteration terminates, and the result of the `for/and` expression is `#f`. If the *body* is never evaluated, then the result of the `for/and` expression is `#t`. Otherwise, the result is the (single) result from the last evaluation of *body*.

Examples:

```
> (for/and ([i '(1 2 3 "x")])
    (i . < . 3))
#f
> (for/and ([i '(1 2 3 4)])
    i)
4
> (for/and ([i '()])
    (error "doesn't get here"))
#t
```

```
(for/or (for-clause ...) body ...+)
```

Iterates like `for`, but when last expression of *body* produces a value other than `#f`, then iteration terminates, and the result of the `for/or` expression is the same (single) value. If the *body* is never evaluated, then the result of the `for/or` expression is `#f`. Otherwise, the result is `#f`.

Examples:

```
> (for/or ([i '(1 2 3 "x")])
    (i . < . 3))
#t
```

```
> (for/or ([i '(1 2 3 4)])
    i)
1
> (for/or ([i '()])
    (error "doesn't get here"))
#f
```

```
(for/sum (for-clause ...) body ...+)
```

Iterates like `for`, but each result of the last *body* is accumulated into a result with `+`.

Example:

```
> (for/sum ([i '(1 2 3 4)]) i)
10
```

```
(for/product (for-clause ...) body ...+)
```

Iterates like `for`, but each result of the last *body* is accumulated into a result with `*`.

Example:

```
> (for/product ([i '(1 2 3 4)]) i)
24
```

```
(for/lists (id ...) (for-clause ...) body ...+)
```

Similar to `for/list`, but the last *body* expression should produce as many values as given *id*s, and the result is as many lists as supplied *id*s. The *id*s are bound to the lists accumulated so far in the *for-clause*s and *body*s.

```
(for/first (for-clause ...) body ...+)
```

Iterates like `for`, but after *body* is evaluated the first time, then the iteration terminates, and the `for/first` result is the (single) result of *body*. If the *body* is never evaluated, then the result of the `for/first` expression is `#f`.

Examples:

```
> (for/first ([i '(1 2 3 "x")]
              #:when (even? i))
    (number->string i))
"2"
> (for/first ([i '()])
    (error "doesn't get here"))
#f
```

```
(for/last (for-clause ...) body ...+)
```

Iterates like `for`, but the `for/last` result is the (single) result of the last evaluation of *body*. If the *body* is never evaluated, then the result of the `for/last` expression is `#f`.

Examples:

```
> (for/last ([i '(1 2 3 4 5)]
               #:when (even? i))
    (number->string i))
"4"
> (for/last ([i '()])
    (error "doesn't get here"))
#f
```

```
(for/fold ([accum-id init-expr] ...) (for-clause ...) . body)
```

Iterates like `for`. Before iteration starts, the *init-expr*s are evaluated to produce initial accumulator values. At the start of each iteration, a location is generated for each *accum-id*, and the corresponding current accumulator value is placed into the location. The last expression in *body* must produce as many values as *accum-id*s, and those values become the current accumulator values. When iteration terminates, the results of the `fold/for` expression are the accumulator values.

Example:

```
> (for/fold ([sum 0]
             [rev-roots null])
            ([i '(1 2 3 4)])
    (values (+ sum i) (cons (sqrt i) rev-roots)))
10
'(2 1.7320508075688772 1.4142135623730951 1)
```

```
(for* (for-clause ...) body ...+)
```

Like `for`, but with an implicit `#:when #t` between each pair of `for-clauses`, so that all sequence iterations are nested.

Example:

```
> (for* ([i '(1 2)]
         [j "ab"])
    (display (list i j)))
(1 a)(1 b)(2 a)(2 b)
```

```
(for*/list (for-clause ...) body ...+)
(for*/lists (id ...) (for-clause ...) body ...+)
(for*/vector (for-clause ...) body ...+)
(for*/vector #:length length-expr (for-clause ...) body ...+)
(for*/hash (for-clause ...) body ...+)
(for*/hasheq (for-clause ...) body ...+)
(for*/hasheqv (for-clause ...) body ...+)
(for*/and (for-clause ...) body ...+)
(for*/or (for-clause ...) body ...+)
(for*/sum (for-clause ...) body ...+)
(for*/product (for-clause ...) body ...+)
(for*/first (for-clause ...) body ...+)
(for*/last (for-clause ...) body ...+)
(for*/fold ([accum-id init-expr] ...) (for-clause ...) body ...+)
```

Like `for/list`, etc., but with the implicit nesting of `for*`.

Example:

```
> (for*/list ([i '(1 2)]
              [j "ab"])
    (list i j))
'((1 #\a) (1 #\b) (2 #\a) (2 #\b))
```

2.18.2 Deriving New Iteration Forms

```
(for/fold/derived orig-datum
  ([accum-id init-expr] ...) (for-clause ...) body ...+)
```

Like `for/fold`, but the extra *orig-datum* is used as the source for all syntax errors.

Examples:

```
> (define-syntax (for/digits stx)
    (syntax-case stx ()
      [(_ clauses . defs+exprs)
       (with-syntax ([original stx])
         #'(let-values
             ([(n k)
               (for/fold/derived original ([n 0] [k 1]) clauses
                 (define d (let () . defs+exprs))
                 (values (+ n (* d k)) (* k 10)))])
             n))]))
```

```
> (for/digits
      [a (in-list '(1 2 3))]
      [b (in-list '(4 5 6))]
    (+ a b))
eval:3:0: for/digits: bad sequence binding clause
  at: a
  in: (for/digits (a (in-list (quote (1 2 3)))) (b (in-list
(quote (4 5 6)))) (+ a b))
> (for/digits
      ([a (in-list '(1 2 3))]
       [b (in-list '(2 4 6))])
    (+ a b))
963
```

```
(for*/fold/derived orig-datum
  ([accum-id init-expr] ...) (for-clause ...) body ...+)
```

Like `for*/fold`, but the extra *orig-datum* is used as the source for all syntax errors.

Examples:

```
> (define-syntax (for*/digits stx)
    (syntax-case stx ()
      [(_ clauses . defs+exprs)
       (with-syntax ([original stx])
         #'(let-values
             ([(n k)
               (for*/fold/derived original ([n 0] [k 1]) clauses
                 (define d (let () . defs+exprs))
                 (values (+ n (* d k)) (* k 10)))])
             n))]))

> (for*/digits
      [ds (in-list '((8 3) (1 1)))]
      [d (in-list ds)]
    d)
eval:6:0: for*/digits: bad sequence binding clause
  at: ds
  in: (for*/digits (ds (in-list (quote ((8 3) (1 1))))) (d
(in-list ds)) d)
> (for*/digits
      ([ds (in-list '((8 3) (1 1)))]
       [d (in-list ds)])
    d)
1138
```

```
(define-sequence-syntax id
  expr-transform-expr
  clause-transform-expr)

                          (or/c (-> identifier?)
  expr-transform-expr :           (syntax? . -> . syntax?))

  clause-transform-expr : (syntax? . -> . syntax?)
```

Defines *id* as syntax. An (*id* . *rest*) form is treated specially when used to generate a sequence in a *clause* of `for` (or one of its variants). In that case, the procedure result of *clause-transform-expr* is called to transform the clause.

When *id* is used in any other expression position, the result of *expr-transform-expr* is used. If it is a procedure of zero arguments, then the result must be an identifier *other-id*, and any use of *id* is converted to a use of *other-id*. Otherwise,*expr-transform-expr* must produce a procedure (of one argument) that is used as a macro transformer.

When the *clause-transform-expr* transformer is used, it is given a *clause* as an argument, where the clause's form is normalized so that the left-hand side is a parenthesized sequence of identifiers. The right-hand side is of the form (*id* . *rest*). The result can be either `#f`, to indicate that the forms should not be treated specially (perhaps because the number of bound identifiers is inconsistent with the (*id* . *rest*) form), or a new *clause* to replace the given one. The new clause might use `:do-in`. To protect identifiers in the result of *clause-transform-expr*, use `for-clause-syntax-protect` instead of `syntax-protect`.

Examples:

```
> (define-sequence-syntax in-digits
    (lambda () #'in-digits/proc)
    (lambda (stx)
      (syntax-case stx ()
        [[(d) (_ nat)]
         #'[(d)
            (:do-in
              ([(n) nat])
              (unless (exact-nonnegative-integer? n)
                (raise-type-error 'in-digits "exact non-negative
integer" n))
              ([i n])
              (not (zero? i))
              ([(j d) (quotient/remainder i 10)])
              #true
              #true
              [j])]])))
```

```
> (define (in-digits/proc n [b 10])
    (for/list ([d (in-digits n)]) d))

> (for/list ([d (in-digits 1138)]) d)
'(8 3 1 1)
> (map in-digits (list 137 216))
'((7 3 1) (6 1 2))
```

```
(:do-in ([(outer-id ...) outer-expr] ...)
        outer-check
        ([loop-id loop-expr] ...)
        pos-guard
        ([(inner-id ...) inner-expr] ...)
        pre-guard
        post-guard
        (loop-arg ...))
```

A form that can only be used as a *seq-expr* in a *clause* of `for` (or one of its variants).

Within a `for`, the pieces of the `:do-in` form are spliced into the iteration essentially as follows:

```
(let-values ([(outer-id ...) outer-expr] ...)
  outer-check
  (let loop ([loop-id loop-expr] ...)
    (if pos-guard
        (let-values ([(inner-id ...) inner-expr] ...)
          (if pre-guard
              (let body-bindings
                   (if post-guard
                       (loop loop-arg ...)
                       done-expr))
              done-expr))
        done-expr)))
```

where *body-bindings* and *done-expr* are from the context of the `:do-in` use. The identifiers bound by the `for` clause are typically part of the `([(`*inner-id* `...)` *inner-expr*`] ...)` section.

The actual `loop` binding and call has additional loop arguments to support iterations in parallel with the `:do-in` form, and the other pieces are similarly accompanied by pieces from parallel iterations.

For an example of `:do-in`, see `define-sequence-syntax`.

```
(for-clause-syntax-protect stx) → syntax?
  stx : syntax?
```

Provided `for-syntax`: Like `syntax-protect`, but allows the `for` expander to disarm the result syntax object, and arms the pieces of a clause instead of the entire syntax object.

Use this function to protect the result of a *clause-transform-expr* that is bound by `define-sequence-syntax`.

2.18.3 Do Loops

```
(do ([id init-expr step-expr-maybe] ...)
    (stop?-expr finish-expr ...)
  expr ...)

step-expr-maybe =
                | step-expr
```

Iteratively evaluates the *expr*s for as long as *stop?-expr* returns `#f`.

To initialize the loop, the *init-expr*s are evaluated in order and bound to the corresponding *id*s. The *id*s are bound in all expressions within the form other than the *init-expr*s.

After the *id*s have been bound, the *stop?-expr* is evaluated. If it produces `#f`, each *expr* is evaluated for its side-effect. The *id*s are then effectively updated with the values of the *step-expr*s, where the default *step-expr* for *id* is just *id*; more precisely, iteration continues with fresh locations for the *id*s that are initialized with the values of the corresponding *step-expr*s.

When *stop?-expr* produces a true value, then the *finish-expr*s are evaluated in order, and the last one is evaluated in tail position to produce the overall value for the `do` form. If no *finish-expr* is provided, the value of the `do` form is `#<void>`.

2.19 Continuation Marks: `with-continuation-mark`

```
(with-continuation-mark key-expr val-expr result-expr)
```

The *key-expr*, `mark-expr`, and *result-expr* expressions are evaluated in order. After *key-expr* is evaluated to obtain a key and `mark-expr` is evaluated to obtain a mark, the key is mapped to the mark in the current continuation's initial frame. If the frame already has a mark for the key, it is replaced. Finally, the *result-expr* is evaluated; the continuation for evaluating *result-expr* is the continuation of the `with-continuation-mark`

expression (so the result of the *result-expr* is the result of the with-continuation-mark expression, and *result-expr* is in tail position for the with-continuation-mark expression).

§9.5 “Continuation Marks” provides more information on continuation marks.

2.20 Quasiquoting: quasiquote, unquote, and unquote-splicing

```
(quasiquote datum)
```

§4.11 “Quasiquoting: quasiquote and ‘” in ***The Racket Guide*** introduces quasiquote.

The same as 'datum if *datum* does not include (unquote *expr*) or (unquote-splicing *expr*). An (unquote *expr*) form escapes from the quote, however, and the result of the *expr* takes the place of the (unquote *expr*) form in the quasiquote result. An (unquote-splicing *expr*) similarly escapes, but the *expr* must produce a list, and its elements are spliced as multiple values place of the (unquote-splicing *expr*), which must appear as the car or a quoted pair, as an element of a quoted vector, or as an element of a quoted prefab structure; in the case of a pair, if the cdr of the relevant quoted pair is empty, then *expr* need not produce a list, and its result is used directly in place of the quoted pair (in the same way that append accepts a non-list final argument). In a quoted hash table, an (unquote *expr*) or (unquote-splicing *expr*) expression escapes only in the second element of an entry pair (i.e., the value), while entry keys are always implicitly quoted. If unquote or unquote-splicing appears within quasiquote in any other way than as (unquote *expr*) or (unquote-splicing *expr*), a syntax error is reported.

Examples:

```
> (quasiquote (0 1 2))
'(0 1 2)
> (quasiquote (0 (unquote (+ 1 2)) 4))
'(0 3 4)
> (quasiquote (0 (unquote-splicing (list 1 2)) 4))
'(0 1 2 4)
> (quasiquote (0 (unquote-splicing 1) 4))
unquote-splicing: contract violation
  expected: list?
  given: 1
> (quasiquote (0 (unquote-splicing 1)))
'(0 . 1)
```

A quasiquote, unquote, or unquote-splicing form is typically abbreviated with `, ,, or ,@, respectively. See also §1.3.7 “Reading Quotes”.

Examples:

```
> `(0 1 2)
'(0 1 2)
> `(1 ,(+ 1 2) 4)
```

```
'(1 3 4)
> `#s(stuff 1 ,(+ 1 2) 4)
'#s(stuff 1 3 4)
> `#hash(("a" . ,(+ 1 2)))
'#hash(("a" . 3))
> `#hash((,(+ 1 2) . "a"))
'#hash((,(+ 1 2) . "a"))
> `(1 ,@(list 1 2) 4)
'(1 1 2 4)
> `#(1 ,@(list 1 2) 4)
'#(1 1 2 4)
```

A `quasiquote` form within the original *datum* increments the level of quasiquotation: within the `quasiquote` form, each `unquote` or `unquote-splicing` is preserved, but a further nested `unquote` or `unquote-splicing` escapes. Multiple nestings of `quasiquote` require multiple nestings of `unquote` or `unquote-splicing` to escape.

Examples:

```
> `(1 `,(+ 1 ,(+ 2 3)) 4)
'(1 `,(+ 1 5) 4)
> `(1 ```,,@,,@(list (+ 1 2)) 4)
'(1 ```,,@,3 4)
```

The `quasiquote` form allocates only as many fresh cons cells, vectors, and boxes as are needed without analyzing `unquote` and `unquote-splicing` expressions. For example, in

```
`(,1 2 3)
```

a single tail `'(2 3)` is used for every evaluation of the `quasiquote` expression.

`unquote`

See `quasiquote`, where `unquote` is recognized as an escape. An `unquote` form as an expression is a syntax error.

`unquote-splicing`

See `quasiquote`, where `unquote-splicing` is recognized as an escape. An `unquote-splicing` form as an expression is a syntax error.

2.21 Syntax Quoting: `quote-syntax`

`(quote-syntax datum)`

Produces a syntax object that preserves the lexical information and source-location information attached to *datum* at expansion time.

Example:

```
> (syntax? (quote-syntax x))
#t
```

2.22 Interaction Wrapper: #%top-interaction

```
(#%top-interaction . form)
```

Expands to simply *form*. The #%top-interaction form is similar to #%app and #%module-begin, in that it provides a hook to control interactive evaluation through load (more precisely, the default load handler) or read-eval-print-loop.

2.23 Limiting Scope: define-package, open-package, ...

```
(require racket/package)
```

The bindings documented in this section are provided by the racket/package library, not racket/base or racket.

```
(define-package package-id exports form ...)
(open-package package-id)

exports = (id ...)
        | #:only (id ...)
        | #:all-defined
        | #:all-defined-except (id ...)
```

The define-package form is based on the module form of Chez Scheme [Waddell99].

The define-package form is similar to module, except that it can appear in any definition context. The *form*s within a define-package form can be definitions or expressions; definitions are not visible outside the define-package form, but *exports* determines a subset of the bindings that can be made visible outside the package using the definition form (open-package *package-id*).

The (*id* ...) and #:only (*id* ...) *exports* forms are equivalent: exactly the listed *id*s are exported. The #:all-defined form exports all definitions from the package body, and #:all-defined-except (*id* ...) exports all definitions except the listed *id*s.

All of the usual definition forms work within a define-package body, and such definitions are visible to all expressions within the body (and, in particular, the definitions can refer to

each other). However, `define-package` handles `define*`, `define*-syntax`, `define*-values`, `define*-syntaxes`, and `open*-package` specially: the bindings introduced by those forms within a `define-package` body are visible only to *`form`*s that appear later in the body, and they can shadow any binding from preceding *`form`*s (even if the preceding binding did not use one of the special `*` definition forms). If an exported identifier is defined multiple times, the last definition is the exported one.

Examples:

```
> (define-package presents (doll)
    (define doll "Molly Coddle")
    (define robot "Destructo"))

> doll
doll: undefined;
 cannot reference undefined identifier
> robot
robot: undefined;
 cannot reference undefined identifier
> (open-package presents)

> doll
"Molly Coddle"
> robot
robot: undefined;
 cannot reference undefined identifier
> (define-package big-russian-doll (middle-russian-doll)
    (define-package middle-russian-doll (little-russian-doll)
      (define little-russian-doll "Anastasia")))

> (open-package big-russian-doll)

> (open-package middle-russian-doll)

> little-russian-doll
"Anastasia"
```

(package-begin *form* ...)

Similar to `define-package`, but it only limits the visible of definitions without binding a package name. If the last *`form`* is an expression, then the expression is in tail position for the `package-begin` form, so that its result is the `package-begin` result.

A `package-begin` form can be used as an expression, but if it is used in a context where definitions are allowed, then the definitions are essentially spliced into the enclosing context (though the defined bindings remain hidden outside the `package-begin`).

Examples:

```
> (package-begin
    (define secret "mimi")
    (list secret))
'("mimi")
> secret
secret: undefined;
 cannot reference undefined identifier
```

```
define*
define*-values
define*-syntax
define*-syntaxes
open*-package
```

Equivalent to `define`, `define-values`, `define-syntax`, `define-syntaxes`, and `open-package`, except within a `define-package` or `package-begin` form, where they create bindings that are visible only to later body forms.

Examples:

```
> (define-package mail (cookies)
    (define* cookies (list 'sugar))
    (define* cookies (cons 'chocolate-chip cookies)))

> (open-package mail)

> cookies
'(chocolate-chip sugar)
> (define-syntax-rule (define-seven id) (define id 7))

> (define-syntax-rule (define*-seven id)
    (begin
      (define-package p (id) (define-seven id))
      (open*-package p)))

> (package-begin
    (define vii 8)
    (define*-seven vii)
    vii)
7
```

```
(package? v) → boolean?
  v : any/c
(package-exported-identifiers id) → (listof identifier?)
  id : identifier?
```

```
(package-original-identifiers id) → (listof identifier?)
  id : identifier?
```

The package?, package-exported-identifiers, and package-original-identifiers functions are exported for-syntax by racket/package.

The package? predicate returns #t if v is a package value as obtained by syntax-local-value on an identifier that is bound to a package.

Given such an identifier, the package-exported-identifiers function returns a list of identifiers that correspond to the bindings that would be introduced by opening the package in the lexical context being expanded. The package-original-identifiers function returns a parallel list of identifiers for existing bindings of package's exports.

2.24 Blocks: block

```
(require racket/block)
```

The bindings documented in this section are provided by the racket/block library, not racket/base or racket.

```
(block defn-or-expr ...)
```

Supports a mixture of expressions and mutually recursive definitions, as in a module body. Unlike an internal-definition context, the last *defn-or-expr* need not be an expression.

The result of the block form is the result of the last *defn-or-expr* if it is an expression, #<void> otherwise. If no *defn-or-expr* is provided (after flattening begin forms), the result is #<void>.

The final *defn-or-expr* is executed in tail position, if it is an expression.

Examples:

```
> (define (f x)
    (block
      (define y (add1 x))
      (displayln y)
      (define z (* 2 y))
      (+ 3 z)))

> (f 12)
13
29
```

2.25 Internal-Definition Limiting: #%stratified-body

```
(#%stratified-body defn-or-expr ...)
```

Like (let () defn-or-expr ...) for an internal-definition context sequence, except that an expression is not allowed to precede a definition, and all definitions are treated as referring to all other definitions (i.e., locations for variables are all allocated first, like letrec and unlike letrec-syntaxes+values).

The #%stratified-body form is useful for implementing syntactic forms or languages that supply a more limited kind of internal-definition context.

2.26 Performance Hints: begin-encourage-inline

```
(require racket/performance-hint)
```

The bindings documented in this section are provided by the racket/performance-hint library, not racket/base or racket.

```
(begin-encourage-inline form ...)
```

Attaches a 'compiler-hint:cross-module-inline syntax property to each form, which is useful when a form is a function definition. See define-values.

3 Datatypes

§3 “Built-In Datatypes” in ***The Racket Guide*** introduces Datatypes.

Each pre-defined datatype comes with a set of procedures for manipulating instances of the datatype.

3.1 Booleans and Equality

True and false *booleans* are represented by the values `#t` and `#f`, respectively, though operations that depend on a boolean value typically treat anything other than `#f` as true. The `#t` value is always `eq?` to itself, and `#f` is always `eq?` to itself.

See §1.3.4 “Reading Booleans” for information on `read`ing booleans and §1.4.3 “Printing Booleans” for information on `print`ing booleans.

See also `and`, `or`, `andmap`, and `ormap`.

```
(boolean? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is `#t` or `#f`, `#f` otherwise.

Examples:

```
> (boolean? #f)
#t
> (boolean? #t)
#t
> (boolean? 'true)
#f
```

```
(not v) → boolean?
  v : any/c
```

Returns `#t` if `v` is `#f`, `#f` otherwise.

Examples:

```
> (not #f)
#t
> (not #t)
#f
> (not 'we-have-no-bananas)
#f
```

```
(equal? v1 v2) → boolean?
  v1 : any/c
  v2 : any/c
```

Two values are `equal?` if and only if they are `eqv?`, unless otherwise specified for a particular datatype.

Datatypes with further specification of `equal?` include strings, byte strings, pairs, mutable pairs, vectors, boxes, hash tables, and inspectable structures. In the last six cases, equality is recursively defined; if both *v1* and *v2* contain reference cycles, they are equal when the infinite unfoldings of the values would be equal. See also `gen:equal+hash` and `prop:impersonator-of`.

Examples:

```
> (equal? 'yes 'yes)
#t
> (equal? 'yes 'no)
#f
> (equal? (expt 2 100) (expt 2 100))
#t
> (equal? 2 2.0)
#f
> (equal? (make-string 3 #\z) (make-string 3 #\z))
#t
```

```
(eqv? v1 v2) → boolean?
  v1 : any/c
  v2 : any/c
```

Two values are `eqv?` if and only if they are `eq?`, unless otherwise specified for a particular datatype.

The number and character datatypes are the only ones for which `eqv?` differs from `eq?`.

Examples:

```
> (eqv? 'yes 'yes)
#t
> (eqv? 'yes 'no)
#f
> (eqv? (expt 2 100) (expt 2 100))
#t
> (eqv? 2 2.0)
#f
> (eqv? (integer->char 955) (integer->char 955))
```

```
#t
> (eqv? (make-string 3 #\z) (make-string 3 #\z))
#f
```

```
(eq? v1 v2) → boolean?
  v1 : any/c
  v2 : any/c
```

Return `#t` if *v1* and *v2* refer to the same object, `#f` otherwise. See also §1.1.6 “Object Identity and Comparisons”.

Examples:

```
> (eq? 'yes 'yes)
#t
> (eq? 'yes 'no)
#f
> (let ([v (mcons 1 2)]) (eq? v v))
#t
> (eq? (mcons 1 2) (mcons 1 2))
#f
> (eq? (make-string 3 #\z) (make-string 3 #\z))
#f
```

```
(equal?/recur v1 v2 recur-proc) → boolean?
  v1 : any/c
  v2 : any/c
  recur-proc : (any/c any/c -> any/c)
```

Like `equal?`, but using *recur-proc* for recursive comparisons (which means that reference cycles are not handled automatically). Non-`#f` results from *recur-proc* are converted to `#t` before being returned by `equal?/recur`.

Examples:

```
> (equal?/recur 1 1 (lambda (a b) #f))
#t
> (equal?/recur '(1) '(1) (lambda (a b) #f))
#f
> (equal?/recur '#(1 1 1) '#(1 1.2 3/4)
                (lambda (a b) (<= (abs (- a b)) 0.25)))
#t
```

```
(immutable? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is an immutable string, byte string, vector, hash table, or box, `#f` otherwise.

Examples:

```
> (immutable? 'hello)
#f
> (immutable? "a string")
#t
> (immutable? (box 5))
#f
> (immutable? #(0 1 2 3))
#t
> (immutable? (make-hash))
#f
> (immutable? (make-immutable-hash '([a b])))
#t
```

```
gen:equal+hash : any/c
```

A generic interface (see §4.4 “Generic Interfaces”) that supplies an equality predicate and hashing functions for a structure type. The following methods must be implemented:

- *equal-proc* `: (-> any/c any/c (-> any/c any/c boolean?) any/c)` — tests whether the first two arguments are equal, where both values are instances of the structure type to which the generic interface is associated (or a subtype of the structure type).

  The third argument is an `equal?` predicate to use for recursive equality checks; use the given predicate instead of `equal?` to ensure that data cycles are handled properly and to work with `equal?/recur` (but beware that an arbitrary function can be provided to `equal?/recur` for recursive checks, which means that arguments provided to the predicate might be exposed to arbitrary code).

  The *equal-proc* is called for a pair of structures only when they are not `eq?`, and only when they both have a `gen:equal+hash` value inherited from the same structure type. With this strategy, the order in which `equal?` receives two structures does not matter. It also means that, by default, a structure sub-type inherits the equality predicate of its parent, if any.

- *hash-proc* `: (-> any/c (-> any/c exact-integer?) exact-integer?)` — computes a hash code for the given structure, like `equal-hash-code`. The first argument is an instance of the structure type (or one of its subtypes) to which the generic interface is associated.

  The second argument is an `equal-hash-code`-like procedure to use for recursive hash-code computation; use the given procedure instead of `equal-hash-code` to ensure that data cycles are handled properly.

- *hash2-proc* `: (-> any/c (-> any/c exact-integer?) exact-integer?)` — computes a secondary hash code for the given structure. This procedure is like *hash-proc*, but analogous to `equal-secondary-hash-code`.

Take care to ensure that *hash-proc* and *hash2-proc* are consistent with *equal-proc*. Specifically, *hash-proc* and *hash2-proc* should produce the same value for any two structures for which *equal-proc* produces a true value.

When a structure type has no gen:equal+hash implementation, then transparent structures (i.e., structures with an inspector that is controlled by the current inspector) are equal? when they are instances of the same structure type (not counting sub-types), and when they have equal? field values. For transparent structures, equal-hash-code and equal-secondary-hash-code derive hash code using the field values. For opaque structure types, equal? is the same as eq?, and equal-hash-code and equal-secondary-hash-code results are based only on eq-hash-code. If a structure has a prop:impersonator-of property, then the prop:impersonator-of property takes precedence over gen:equal+hash if the property value's procedure returns a non-#f value when applied to the structure.

Examples:

```
> (define (farm=? farm1 farm2 recursive-equal?)
    (and (= (farm-apples farm1)
            (farm-apples farm2))
         (= (farm-oranges farm1)
            (farm-oranges farm2))
         (= (farm-sheep farm1)
            (farm-sheep farm2))))

> (define (farm-hash-1 farm recursive-equal-hash)
    (+ (* 10000 (farm-apples farm))
       (* 100 (farm-oranges farm))
       (* 1 (farm-sheep farm))))

> (define (farm-hash-2 farm recursive-equal-hash)
    (+ (* 10000 (farm-sheep farm))
       (* 100 (farm-apples farm))
       (* 1 (farm-oranges farm))))

> (define-struct farm (apples oranges sheep)
                 #:methods gen:equal+hash
                 [(define equal-proc farm=?)
                  (define hash-proc  farm-hash-1)
                  (define hash2-proc farm-hash-2)])

> (define east (make-farm 5 2 20))

> (define west (make-farm 18 6 14))

> (define north (make-farm 5 20 20))
```

```
> (define south (make-farm 18 6 14))

> (equal? east west)
#f
> (equal? east north)
#f
> (equal? west south)
#t
```

prop:equal+hash : struct-type-property?

A deprecated structure type property (see §4.3 “Structure Type Properties”) that supplies an equality predicate and hashing functions for a structure type. gen:equal+hash should be used instead. Accepts a list of three procedures that correspond to the methods of gen:equal+hash.

3.1.1 Boolean Aliases

```
(require racket/bool)
```

The bindings documented in this section are provided by the racket/bool and racket libraries, but not racket/base.

true : boolean?

An alias for #t.

false : boolean?

An alias for #f.

(symbol=? a b) → boolean?
  a : symbol?
  b : symbol?

Returns (equal? a b) (if a and b are symbols).

(boolean=? a b) → boolean?
  a : boolean?
  b : boolean?

Returns (equal? a b) (if a and b are booleans).

(false? v) → boolean?
  v : any/c

Returns (not v).

(nand expr ...)

Same as (not (and expr ...)).

Examples:

```
> (nand #f #t)
#t
> (nand #f (error 'ack "we don't get here"))
#t
```

(nor expr ...)

Same as (not (or expr ...)).

In the two argument case, returns #t if neither of the arguments is a true value.

Examples:

```
> (nor #f #t)
#f
> (nor #t (error 'ack "we don't get here"))
#f
```

(implies expr1 expr2)

Checks to be sure that the first expression implies the second.

Same as (if expr1 expr2 #t).

Examples:

```
> (implies #f #t)
#t
> (implies #f #f)
#t
> (implies #t #f)
#f
> (implies #f (error 'ack "we don't get here"))
#t
```

(xor b1 b2) → any
  b1 : any/c
  b2 : any/c

Returns the exclusive or of *b1* and *b2*.

If exactly one of *b1* and *b2* is not `#f`, then return it. Otherwise, returns `#f`.

Examples:

```
> (xor 11 #f)
11
> (xor #f 22)
22
> (xor 11 22)
#f
> (xor #f #f)
#f
```

3.2 Numbers

§3.2 "Numbers" in ***The Racket Guide*** introduces numbers.

All *numbers* are *complex numbers*. Some of them are *real numbers*, and all of the real numbers that can be represented are also *rational numbers*, except for `+inf.0` (positive infinity), `+inf.f` (single-precision variant), `-inf.0` (negative infinity), `-inf.f` (single-precision variant), `+nan.0` (not-a-number), and `+nan.f` (single-precision variant). Among the rational numbers, some are *integers*, because `round` applied to the number produces the same number.

See §1.3.3 "Reading Numbers" for information on the syntax of number literals.

Orthogonal to those categories, each number is also either an *exact number* or an *inexact number*. Unless otherwise specified, computations that involve an inexact number produce inexact results. Certain operations on inexact numbers, however, produce an exact number, such as multiplying an inexact number with an exact `0`. Some operations, which can produce an irrational number for rational arguments (e.g., `sqrt`), may produce inexact results even for exact arguments.

In the case of complex numbers, either the real and imaginary parts are both exact or inexact, or the number has an exact zero real part and an inexact imaginary part; a complex number with an exact zero imaginary part is a real number.

Inexact real numbers are implemented as either single- or double-precision IEEE floating-point numbers—the latter by default, and the former only when a computation starts with numerical constants specified as single-precision numbers. Inexact real numbers that are represented as double-precision floating-point numbers are *flonums*.

The precision and size of exact numbers is limited only by available memory (and the precision of operations that can produce irrational numbers). In particular, adding, multiplying, subtracting, and dividing exact numbers always produces an exact result.

Inexact numbers can be coerced to exact form, except for the inexact numbers `+inf.0`, `+inf.f`, `-inf.0`, `-inf.f`, `+nan.0`, and `+nan.f`, which have no exact form. Dividing a

number by exact zero raises an exception; dividing a non-zero number other than +nan.0 or +nan.f by an inexact zero returns +inf.0, +inf.f, -inf.0 or -inf.f, depending on the sign and precision of the dividend. The +nan.0 value is not = to itself, but +nan.0 is eqv? to itself, and +nan.f is similarly eqv? but not = to itself. Conversely, (= 0.0 -0.0) is #t, but (eqv? 0.0 -0.0) is #f, and the same for 0.0f0 and -0.0f0 (which are single-precision variants). The datum -nan.0 refers to the same constant as +nan.0, and -nan.f is the same as +nan.f.

Calculations with infinites produce results consistent with IEEE double- or single-precision floating point where IEEE specifies the result; in cases where IEEE provides no specification, such as (angle +inf.0+inf.0i), the result corresponds to the limit approaching infinity, or +nan.0 or +nan.f if no such limit exists.

A *fixnum* is an exact integer whose two's complement representation fit into 31 bits on a 32-bit platform or 63 bits on a 64-bit platform; furthermore, no allocation is required when computing with fixnums. See also the racket/fixnum module, below.

Two fixnums that are = are also the same according to eq?. Otherwise, the result of eq? applied to two numbers is undefined, except that numbers produced by the default reader in read-syntax mode are interned and therefore eq? when they are eqv?.

Two numbers are eqv? when they are both inexact with the same precision or both exact, and when they are = (except for +nan.0, +nan.f, 0.0, 0.0f0, -0.0, and -0.0f0, as noted above). Two numbers are equal? when they are eqv?.

See §1.3.3 "Reading Numbers" for information on reading numbers and §1.4.2 "Printing Numbers" for information on printing numbers.

3.2.1 Number Types

```
(number? v) → boolean?
  v : any/c
```

Returns #t if v is a number, #f otherwise.

Examples:

```
> (number? 1)
#t
> (number? 2+3i)
#t
> (number? "hello")
#f
```

```
(complex? v) → boolean?
  v : any/c
```

Returns `(number? v)`, because all numbers are complex numbers.

```
(real? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a real number, `#f` otherwise.

Examples:

```
> (real? 1)
#t
> (real? +inf.0)
#t
> (real? 2+3i)
#f
> (real? 2.0+0.0i)
#f
> (real? "hello")
#f
```

```
(rational? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a rational number, `#f` otherwise.

Examples:

```
> (rational? 1)
#t
> (rational? +inf.0)
#f
> (rational? "hello")
#f
```

```
(integer? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a number that is an integer, `#f` otherwise.

Examples:

```
> (integer? 1)
#t
> (integer? 2.3)
#f
> (integer? 4.0)
#t
```

```
> (integer? +inf.0)
#f
> (integer? 2+3i)
#f
> (integer? "hello")
#f
```

```
(exact-integer? v) → boolean?
  v : any/c
```

Returns `(and (integer? v) (exact? v))`.

Examples:

```
> (exact-integer? 1)
#t
> (exact-integer? 4.0)
#f
```

```
(exact-nonnegative-integer? v) → boolean?
  v : any/c
```

Returns `(and (exact-integer? v) (not (negative? v)))`.

Examples:

```
> (exact-nonnegative-integer? 0)
#t
> (exact-nonnegative-integer? -1)
#f
```

```
(exact-positive-integer? v) → boolean?
  v : any/c
```

Returns `(and (exact-integer? v) (positive? v))`.

Examples:

```
> (exact-positive-integer? 1)
#t
> (exact-positive-integer? 0)
#f
```

```
(inexact-real? v) → boolean?
  v : any/c
```

Returns `(and (real? v) (inexact? v))`.

```
(fixnum? v) → boolean?
  v : any/c
```

Return `#t` if `v` is a fixnum, `#f` otherwise.

Note: the result of this function is platform-dependent, so using it in syntax transformers can lead to platform-dependent bytecode files.

```
(flonum? v) → boolean?
  v : any/c
```

Return `#t` if `v` is a flonum, `#f` otherwise.

```
(double-flonum? v) → boolean?
  v : any/c
```

Identical to `flonum?`.

```
(single-flonum? v) → boolean?
  v : any/c
```

Return `#t` if `v` is a single-precision floating-point number, `#f` otherwise.

```
(zero? z) → boolean?
  z : number?
```

Returns `(= 0 z)`.

Examples:

```
> (zero? 0)
#t
> (zero? -0.0)
#t
```

```
(positive? x) → boolean?
  x : real?
```

Returns `(> x 0)`.

Examples:

```
> (positive? 10)
#t
> (positive? -10)
#f
> (positive? 0.0)
#f
```

```
(negative? x) → boolean?
  x : real?
```

Returns `(< x 0)`.

Examples:

```
> (negative? 10)
#f
> (negative? -10)
#t
> (negative? -0.0)
#f
```

```
(even? n) → boolean?
  n : integer?
```

Returns `(zero? (modulo n 2))`.

Examples:

```
> (even? 10.0)
#t
> (even? 11)
#f
> (even? +inf.0)
even?: contract violation
  expected: integer
  given: +inf.0
```

```
(odd? n) → boolean?
  n : integer?
```

Returns `(not (even? n))`.

Examples:

```
> (odd? 10.0)
#f
> (odd? 11)
#t
> (odd? +inf.0)
odd?: contract violation
  expected: integer
  given: +inf.0
```

```
(exact? z) → boolean?
  z : number?
```

Returns `#t` if `z` is an exact number, `#f` otherwise.

Examples:

```
> (exact? 1)
#t
> (exact? 1.0)
#f
```

```
(inexact? z) → boolean?
  z : number?
```

Returns `#t` if `z` is an inexact number, `#f` otherwise.

Examples:

```
> (inexact? 1)
#f
> (inexact? 1.0)
#t
```

```
(inexact->exact z) → exact?
  z : number?
```

Coerces `z` to an exact number. If `z` is already exact, it is returned. If `z` is `+inf.0`, `-inf.0`, or `+nan.0`, then the `exn:fail:contract` exception is raised.

Examples:

```
> (inexact->exact 1)
1
> (inexact->exact 1.0)
1
```

```
(exact->inexact z) → inexact?
  z : number?
```

Coerces `z` to an inexact number. If `z` is already inexact, it is returned.

Examples:

```
> (exact->inexact 1)
1.0
> (exact->inexact 1.0)
1.0
```

```
(real->single-flonum x) → single-flonum?
  x : real?
```

Coerces `x` to a single-precision floating-point number. If `x` is already a single-precision floating-point number, it is returned.

```
(real->double-flonum x) → flonum?
  x : real?
```

Coerces `x` to a double-precision floating-point number. If `x` is already a double-precision floating-point number, it is returned.

3.2.2 Generic Numerics

Most Racket numeric operations work on any kind of number.

Arithmetic

```
(+ z ...) → number?
  z : number?
```

Returns the sum of the `z`s, adding pairwise from left to right. If no arguments are provided, the result is `0`.

Examples:

```
> (+ 1 2)
3
> (+ 1.0 2+3i 5)
8.0+3.0i
> (+)
0
```

```
(- z) → number?
  z : number?
(- z w ...+) → number?
  z : number?
  w : number?
```

When no `w`s are supplied, returns `(- 0 z)`. Otherwise, returns the subtraction of the `w`s from `z` working pairwise from left to right.

Examples:

```
> (- 5 3.0)
2.0
> (- 1)
-1
> (- 2+7i 1 3)
-2+7i
```

```
(* z ...) → number?
  z : number?
```

Returns the product of the `z`s, multiplying pairwise from left to right. If no arguments are provided, the result is `1`. Multiplying any number by exact `0` produces exact `0`.

Examples:

```
> (* 2 3)
6
> (* 8.0 9)
72.0
> (* 1+2i 3+4i)
-5+10i
```

```
(/ z) → number?
  z : number?
(/ z w ...+) → number?
  z : number?
  w : number?
```

When no `w`s are supplied, returns `(/ 1 z)`. Otherwise, returns the division of `z` by the `w`s working pairwise from left to right.

If `z` is exact `0` and no `w` is exact `0`, then the result is exact `0`. If any `w` is exact `0`, the `exn:fail:contract:divide-by-zero` exception is raised.

Examples:

```
> (/ 3 4)
3/4
> (/ 81 3 3)
9
> (/ 10.0)
0.1
> (/ 1+2i 3+4i)
11/25+2/25i
```

```
(quotient n m) → integer?
  n : integer?
  m : integer?
```

Returns `(truncate (/ n m))`.

Examples:

```
> (quotient 10 3)
3
> (quotient -10.0 3)
-3.0
> (quotient +inf.0 3)
quotient: contract violation
  expected: integer?
  given: +inf.0
  argument position: 1st
  other arguments...:
   3
```

```
(remainder n m) → integer?
  n : integer?
  m : integer?
```

Returns q with the same sign as n such that

- `(abs q)` is between `0` (inclusive) and `(abs m)` (exclusive), and
- `(+ q (* m (quotient n m)))` equals n.

If m is exact `0`, the `exn:fail:contract:divide-by-zero` exception is raised.

Examples:

```
> (remainder 10 3)
1
> (remainder -10.0 3)
-1.0
> (remainder 10.0 -3)
1.0
> (remainder -10 -3)
-1
> (remainder +inf.0 3)
remainder: contract violation
  expected: integer?
```

```
  given: +inf.0
  argument position: 1st
  other arguments...:
   3
```

```
(quotient/remainder n m) → integer? integer?
  n : integer?
  m : integer?
```

Returns `(values (quotient n m) (remainder n m))`, but the combination may be computed more efficiently than separate calls to `quotient` and `remainder`.

Example:

```
> (quotient/remainder 10 3)
3
1
```

```
(modulo n m) → integer?
  n : integer?
  m : integer?
```

Returns `q` with the same sign as `m` where

- `(abs q)` is between `0` (inclusive) and `(abs m)` (exclusive), and
- the difference between `q` and `(- n (* m (quotient n m)))` is a multiple of `m`.

If `m` is exact `0`, the `exn:fail:contract:divide-by-zero` exception is raised.

Examples:

```
> (modulo 10 3)
1
> (modulo -10.0 3)
2.0
> (modulo 10.0 -3)
-2.0
> (modulo -10 -3)
-1
> (modulo +inf.0 3)
modulo: contract violation
  expected: integer?
  given: +inf.0
  argument position: 1st
  other arguments...:
   3
```

```
(add1 z) → number?
  z : number?
```

Returns `(+ z 1)`.

```
(sub1 z) → number?
  z : number?
```

Returns `(- z 1)`.

```
(abs x) → number?
  x : real?
```

Returns the absolute value of `x`.

Examples:

```
> (abs 1.0)
1.0
> (abs -1)
1
```

```
(max x ...+) → real?
  x : real?
```

Returns the largest of the `x`s, or `+nan.0` if any `x` is `+nan.0`. If any `x` is inexact, the result is coerced to inexact.

Examples:

```
> (max 1 3 2)
3
> (max 1 3 2.0)
3.0
```

```
(min x ...+) → real?
  x : real?
```

Returns the smallest of the `x`s, or `+nan.0` if any `x` is `+nan.0`. If any `x` is inexact, the result is coerced to inexact.

Examples:

```
> (min 1 3 2)
1
> (min 1 3 2.0)
1.0
```

```
(gcd n ...) → rational?
  n : rational?
```

Returns the greatest common divisor (a non-negative number) of the *n*s; for non-integer *n*s, the result is the `gcd` of the numerators divided by the `lcm` of the denominators. If no arguments are provided, the result is `0`. If all arguments are zero, the result is zero.

Examples:

```
> (gcd 10)
10
> (gcd 12 81.0)
3.0
> (gcd 1/2 1/3)
1/6
```

```
(lcm n ...) → rational?
  n : rational?
```

Returns the least common multiple (a non-negative number) of the *n*s; non-integer *n*s, the result is the absolute value of the product divided by the `gcd`. If no arguments are provided, the result is `1`. If any argument is zero, the result is zero; furthermore, if any argument is exact `0`, the result is exact `0`.

Examples:

```
> (lcm 10)
10
> (lcm 3 4.0)
12.0
> (lcm 1/2 2/3)
2
```

```
(round x) → (or/c integer? +inf.0 -inf.0 +nan.0)
  x : real?
```

Returns the integer closest to `x`, resolving ties in favor of an even number, but `+inf.0`, `-inf.0`, and `+nan.0` round to themselves.

Examples:

```
> (round 17/4)
4
> (round -17/4)
-4
```

```
> (round 2.5)
2.0
> (round -2.5)
-2.0
> (round +inf.0)
+inf.0
```

```
(floor x) → (or/c integer? +inf.0 -inf.0 +nan.0)
  x : real?
```

Returns the largest integer that is no more than `x`, but `+inf.0`, `-inf.0`, and `+nan.0` floor to themselves.

Examples:

```
> (floor 17/4)
4
> (floor -17/4)
-5
> (floor 2.5)
2.0
> (floor -2.5)
-3.0
> (floor +inf.0)
+inf.0
```

```
(ceiling x) → (or/c integer? +inf.0 -inf.0 +nan.0)
  x : real?
```

Returns the smallest integer that is at least as large as `x`, but `+inf.0`, `-inf.0`, and `+nan.0` ceiling to themselves.

Examples:

```
> (ceiling 17/4)
5
> (ceiling -17/4)
-4
> (ceiling 2.5)
3.0
> (ceiling -2.5)
-2.0
> (ceiling +inf.0)
+inf.0
```

```
(truncate x) → (or/c integer? +inf.0 -inf.0 +nan.0)
  x : real?
```

Returns the integer farthest from 0 that is not farther from 0 than *x*, but +inf.0, -inf.0, and +nan.0 truncate to themselves.

Examples:

```
> (truncate 17/4)
4
> (truncate -17/4)
-4
> (truncate 2.5)
2.0
> (truncate -2.5)
-2.0
> (truncate +inf.0)
+inf.0
```

```
(numerator q) → integer?
  q : rational?
```

Coerces *q* to an exact number, finds the numerator of the number expressed in its simplest fractional form, and returns this number coerced to the exactness of *q*.

Examples:

```
> (numerator 5)
5
> (numerator 17/4)
17
> (numerator 2.3)
2589569785738035.0
```

```
(denominator q) → integer?
  q : rational?
```

Coerces *q* to an exact number, finds the numerator of the number expressed in its simplest fractional form, and returns this number coerced to the exactness of *q*.

Examples:

```
> (denominator 5)
1
> (denominator 17/4)
4
> (denominator 2.3)
1125899906842624.0
```

```
(rationalize x tolerance) → real?
  x : real?
  tolerance : real?
```

Among the real numbers within (abs *tolerance*) of *x*, returns the one corresponding to an exact number whose denominator is the smallest. If multiple integers are within *tolerance* of *x*, the one closest to 0 is used.

Examples:

```
> (rationalize 1/4 1/10)
1/3
> (rationalize -1/4 1/10)
-1/3
> (rationalize 1/4 1/4)
0
> (rationalize 11/40 1/4)
1/2
```

Number Comparison

```
(= z w ...+) → boolean?
  z : number?
  w : number?
```

Returns #t if all of the arguments are numerically equal, #f otherwise. An inexact number is numerically equal to an exact number when the exact coercion of the inexact number is the exact number. Also, 0.0 and -0.0 are numerically equal, but +nan.0 is not numerically equal to itself.

Examples:

```
> (= 1 1.0)
#t
> (= 1 2)
#f
> (= 2+3i 2+3i 2+3i)
#t
```

```
(< x y ...+) → boolean?
  x : real?
  y : real?
```

Returns #t if the arguments in the given order are strictly increasing, #f otherwise.

Examples:

```
> (< 1 1)
#f
```

```
> (< 1 2 3)
#t
> (< 1 +inf.0)
#t
> (< 1 +nan.0)
#f
```

```
(<= x y ...+) → boolean?
  x : real?
  y : real?
```

Returns `#t` if the arguments in the given order are non-decreasing, `#f` otherwise.

Examples:

```
> (<= 1 1)
#t
> (<= 1 2 1)
#f
```

```
(> x y ...+) → boolean?
  x : real?
  y : real?
```

Returns `#t` if the arguments in the given order are strictly decreasing, `#f` otherwise.

Examples:

```
> (> 1 1)
#f
> (> 3 2 1)
#t
> (> +inf.0 1)
#t
> (> +nan.0 1)
#f
```

```
(>= x y ...+) → boolean?
  x : real?
  y : real?
```

Returns `#t` if the arguments in the given order are non-increasing, `#f` otherwise.

Examples:

```
> (>= 1 1)
#t
> (>= 1 2 1)
#f
```

Powers and Roots

```
(sqrt z) → number?
  z : number?
```

Returns the principal square root of `z`. The result is exact if `z` is exact and `z`'s square root is rational. See also `integer-sqrt`.

Examples:

```
> (sqrt 4/9)
2/3
> (sqrt 2)
1.4142135623730951
> (sqrt -1)
0+1i
```

```
(integer-sqrt n) → complex?
  n : integer?
```

Returns `(floor (sqrt n))` for positive `n`. For negative `n`, the result is `(* (integer-sqrt (- n)) 0+1i)`.

Examples:

```
> (integer-sqrt 4.0)
2.0
> (integer-sqrt 5)
2
```

```
(integer-sqrt/remainder n) → integer? integer?
  n : integer?
```

Returns `(integer-sqrt n)` and `(- n (expt (integer-sqrt n) 2))`.

Examples:

```
> (integer-sqrt/remainder 4.0)
2.0
0.0
> (integer-sqrt/remainder 5)
2
1
```

```
(expt z w) → number?
  z : number?
  w : number?
```

Returns `z` raised to the power of `w`. If `w` is exact `0`, the result is exact `1`. If `z` is exact `0` and `w` is negative, the `exn:fail:contract:divide-by-zero` exception is raised.

Examples:

```
> (expt 2 3)
8
> (expt 4 0.5)
2.0
> (expt +inf.0 0)
1
```

```
(exp z) → number?
  z : number?
```

Returns Euler's number raised to the power of `z`. The result is normally inexact, but it is exact `1` when `z` is an exact `0`.

Examples:

```
> (exp 1)
2.718281828459045
> (exp 2+3i)
-7.315110094901103+1.0427436562359045i
> (exp 0)
1
```

```
(log z) → number?
  z : number?
```

Returns the natural logarithm of `z`. The result is normally inexact, but it is exact `0` when `z` is an exact `1`. When `z` is exact `0`, `exn:fail:contract:divide-by-zero` exception is raised.

Examples:

```
> (log (exp 1))
1.0
> (log 2+3i)
1.2824746787307684+0.982793723247329i
> (log 1)
0
```

Trigonometric Functions

```
(sin z) → number?
  z : number?
```

Returns the sine of z, where z is in radians. The result is normally inexact, but it is exact 0 if z is exact 0.

Examples:

```
> (sin 3.14159)
2.65358979335273e-06
> (sin 1.0+5.0i)
62.44551846769653+40.0921657779984i
```

```
(cos z) → number?
  z : number?
```

Returns the cosine of z, where z is in radians.

Examples:

```
> (cos 3.14159)
-0.9999999999964793
> (cos 1.0+5.0i)
40.095806306298826-62.43984868079963i
```

```
(tan z) → number?
  z : number?
```

Returns the tangent of z, where z is in radians. The result is normally inexact, but it is exact 0 if z is exact 0.

Examples:

```
> (tan 0.7854)
1.0000036732118496
> (tan 1.0+5.0i)
8.256719834227411e-05+1.0000377833796008i
```

```
(asin z) → number?
  z : number?
```

Returns the arcsine in radians of z. The result is normally inexact, but it is exact 0 if z is exact 0.

Examples:

```
> (asin 0.25)
0.25268025514207865
> (asin 1.0+5.0i)
0.1937931365549321+2.3309746530493123i
```

```
(acos z) → number?
  z : number?
```

Returns the arccosine in radians of `z`.

Examples:

```
> (acos 0.25)
1.318116071652818
> (acos 1.0+5.0i)
1.3770031902399644-2.3309746530493123i
```

```
(atan z) → number?
  z : number?
(atan y x) → number?
  y : real?
  x : real?
```

In the one-argument case, returns the arctangent of the inexact approximation of `z`, except that the result is an exact `0` for an exact `0` argument.

In the two-argument case, the result is roughly the same as `(atan (/ (exact->inexact y)) (exact->inexact x))`, but the signs of `y` and `x` determine the quadrant of the result. Moreover, a suitable angle is returned when `y` divided by `x` produces `+nan.0` in the case that neither `y` nor `x` is `+nan.0`. Finally, if `y` is exact `0` and `x` is an exact positive number, the result is exact `0`. If both `x` and `y` are exact `0`, the `exn:fail:contract:divide-by-zero` exception is raised.

Examples:

```
> (atan 0.5)
0.4636476090008061
> (atan 2 1)
1.1071487177940904
> (atan -2 -1)
-2.0344439357957027
> (atan 1.0+5.0i)
1.530881333938778+0.19442614214700213i
> (atan +inf.0 -inf.0)
2.356194490192345
```

Complex Numbers

```
(make-rectangular x y) → number?
  x : real?
  y : real?
```

Returns `(+ x (* y 0+1i))`.

Example:

```
> (make-rectangular 3 4.0)
3.0+4.0i
```

```
(make-polar magnitude angle) → number?
  magnitude : real?
  angle : real?
```

Returns `(+ (* magnitude (cos angle)) (* magnitude (sin angle) 0+1i))`.

Examples:

```
> (make-polar 10 (* pi 1/2))
6.123233995736766e-16+10.0i
> (make-polar 10 (* pi 1/4))
7.0710678118654755+7.071067811865475i
```

```
(real-part z) → real?
  z : number?
```

Returns the real part of the complex number `z` in rectangle coordinates.

Examples:

```
> (real-part 3+4i)
3
> (real-part 5.0)
5.0
```

```
(imag-part z) → real?
  z : number?
```

Returns the imaginary part of the complex number `z` in rectangle coordinates.

Examples:

```
> (imag-part 3+4i)
4
> (imag-part 5.0)
0
> (imag-part 5.0+0.0i)
0.0
```

```
(magnitude z) → (and/c real? (not/c negative?))
  z : number?
```

Returns the magnitude of the complex number *z* in polar coordinates.

Examples:

```
> (magnitude -3)
3
> (magnitude 3.0)
3.0
> (magnitude 3+4i)
5
```

```
(angle z) → real?
  z : number?
```

Returns the angle of the complex number *z* in polar coordinates.

The result is guaranteed to be between `(- pi)` and `pi`, possibly equal to `pi` (but never equal to `(- pi)`).

Examples:

```
> (angle -3)
3.141592653589793
> (angle 3.0)
0
> (angle 3+4i)
0.9272952180016122
> (angle +inf.0+inf.0i)
0.7853981633974483
> (angle -1)
3.141592653589793
```

Bitwise Operations

```
(bitwise-ior n ...) → exact-integer?
  n : exact-integer?
```

Returns the bitwise "inclusive or" of the *n*s in their (semi-infinite) two's complement representation. If no arguments are provided, the result is `0`.

Examples:

```
> (bitwise-ior 1 2)
3
> (bitwise-ior -32 1)
-31
```

```
(bitwise-and n ...) → exact-integer?
  n : exact-integer?
```

Returns the bitwise “and” of the `n`s in their (semi-infinite) two’s complement representation. If no arguments are provided, the result is `-1`.

Examples:

```
> (bitwise-and 1 2)
0
> (bitwise-and -32 -1)
-32
```

```
(bitwise-xor n ...) → exact-integer?
  n : exact-integer?
```

Returns the bitwise “exclusive or” of the `n`s in their (semi-infinite) two’s complement representation. If no arguments are provided, the result is `0`.

Examples:

```
> (bitwise-xor 1 5)
4
> (bitwise-xor -32 -1)
31
```

```
(bitwise-not n) → exact-integer?
  n : exact-integer?
```

Returns the bitwise “not” of `n` in its (semi-infinite) two’s complement representation.

Examples:

```
> (bitwise-not 5)
-6
> (bitwise-not -1)
0
```

```
(bitwise-bit-set? n m) → boolean?
  n : exact-integer?
  m : exact-nonnegative-integer?
```

Returns `#t` when the `m`th bit of `n` is set in `n`’s (semi-infinite) two’s complement representation.

This operation is equivalent to `(not (zero? (bitwise-and n (arithmetic-shift 1 m))))`, but it is faster and runs in constant time when `n` is positive.

Examples:

```
> (bitwise-bit-set? 5 0)
#t
> (bitwise-bit-set? 5 2)
#t
> (bitwise-bit-set? -5 (expt 2 700))
#t
```

```
(bitwise-bit-field n start end) → exact-integer?
  n : exact-integer?
  start : exact-nonnegative-integer?
  end : (and/c exact-nonnegative-integer?
               (start . <= . end))
```

Extracts the bits between position *start* and `(- end 1)` (inclusive) from *n* and shifts them down to the least significant portion of the number.

This operation is equivalent to the computation

```
(bitwise-and (sub1 (arithmetic-shift 1 (- end start)))
             (arithmetic-shift n (- start)))
```

but it runs in constant time when *n* is positive, *start* and *end* are fixnums, and `(- end start)` is no more than the maximum width of a fixnum.

Each pair of examples below uses the same numbers, showing the result both in binary and as integers.

Examples:

```
> (format "~b" (bitwise-bit-field (string->number "1101" 2) 1 1))
"0"
> (bitwise-bit-field 13 1 1)
0
> (format "~b" (bitwise-bit-field (string->number "1101" 2) 1 3))
"10"
> (bitwise-bit-field 13 1 3)
2
> (format "~b" (bitwise-bit-field (string->number "1101" 2) 1 4))
"110"
> (bitwise-bit-field 13 1 4)
6
```

```
(arithmetic-shift n m) → exact-integer?
  n : exact-integer?
  m : exact-integer?
```

Returns the bitwise "shift" of *n* in its (semi-infinite) two's complement representation. If *m* is non-negative, the integer *n* is shifted left by *m* bits; i.e., *m* new zeros are introduced as rightmost digits. If *m* is negative, *n* is shifted right by `(- m)` bits; i.e., the rightmost *m* digits are dropped.

Examples:

```
> (arithmetic-shift 1 10)
1024
> (arithmetic-shift 255 -3)
31
```

```
(integer-length n) → exact-integer?
  n : exact-integer?
```

Returns the number of bits in the (semi-infinite) two's complement representation of *n* after removing all leading zeros (for non-negative *n*) or ones (for negative *n*).

Examples:

```
> (integer-length 8)
4
> (integer-length -8)
3
```

Random Numbers

```
(random k [generator]) → exact-nonnegative-integer?
  k : (integer-in 1 4294967087)
  generator : pseudo-random-generator?
            = (current-pseudo-random-generator)
(random [generator]) → (and/c real? inexact? (>/c 0) (</c 1))
  generator : pseudo-random-generator?
            = (current-pseudo-random-generator)
```

When called with an integer argument *k*, returns a random exact integer in the range `0` to *k*-1. When called with zero arguments, returns a random inexact number between `0` and `1`, exclusive.

In each case, the number is provided by the given pseudo-random number generator (which defaults to the current one, as produced by `current-pseudo-random-generator`). The generator maintains an internal state for generating numbers. The random number generator uses a 54-bit version of L'Ecuyer's MRG32k3a algorithm [L'Ecuyer02].

```
(random-seed k) → void?
  k : (integer-in 1 (sub1 (expt 2 31)))
```

Seeds the current pseudo-random number generator with *k*. Seeding a generator sets its internal state deterministically; that is, seeding a generator with a particular number forces it to produce a sequence of pseudo-random numbers that is the same across runs and across platforms.

The random-seed function is convenient for some purposes, but note that the space of states for a pseudo-random number generator is much larger that the space of allowed values for *k*. Use vector->pseudo-random-generator! to set a pseudo-random number generator to any of its possible states.

```
(make-pseudo-random-generator) → pseudo-random-generator?
```

Returns a new pseudo-random number generator. The new generator is seeded with a number derived from (current-milliseconds).

```
(pseudo-random-generator? v) → boolean?
  v : any/c
```

Returns #t if *v* is a pseudo-random number generator, #f otherwise.

```
(current-pseudo-random-generator) → pseudo-random-generator?
(current-pseudo-random-generator generator) → void?
  generator : pseudo-random-generator?
```

A parameter that determines the pseudo-random number generator used by random.

```
(pseudo-random-generator->vector generator)
 → pseudo-random-generator-vector?
  generator : pseudo-random-generator?
```

Produces a vector that represents the complete internal state of *generator*. The vector is suitable as an argument to vector->pseudo-random-generator to recreate the generator in its current state (across runs and across platforms).

```
(vector->pseudo-random-generator vec)
 → pseudo-random-generator?
  vec : pseudo-random-generator-vector?
```

Produces a pseudo-random number generator whose internal state corresponds to *vec*.

```
(vector->pseudo-random-generator! generator
                                  vec)      → void?
  generator : pseudo-random-generator?
  vec : pseudo-random-generator-vector?
```

Like vector->pseudo-random-generator, but changes *generator* to the given state, instead of creating a new generator.

```
(pseudo-random-generator-vector? v) → boolean?
  v : any/c
```

Returns #t if *v* is a vector of six exact integers, where the first three integers are in the range 0 to 4294967086, inclusive; the last three integers are in the range 0 to 4294944442, inclusive; at least one of the first three integers is non-zero; and at least one of the last three integers is non-zero. Otherwise, the result is #f.

Number–String Conversions

```
(number->string z [radix]) → string?
  z : number?
  radix : (or/c 2 8 10 16) = 10
```

Returns a string that is the printed form of *z* in the base specified by *radix*. If *z* is inexact, *radix* must be 10, otherwise the exn:fail:contract exception is raised.

Examples:

```
> (number->string 3.0)
"3.0"
> (number->string 255 8)
"377"
```

```
(string->number s [radix]) → (or/c number? #f)
  s : string?
  radix : (integer-in 2 16) = 10
```

Reads and returns a number datum from *s* (see §1.3.3 “Reading Numbers”), returning #f if *s* does not parse exactly as a number datum (with no whitespace). The optional *radix* argument specifies the default base for the number, which can be overridden by #b, #o, #d, or #x in the string. The read-decimal-as-inexact parameter affects string->number in the same as way as read.

Examples:

```
> (string->number "3.0+2.5i")
3.0+2.5i
> (string->number "hello")
#f
> (string->number "111" 7)
57
> (string->number "#b111" 7)
7
```

```
(real->decimal-string n [decimal-digits]) → string?
  n : real?
  decimal-digits : exact-nonnegative-integer? = 2
```

Prints *n* into a string and returns the string. The printed form of *n* shows exactly *decimal-digits* digits after the decimal point. The printed form uses a minus sign if *n* is negative, and it does not use a plus sign if *n* is positive.

Before printing, *n* is converted to an exact number, multiplied by (expt 10 *decimal-digits*), rounded, and then divided again by (expt 10 *decimal-digits*). The result of this process is an exact number whose decimal representation has no more than *decimal-digits* digits after the decimal (and it is padded with trailing zeros if necessary).

Examples:

```
> (real->decimal-string pi)
"3.14"
> (real->decimal-string pi 5)
"3.14159"
```

```
(integer-bytes->integer  bstr
                         signed?
                        [big-endian?
                         start
                         end])        → exact-integer?
  bstr : bytes?
  signed? : any/c
  big-endian? : any/c = (system-big-endian?)
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (bytes-length bstr)
```

Converts the machine-format number encoded in *bstr* to an exact integer. The *start* and *end* arguments specify the substring to decode, where (- *end start*) must be 2, 4, or 8. If *signed?* is true, then the bytes are decoded as a two's-complement number, otherwise it is decoded as an unsigned integer. If *big-endian?* is true, then the first character's ASCII value provides the most significant eight bits of the number, otherwise the first character provides the least-significant eight bits, and so on.

```
(integer->integer-bytes  n
                         size-n
                         signed?
                        [big-endian?
                         dest-bstr
                         start])      → bytes?
  n : exact-integer?
```

```
  size-n : (or/c 2 4 8)
  signed? : any/c
  big-endian? : any/c = (system-big-endian?)
  dest-bstr : (and/c bytes? (not/c immutable?))
            = (make-bytes size-n)
  start : exact-nonnegative-integer? = 0
```

Converts the exact integer *n* to a machine-format number encoded in a byte string of length *size-n*, which must be 2, 4, or 8. If *signed?* is true, then the number is encoded as two's complement, otherwise it is encoded as an unsigned bit stream. If *big-endian?* is true, then the most significant eight bits of the number are encoded in the first character of the resulting byte string, otherwise the least-significant bits are encoded in the first byte, and so on.

The *dest-bstr* argument must be a mutable byte string of length *size-n*. The encoding of *n* is written into *dest-bstr* starting at offset *start*, and *dest-bstr* is returned as the result.

If *n* cannot be encoded in a string of the requested size and format, the `exn:fail:contract` exception is raised. If *dest-bstr* is not of length *size-n*, the `exn:fail:contract` exception is raised.

```
(floating-point-bytes->real bstr
                            [big-endian?
                             start
                             end])       → flonum?
  bstr : bytes?
  big-endian? : any/c = (system-big-endian?)
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (bytes-length bstr)
```

Converts the IEEE floating-point number encoded in *bstr* from position *start* (inclusive) to *end* (exclusive) to an inexact real number. The difference between *start* an *end* must be either 4 or 8 bytes. If *big-endian?* is true, then the first byte's ASCII value provides the most significant eight bits of the IEEE representation, otherwise the first byte provides the least-significant eight bits, and so on.

```
(real->floating-point-bytes x
                            size-n
                            [big-endian?
                             dest-bstr
                             start])     → bytes?
  x : real?
  size-n : (or/c 4 8)
  big-endian? : any/c = (system-big-endian?)
  dest-bstr : (and/c bytes? (not/c immutable?))
            = (make-bytes size-n)
```

```
  start : exact-nonnegative-integer? = 0
```

Converts the real number *x* to its IEEE representation in a byte string of length *size-n*, which must be 4 or 8. If *big-endian?* is true, then the most significant eight bits of the number are encoded in the first byte of the resulting byte string, otherwise the least-significant bits are encoded in the first character, and so on.

The *dest-bstr* argument must be a mutable byte string of length *size-n*. The encoding of n is written into *dest-bstr* starting with byte *start*, and *dest-bstr* is returned as the result.

If *dest-bstr* is provided and it has less than *start* plus *size-n* bytes, the exn:fail:contract exception is raised.

```
(system-big-endian?) → boolean?
```

Returns #t if the native encoding of numbers is big-endian for the machine running Racket, #f if the native encoding is little-endian.

Extra Constants and Functions

```
(require racket/math)
```

The bindings documented in this section are provided by the racket/math and racket libraries, but not racket/base.

```
pi : flonum?
```

An approximation of π, the ratio of a circle's circumference to its diameter.

Examples:

```
> pi
3.141592653589793
> (cos pi)
-1.0
```

```
pi.f : single-flonum?
```

Like pi, but in single precision.

Examples:

```
> pi.f
```

```
3.1415927f0
> (* 2.0f0 pi)
6.283185307179586
> (* 2.0f0 pi.f)
6.2831855f0
```

```
(degrees->radians x) → real?
  x : real?
```

Converts an x-degree angle to radians.

Examples:

```
> (degrees->radians 180)
3.141592653589793
> (sin (degrees->radians 45))
0.7071067811865475
```

```
(radians->degrees x) → real?
  x : real?
```

Converts x radians to degrees.

Examples:

```
> (radians->degrees pi)
180.0
> (radians->degrees (* 1/4 pi))
45.0
```

```
(sqr z) → number?
  z : number?
```

Returns `(* z z)`.

```
(sgn x) → (or/c 1 0 -1 1.0 0.0 -1.0)
  x : real?
```

Returns the sign of x as either -1, 0, or 1.

Examples:

```
> (sgn 10)
1
> (sgn -10.0)
-1.0
> (sgn 0)
0
```

```
(conjugate z) → number?
  z : number?
```

Returns the complex conjugate of `z`.

Examples:

```
> (conjugate 1)
1
> (conjugate 3+4i)
3-4i
```

```
(sinh z) → number?
  z : number?
```

Returns the hyperbolic sine of `z`.

```
(cosh z) → number?
  z : number?
```

Returns the hyperbolic cosine of `z`.

```
(tanh z) → number?
  z : number?
```

Returns the hyperbolic tangent of `z`.

```
(exact-round x) → exact-integer?
  x : rational?
```

Equivalent to `(inexact->exact (round x))`.

```
(exact-floor x) → exact-integer?
  x : rational?
```

Equivalent to `(inexact->exact (floor x))`.

```
(exact-ceiling x) → exact-integer?
  x : rational?
```

Equivalent to `(inexact->exact (ceiling x))`.

```
(exact-truncate x) → exact-integer?
  x : rational?
```

Equivalent to `(inexact->exact (truncate x))`.

```
(order-of-magnitude r) → (and/c exact? integer?)
  r : (and/c real? positive?)
```

Computes the greatest exact integer `m` such that:

```
(<= (expt 10 m)
    (inexact->exact r))
```

Hence also:

```
(< (inexact->exact r)
   (expt 10 (add1 m)))
```

Examples:

```
> (order-of-magnitude 999)
2
> (order-of-magnitude 1000)
3
> (order-of-magnitude 1/100)
-2
> (order-of-magnitude 1/101)
-3
```

```
(nan? x) → boolean?
  x : real?
```

Returns `#t` if `x` is `eqv?` to `+nan.0` or `+nan.f`; otherwise `#f`.

```
(infinite? x) → boolean?
  x : real?
```

Returns `#t` if `z` is `+inf.0`, `-inf.0`, `+inf.f`, `-inf.f`; otherwise `#f`.

3.2.3 Flonums

```
(require racket/flonum)
```

The `racket/flonum` library provides operations like `fl+` that consume and produce only flonums. Flonum-specific operations can provide better performance when used consistently, and they are as safe as generic operations like `+`.

See also §18.7 "Fixnum and Flonum Optimizations" in ***The Racket Guide***.

Flonum Arithmetic

```
(fl+ a b) → flonum?
  a : flonum?
  b : flonum?
(fl- a b) → flonum?
  a : flonum?
  b : flonum?
(fl* a b) → flonum?
  a : flonum?
  b : flonum?
(fl/ a b) → flonum?
  a : flonum?
  b : flonum?
(flabs a) → flonum?
  a : flonum?
```

Like `+`, `-`, `*`, `/`, and `abs`, but constrained to consume flonums. The result is always a flonum.

```
(fl= a b) → boolean?
  a : flonum?
  b : flonum?
(fl< a b) → boolean?
  a : flonum?
  b : flonum?
(fl> a b) → boolean?
  a : flonum?
  b : flonum?
(fl<= a b) → boolean?
  a : flonum?
  b : flonum?
(fl>= a b) → boolean?
  a : flonum?
  b : flonum?
(flmin a b) → flonum?
  a : flonum?
  b : flonum?
(flmax a b) → flonum?
  a : flonum?
  b : flonum?
```

Like `=`, `<`, `>`, `<=`, `>=`, `min`, and `max`, but constrained to consume flonums.

```
(flround a) → flonum?
  a : flonum?
```

```
(flfloor a) → flonum?
  a : flonum?
(flceiling a) → flonum?
  a : flonum?
(fltruncate a) → flonum?
  a : flonum?
```

Like `round`, `floor`, `ceiling`, and `truncate`, but constrained to consume flonums.

```
(flsin a) → flonum?
  a : flonum?
(flcos a) → flonum?
  a : flonum?
(fltan a) → flonum?
  a : flonum?
(flasin a) → flonum?
  a : flonum?
(flacos a) → flonum?
  a : flonum?
(flatan a) → flonum?
  a : flonum?
(fllog a) → flonum?
  a : flonum?
(flexp a) → flonum?
  a : flonum?
(flsqrt a) → flonum?
  a : flonum?
```

Like `sin`, `cos`, `tan`, `asin`, `acos`, `atan`, `log`, `exp`, and `sqrt`, but constrained to consume and produce flonums. The result is `+nan.0` when a number outside the range `-1.0` to `1.0` is given to `flasin` or `flacos`, or when a negative number is given to `fllog` or `flsqrt`.

```
(flexpt a b) → flonum?
  a : flonum?
  b : flonum?
```

Like `expt`, but constrained to consume and produce flonums. The result is `+nan.0` when `a` is negative and `b` is not an integer or when `a` is zero and `b` is not positive.

```
(->fl a) → flonum?
  a : exact-integer?
```

Like `exact->inexact`, but constrained to consume exact integers, so the result is always a flonum.

```
(fl->exact-integer a) → exact-integer?
  a : flonum?
```

Like `inexact->exact`, but constrained to consume an integer flonum, so the result is always an exact integer.

```
(make-flrectangular a b)
 → (and/c complex? inexact? (not/c real?))
  a : flonum?
  b : flonum?
(flreal-part a) → flonum?
  a : (and/c complex? inexact? (not/c real?))
(flimag-part a) → flonum?
  a : (and/c complex? inexact? (not/c real?))
```

Like `make-rectangular`, `real-part`, and `imag-part`, but both parts of the complex number must be inexact.

Flonum Vectors

A *flvector* is like a vector, but it holds only inexact real numbers. This representation can be more compact, and unsafe operations on flvectors (see `racket/unsafe/ops`) can execute more efficiently than unsafe operations on vectors of inexact reals.

An f64vector as provided by `ffi/vector` stores the same kinds of values as a flvector, but with extra indirections that make f64vectors more convenient for working with foreign libraries. The lack of indirections makes unsafe flvector access more efficient.

Two flvectors are `equal?` if they have the same length, and if the values in corresponding slots of the flvectors are `equal?`.

```
(flvector? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a flvector, `#f` otherwise.

```
(flvector x ...) → flvector?
  x : flonum?
```

Creates a flvector containing the given inexact real numbers.

Example:

```
> (flvector 2.0 3.0 4.0 5.0)
#<flvector>
```

```
(make-flvector size [x]) → flvector?
  size : exact-nonnegative-integer?
  x : flonum? = 0.0
```

Creates a flvector with *size* elements, where every slot in the flvector is filled with *x*.

Example:

```
> (make-flvector 4 3.0)
#<flvector>
```

```
(flvector-length vec) → exact-nonnegative-integer?
  vec : flvector?
```

Returns the length of *vec* (i.e., the number of slots in the flvector).

```
(flvector-ref vec pos) → flonum?
  vec : flvector?
  pos : exact-nonnegative-integer?
```

Returns the inexact real number in slot *pos* of *vec*. The first slot is position 0, and the last slot is one less than (flvector-length *vec*).

```
(flvector-set! vec pos x) → flonum?
  vec : flvector?
  pos : exact-nonnegative-integer?
  x : flonum?
```

Sets the inexact real number in slot *pos* of *vec*. The first slot is position 0, and the last slot is one less than (flvector-length *vec*).

```
(flvector-copy vec [start end]) → flvector?
  vec : flvector?
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (vector-length v)
```

Creates a fresh flvector of size (- *end start*), with all of the elements of *vec* from *start* (inclusive) to *end* (exclusive).

```
(in-flvector vec [start stop step]) → sequence?
  vec : flvector?
  start : exact-nonnegative-integer? = 0
  stop : (or/c exact-integer? #f) = #f
  step : (and/c exact-integer? (not/c zero?)) = 1
```

Returns a sequence equivalent to *vec* when no optional arguments are supplied.

The optional arguments *start*, *stop*, and *step* are as in in-vector.

A in-flvector application can provide better performance for flvector iteration when it appears directly in a for clause.

```
(for/flvector (for-clause ...) body ...)
(for/flvector #:length length-expr (for-clause ...) body ...)
(for*/flvector (for-clause ...) body ...)
(for*/flvector #:length length-expr (for-clause ...) body ...)
```

Like for/vector or for*/vector, but for flvectors.

```
(shared-flvector x ...) → flvector?
  x : flonum?
```

Creates a flvector containing the given inexact real numbers. For communication among places, the new flvector is allocated in the shared memory space.

Example:

```
> (shared-flvector 2.0 3.0 4.0 5.0)
#<flvector>
```

```
(make-shared-flvector size [x]) → flvector?
  size : exact-nonnegative-integer?
  x : flonum? = 0.0
```

Creates a flvector with *size* elements, where every slot in the flvector is filled with *x*. For communication among places, the new flvector is allocated in the shared memory space.

Example:

```
> (make-shared-flvector 4 3.0)
#<flvector>
```

3.2.4 Fixnums

```
(require racket/fixnum)
```

The racket/fixnum library provides operations like fx+ that consume and produce only fixnums. The operations in this library are meant to be safe versions of unsafe operations like unsafe-fx+. These safe operations are generally no faster than using generic primitives like +.

The expected use of the racket/fixnum library is for code where the require of racket/fixnum is replaced with

```
(require (filtered-in
          (λ (name) (regexp-replace #rx"unsafe-" name ""))
          racket/unsafe/ops))
```

to drop in unsafe versions of the library. Alternately, when encountering crashes with code that uses unsafe fixnum operations, use the `racket/fixnum` library to help debug the problems.

Fixnum Arithmetic

```
(fx+ a b) → fixnum?
  a : fixnum?
  b : fixnum?
(fx- a b) → fixnum?
  a : fixnum?
  b : fixnum?
(fx* a b) → fixnum?
  a : fixnum?
  b : fixnum?
(fxquotient a b) → fixnum?
  a : fixnum?
  b : fixnum?
(fxremainder a b) → fixnum?
  a : fixnum?
  b : fixnum?
(fxmodulo a b) → fixnum?
  a : fixnum?
  b : fixnum?
(fxabs a) → fixnum?
  a : fixnum?
```

Safe versions of `unsafe-fx+`, `unsafe-fx-`, `unsafe-fx*`, `unsafe-fxquotient`, `unsafe-fxremainder`, `unsafe-fxmodulo`, and `unsafe-fxabs`. The `exn:fail:contract:non-fixnum-result` exception is raised if the arithmetic result would not be a fixnum.

```
(fxand a b) → fixnum?
  a : fixnum?
  b : fixnum?
(fxior a b) → fixnum?
  a : fixnum?
  b : fixnum?
(fxxor a b) → fixnum?
  a : fixnum?
  b : fixnum?
```

```
(fxnot a) → fixnum?
  a : fixnum?
(fxlshift a b) → fixnum?
  a : fixnum?
  b : fixnum?
(fxrshift a b) → fixnum?
  a : fixnum?
  b : fixnum?
```

Safe versions of `unsafe-fxand`, `unsafe-fxior`, `unsafe-fxxor`, `unsafe-fxnot`, `unsafe-fxlshift`, and `unsafe-fxrshift`. The `exn:fail:contract:non-fixnum-result` exception is raised if the arithmetic result would not be a fixnum.

```
(fx= a b) → boolean?
  a : fixnum?
  b : fixnum?
(fx< a b) → boolean?
  a : fixnum?
  b : fixnum?
(fx> a b) → boolean?
  a : fixnum?
  b : fixnum?
(fx<= a b) → boolean?
  a : fixnum?
  b : fixnum?
(fx>= a b) → boolean?
  a : fixnum?
  b : fixnum?
(fxmin a b) → fixnum?
  a : fixnum?
  b : fixnum?
(fxmax a b) → fixnum?
  a : fixnum?
  b : fixnum?
```

Safe versions of `unsafe-fx=`, `unsafe-fx<`, `unsafe-fx>`, `unsafe-fx<=`, `unsafe-fx>=`, `unsafe-fxmin`, and `unsafe-fxmax`.

```
(fx->fl a) → flonum?
  a : fixnum?
(fl->fx a) → fixnum?
  a : flonum?
```

Safe versions of `unsafe-fx->fl` and `unsafe-fl->fx`.

Fixnum Vectors

A *fxvector* is like a vector, but it holds only fixnums. The only advantage of a fxvector over a vector is that a shared version can be created with functions like `shared-fxvector`.

Two fxvectors are `equal?` if they have the same length, and if the values in corresponding slots of the fxvectors are `equal?`.

```
(fxvector? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a fxvector, `#f` otherwise.

```
(fxvector x ...) → fxvector?
  x : fixnum?
```

Creates a fxvector containing the given fixnums.

Example:

```
> (fxvector 2 3 4 5)
#<fxvector>
```

```
(make-fxvector size [x]) → fxvector?
  size : exact-nonnegative-integer?
  x : fixnum? = 0
```

Creates a fxvector with *size* elements, where every slot in the fxvector is filled with *x*.

Example:

```
> (make-fxvector 4 3)
#<fxvector>
```

```
(fxvector-length vec) → exact-nonnegative-integer?
  vec : fxvector?
```

Returns the length of *vec* (i.e., the number of slots in the fxvector).

```
(fxvector-ref vec pos) → fixnum?
  vec : fxvector?
  pos : exact-nonnegative-integer?
```

Returns the fixnum in slot *pos* of *vec*. The first slot is position `0`, and the last slot is one less than `(fxvector-length vec)`.

```
(fxvector-set! vec pos x) → fixnum?
  vec : fxvector?
  pos : exact-nonnegative-integer?
  x : fixnum?
```

Sets the fixnum in slot *pos* of *vec*. The first slot is position 0, and the last slot is one less than (fxvector-length *vec*).

```
(fxvector-copy vec [start end]) → fxvector?
  vec : fxvector?
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (vector-length v)
```

Creates a fresh fxvector of size (- *end* *start*), with all of the elements of *vec* from *start* (inclusive) to *end* (exclusive).

```
(in-fxvector vec [start stop step]) → sequence?
  vec : fxvector?
  start : exact-nonnegative-integer? = 0
  stop : (or/c exact-integer? #f) = #f
  step : (and/c exact-integer? (not/c zero?)) = 1
```

Returns a sequence equivalent to *vec* when no optional arguments are supplied.

The optional arguments *start*, *stop*, and *step* are as in in-vector.

An in-fxvector application can provide better performance for fxvector iteration when it appears directly in a for clause.

```
(for/fxvector (for-clause ...) body ...)
(for/fxvector #:length length-expr (for-clause ...) body ...)
(for*/fxvector (for-clause ...) body ...)
(for*/fxvector #:length length-expr (for-clause ...) body ...)
```

Like for/vector or for*/vector, but for fxvectors.

```
(shared-fxvector x ...) → fxvector?
  x : fixnum?
```

Creates a fxvector containing the given fixnums. For communication among places, the new fxvector is allocated in the shared memory space.

Example:

```
> (shared-fxvector 2 3 4 5)
#<fxvector>
```

```
(make-shared-fxvector size [x]) → fxvector?
  size : exact-nonnegative-integer?
  x : fixnum? = 0
```

Creates a fxvector with *size* elements, where every slot in the fxvector is filled with *x*. For communication among places, the new fxvector is allocated in the shared memory space.

Example:

```
> (make-shared-fxvector 4 3)
#<fxvector>
```

3.3 Strings

§3.4 “Strings (Unicode)” in ***The Racket Guide*** introduces strings.

A *string* is a fixed-length array of characters.

A string can be *mutable* or *immutable*. When an immutable string is provided to a procedure like `string-set!`, the `exn:fail:contract` exception is raised. String constants generated by the default reader (see §1.3.6 “Reading Strings”) are immutable, and they are interned in `read-syntax` mode.

Two strings are `equal?` when they have the same length and contain the same sequence of characters.

A string can be used as a single-valued sequence (see §3.14.1 “Sequences”). The characters of the string serve as elements of the sequence. See also `in-string`.

See §1.3.6 “Reading Strings” for information on `read`ing strings and §1.4.5 “Printing Strings” for information on `print`ing strings.

See also: `immutable?`, `symbol->string`, `bytes->string/utf-8`.

3.3.1 String Constructors, Selectors, and Mutators

```
(string? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a string, `#f` otherwise.

Examples:

```
> (string? "Apple")
#t
> (string? 'apple)
#f
```

```
(make-string k [char]) → string?
  k : exact-nonnegative-integer?
  char : char? = #\nul
```

Returns a new mutable string of length *k* where each position in the string is initialized with the character *char*.

Example:

```
> (make-string 5 #\z)
"zzzzz"
```

```
(string char ...) → string?
  char : char?
```

Returns a new mutable string whose length is the number of provided *char*s, and whose positions are initialized with the given *char*s.

Example:

```
> (string #\A #\p #\p #\l #\e)
"Apple"
```

```
(string->immutable-string str) → (and/c string? immutable?)
  str : string?
```

Returns an immutable string with the same content as *str*, returning *str* itself if *str* is immutable.

```
(string-length str) → exact-nonnegative-integer?
  str : string?
```

Returns the length of *str*.

Example:

```
> (string-length "Apple")
5
```

```
(string-ref str k) → char?
  str : string?
  k : exact-nonnegative-integer?
```

Returns the character at position *k* in *str*. The first position in the string corresponds to 0, so the position *k* must be less than the length of the string, otherwise the exn:fail:contract exception is raised.

Example:

```
> (string-ref "Apple" 0)
#\A
```

```
(string-set! str k char) → void?
  str : (and/c string? (not/c immutable?))
  k : exact-nonnegative-integer?
  char : char?
```

Changes the character position *k* in *str* to *char*. The first position in the string corresponds to 0, so the position *k* must be less than the length of the string, otherwise the exn:fail:contract exception is raised.

Examples:

```
> (define s (string #\A #\p #\p #\l #\e))

> (string-set! s 4 #\y)

> s
"Apply"
```

```
(substring str start [end]) → string?
  str : string?
  start : exact-nonnegative-integer?
  end : exact-nonnegative-integer? = (string-length str)
```

Returns a new mutable string that is (- *end* *start*) characters long, and that contains the same characters as *str* from *start* inclusive to *end* exclusive. The first position in a string corresponds to 0, so the *start* and *end* arguments so they must be less than or equal to the length of *str*, and *end* must be greater than or equal to *start*, otherwise the exn:fail:contract exception is raised.

Examples:

```
> (substring "Apple" 1 3)
"pp"
> (substring "Apple" 1)
"pple"
```

```
(string-copy str) → string?
  str : string?
```

Returns (substring str 0).

```
(string-copy! dest
              dest-start
              src
             [src-start
              src-end])  → void?
  dest : (and/c string? (not/c immutable?))
  dest-start : exact-nonnegative-integer?
  src : string?
  src-start : exact-nonnegative-integer? = 0
  src-end : exact-nonnegative-integer? = (string-length src)
```

Changes the characters of *dest* starting at position *dest-start* to match the characters in *src* from *src-start* (inclusive) to *src-end* (exclusive), where the first position in a string corresponds to 0. The strings *dest* and *src* can be the same string, and in that case the destination region can overlap with the source region; the destination characters after the copy match the source characters from before the copy. If any of *dest-start*, *src-start*, or *src-end* are out of range (taking into account the sizes of the strings and the source and destination regions), the exn:fail:contract exception is raised.

Examples:

```
> (define s (string #\A #\p #\p #\l #\e))

> (string-copy! s 4 "y")

> (string-copy! s 0 s 3 4)

> s
"lpply"
```

```
(string-fill! dest char) → void?
  dest : (and/c string? (not/c immutable?))
  char : char?
```

Changes *dest* so that every position in the string is filled with *char*.

Examples:

```
> (define s (string #\A #\p #\p #\l #\e))

> (string-fill! s #\q)

> s
"qqqqq"
```

```
(string-append str ...) → string?
  str : string?
```

Returns a new mutable string that is as long as the sum of the given *str*s' lengths, and that contains the concatenated characters of the given *str*s. If no *str*s are provided, the result is a zero-length string.

Example:

```
> (string-append "Apple" "Banana")
"AppleBanana"
```

```
(string->list str) → (listof char?)
  str : string?
```

Returns a new list of characters corresponding to the content of *str*. That is, the length of the list is `(string-length str)`, and the sequence of characters in *str* is the same sequence in the result list.

Example:

```
> (string->list "Apple")
'(#\A #\p #\p #\l #\e)
```

```
(list->string lst) → string?
  lst : (listof char?)
```

Returns a new mutable string whose content is the list of characters in *lst*. That is, the length of the string is `(length lst)`, and the sequence of characters in *lst* is the same sequence in the result string.

Example:

```
> (list->string (list #\A #\p #\p #\l #\e))
"Apple"
```

```
(build-string n proc) → string?
  n : exact-nonnegative-integer?
  proc : (exact-nonnegative-integer? . -> . char?)
```

Creates a string of *n* characters by applying *proc* to the integers from `0` to `(sub1 n)` in order. If *str* is the resulting string, then `(string-ref str i)` is the character produced by `(proc i)`.

Example:

```
> (build-string 5 (lambda (i) (integer->char (+ i 97))))
"abcde"
```

3.3.2 String Comparisons

```
(string=? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Returns `#t` if all of the arguments are `equal?`.

Examples:

```
> (string=? "Apple" "apple")
#f
> (string=? "a" "as" "a")
#f
```

```
(string<? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Returns `#t` if the arguments are lexicographically sorted increasing, where individual characters are ordered by `char<?`, `#f` otherwise.

Examples:

```
> (string<? "Apple" "apple")
#t
> (string<? "apple" "Apple")
#f
> (string<? "a" "b" "c")
#t
```

```
(string<=? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like `string<?`, but checks whether the arguments are nondecreasing.

Examples:

```
> (string<=? "Apple" "apple")
#t
> (string<=? "apple" "Apple")
#f
> (string<=? "a" "b" "b")
#t
```

```
(string>? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like string<?, but checks whether the arguments are decreasing.

Examples:

```
> (string>? "Apple" "apple")
#f
> (string>? "apple" "Apple")
#t
> (string>? "c" "b" "a")
#t
```

```
(string>=? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like string<?, but checks whether the arguments are nonincreasing.

Examples:

```
> (string>=? "Apple" "apple")
#f
> (string>=? "apple" "Apple")
#t
> (string>=? "c" "b" "b")
#t
```

```
(string-ci=? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Returns #t if all of the arguments are eqv? after locale-insensitive case-folding via string-foldcase.

Examples:

```
> (string-ci=? "Apple" "apple")
#t
> (string-ci=? "a" "a" "a")
#t
```

```
(string-ci<? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like `string<?`, but checks whether the arguments would be in increasing order if each was first case-folded using `string-foldcase` (which is locale-insensitive).

Examples:

```
> (string-ci<? "Apple" "apple")
#f
> (string-ci<? "apple" "banana")
#t
> (string-ci<? "a" "b" "c")
#t
```

```
(string-ci<=? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like `string-ci<?`, but checks whether the arguments would be nondecreasing after case-folding.

Examples:

```
> (string-ci<=? "Apple" "apple")
#t
> (string-ci<=? "apple" "Apple")
#t
> (string-ci<=? "a" "b" "b")
#t
```

```
(string-ci>? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like `string-ci<?`, but checks whether the arguments would be decreasing after case-folding.

Examples:

```
> (string-ci>? "Apple" "apple")
#f
> (string-ci>? "banana" "Apple")
#t
> (string-ci>? "c" "b" "a")
#t
```

```
(string-ci>=? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like string-ci<?, but checks whether the arguments would be nonincreasing after case-folding.

Examples:

```
> (string-ci>=? "Apple" "apple")
#t
> (string-ci>=? "apple" "Apple")
#t
> (string-ci>=? "c" "b" "b")
#t
```

3.3.3 String Conversions

(string-upcase *str*) → string?
  str : string?

Returns a string whose characters are the upcase conversion of the characters in *str*. The conversion uses Unicode's locale-independent conversion rules that map code-point sequences to code-point sequences (instead of simply mapping a 1-to-1 function on code points over the string), so the string produced by the conversion can be longer than the input string.

Examples:

```
> (string-upcase "abc!")
"ABC!"
> (string-upcase "Straße")
"STRASSE"
```

(string-downcase *string*) → string?
  string : string?

Like string-upcase, but the downcase conversion.

Examples:

```
> (string-downcase "aBC!")
"abc!"
> (string-downcase "Straße")
"straße"
> (string-downcase "KAOΣ")
"καοσ"
> (string-downcase "Σ")
"σ"
```

(string-titlecase *string*) → string?
  string : string?

Like string-upcase, but the titlecase conversion only for the first character in each sequence of cased characters in str (ignoring case-ignorable characters).

Examples:

```
> (string-titlecase "aBC  tw0")
"Abc  Two"
> (string-titlecase "y2k")
"Y2K"
> (string-titlecase "main straße")
"Main Straße"
> (string-titlecase "stra ße")
"Stra Sse"
```

(string-foldcase *string*) → string?
  string : string?

Like string-upcase, but the case-folding conversion.

Examples:

```
> (string-foldcase "aBC!")
"abc!"
> (string-foldcase "Straße")
"strasse"
> (string-foldcase "KA0Σ")
"κα0σ"
```

(string-normalize-nfd *string*) → string?
  string : string?

Returns a string that is the Unicode normalized form D of *string*. If the given string is already in the corresponding Unicode normal form, the string may be returned directly as the result (instead of a newly allocated string).

(string-normalize-nfkd *string*) → string?
  string : string?

Like string-normalize-nfd, but for normalized form KD.

(string-normalize-nfc *string*) → string?
  string : string?

Like string-normalize-nfd, but for normalized form C.

(string-normalize-nfkc *string*) → string?
  string : string?

Like string-normalize-nfd, but for normalized form KC.

3.3.4 Locale-Specific String Operations

```
(string-locale=? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like `string=?`, but the strings are compared in a locale-specific way, based on the value of `current-locale`. See §12.1.1 “Encodings and Locales” for more information on locales.

```
(string-locale<? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like `string<?`, but the sort order compares strings in a locale-specific way, based on the value of `current-locale`. In particular, the sort order may not be simply a lexicographic extension of character ordering.

```
(string-locale>? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like `string>?`, but locale-specific like `string-locale<?`.

```
(string-locale-ci=? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like `string-locale=?`, but strings are compared using rules that are both locale-specific and case-insensitive (depending on what “case-insensitive” means for the current locale).

```
(string-locale-ci<? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like `string<?`, but both locale-sensitive and case-insensitive like `string-locale-ci=?`.

```
(string-locale-ci>? str1 str2 ...+) → boolean?
  str1 : string?
  str2 : string?
```

Like `string>?`, but both locale-sensitive and case-insensitive like `string-locale-ci=?`.

```
(string-locale-upcase string) → string?
  string : string?
```

Like string-upcase, but using locale-specific case-conversion rules based on the value of current-locale.

```
(string-locale-downcase string) → string?
  string : string?
```

Like string-downcase, but using locale-specific case-conversion rules based on the value of current-locale.

3.3.5 Additional String Functions

```
(require racket/string)
```

The bindings documented in this section are provided by the racket/string and racket libraries, but not racket/base.

```
(string-append* str ... strs) → string?
  str : string?
  strs : (listof string?)
```

Like string-append, but the last argument is used as a list of arguments for string-append, so (string-append* str ... strs) is the same as (apply string-append str ... strs). In other words, the relationship between string-append and string-append* is similar to the one between list and list*.

Examples:

```
> (string-append* "a" "b" '("c" "d"))
"abcd"
> (string-append* (cdr (append* (map (lambda (x) (list ", " x))
                                     '("Alpha" "Beta" "Gamma")))))
"Alpha, Beta, Gamma"
```

```
(string-join strs
             [sep
              #:before-first before-first
              #:before-last before-last
              #:after-last after-last])    → string?
  strs : (listof string?)
  sep : string? = " "
  before-first : string? = ""
  before-last : string? = sep
  after-last : string? = ""
```

Appends the strings in strs, inserting sep between each pair of strings in strs. before-last, before-first, and after-last are analogous to the inputs of add-between: they

specify an alternate separator between the last two strings, a prefix string, and a suffix string respectively.

Examples:

```
> (string-join '("one" "two" "three" "four"))
"one two three four"
> (string-join '("one" "two" "three" "four") ", ")
"one, two, three, four"
> (string-join '("one" "two" "three" "four") " potato ")
"one potato two potato three potato four"
> (string-join '("x" "y" "z") ", "
               #:before-first "Todo: "
               #:before-last " and "
               #:after-last ".")
"Todo: x, y and z."
```

```
(string-normalize-spaces str
                         [sep
                          space
                          #:trim? trim?
                          #:repeat? repeat?]) → string?
  str : string?
  sep : (or/c string? regexp?) = #px"\\s+"
  space : string? = " "
  trim? : any/c = #t
  repeat? : any/c = #f
```

Normalizes spaces in the input *str* by trimming it (using string-trim and *sep*) and replacing all whitespace sequences in the result with *space*, which defaults to a single space.

Example:

```
> (string-normalize-spaces "  foo bar  baz \r\n\t")
"foo bar baz"
```

The result of (string-normalize-spaces *str sep space*) is the same as (string-join (string-split *str sep* ....) *space*).

```
(string-replace str from to [#:all all?]) → string?
  str : string?
  from : (or/c string? regexp?)
  to : string?
  all? : any/c = #t
```

Returns *str* with all occurrences of *from* replaced with by *to*. If *from* is a string, it is matched literally (as opposed to being used as a regular expression).

By default, all occurrences are replaced, but only the first match is replaced if *all?* is #f.

Example:

```
> (string-replace "foo bar baz" "bar" "blah")
"foo blah baz"
```

```
(string-split  str
               [sep
                #:trim? trim?
                #:repeat? repeat?]) → (listof string?)
  str : string?
  sep : (or/c string? regexp?) = #px"\\s+"
  trim? : any/c = #t
  repeat? : any/c = #f
```

Splits the input *str* on whitespaces, returning a list of substrings of *str* that are separated by *sep*. The input is first trimmed using *sep* (see string-trim), unless *trim?* is #f. Empty matches are handled in the same way as for regexp-split. As a special case, if *str* is the empty string after trimming, the result is '() instead of '("").

Like string-trim, provide *sep* to use a different separator, and *repeat?* controls matching repeated sequences.

Examples:

```
> (string-split "  foo bar  baz \r\n\t")
'("foo" "bar" "baz")
> (string-split "  ")
'()
> (string-split "  " #:trim? #f)
'("" "")
```

```
(string-trim  str
              [sep
               #:left? left?
               #:right? right?
               #:repeat? repeat?]) → string?
  str : string?
  sep : (or/c string? regexp?) = #px"\\s+"
  left? : any/c = #t
  right? : any/c = #t
  repeat? : any/c = #f
```

Trims the input *str* by removing prefix and suffix *sep*, which defaults to whitespace. A string *sep* is matched literally (as opposed to being used as a regular expression).

Use `#:left? #f` or `#:right? #f` to suppress trimming the corresponding side. When *repeat?* is `#f` (the default), only one match is removed from each side; when *repeat?* is true, all initial or trailing matches are trimmed (which is an alternative to using a regular expression *sep* that contains `+`).

Examples:

```
> (string-trim "  foo bar  baz \r\n\t")
"foo bar  baz"
> (string-trim "  foo bar  baz \r\n\t" " " #:repeat? #t)
"foo bar  baz \r\n\t"
> (string-trim "aaaxaayaa" "aa")
"axaay"
```

3.4 Byte Strings

§3.5 “Bytes and Byte Strings” in ***The Racket Guide*** introduces byte strings.

A *byte string* is a fixed-length array of bytes. A *byte* is an exact integer between `0` and `255` inclusive.

A byte string can be *mutable* or *immutable*. When an immutable byte string is provided to a procedure like `bytes-set!`, the `exn:fail:contract` exception is raised. Byte-string constants generated by the default reader (see §1.3.6 “Reading Strings”) are immutable, and they are interned in `read-syntax` mode.

Two byte strings are `equal?` when they have the same length and contain the same sequence of bytes.

A byte string can be used as a single-valued sequence (see §3.14.1 “Sequences”). The bytes of the string serve as elements of the sequence. See also `in-bytes`.

See §1.3.6 “Reading Strings” for information on `read`ing byte strings and §1.4.5 “Printing Strings” for information on `print`ing byte strings.

See also: `immutable?`.

3.4.1 Byte String Constructors, Selectors, and Mutators

```
(bytes? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a byte string, `#f` otherwise.

Examples:

```
> (bytes? #"Apple")
#t
> (bytes? "Apple")
#f
```

```
(make-bytes k [b]) → bytes?
  k : exact-nonnegative-integer?
  b : byte? = 0
```

Returns a new mutable byte string of length *k* where each position in the byte string is initialized with the byte *b*.

Example:

```
> (make-bytes 5 65)
#"AAAAA"
```

```
(bytes b ...) → bytes?
  b : byte?
```

Returns a new mutable byte string whose length is the number of provided *b*s, and whose positions are initialized with the given *b*s.

Example:

```
> (bytes 65 112 112 108 101)
#"Apple"
```

```
(bytes->immutable-bytes bstr) → (and/c bytes? immutable?)
  bstr : bytes?
```

Returns an immutable byte string with the same content as *bstr*, returning *bstr* itself if *bstr* is immutable.

Examples:

```
> (bytes->immutable-bytes (bytes 65 65 65))
#"AAA"
> (define b (bytes->immutable-bytes (make-bytes 5 65)))

> (bytes->immutable-bytes b)
#"AAAAA"
> (eq? (bytes->immutable-bytes b) b)
#t
```

```
(byte? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a byte (i.e., an exact integer between `0` and `255` inclusive), `#f` otherwise.

Examples:

```
> (byte? 65)
#t
> (byte? 0)
#t
> (byte? 256)
#f
> (byte? -1)
#f
```

```
(bytes-length bstr) → exact-nonnegative-integer?
  bstr : bytes?
```

Returns the length of *bstr*.

Example:

```
> (bytes-length #"Apple")
5
```

```
(bytes-ref bstr k) → byte?
  bstr : bytes?
  k : exact-nonnegative-integer?
```

Returns the character at position *k* in *bstr*. The first position in the bytes coore-sponds to `0`, so the position *k* must be less than the length of the bytes, otherwise the `exn:fail:contract` exception is raised.

Example:

```
> (bytes-ref #"Apple" 0)
65
```

```
(bytes-set! bstr k b) → void?
  bstr : (and/c bytes? (not/c immutable?))
  k : exact-nonnegative-integer?
  b : byte?
```

Changes the character position *k* in *bstr* to *b*. The first position in the byte string coore-sponds to `0`, so the position *k* must be less than the length of the bytes, otherwise the `exn:fail:contract` exception is raised.

Examples:

```
> (define s (bytes 65 112 112 108 101))
```

```
> (bytes-set! s 4 121)

> s
#"Apply"
```

```
(subbytes bstr start [end]) → bytes?
  bstr : bytes?
  start : exact-nonnegative-integer?
  end : exact-nonnegative-integer? = (bytes-length str)
```

Returns a new mutable byte string that is `(- end start)` bytes long, and that contains the same bytes as *bstr* from *start* inclusive to *end* exclusive. The *start* and *end* arguments must be less than or equal to the length of *bstr*, and *end* must be greater than or equal to *start*, otherwise the `exn:fail:contract` exception is raised.

Examples:

```
> (subbytes #"Apple" 1 3)
#"pp"
> (subbytes #"Apple" 1)
#"pple"
```

```
(bytes-copy bstr) → bytes?
  bstr : bytes?
```

Returns `(subbytes str 0)`.

```
(bytes-copy!  dest
              dest-start
              src
             [src-start
              src-end])   → void?
  dest : (and/c bytes? (not/c immutable?))
  dest-start : exact-nonnegative-integer?
  src : bytes?
  src-start : exact-nonnegative-integer? = 0
  src-end : exact-nonnegative-integer? = (bytes-length src)
```

Changes the bytes of *dest* starting at position *dest-start* to match the bytes in *src* from *src-start* (inclusive) to *src-end* (exclusive). The bytes strings *dest* and *src* can be the same byte string, and in that case the destination region can overlap with the source region; the destination bytes after the copy match the source bytes from before the copy. If any of *dest-start*, *src-start*, or *src-end* are out of range (taking into account the sizes of the bytes strings and the source and destination regions), the `exn:fail:contract` exception is raised.

Examples:

```
> (define s (bytes 65 112 112 108 101))

> (bytes-copy! s 4 #"y")

> (bytes-copy! s 0 s 3 4)

> s
#"lpply"
```

(bytes-fill! *dest b*) → void?
  dest : (and/c bytes? (not/c immutable?))
  b : byte?

Changes *dest* so that every position in the bytes is filled with *b*.

Examples:

```
> (define s (bytes 65 112 112 108 101))

> (bytes-fill! s 113)

> s
#"qqqqq"
```

(bytes-append *bstr* ...) → bytes?
  bstr : bytes?

Returns a new mutable byte string that is as long as the sum of the given *bstr*s' lengths, and that contains the concatenated bytes of the given *bstr*s. If no *bstr*s are provided, the result is a zero-length byte string.

Example:

```
> (bytes-append #"Apple" #"Banana")
#"AppleBanana"
```

(bytes->list *bstr*) → (listof byte?)
  bstr : bytes?

Returns a new list of bytes corresponding to the content of *bstr*. That is, the length of the list is (bytes-length *bstr*), and the sequence of bytes in *bstr* is the same sequence in the result list.

Example:

```
> (bytes->list #"Apple")
'(65 112 112 108 101)
```

```
(list->bytes lst) → bytes?
  lst : (listof byte?)
```

Returns a new mutable byte string whose content is the list of bytes in `lst`. That is, the length of the byte string is `(length lst)`, and the sequence of bytes in `lst` is the same sequence in the result byte string.

Example:

```
> (list->bytes (list 65 112 112 108 101))
#"Apple"
```

```
(make-shared-bytes k [b]) → bytes?
  k : exact-nonnegative-integer?
  b : byte? = 0
```

Returns a new mutable byte string of length `k` where each position in the byte string is initialized with the byte `b`. For communication among places, the new byte string is allocated in the shared memory space.

Example:

```
> (make-shared-bytes 5 65)
#"AAAAA"
```

```
(shared-bytes b ...) → bytes?
  b : byte?
```

Returns a new mutable byte string whose length is the number of provided `b`s, and whose positions are initialized with the given `b`s. For communication among places, the new byte string is allocated in the shared memory space.

Example:

```
> (shared-bytes 65 112 112 108 101)
#"Apple"
```

3.4.2 Byte String Comparisons

```
(bytes=? bstr1 bstr2 ...+) → boolean?
  bstr1 : bytes?
  bstr2 : bytes?
```

Returns #t if all of the arguments are eqv?.

Examples:

```
> (bytes=? #"Apple" #"apple")
#f
> (bytes=? #"a" #"as" #"a")
#f
```

```
(bytes<? bstr1 bstr2 ...+) → boolean?
  bstr1 : bytes?
  bstr2 : bytes?
```

Returns #t if the arguments are lexicographically sorted increasing, where individual bytes are ordered by <, #f otherwise.

Examples:

```
> (bytes<? #"Apple" #"apple")
#t
> (bytes<? #"apple" #"Apple")
#f
> (bytes<? #"a" #"b" #"c")
#t
```

```
(bytes>? bstr1 bstr2 ...+) → boolean?
  bstr1 : bytes?
  bstr2 : bytes?
```

Like bytes<?, but checks whether the arguments are decreasing.

Examples:

```
> (bytes>? #"Apple" #"apple")
#f
> (bytes>? #"apple" #"Apple")
#t
> (bytes>? #"c" #"b" #"a")
#t
```

3.4.3 Bytes to/from Characters, Decoding and Encoding

```
(bytes->string/utf-8 bstr [err-char start end]) → string?
  bstr : bytes?
  err-char : (or/c #f char?) = #f
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (bytes-length bstr)
```

Produces a string by decoding the *start* to *end* substring of *bstr* as a UTF-8 encoding of Unicode code points. If *err-char* is not `#f`, then it is used for bytes that fall in the range `128` to `255` but are not part of a valid encoding sequence. (This rule is consistent with reading characters from a port; see §12.1.1 “Encodings and Locales” for more details.) If *err-char* is `#f`, and if the *start* to *end* substring of *bstr* is not a valid UTF-8 encoding overall, then the `exn:fail:contract` exception is raised.

Example:

```
> (bytes->string/utf-8 (bytes 195 167 195 176 195 182 194 163))
"çðö£"
```

```
(bytes->string/locale  bstr
                       [err-char
                        start
                        end])     → string?
  bstr : bytes?
  err-char : (or/c #f char?) = #f
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (bytes-length bstr)
```

Produces a string by decoding the *start* to *end* substring of *bstr* using the current locale’s encoding (see also §12.1.1 “Encodings and Locales”). If *err-char* is not `#f`, it is used for each byte in *bstr* that is not part of a valid encoding; if *err-char* is `#f`, and if the *start* to *end* substring of *bstr* is not a valid encoding overall, then the `exn:fail:contract` exception is raised.

```
(bytes->string/latin-1  bstr
                        [err-char
                         start
                         end])     → string?
  bstr : bytes?
  err-char : (or/c #f char?) = #f
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (bytes-length bstr)
```

Produces a string by decoding the *start* to *end* substring of *bstr* as a Latin-1 encoding of Unicode code points; i.e., each byte is translated directly to a character using `integer->char`, so the decoding always succeeds. The *err-char* argument is ignored, but present for consistency with the other operations.

Example:

```
> (bytes->string/latin-1 (bytes 254 211 209 165))
"þÓÑ¥"
```

```
(string->bytes/utf-8 str [err-byte start end]) → bytes?
  str : string?
  err-byte : (or/c #f byte?) = #f
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (string-length str)
```

Produces a byte string by encoding the *start* to *end* substring of *str* via UTF-8 (always succeeding). The *err-byte* argument is ignored, but included for consistency with the other operations.

Examples:

```
> (define b
    (bytes->string/utf-8
     (bytes 195 167 195 176 195 182 194 163)))

> (string->bytes/utf-8 b)
#"\303\247\303\260\303\266\302\243"
> (bytes->string/utf-8 (string->bytes/utf-8 b))
"çðö£"
```

```
(string->bytes/locale str [err-byte start end]) → bytes?
  str : string?
  err-byte : (or/c #f byte?) = #f
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (string-length str)
```

Produces a string by encoding the *start* to *end* substring of *str* using the current locale's encoding (see also §12.1.1 "Encodings and Locales"). If *err-byte* is not `#f`, it is used for each character in *str* that cannot be encoded for the current locale; if *err-byte* is `#f`, and if the *start* to *end* substring of *str* cannot be encoded, then the `exn:fail:contract` exception is raised.

```
(string->bytes/latin-1  str
                        [err-byte
                         start
                         end])     → bytes?
  str : string?
  err-byte : (or/c #f byte?) = #f
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (string-length str)
```

Produces a string by encoding the *start* to *end* substring of *str* using Latin-1; i.e., each character is translated directly to a byte using `char->integer`. If *err-byte* is not `#f`, it is

used for each character in *str* whose value is greater than 255. If *err-byte* is #f, and if the *start* to *end* substring of *str* has a character with a value greater than 255, then the exn:fail:contract exception is raised.

Examples:

```
> (define b
    (bytes->string/latin-1 (bytes 254 211 209 165)))

> (string->bytes/latin-1 b)
#"\376\323\321\245"
> (bytes->string/latin-1 (string->bytes/latin-1 b))
"þÓÑ¥"
```

```
(string-utf-8-length str [start end]) → exact-nonnegative-integer?
  str : string?
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (string-length str)
```

Returns the length in bytes of the UTF-8 encoding of *str*'s substring from *start* to *end*, but without actually generating the encoded bytes.

Examples:

```
> (string-utf-8-length
    (bytes->string/utf-8 (bytes 195 167 195 176 195 182 194 163)))
8
> (string-utf-8-length "hello")
5
```

```
(bytes-utf-8-length bstr [err-char start end])
 → exact-nonnegative-integer?
  bstr : bytes?
  err-char : (or/c #f char?) = #f
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (bytes-length bstr)
```

Returns the length in characters of the UTF-8 decoding of *bstr*'s substring from *start* to *end*, but without actually generating the decoded characters. If *err-char* is #f and the substring is not a UTF-8 encoding overall, the result is #f. Otherwise, *err-char* is used to resolve decoding errors as in bytes->string/utf-8.

Examples:

```
> (bytes-utf-8-length (bytes 195 167 195 176 195 182 194 163))
4
> (bytes-utf-8-length (make-bytes 5 65))
5
```

```
(bytes-utf-8-ref bstr [skip err-char start end]) → char?
  bstr : bytes?
  skip : exact-nonnegative-integer? = 0
  err-char : (or/c #f char?) = #f
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (bytes-length bstr)
```

Returns the *skip*th character in the UTF-8 decoding of *bstr*'s substring from *start* to *end*, but without actually generating the other decoded characters. If the substring is not a UTF-8 encoding up to the *skip*th character (when *err-char* is `#f`), or if the substring decoding produces fewer than *skip* characters, the result is `#f`. If *err-char* is not `#f`, it is used to resolve decoding errors as in `bytes->string/utf-8`.

Examples:

```
> (bytes-utf-8-ref (bytes 195 167 195 176 195 182 194 163) 0)
#\ç
> (bytes-utf-8-ref (bytes 195 167 195 176 195 182 194 163) 1)
#\ð
> (bytes-utf-8-ref (bytes 195 167 195 176 195 182 194 163) 2)
#\ö
> (bytes-utf-8-ref (bytes 65 66 67 68) 0)
#\A
> (bytes-utf-8-ref (bytes 65 66 67 68) 1)
#\B
> (bytes-utf-8-ref (bytes 65 66 67 68) 2)
#\C
```

```
(bytes-utf-8-index  bstr
                   [skip
                    err-char
                    start
                    end])     → exact-nonnegative-integer?
  bstr : bytes?
  skip : exact-nonnegative-integer? = 0
  err-char : (or/c #f char?) = #f
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (bytes-length bstr)
```

Returns the offset in bytes into *bstr* at which the *skip*th character's encoding starts in the UTF-8 decoding of *bstr*'s substring from *start* to *end* (but without actually generating the other decoded characters). The result is relative to the start of *bstr*, not to *start*. If the substring is not a UTF-8 encoding up to the *skip*th character (when *err-char* is `#f`), or if the substring decoding produces fewer than *skip* characters, the result is `#f`. If *err-char* is not `#f`, it is used to resolve decoding errors as in `bytes->string/utf-8`.

Examples:

```
> (bytes-utf-8-index (bytes 195 167 195 176 195 182 194 163) 0)
0
> (bytes-utf-8-index (bytes 195 167 195 176 195 182 194 163) 1)
2
> (bytes-utf-8-index (bytes 195 167 195 176 195 182 194 163) 2)
4
> (bytes-utf-8-index (bytes 65 66 67 68) 0)
0
> (bytes-utf-8-index (bytes 65 66 67 68) 1)
1
> (bytes-utf-8-index (bytes 65 66 67 68) 2)
2
```

3.4.4 Bytes to Bytes Encoding Conversion

```
(bytes-open-converter from-name to-name) → bytes-converter?
  from-name : string?
  to-name : string?
```

Produces a *byte converter* to go from the encoding named by *from-name* to the encoding named by *to-name*. If the requested conversion pair is not available, `#f` is returned instead of a converter.

Certain encoding combinations are always available:

- `(bytes-open-converter "UTF-8" "UTF-8")` — the identity conversion, except that encoding errors in the input lead to a decoding failure.
- `(bytes-open-converter "UTF-8-permissive" "UTF-8")` — the identity conversion, except that any input byte that is not part of a valid encoding sequence is effectively replaced by the UTF-8 encoding sequence for `#\uFFFD`. (This handling of invalid sequences is consistent with the interpretation of port bytes streams into characters; see §12.1 “Ports”.)
- `(bytes-open-converter "" "UTF-8")` — converts from the current locale’s default encoding (see §12.1.1 “Encodings and Locales”) to UTF-8.
- `(bytes-open-converter "UTF-8" "")` — converts from UTF-8 to the current locale’s default encoding (see §12.1.1 “Encodings and Locales”).
- `(bytes-open-converter "platform-UTF-8" "platform-UTF-16")` — converts UTF-8 to UTF-16 on Unix and Mac OS X, where each UTF-16 code unit is a sequence of two bytes ordered by the current platform’s endianness. On Windows,

the input can include encodings that are not valid UTF-8, but which naturally extend the UTF-8 encoding to support unpaired surrogate code units, and the output is a sequence of UTF-16 code units (as little-endian byte pairs), potentially including unpaired surrogates.

- `(bytes-open-converter "platform-UTF-8-permissive" "platform-UTF-16")` — like `(bytes-open-converter "platform-UTF-8" "platform-UTF-16")`, but an input byte that is not part of a valid UTF-8 encoding sequence (or valid for the unpaired-surrogate extension on Windows) is effectively replaced with `(char->integer #\?)`.

- `(bytes-open-converter "platform-UTF-16" "platform-UTF-8")` — converts UTF-16 (bytes ordered by the current platform's endianness) to UTF-8 on Unix and Mac OS X. On Windows, the input can include UTF-16 code units that are unpaired surrogates, and the corresponding output includes an encoding of each surrogate in a natural extension of UTF-8. On Unix and Mac OS X, surrogates are assumed to be paired: a pair of bytes with the bits `55296` starts a surrogate pair, and the `1023` bits are used from the pair and following pair (independent of the value of the `56320` bits). On all platforms, performance may be poor when decoding from an odd offset within an input byte string.

A newly opened byte converter is registered with the current custodian (see §13.7 "Custodians"), so that the converter is closed when the custodian is shut down. A converter is not registered with a custodian (and does not need to be closed) if it is one of the guaranteed combinations not involving `""` on Unix, or if it is any of the guaranteed combinations (including `""`) on Windows and Mac OS X.

The set of available encodings and combinations varies by platform, depending on the `iconv` library that is installed; the *from-name* and *to-name* arguments are passed on to `iconv_open`. On Windows, `"iconv.dll"` or `"libiconv.dll"` must be in the same directory as `"libmzsch`*VERS*`.dll"` (where *VERS* is a version number), in the user's path, in the system directory, or in the current executable's directory at run time, and the DLL must either supply `_errno` or link to `"msvcrt.dll"` for `_errno`; otherwise, only the guaranteed combinations are available.

In the Racket software distributions for Windows, a suitable `"iconv.dll"` is included with `"libmzsch`*VERS*`.dll"`.

Use `bytes-convert` with the result to convert byte strings.

```
(bytes-close-converter converter) → void
  converter : bytes-converter?
```

Closes the given converter, so that it can no longer be used with `bytes-convert` or `bytes-convert-end`.

```
(bytes-convert  converter
                src-bstr
               [src-start-pos
                src-end-pos
                dest-bstr
                dest-start-pos
                dest-end-pos])
   (or/c bytes? exact-nonnegative-integer?)
 → exact-nonnegative-integer?
   (or/c 'complete 'continues 'aborts 'error)
  converter : bytes-converter?
  src-bstr : bytes?
  src-start-pos : exact-nonnegative-integer? = 0
  src-end-pos :  exact-nonnegative-integer?
              = (bytes-length src-bstr)
  dest-bstr : (or/c bytes? #f) = #f
  dest-start-pos : exact-nonnegative-integer? = 0
  dest-end-pos :  (or/c exact-nonnegative-integer? #f)
                = (and dest-bstr
                       (bytes-length dest-bstr))
```

Converts the bytes from *src-start-pos* to *src-end-pos* in *src-bstr*.

If *dest-bstr* is not `#f`, the converted bytes are written into *dest-bstr* from *dest-start-pos* to *dest-end-pos*. If *dest-bstr* is `#f`, then a newly allocated byte string holds the conversion results, and if *dest-end-pos* is not `#f`, the size of the result byte string is no more than `(- dest-end-pos dest-start-pos)`.

The result of `bytes-convert` is three values:

- *result-bstr* or *dest-wrote-amt* — a byte string if *dest-bstr* is `#f` or not provided, or the number of bytes written into *dest-bstr* otherwise.
- *src-read-amt* — the number of bytes successfully converted from *src-bstr*.
- `'complete`, `'continues`, `'aborts`, or `'error` — indicates how conversion terminated:
  - `'complete`: The entire input was processed, and *src-read-amt* will be equal to `(- src-end-pos src-start-pos)`.
  - `'continues`: Conversion stopped due to the limit on the result size or the space in *dest-bstr*; in this case, fewer than `(- dest-end-pos dest-start-pos)` bytes may be returned if more space is needed to process the next complete encoding sequence in *src-bstr*.
  - `'aborts`: The input stopped part-way through an encoding sequence, and more input bytes are necessary to continue. For example, if the last byte of input is `195`

for a "UTF-8-permissive" decoding, the result is 'aborts, because another byte is needed to determine how to use the 195 byte.

– 'error: The bytes starting at (+ *src-start-pos* *src-read-amt*) bytes in *src-bstr* do not form a legal encoding sequence. This result is never produced for some encodings, where all byte sequences are valid encodings. For example, since "UTF-8-permissive" handles an invalid UTF-8 sequence by dropping characters or generating "?," every byte sequence is effectively valid.

Applying a converter accumulates state in the converter (even when the third result of bytes-convert is 'complete). This state can affect both further processing of input and further generation of output, but only for conversions that involve "shift sequences" to change modes within a stream. To terminate an input sequence and reset the converter, use bytes-convert-end.

Examples:

```
> (define convert (bytes-open-converter "UTF-8" "UTF-16"))

> (bytes-convert convert (bytes 65 66 67 68))
#"\377\376A\0B\0C\0D\0"
4
'complete
> (bytes 195 167 195 176 195 182 194 163)
#"\303\247\303\260\303\266\302\243"
> (bytes-convert convert (bytes 195 167 195 176 195 182 194 163))
#"\347\0\360\0\366\0\243\0"
8
'complete
> (bytes-close-converter convert)
```

```
(bytes-convert-end  converter
                   [dest-bstr
                    dest-start-pos
                    dest-end-pos])
 → (or/c bytes? exact-nonnegative-integer?)
   (or/c 'complete 'continues)
  converter : bytes-converter?
  dest-bstr : (or/c bytes? #f) = #f
  dest-start-pos : exact-nonnegative-integer? = 0
  dest-end-pos : (or/c exact-nonnegative-integer? #f)
               = (and dest-bstr
                      (bytes-length dest-bstr))
```

Like bytes-convert, but instead of converting bytes, this procedure generates an ending sequence for the conversion (sometimes called a "shift sequence"), if any. Few encodings

use shift sequences, so this function will succeed with no output for most encodings. In any case, successful output of a (possibly empty) shift sequence resets the converter to its initial state.

The result of bytes-convert-end is two values:

- *result-bstr* or *dest-wrote-amt* — a byte string if *dest-bstr* is #f or not provided, or the number of bytes written into *dest-bstr* otherwise.
- 'complete or 'continues — indicates whether conversion completed. If 'complete, then an entire ending sequence was produced. If 'continues, then the conversion could not complete due to the limit on the result size or the space in *dest-bstr*, and the first result is either an empty byte string or 0.

(bytes-converter? *v*) → boolean?
  v : any/c

Returns #t if *v* is a byte converter produced by bytes-open-converter, #f otherwise.

Examples:

```
> (bytes-converter? (bytes-open-converter "UTF-8" "UTF-16"))
#t
> (bytes-converter? (bytes-open-converter "whacky" "not likely"))
#f
> (define b (bytes-open-converter "UTF-8" "UTF-16"))

> (bytes-close-converter b)

> (bytes-converter? b)
#t
```

(locale-string-encoding) → any

Returns a string for the current locale's encoding (i.e., the encoding normally identified by ""). See also system-language+country.

3.4.5 Additional Byte String Functions

```
(require racket/bytes)
```

The bindings documented in this section are provided by the racket/bytes and racket libraries, but not racket/base.

```
(bytes-append* str ... strs) → bytes?
  str : bytes?
  strs : (listof bytes?)
```

Like bytes-append, but the last argument is used as a list of arguments for bytes-append, so (bytes-append* *str* ... *strs*) is the same as (apply bytes-append *str* ... *strs*). In other words, the relationship between bytes-append and bytes-append* is similar to the one between list and list*.

Examples:

```
> (bytes-append* #"a" #"b" '(#"c" #"d"))
#"abcd"
> (bytes-append* (cdr (append* (map (lambda (x) (list #", " x))
                                    '(#"Alpha" #"Beta" #"Gamma")))))
#"Alpha, Beta, Gamma"
```

```
(bytes-join strs sep) → bytes?
  strs : (listof bytes?)
  sep : bytes?
```

Appends the byte strings in *strs*, inserting *sep* between each pair of bytes in *strs*.

Example:

```
> (bytes-join '(#"one" #"two" #"three" #"four") #" potato ")
#"one potato two potato three potato four"
```

3.5 Characters

§3.3 “Characters” in ***The Racket Guide*** introduces characters.

Characters range over Unicode scalar values, which includes characters whose values range from #x0 to #x10FFFF, but not including #xD800 to #xDFFF. The scalar values are a subset of the Unicode code points.

Two characters are eqv? if they correspond to the same scalar value. For each scalar value less than 256, character values that are eqv? are also eq?. Characters produced by the default reader are interned in read-syntax mode.

See §1.3.13 “Reading Characters” for information on reading characters and §1.4.10 “Printing Characters” for information on printing characters.

3.5.1 Characters and Scalar Values

```
(char? v) → boolean?
  v : any/c
```

Return `#t` if *v* is a character, `#f` otherwise.

```
(char->integer char) → exact-integer?
  char : char?
```

Returns a character's code-point number.

Example:

```
> (char->integer #\A)
65
```

```
(integer->char k) → char?
      (and/c exact-integer?
  k :        (or/c (integer-in 0 55295)
                   (integer-in 57344 1114111)))
```

Return the character whose code-point number is *k*. For *k* less than 256, the result is the same object for the same *k*.

Example:

```
> (integer->char 65)
#\A
```

```
(char-utf-8-length char) → (integer-in 1 6)
  char : char?
```

Produces the same result as `(bytes-length (string->bytes/utf-8 (string char)))`.

3.5.2 Character Comparisons

```
(char=? char1 char2 ...+) → boolean?
  char1 : char?
  char2 : char?
```

Returns `#t` if all of the arguments are `eqv?`.

Examples:

```
> (char=? #\a #\a)
#t
> (char=? #\a #\A #\a)
#f
```

```
(char<? char1 char2 ...+) → boolean?
  char1 : char?
  char2 : char?
```

Returns #t if the arguments are sorted increasing, where two characters are ordered by their scalar values, #f otherwise.

Examples:

```
> (char<? #\A #\a)
#t
> (char<? #\a #\A)
#f
> (char<? #\a #\b #\c)
#t
```

```
(char<=? char1 char2 ...+) → boolean?
  char1 : char?
  char2 : char?
```

Like char<?, but checks whether the arguments are nondecreasing.

Examples:

```
> (char<=? #\A #\a)
#t
> (char<=? #\a #\A)
#f
> (char<=? #\a #\b #\b)
#t
```

```
(char>? char1 char2 ...+) → boolean?
  char1 : char?
  char2 : char?
```

Like char<?, but checks whether the arguments are decreasing.

Examples:

```
> (char>? #\A #\a)
#f
> (char>? #\a #\A)
#t
> (char>? #\c #\b #\a)
#t
```

```
(char>=? char1 char2 ...+) → boolean?
  char1 : char?
  char2 : char?
```

Like `char<?`, but checks whether the arguments are nonincreasing.

Examples:

```
> (char>=? #\A #\a)
#f
> (char>=? #\a #\A)
#t
> (char>=? #\c #\b #\b)
#t
```

```
(char-ci=? char1 char2 ...+) → boolean?
  char1 : char?
  char2 : char?
```

Returns `#t` if all of the arguments are `eqv?` after locale-insensitive case-folding via `char-foldcase`.

Examples:

```
> (char-ci=? #\A #\a)
#t
> (char-ci=? #\a #\a #\a)
#t
```

```
(char-ci<? char1 char2 ...+) → boolean?
  char1 : char?
  char2 : char?
```

Like `char<?`, but checks whether the arguments would be in increasing order if each was first case-folded using `char-foldcase` (which is locale-insensitive).

Examples:

```
> (char-ci<? #\A #\a)
#f
> (char-ci<? #\a #\b)
#t
> (char-ci<? #\a #\b #\c)
#t
```

```
(char-ci<=? char1 char2 ...+) → boolean?
  char1 : char?
  char2 : char?
```

Like char-ci<?, but checks whether the arguments would be nondecreasing after case-folding.

Examples:

```
> (char-ci<=? #\A #\a)
#t
> (char-ci<=? #\a #\A)
#t
> (char-ci<=? #\a #\b #\b)
#t
```

```
(char-ci>? char1 char2 ...+) → boolean?
  char1 : char?
  char2 : char?
```

Like char-ci<?, but checks whether the arguments would be decreasing after case-folding.

Examples:

```
> (char-ci>? #\A #\a)
#f
> (char-ci>? #\b #\A)
#t
> (char-ci>? #\c #\b #\a)
#t
```

```
(char-ci>=? char1 char2 ...+) → boolean?
  char1 : char?
  char2 : char?
```

Like char-ci<?, but checks whether the arguments would be nonincreasing after case-folding.

Examples:

```
> (char-ci>=? #\A #\a)
#t
> (char-ci>=? #\a #\A)
#t
> (char-ci>=? #\c #\b #\b)
#t
```

3.5.3 Classifications

```
(char-alphabetic? char) → boolean?
  char : char?
```

Returns `#t` if *char* has the Unicode “Alphabetic” property.

```
(char-lower-case? char) → boolean?
  char : char?
```

Returns `#t` if *char* has the Unicode “Lowercase” property.

```
(char-upper-case? char) → boolean?
  char : char?
```

Returns `#t` if *char* has the Unicode “Uppercase” property.

```
(char-title-case? char) → boolean?
  char : char?
```

Returns `#t` if *char*’s Unicode general category is Lt, `#f` otherwise.

```
(char-numeric? char) → boolean?
  char : char?
```

Returns `#t` if *char* has the Unicode “Numeric” property.

```
(char-symbolic? char) → boolean?
  char : char?
```

Returns `#t` if *char*’s Unicode general category is Sm, Sc, Sk, or So, `#f` otherwise.

```
(char-punctuation? char) → boolean?
  char : char?
```

Returns `#t` if *char*’s Unicode general category is Pc, Pd, Ps, Pe, Pi, Pf, or Po, `#f` otherwise.

```
(char-graphic? char) → boolean?
  char : char?
```

Returns `#t` if *char*’s Unicode general category is Ll, Lm, Lo, Lt, Lu, Nd, Nl, No, Mn, Mc, or Me, or if one of the following produces `#t` when applied to *char*: `char-alphabetic?`, `char-numeric?`, `char-symbolic?`, or `char-punctuation?`.

```
(char-whitespace? char) → boolean?
  char : char?
```

Returns `#t` if *char* has the Unicode “White_Space” property.

```
(char-blank? char) → boolean?
  char : char?
```

Returns `#t` if *char*'s Unicode general category is Zs or if *char* is `#\tab`. (These correspond to horizontal whitespace.)

```
(char-iso-control? char) → boolean?
  char : char?
```

Return `#t` if *char* is between `#\nul` and `#\u001F` inclusive or `#\rubout` and `#\u009F` inclusive.

```
(char-general-category char) → symbol?
  char : char?
```

Returns a symbol representing the character's Unicode general category, which is `'lu`, `'ll`, `'lt`, `'lm`, `'lo`, `'mn`, `'mc`, `'me`, `'nd`, `'nl`, `'no`, `'ps`, `'pe`, `'pi`, `'pf`, `'pd`, `'pc`, `'po`, `'sc`, `'sm`, `'sk`, `'so`, `'zs`, `'zp`, `'zl`, `'cc`, `'cf`, `'cs`, `'co`, or `'cn`.

```
(make-known-char-range-list)
    (listof (list/c exact-nonnegative-integer?
 →                  exact-nonnegative-integer?
                    boolean?))
```

Produces a list of three-element lists, where each three-element list represents a set of consecutive code points for which the Unicode standard specifies character properties. Each three-element list contains two integers and a boolean; the first integer is a starting code-point value (inclusive), the second integer is an ending code-point value (inclusive), and the boolean is `#t` when all characters in the code-point range have identical results for all of the character predicates above. The three-element lists are ordered in the overall result list such that later lists represent larger code-point values, and all three-element lists are separated from every other by at least one code-point value that is not specified by Unicode.

3.5.4 Character Conversions

```
(char-upcase char) → char?
  char : char?
```

Produces a character consistent with the 1-to-1 code point mapping defined by Unicode. If *char* has no upcase mapping, `char-upcase` produces *char*.

String procedures, such as `string-upcase`, handle the case where Unicode defines a locale-independent mapping from the code point to a code-point sequence (in addition to the 1-1 mapping on scalar values).

Examples:

```
> (char-upcase #\a)
#\A
> (char-upcase #\λ)
#\Λ
> (char-upcase #\space)
#\space
```

```
(char-downcase char) → char?
  char : char?
```

Like `char-upcase`, but for the Unicode downcase mapping.

Examples:

```
> (char-downcase #\A)
#\a
> (char-downcase #\Λ)
#\λ
> (char-downcase #\space)
#\space
```

```
(char-titlecase char) → char?
  char : char?
```

Like `char-upcase`, but for the Unicode titlecase mapping.

Examples:

```
> (char-upcase #\a)
#\A
> (char-upcase #\λ)
#\Λ
> (char-upcase #\space)
#\space
```

```
(char-foldcase char) → char?
  char : char?
```

Like `char-upcase`, but for the Unicode case-folding mapping.

Examples:

```
> (char-foldcase #\A)
#\a
> (char-foldcase #\Σ)
#\σ
> (char-foldcase #\ς)
#\σ
> (char-foldcase #\space)
#\space
```

3.6 Symbols

§3.6 “Symbols” in ***The Racket Guide*** introduces symbols.

A *symbol* is like an immutable string, but symbols are normally interned, so that two symbols with the same character content are normally `eq?`. All symbols produced by the default reader (see §1.3.2 “Reading Symbols”) are interned.

The two procedures `string->uninterned-symbol` and `gensym` generate *uninterned* symbols, i.e., symbols that are not `eq?`, `eqv?`, or `equal?` to any other symbol, although they may print the same as other symbols.

The procedure `string->unreadable-symbol` returns an *unreadable symbol* that is partially interned. The default reader (see §1.3.2 “Reading Symbols”) never produces a unreadable symbol, but two calls to `string->unreadable-symbol` with `equal?` strings produce `eq?` results. An unreadable symbol can print the same as an interned or uninterned symbol. Unreadable symbols are useful in expansion and compilation to avoid collisions with symbols that appear in the source; they are usually not generated directly, but they can appear in the result of functions like `identifier-binding`.

Interned and unreadable symbols are only weakly held by the internal symbol table. This weakness can never affect the result of an `eq?`, `eqv?`, or `equal?` test, but a symbol may disappear when placed into a weak box (see §15.1 “Weak Boxes”) used as the key in a weak hash table (see §3.13 “Hash Tables”), or used as an ephemeron key (see §15.2 “Ephemerons”).

See §1.3.2 “Reading Symbols” for information on `read`ing symbols and §1.4.1 “Printing Symbols” for information on `print`ing symbols.

```
(symbol? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a symbol, `#f` otherwise.

Examples:

```
> (symbol? 'Apple)
#t
> (symbol? 10)
#f
```

```
(symbol-interned? sym) → boolean?
  sym : symbol?
```

Returns `#t` if *sym* is interned, `#f` otherwise.

Examples:

```
> (symbol-interned? 'Apple)
```

```
#t
> (symbol-interned? (gensym))
#f
> (symbol-interned? (string->unreadable-symbol "Apple"))
#f
```

```
(symbol-unreadable? sym) → boolean?
  sym : symbol?
```

Returns `#t` if *sym* is an unreadable symbol, `#f` otherwise.

Examples:

```
> (symbol-unreadable? 'Apple)
#f
> (symbol-unreadable? (gensym))
#f
> (symbol-unreadable? (string->unreadable-symbol "Apple"))
#t
```

```
(symbol->string sym) → string?
  sym : symbol?
```

Returns a freshly allocated mutable string whose characters are the same as in *sym*.

Example:

```
> (symbol->string 'Apple)
"Apple"
```

```
(string->symbol str) → symbol?
  str : string?
```

Returns an interned symbol whose characters are the same as in *str*.

Examples:

```
> (string->symbol "Apple")
'Apple
> (string->symbol "1")
'|1|
```

```
(string->uninterned-symbol str) → symbol?
  str : string?
```

Like `(string->symbol str)`, but the resulting symbol is a new uninterned symbol. Calling `string->uninterned-symbol` twice with the same *str* returns two distinct symbols.

Examples:

```
> (string->uninterned-symbol "Apple")
'Apple
> (eq? 'a (string->uninterned-symbol "a"))
#f
> (eq? (string->uninterned-symbol "a")
       (string->uninterned-symbol "a"))
#f
```

```
(string->unreadable-symbol str) → symbol?
  str : string?
```

Like `(string->symbol str)`, but the resulting symbol is a new unreadable symbol. Calling `string->unreadable-symbol` twice with equivalent *str*s returns the same symbol, but `read` never produces the symbol.

Examples:

```
> (string->unreadable-symbol "Apple")
'Apple
> (eq? 'a (string->unreadable-symbol "a"))
#f
> (eq? (string->unreadable-symbol "a")
       (string->unreadable-symbol "a"))
#t
```

```
(gensym [base]) → symbol?
  base : (or/c string? symbol?) = "g"
```

Returns a new uninterned symbol with an automatically-generated name. The optional *base* argument is a prefix symbol or string.

Example:

```
> (gensym "apple")
'apple239220
```

3.7 Regular Expressions

§9 “Regular Expressions” in ***The Racket Guide*** introduces regular expressions.

Regular expressions are specified as strings or byte strings, using the same pattern language as either the Unix utility `egrep` or Perl. A string-specified pattern produces a character regexp matcher, and a byte-string pattern produces a byte regexp matcher. If a character regexp is used with a byte string or input port, it matches UTF-8 encodings (see §12.1.1 “Encodings and Locales”) of matching character streams; if a byte regexp is used with a character string, it matches bytes in the UTF-8 encoding of the string.

Regular expressions can be compiled into a *regexp value* for repeated matches. The `regexp` and `byte-regexp` procedures convert a string or byte string (respectively) into a regexp value using one syntax of regular expressions that is most compatible to `egrep`. The `pregexp` and `byte-pregexp` procedures produce a regexp value using a slightly different syntax of regular expressions that is more compatible with Perl.

Two regular expressions are `equal?` if they have the same source, use the same pattern language, and are both character regexps or both byte regexps.

A literal or printed regular expression starts with `#rx` or `#px`. See §1.3.15 "Reading Regular Expressions" for information on `read`ing regular expressions and §1.4.12 "Printing Regular Expressions" for information on `print`ing regular expressions. Regular expressions produced by the default reader are interned in `read-syntax` mode.

The internal size of a regexp value is limited to 32 kilobytes; this limit roughly corresponds to a source string with 32,000 literal characters or 5,000 operators.

3.7.1 Regexp Syntax

The following syntax specifications describe the content of a string that represents a regular expression. The syntax of the corresponding string may involve extra escape characters. For example, the regular expression `(.*)\1` can be represented with the string `"(.*)\\1"` or the regexp constant `#rx"(.*)\\1"`; the `\` in the regular expression must be escaped to include it in a string or regexp constant.

The `regexp` and `pregexp` syntaxes share a common core:

⟨*regexp*⟩	::=	⟨*pces*⟩	Match ⟨*pces*⟩	
	|	⟨*regexp*⟩`|`⟨*regexp*⟩	Match either ⟨*regexp*⟩, try left first	ex1
⟨*pces*⟩	::=	⟨*pce*⟩	Match ⟨*pce*⟩	
	|	⟨*pce*⟩⟨*pces*⟩	Match ⟨*pce*⟩ followed by ⟨*pces*⟩	
⟨*pce*⟩	::=	⟨*repeat*⟩	Match ⟨*repeat*⟩, longest possible	ex3
	|	⟨*repeat*⟩`?`	Match ⟨*repeat*⟩, shortest possible	ex6
	|	⟨*atom*⟩	Match ⟨*atom*⟩ exactly once	
⟨*repeat*⟩	::=	⟨*atom*⟩`*`	Match ⟨*atom*⟩ 0 or more times	ex3
	|	⟨*atom*⟩`+`	Match ⟨*atom*⟩ 1 or more times	ex4
	|	⟨*atom*⟩`?`	Match ⟨*atom*⟩ 0 or 1 times	ex5
⟨*atom*⟩	::=	`(`⟨*regexp*⟩`)`	Match sub-expression ⟨*regexp*⟩ and report	ex11
	|	`[`⟨*rng*⟩`]`	Match any character in ⟨*rng*⟩	ex2
	|	`[^`⟨*rng*⟩`]`	Match any character not in ⟨*rng*⟩	ex12
	|	`.`	Match any (except newline in multi mode)	ex13
	|	`^`	Match start (or after newline in multi mode)	ex14
	|	`$`	Match end (or before newline in multi mode)	ex15
	|	⟨*literal*⟩	Match a single literal character	ex1
	|	`(?`⟨*mode*⟩`:`⟨*regexp*⟩`)`	Match ⟨*regexp*⟩ using ⟨*mode*⟩	ex35

	`(?>`⟨*regexp*⟩`)`	Match ⟨*regexp*⟩, only first possible	
	⟨*look*⟩	Match empty if ⟨*look*⟩ matches	
	`(?`⟨*tst*⟩⟨*pces*⟩`|`⟨*pces*⟩`)`	Match 1st ⟨*pces*⟩ if ⟨*tst*⟩, else 2nd ⟨*pces*⟩	ex36
	`(?`⟨*tst*⟩⟨*pces*⟩`)`	Match ⟨*pces*⟩ if ⟨*tst*⟩, empty if not ⟨*tst*⟩	
⟨*rng*⟩ ::=	`]`	⟨*rng*⟩ contains `]` only	ex27
	`-`	⟨*rng*⟩ contains `-` only	ex28
	⟨*mrng*⟩	⟨*rng*⟩ contains everything in ⟨*mrng*⟩	
	⟨*mrng*⟩`-`	⟨*rng*⟩ contains `-` and everything in ⟨*mrng*⟩	
⟨*mrng*⟩ ::=	`]`⟨*lrng*⟩	⟨*mrng*⟩ contains `]` and everything in ⟨*lrng*⟩	ex29
	`-`⟨*lrng*⟩	⟨*mrng*⟩ contains `-` and everything in ⟨*lrng*⟩	ex29
	⟨*lirng*⟩	⟨*mrng*⟩ contains everything in ⟨*lirng*⟩	
⟨*lirng*⟩ ::=	⟨*riliteral*⟩	⟨*lirng*⟩ contains a literal character	
	⟨*riliteral*⟩`-`⟨*rliteral*⟩	⟨*lirng*⟩ contains Unicode range inclusive	ex22
	⟨*lirng*⟩⟨*lrng*⟩	⟨*lirng*⟩ contains everything in both	
⟨*lrng*⟩ ::=	`^`	⟨*lrng*⟩ contains `^`	ex30
	⟨*rliteral*⟩`-`⟨*rliteral*⟩	⟨*lrng*⟩ contains Unicode range inclusive	
	`^`⟨*lrng*⟩	⟨*lrng*⟩ contains `^` and more	
	⟨*lirng*⟩	⟨*lrng*⟩ contains everything in ⟨*lirng*⟩	
⟨*look*⟩ ::=	`(?=`⟨*regexp*⟩`)`	Match if ⟨*regexp*⟩ matches	ex31
	`(?!`⟨*regexp*⟩`)`	Match if ⟨*regexp*⟩ doesn't match	ex32
	`(?<=`⟨*regexp*⟩`)`	Match if ⟨*regexp*⟩ matches preceding	ex33
	`(?<!`⟨*regexp*⟩`)`	Match if ⟨*regexp*⟩ doesn't match preceding	ex34
⟨*tst*⟩ ::=	`(`⟨*n*⟩`)`	True if Nth `(` has a match	
	⟨*look*⟩	True if ⟨*look*⟩ matches	ex36
⟨*mode*⟩ ::=		Like the enclosing mode	
	⟨*mode*⟩`i`	Like ⟨*mode*⟩, but case-insensitive	ex35
	⟨*mode*⟩`-i`	Like ⟨*mode*⟩, but sensitive	
	⟨*mode*⟩`s`	Like ⟨*mode*⟩, but not in multi mode	
	⟨*mode*⟩`-s`	Like ⟨*mode*⟩, but in multi mode	
	⟨*mode*⟩`m`	Like ⟨*mode*⟩, but in multi mode	
	⟨*mode*⟩`-m`	Like ⟨*mode*⟩, but not in multi mode	

The following completes the grammar for `regexp`, which treats `{` and `}` as literals, `\` as a literal within ranges, and `\` as a literal producer outside of ranges.

⟨*literal*⟩ ::=	Any character except `(`, `)`, `*`, `+`, `?`, `[`, `.`, `^`, `\`, or `|`		
	`\`⟨*aliteral*⟩	Match ⟨*aliteral*⟩	ex21
⟨*aliteral*⟩ ::=	Any character		
⟨*riliteral*⟩ ::=	Any character except `]`, `-`, or `^`		
⟨*rliteral*⟩ ::=	Any character except `]` or `-`		

The following completes the grammar for `pregexp`, which uses `{` and `}` bounded repetition and uses `\` for meta-characters both inside and outside of ranges.

⟨*repeat*⟩ ::=	...	...	
	⟨*atom*⟩`{`⟨*n*⟩`}`	Match ⟨*atom*⟩ exactly ⟨*n*⟩ times	ex7

	| ⟨*atom*⟩`{`⟨*n*⟩`,}`	Match ⟨*atom*⟩ ⟨*n*⟩ or more times	ex8
	| ⟨*atom*⟩`{,`⟨*m*⟩`}`	Match ⟨*atom*⟩ between 0 and ⟨*m*⟩ times	ex9
	| ⟨*atom*⟩`{`⟨*n*⟩`,`⟨*m*⟩`}`	Match ⟨*atom*⟩ between ⟨*n*⟩ and ⟨*m*⟩ times	ex10
⟨*atom*⟩	::= ...	...	
	| `\`⟨*n*⟩	Match latest reported match for ⟨*n*⟩th `(`	ex16
	| ⟨*class*⟩	Match any character in ⟨*class*⟩	
	| `\b`	Match `\w*` boundary	ex17
	| `\B`	Match where `\b` does not	ex18
	| `\p{`⟨*property*⟩`}`	Match (UTF-8 encoded) in ⟨*property*⟩	ex19
	| `\P{`⟨*property*⟩`}`	Match (UTF-8 encoded) not in ⟨*property*⟩	ex20
⟨*literal*⟩	::= Any character except `(`, `)`, `*`, `+`, `?`, `[`, `]`, `{`, `}`, `.`, `^`, `\`, or `|`		
	| `\`⟨*aliteral*⟩	Match ⟨*aliteral*⟩	ex21
⟨*aliteral*⟩	::= Any character except `a-z`, `A-Z`, `0-9`		
⟨*lirng*⟩	::= ...	...	
	| ⟨*class*⟩	⟨*lirng*⟩ contains all characters in ⟨*class*⟩	
	| ⟨*posix*⟩	⟨*lirng*⟩ contains all characters in ⟨*posix*⟩	ex26
	| `\`⟨*eliteral*⟩	⟨*lirng*⟩ contains ⟨*eliteral*⟩	
⟨*riliteral*⟩	::= Any character except `]`, `\`, `-`, or `^`		
⟨*rliteral*⟩	::= Any character except `]`, `\`, or `-`		
⟨*eliteral*⟩	::= Any character except `a-z`, `A-Z`		
⟨*class*⟩	::= `\d`	Contains `0-9`	ex23
	| `\D`	Contains ASCII other than those in `\d`	
	| `\w`	Contains `a-z`, `A-Z`, `0-9`, `_`	ex24
	| `\W`	Contains ASCII other than those in `\w`	
	| `\s`	Contains space, tab, newline, formfeed, return	ex25
	| `\S`	Contains ASCII other than those in `\s`	
⟨*posix*⟩	::= `[:alpha:]`	Contains `a-z`, `A-Z`	
	| `[:upper:]`	Contains `A-Z`	
	| `[:lower:]`	Contains `a-z`	ex26
	| `[:digit:]`	Contains `0-9`	
	| `[:xdigit:]`	Contains `0-9`, `a-f`, `A-F`	
	| `[:alnum:]`	Contains `a-z`, `A-Z`, `0-9`	
	| `[:word:]`	Contains `a-z`, `A-Z`, `0-9`, `_`	
	| `[:blank:]`	Contains space and tab	
	| `[:space:]`	Contains space, tab, newline, formfeed, return	
	| `[:graph:]`	Contains all ASCII characters that use ink	
	| `[:print:]`	Contains space, tab, and ASCII ink users	
	| `[:cntrl:]`	Contains all characters with scalar value < 32	
	| `[:ascii:]`	Contains all ASCII characters	
⟨*property*⟩	::= ⟨*category*⟩	Includes all characters in ⟨*category*⟩	
	| `^`⟨*category*⟩	Includes all characters not in ⟨*category*⟩	

The Unicode categories follow.

⟨*category*⟩	::= `Ll` Letter, lowercase		ex19
	| `Lu` Letter, uppercase		

`Lt`	Letter, titlecase
`Lm`	Letter, modifier
`L&`	Union of `Ll`, `Lu`, `Lt`, and `Lm`
`Lo`	Letter, other
`L`	Union of `L&` and `Lo`
`Nd`	Number, decimal digit
`Nl`	Number, letter
`No`	Number, other
`N`	Union of `Nd`, `Nl`, and `No`
`Ps`	Punctuation, open
`Pe`	Punctuation, close
`Pi`	Punctuation, initial quote
`Pf`	Punctuation, final quote
`Pc`	Punctuation, connector
`Pd`	Punctuation, dash
`Po`	Punctuation, other
`P`	Union of `Ps`, `Pe`, `Pi`, `Pf`, `Pc`, `Pd`, and `Po`
`Mn`	Mark, non-spacing
`Mc`	Mark, spacing combining
`Me`	Mark, enclosing
`M`	Union of `Mn`, `Mc`, and `Me`
`Sc`	Symbol, currency
`Sk`	Symbol, modifier
`Sm`	Symbol, math
`So`	Symbol, other
`S`	Union of `Sc`, `Sk`, `Sm`, and `So`
`Zl`	Separator, line
`Zp`	Separator, paragraph
`Zs`	Separator, space
`Z`	Union of `Zl`, `Zp`, and `Zs`
`Cc`	Other, control
`Cf`	Other, format
`Cs`	Other, surrogate
`Cn`	Other, not assigned
`Co`	Other, private use
`C`	Union of `Cc`, `Cf`, `Cs`, `Cn`, and `Co`
`.`	Union of all Unicode categories

Examples:

```
> (regexp-match #rx"a|b" "cat") ; ex1
'("a")
> (regexp-match #rx"[at]" "cat") ; ex2
'("a")
> (regexp-match #rx"ca*[at]" "caaat") ; ex3
'("caaat")
```

```
> (regexp-match #rx"ca+[at]" "caaat") ; ex4
'("caaat")
> (regexp-match #rx"ca?t?" "ct") ; ex5
'("ct")
> (regexp-match #rx"ca*?[at]" "caaat") ; ex6
'("ca")
> (regexp-match #px"ca{2}" "caaat") ; ex7, uses #px
'("caa")
> (regexp-match #px"ca{2,}t" "catcaat") ; ex8, uses #px
'("caat")
> (regexp-match #px"ca{,2}t" "caaatcat") ; ex9, uses #px
'("cat")
> (regexp-match #px"ca{1,2}t" "caaatcat") ; ex10, uses #px
'("cat")
> (regexp-match #rx"(c*)(a*)" "caat") ; ex11
'("caa" "c" "aa")
> (regexp-match #rx"[^ca]" "caat") ; ex12
'("t")
> (regexp-match #rx".(.)." "cat") ; ex13
'("cat" "a")
> (regexp-match #rx"^a|^c" "cat") ; ex14
'("c")
> (regexp-match #rx"a$|t$" "cat") ; ex15
'("t")
> (regexp-match #px"c(.)\\1t" "caat") ; ex16, uses #px
'("caat" "a")
> (regexp-match #px".\\b." "cat in hat") ; ex17, uses #px
'("t ")
> (regexp-match #px".\\B." "cat in hat") ; ex18, uses #px
'("ca")
> (regexp-match #px"\\p{Ll}" "Cat") ; ex19, uses #px
'("a")
> (regexp-match #px"\\P{Ll}" "cat!") ; ex20, uses #px
'("!")
> (regexp-match #rx"\\|" "c|t") ; ex21
'("|")
> (regexp-match #rx"[a-f]*" "cat") ; ex22
'("ca")
> (regexp-match #px"[a-f\\d]*" "1cat") ; ex23, uses #px
'("1ca")
> (regexp-match #px" [\\w]" "cat hat") ; ex24, uses #px
'(" h")
> (regexp-match #px"t[\\s]" "cat\nhat") ; ex25, uses #px
'("t\n")
> (regexp-match #px"[[:lower:]]+" "Cat") ; ex26, uses #px
'("at")
```

```
> (regexp-match #rx"[]]" "c]t") ; ex27
'("]")
> (regexp-match #rx"[-]" "c-t") ; ex28
'("-")
> (regexp-match #rx"[]a[]+" "c[a]t") ; ex29
'("[a]")
> (regexp-match #rx"[a^]+" "ca^t") ; ex30
'("a^")
> (regexp-match #rx".a(?=p)" "cat nap") ; ex31
'("na")
> (regexp-match #rx".a(?!t)" "cat nap") ; ex32
'("na")
> (regexp-match #rx"(?<=n)a." "cat nap") ; ex33
'("ap")
> (regexp-match #rx"(?<!c)a." "cat nap") ; ex34
'("ap")
> (regexp-match #rx"(?i:a)[tp]" "cAT nAp") ; ex35
'("Ap")
> (regexp-match #rx"(?(?<=c)a|b)+" "cabal") ; ex36
'("ab")
```

3.7.2 Additional Syntactic Constraints

In addition to matching a grammar, regular expressions must meet two syntactic restrictions:

- In a ⟨*repeat*⟩ other than ⟨*atom*⟩`?`, the ⟨*atom*⟩ must not match an empty sequence.
- In a `(?<=`⟨*regexp*⟩`)` or `(?<!`⟨*regexp*⟩`)`, the ⟨*regexp*⟩ must match a bounded sequence only.

These constraints are checked syntactically by the following type system. A type $[n, m]$ corresponds to an expression that matches between n and m characters. In the rule for `(`⟨*Regexp*⟩`)`, N means the number such that the opening parenthesis is the Nth opening parenthesis for collecting match reports. Non-emptiness is inferred for a backreference pattern, `\`⟨N⟩, so that a backreference can be used for repetition patterns; in the case of mutual dependencies among backreferences, the inference chooses the fixpoint that maximizes non-emptiness. Finiteness is not inferred for backreferences (i.e., a backreference is assumed to match an arbitrarily large sequence).

$$\frac{\langle regexp\rangle_1 \ : \ [n_1, m_1] \qquad \langle regexp\rangle_2 \ : \ [n_2, m_2]}{\langle regexp\rangle_1\,|\,\langle regexp\rangle_2 \ : \ [\min(n_1, n_2), \max(m_1, m_2)]}$$

$$\frac{\langle pce\rangle \ : \ [n_1, m_1] \qquad \langle pces\rangle \ : \ [n_2, m_2]}{\langle pce\rangle\langle pces\rangle \ : \ [n_1{+}n_2, m_1{+}m_2]}$$

$$\frac{\langle repeat\rangle \;:\; [n,m]}{\langle repeat\rangle\texttt{?} \;:\; [0,m]} \qquad \frac{\langle atom\rangle \;:\; [n,m] \quad n>0}{\langle atom\rangle\texttt{*} \;:\; [0,\infty]}$$

$$\frac{\langle atom\rangle \;:\; [n,m] \quad n>0}{\langle atom\rangle\texttt{+} \;:\; [1,\infty]} \qquad \frac{\langle atom\rangle \;:\; [n,m]}{\langle atom\rangle\texttt{?} \;:\; [0,m]}$$

$$\frac{\langle atom\rangle \;:\; [n,m] \quad n>0}{\langle atom\rangle\texttt{\{}\langle n\rangle\texttt{\}} \;:\; [n\text{*}\langle n\rangle, m\text{*}\langle n\rangle]}$$

$$\frac{\langle atom\rangle \;:\; [n,m] \quad n>0}{\langle atom\rangle\texttt{\{}\langle n\rangle\texttt{,\}} \;:\; [n\text{*}\langle n\rangle, \infty]}$$

$$\frac{\langle atom\rangle \;:\; [n,m] \quad n>0}{\langle atom\rangle\texttt{\{,}\langle m\rangle\texttt{\}} \;:\; [0, m\text{*}\langle m\rangle]}$$

$$\frac{\langle atom\rangle \;:\; [n,m] \quad n>0}{\langle atom\rangle\texttt{\{}\langle n\rangle\texttt{,}\langle m\rangle\texttt{\}} \;:\; [n\text{*}\langle n\rangle, m\text{*}\langle m\rangle]}$$

$$\frac{\langle regexp\rangle \;:\; [n,m]}{\texttt{(}\langle regexp\rangle\texttt{)} \;:\; [n,m] \quad \alpha_{\mathrm{N}}{=}n}$$

$$\frac{\langle regexp\rangle \;:\; [n,m]}{\texttt{(?}\langle mode\rangle\texttt{:}\langle regexp\rangle\texttt{)} \;:\; [n,m]}$$

$$\frac{\langle regexp\rangle \;:\; [n,m]}{\texttt{(?=}\langle regexp\rangle\texttt{)} \;:\; [0,0]} \qquad \frac{\langle regexp\rangle \;:\; [n,m]}{\texttt{(?!}\langle regexp\rangle\texttt{)} \;:\; [0,0]}$$

$$\frac{\langle regexp\rangle \;:\; [n,m] \quad m<\infty}{\texttt{(?<=}\langle regexp\rangle\texttt{)} \;:\; [0,0]} \qquad \frac{\langle regexp\rangle \;:\; [n,m] \quad m<\infty}{\texttt{(?<!}\langle regexp\rangle\texttt{)} \;:\; [0,0]}$$

$$\frac{\langle regexp\rangle \;:\; [n,m]}{\texttt{(?>}\langle regexp\rangle\texttt{)} \;:\; [n,m]}$$

$$\frac{\langle tst\rangle \;:\; [n_0,m_0] \quad \langle pces\rangle_1 \;:\; [n_1,m_1] \quad \langle pces\rangle_2 \;:\; [n_2,m_2]}{\texttt{(?}\langle tst\rangle\langle pces\rangle_1\texttt{|}\langle pces\rangle_2\texttt{)} \;:\; [\min(n_1,n_2), \max(m_1,m_2)]}$$

$$\frac{\langle tst\rangle \;:\; [n_0,m_0] \quad \langle pces\rangle \;:\; [n_1,m_1]}{\texttt{(?}\langle tst\rangle\langle pces\rangle\texttt{)} \;:\; [0,m_1]}$$

$$\texttt{(}\langle n\rangle\texttt{)} \;:\; [\alpha_{\mathrm{N}},\infty] \qquad \texttt{[}\langle rng\rangle\texttt{]} \;:\; [1,1] \qquad \texttt{[\^{}}\langle rng\rangle\texttt{]} \;:\; [1,1]$$

$$\texttt{.} \;:\; [1,1] \qquad \texttt{\^{}} \;:\; [0,0] \qquad \texttt{\$} \;:\; [0,0]$$

$$\langle literal\rangle \;:\; [1,1] \qquad \texttt{\textbackslash}\langle n\rangle \;:\; [\alpha_{\mathrm{N}},\infty] \qquad \langle class\rangle \;:\; [1,1]$$

`\b` : [0, 0] `\B` : [0, 0]

`\p{`⟨*property*⟩`}` : [1, 6] `\P{`⟨*property*⟩`}` : [1, 6]

3.7.3 Regexp Constructors

```
(regexp? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a regexp value created by `regexp` or `pregexp`, `#f` otherwise.

```
(pregexp? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a regexp value created by `pregexp` (not `regexp`), `#f` otherwise.

```
(byte-regexp? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a regexp value created by `byte-regexp` or `byte-pregexp`, `#f` otherwise.

```
(byte-pregexp? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a regexp value created by `byte-pregexp` (not `byte-regexp`), `#f` otherwise.

```
(regexp str) → regexp?
  str : string?
```

Takes a string representation of a regular expression (using the syntax in §3.7.1 “Regexp Syntax”) and compiles it into a regexp value. Other regular expression procedures accept either a string or a regexp value as the matching pattern. If a regular expression string is used multiple times, it is faster to compile the string once to a regexp value and use it for repeated matches instead of using the string each time.

The `object-name` procedure returns the source string for a regexp value.

Examples:

```
> (regexp "ap*le")
#rx"ap*le"
> (object-name #rx"ap*le")
"ap*le"
```

```
(pregexp string) → pregexp?
  string : string?
```

Like regexp, except that it uses a slightly different syntax (see §3.7.1 “Regexp Syntax”). The result can be used with regexp-match, etc., just like the result from regexp.

Examples:

```
> (pregexp "ap*le")
#px"ap*le"
> (regexp? #px"ap*le")
#t
```

```
(byte-regexp bstr) → byte-regexp?
  bstr : bytes?
```

Takes a byte-string representation of a regular expression (using the syntax in §3.7.1 “Regexp Syntax”) and compiles it into a byte-regexp value.

The object-name procedure returns the source byte string for a regexp value.

Examples:

```
> (byte-regexp #"ap*le")
#rx#"ap*le"
> (object-name #rx#"ap*le")
#"ap*le"
> (byte-regexp "ap*le")
byte-regexp: contract violation
  expected: byte?
  given: "ap*le"
```

```
(byte-pregexp bstr) → byte-pregexp?
  bstr : bytes?
```

Like byte-regexp, except that it uses a slightly different syntax (see §3.7.1 “Regexp Syntax”). The result can be used with regexp-match, etc., just like the result from byte-regexp.

Example:

```
> (byte-pregexp #"ap*le")
#px#"ap*le"
```

```
(regexp-quote str [case-sensitive?]) → string?
  str : string?
```

```
  case-sensitive? : any/c = #t
(regexp-quote bstr [case-sensitive?]) → bytes?
  bstr : bytes?
  case-sensitive? : any/c = #t
```

Produces a string or byte string suitable for use with `regexp` to match the literal sequence of characters in *str* or sequence of bytes in *bstr*. If *case-sensitive?* is true (the default), the resulting regexp matches letters in *str* or `bytes` case-sensitively, otherwise it matches case-insensitively.

Examples:

```
> (regexp-match "." "apple.scm")
'("a")
> (regexp-match (regexp-quote ".") "apple.scm")
'(".")
```

```
(regexp-max-lookbehind pattern) → exact-nonnegative-integer?
  pattern : (or/c regexp? byte-regexp?)
```

Returns the maximum number of bytes that *pattern* may consult before the starting position of a match to determine the match. For example, the pattern `(?<=abc)d` consults three bytes preceding a matching `d`, while `e(?<=a..)d` consults two bytes before a matching `ed`. A `^` pattern may consult a preceding byte to determine whether the current position is the start of the input or of a line.

3.7.4 Regexp Matching

```
(regexp-match  pattern
               input
              [start-pos
               end-pos
               output-port
               input-prefix])
   (if (and (or (string? pattern) (regexp? pattern))
            (or (string? input) (path? input)))
→      (or/c #f (cons/c string? (listof (or/c string? #f))))
       (or/c #f (cons/c bytes?  (listof (or/c bytes?  #f)))))
  pattern : (or/c string? bytes? regexp? byte-regexp?)
  input : (or/c string? bytes? path? input-port?)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : (or/c exact-nonnegative-integer? #f) = #f
  output-port : (or/c output-port? #f) = #f
  input-prefix : bytes? = #""
```

Attempts to match *pattern* (a string, byte string, regexp value, or byte-regexp value) once to a portion of *input*. The matcher finds a portion of *input* that matches and is closest to the start of the input (after *start-pos*).

If *input* is a path, it is converted to a byte string with path->bytes if *pattern* is a byte string or a byte-based regexp. Otherwise, *input* is converted to a string with path->string.

The optional *start-pos* and *end-pos* arguments select a portion of *input* for matching; the default is the entire string or the stream up to an end-of-file. When *input* is a string, *start-pos* is a character position; when *input* is a byte string, then *start-pos* is a byte position; and when *input* is an input port, *start-pos* is the number of bytes to skip before starting to match. The *end-pos* argument can be #f, which corresponds to the end of the string or an end-of-file in the stream; otherwise, it is a character or byte position, like *start-pos*. If *input* is an input port, and if an end-of-file is reached before *start-pos* bytes are skipped, then the match fails.

In *pattern*, a start-of-string ^ refers to the first position of *input* after *start-pos*, assuming that *input-prefix* is #"". The end-of-input $ refers to the *end-pos*th position or (in the case of an input port) an end-of-file, whichever comes first.

The *input-prefix* specifies bytes that effectively precede *input* for the purposes of ^ and other look-behind matching. For example, a #"" prefix means that ^ matches at the beginning of the stream, while a #"\n" *input-prefix* means that a start-of-line ^ can match the beginning of the input, while a start-of-file ^ cannot.

If the match fails, #f is returned. If the match succeeds, a list containing strings or byte string, and possibly #f, is returned. The list contains strings only if *input* is a string and *pattern* is not a byte regexp. Otherwise, the list contains byte strings (substrings of the UTF-8 encoding of *input*, if *input* is a string).

The first [byte] string in a result list is the portion of *input* that matched *pattern*. If two portions of *input* can match *pattern*, then the match that starts earliest is found.

Additional [byte] strings are returned in the list if *pattern* contains parenthesized sub-expressions (but not when the opening parenthesis is followed by ?). Matches for the sub-expressions are provided in the order of the opening parentheses in *pattern*. When sub-expressions occur in branches of an | "or" pattern, in a * "zero or more" pattern, or other places where the overall pattern can succeed without a match for the sub-expression, then a #f is returned for the sub-expression if it did not contribute to the final match. When a single sub-expression occurs within a * "zero or more" pattern or other multiple-match positions, then the rightmost match associated with the sub-expression is returned in the list.

If the optional *output-port* is provided as an output port, the part of *input* from its beginning (not *start-pos*) that precedes the match is written to the port. All of *input* up to *end-pos* is written to the port if no match is found. This functionality is most useful when *input* is an input port.

When matching an input port, a match failure reads up to *end-pos* bytes (or end-of-file), even if *pattern* begins with a start-of-string `^`; see also `regexp-try-match`. On success, all bytes up to and including the match are eventually read from the port, but matching proceeds by first peeking bytes from the port (using `peek-bytes-avail!`), and then (re-)reading matching bytes to discard them after the match result is determined. Non-matching bytes may be read and discarded before the match is determined. The matcher peeks in blocking mode only as far as necessary to determine a match, but it may peek extra bytes to fill an internal buffer if immediately available (i.e., without blocking). Greedy repeat operators in *pattern*, such as `*` or `+`, tend to force reading the entire content of the port (up to *end-pos*) to determine a match.

If the input port is read simultaneously by another thread, or if the port is a custom port with inconsistent reading and peeking procedures (see §12.1.9 “Custom Ports”), then the bytes that are peeked and used for matching may be different than the bytes read and discarded after the match completes; the matcher inspects only the peeked bytes. To avoid such interleaving, use `regexp-match-peek` (with a `progress-evt` argument) followed by `port-commit-peeked`.

Examples:

```
> (regexp-match #rx"x." "12x4x6")
'("x4")
> (regexp-match #rx"y." "12x4x6")
#f
> (regexp-match #rx"x." "12x4x6" 3)
'("x6")
> (regexp-match #rx"x." "12x4x6" 3 4)
#f
> (regexp-match #rx#"x." "12x4x6")
'(#"x4")
> (regexp-match #rx"x." "12x4x6" 0 #f (current-output-port))
12
'("x4")
> (regexp-match #rx"(-[0-9]*)+" "a-12--345b")
'("-12--345" "-345")
```

```
(regexp-match* pattern
               input
              [start-pos
               end-pos
               input-prefix
               #:match-select match-select
               #:gap-select? gap-select])
  (if (and (or (string? pattern) (regexp? pattern))
           (or (string? input) (path? input)))
→     (listof (or/c string? (listof (or/c #f string?))))
      (listof (or/c bytes? (listof (or/c #f bytes?)))))
```

```
pattern : (or/c string? bytes? regexp? byte-regexp?)
input : (or/c string? bytes? path? input-port?)
start-pos : exact-nonnegative-integer? = 0
end-pos : (or/c exact-nonnegative-integer? #f) = #f
input-prefix : bytes? = #""
match-select : (or/c (list? . -> . (or/c any/c list?)) #f) = car
gap-select : any/c = #f
```

Like `regexp-match`, but the result is a list of strings or byte strings corresponding to a sequence of matches of *pattern* in *input*.

The *pattern* is used in order to find matches, where each match attempt starts at the end of the last match, and `^` is allowed to match the beginning of the input (if *input-prefix* is `#""`) only for the first match. Empty matches are handled like other matches, returning a zero-length string or byte sequence (they are more useful in making this a complement of `regexp-split`), but *pattern* is restricted from matching an empty sequence immediately after an empty match.

If *input* contains no matches (in the range *start-pos* to *end-pos*), `null` is returned. Otherwise, each item in the resulting list is a distinct substring or byte sequence from *input* that matches *pattern*. The *end-pos* argument can be `#f` to match to the end of *input* (which corresponds to an end-of-file if *input* is an input port).

Examples:

```
> (regexp-match* #rx"x." "12x4x6")
'("x4" "x6")
> (regexp-match* #rx"x*" "12x4x6")
'("" "" "x" "" "x" "" "")
```

match-select specifies the collected results. The default of `car` means that the result is the list of matches without returning parenthesized sub-patterns. It can be given as a 'selector' function which chooses an item from a list, or it can choose a list of items. For example, you can use `cdr` to get a list of lists of parenthesized sub-patterns matches, or `values` (as an identity function) to get the full matches as well. (Note that the selector must choose an element of its input list or a list of elements, but it must not inspect its input as they can be either a list of strings or a list of position pairs. Furthermore, the selector must be consistent in its choice(s).)

Examples:

```
> (regexp-match* #rx"x(.)" "12x4x6" #:match-select cadr)
'("4" "6")
> (regexp-match* #rx"x(.)" "12x4x6" #:match-select values)
'(("x4" "4") ("x6" "6"))
```

In addition, specifying *gap-select* as a non-#f value will make the result an interleaved list of the matches as well as the separators between them matches, starting and ending with a separator. In this case, *match-select* can be given as #f to return *only* the separators, making such uses equivalent to regexp-split.

Examples:

```
> (regexp-match* #rx"x(.)" "12x4x6" #:match-select cadr #:gap-
select? #t)
'("12" "4" "" "6" "")
> (regexp-match* #rx"x(.)" "12x4x6" #:match-select #f #:gap-
select? #t)
'("12" "" "")
```

```
(regexp-try-match  pattern
                   input
                  [start-pos
                   end-pos
                   output-port
                   input-prefix])
   (if (and (or (string? pattern) (regexp? pattern))
            (string? input))
→      (or/c #f (cons/c string? (listof (or/c string? #f))))
       (or/c #f (cons/c bytes?  (listof (or/c bytes?  #f)))))
 pattern : (or/c string? bytes? regexp? byte-regexp?)
 input : input-port?
 start-pos : exact-nonnegative-integer? = 0
 end-pos : (or/c exact-nonnegative-integer? #f) = #f
 output-port : (or/c output-port? #f) = #f
 input-prefix : bytes? = #""
```

Like regexp-match on input ports, except that if the match fails, no characters are read and discarded from in.

This procedure is especially useful with a *pattern* that begins with a start-of-string ^ or with a non-#f *end-pos*, since each limits the amount of peeking into the port. Otherwise, beware that a large portion of the stream may be peeked (and therefore pulled into memory) before the match succeeds or fails.

```
(regexp-match-positions  pattern
                         input
                        [start-pos
                         end-pos
                         output-port
                         input-prefix])
```

```
   (or/c (cons/c (cons/c exact-nonnegative-integer?
                         exact-nonnegative-integer?)
                 (listof (or/c (cons/c exact-nonnegative-integer?
→                                       exact-nonnegative-integer?)
                               #f)))
         #f)
 pattern : (or/c string? bytes? regexp? byte-regexp?)
 input : (or/c string? bytes? path? input-port?)
 start-pos : exact-nonnegative-integer? = 0
 end-pos : (or/c exact-nonnegative-integer? #f) = #f
 output-port : (or/c output-port? #f) = #f
 input-prefix : bytes? = #""
```

Like `regexp-match`, but returns a list of number pairs (and `#f`) instead of a list of strings. Each pair of numbers refers to a range of characters or bytes in *input*. If the result for the same arguments with `regexp-match` would be a list of byte strings, the resulting ranges correspond to byte ranges; in that case, if *input* is a character string, the byte ranges correspond to bytes in the UTF-8 encoding of the string.

Range results are returned in a `substring`- and `subbytes`-compatible manner, independent of *start-pos*. In the case of an input port, the returned positions indicate the number of bytes that were read, including *start-pos*, before the first matching byte.

Examples:

```
> (regexp-match-positions #rx"x." "12x4x6")
'((2 . 4))
> (regexp-match-positions #rx"x." "12x4x6" 3)
'((4 . 6))
> (regexp-match-positions #rx"(-[0-9]*)+" "a-12--345b")
'((1 . 9) (5 . 9))
```

```
(regexp-match-positions* pattern
                         input
                        [start-pos
                         end-pos
                         input-prefix
                         #:match-select match-select])
   (or/c (listof (cons/c exact-nonnegative-integer?
                         exact-nonnegative-integer?))
→        (listof (listof (or/c #f (cons/c exact-nonnegative-integer?
                                          exact-nonnegative-integer?)))))
 pattern : (or/c string? bytes? regexp? byte-regexp?)
 input : (or/c string? bytes? path? input-port?)
 start-pos : exact-nonnegative-integer? = 0
 end-pos : (or/c exact-nonnegative-integer? #f) = #f
 input-prefix : bytes? = #""
```

```
  match-select : (list? . -> . (or/c any/c list?)) = car
```

Like `regexp-match-positions`, but returns multiple matches like `regexp-match*`.

Examples:

```
> (regexp-match-positions* #rx"x." "12x4x6")
'((2 . 4) (4 . 6))
> (regexp-match-positions* #rx"x(.)" "12x4x6" #:match-select cadr)
'((3 . 4) (5 . 6))
```

Note that unlike `regexp-match*`, there is no `#:gap-select?` input keyword, as this information can be easily inferred from the resulting matches.

```
(regexp-match?  pattern
                input
               [start-pos
                end-pos
                output-port
                input-prefix]) → boolean?
  pattern : (or/c string? bytes? regexp? byte-regexp?)
  input : (or/c string? bytes? path? input-port?)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : (or/c exact-nonnegative-integer? #f) = #f
  output-port : (or/c output-port? #f) = #f
  input-prefix : bytes? = #""
```

Like `regexp-match`, but returns merely `#t` when the match succeeds, `#f` otherwise.

Examples:

```
> (regexp-match? #rx"x." "12x4x6")
#t
> (regexp-match? #rx"y." "12x4x6")
#f
```

```
(regexp-match-exact? pattern input) → boolean?
  pattern : (or/c string? bytes? regexp? byte-regexp?)
  input : (or/c string? bytes? path?)
```

Like `regexp-match?`, but `#t` is only returned when the entire content of *input* matches *pattern*.

Examples:

```
> (regexp-match-exact? #rx"x." "12x4x6")
#f
> (regexp-match-exact? #rx"1.*x." "12x4x6")
#t
```

```
(regexp-match-peek  pattern
                    input
                   [start-pos
                    end-pos
                    progress
                    input-prefix])
 → (or/c (cons/c bytes? (listof (or/c bytes? #f)))
         #f)
  pattern : (or/c string? bytes? regexp? byte-regexp?)
  input : input-port?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : (or/c exact-nonnegative-integer? #f) = #f
  progress : (or/c evt #f) = #f
  input-prefix : bytes? = #""
```

Like `regexp-match` on input ports, but only peeks bytes from *input* instead of reading them. Furthermore, instead of an output port, the last optional argument is a progress event for *input* (see `port-progress-evt`). If *progress* becomes ready, then the match stops peeking from *input* and returns `#f`. The *progress* argument can be `#f`, in which case the peek may continue with inconsistent information if another process meanwhile reads from *input*.

Examples:

```
> (define p (open-input-string "a abcd"))

> (regexp-match-peek ".*bc" p)
'(#"a abc")
> (regexp-match-peek ".*bc" p 2)
'(#"abc")
> (regexp-match ".*bc" p 2)
'(#"abc")
> (peek-char p)
#\d
> (regexp-match ".*bc" p)
#f
> (peek-char p)
#<eof>
```

```
(regexp-match-peek-positions  pattern
                              input
                             [start-pos
                              end-pos
                              progress
                              input-prefix])
```

```
   (or/c (cons/c (cons/c exact-nonnegative-integer?
                         exact-nonnegative-integer?)
                 (listof (or/c (cons/c exact-nonnegative-integer?
→                                       exact-nonnegative-integer?)
                               #f)))
         #f)
 pattern : (or/c string? bytes? regexp? byte-regexp?)
 input : input-port?
 start-pos : exact-nonnegative-integer? = 0
 end-pos : (or/c exact-nonnegative-integer? #f) = #f
 progress : (or/c evt #f) = #f
 input-prefix : bytes? = #""
```

Like `regexp-match-positions` on input ports, but only peeks bytes from *input* instead of reading them, and with a *progress* argument like `regexp-match-peek`.

```
(regexp-match-peek-immediate  pattern
                              input
                             [start-pos
                              end-pos
                              progress
                              input-prefix])
   (or/c (cons/c bytes? (listof (or/c bytes? #f)))
→        #f)
 pattern : (or/c string? bytes? regexp? byte-regexp?)
 input : input-port?
 start-pos : exact-nonnegative-integer? = 0
 end-pos : (or/c exact-nonnegative-integer? #f) = #f
 progress : (or/c evt #f) = #f
 input-prefix : bytes? = #""
```

Like `regexp-match-peek`, but it attempts to match only bytes that are available from *input* without blocking. The match fails if not-yet-available characters might be used to match *pattern*.

```
(regexp-match-peek-positions-immediate  pattern
                                        input
                                       [start-pos
                                        end-pos
                                        progress
                                        input-prefix])
   (or/c (cons/c (cons/c exact-nonnegative-integer?
                         exact-nonnegative-integer?)
                 (listof (or/c (cons/c exact-nonnegative-integer?
→                                       exact-nonnegative-integer?)
                               #f)))
         #f)
```

```
  pattern : (or/c string? bytes? regexp? byte-regexp?)
  input : input-port?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : (or/c exact-nonnegative-integer? #f) = #f
  progress : (or/c evt #f) = #f
  input-prefix : bytes? = #""
```

Like regexp-match-peek-positions, but it attempts to match only bytes that are available from *input* without blocking. The match fails if not-yet-available characters might be used to match *pattern*.

```
(regexp-match-peek-positions* pattern
                              input
                             [start-pos
                              end-pos
                              input-prefix
                              #:match-select match-select])
  (or/c (listof (cons/c exact-nonnegative-integer?
                        exact-nonnegative-integer?))
→       (listof (listof (or/c #f (cons/c exact-nonnegative-integer?
                                         exact-nonnegative-integer?)))))
 pattern : (or/c string? bytes? regexp? byte-regexp?)
 input : input-port?
 start-pos : exact-nonnegative-integer? = 0
 end-pos : (or/c exact-nonnegative-integer? #f) = #f
 input-prefix : bytes? = #""
 match-select : (list? . -> . (or/c any/c list?)) = car
```

Like regexp-match-peek-positions, but returns multiple matches like regexp-match-positions*.

```
(regexp-match/end pattern
                  input
                 [start-pos
                  end-pos
                  output-port
                  input-prefix
                  count])
  (if (and (or (string? pattern) (regexp? pattern))
           (or/c (string? input) (path? input)))
→     (or/c #f (cons/c string? (listof (or/c string? #f))))
      (or/c #f (cons/c bytes?  (listof (or/c bytes?  #f)))))
  (or/c #f bytes?)
 pattern : (or/c string? bytes? regexp? byte-regexp?)
 input : (or/c string? bytes? path? input-port?)
```

```
  start-pos : exact-nonnegative-integer? = 0
  end-pos : (or/c exact-nonnegative-integer? #f) = #f
  output-port : (or/c output-port? #f) = #f
  input-prefix : bytes? = #""
  count : nonnegative-exact-integer? = 1
```

Like regexp-match, but with a second result: a byte string of up to *count* bytes that correspond to the input (possibly including the *input-prefix*) leading to the end of the match; the second result is #f if no match is found.

The second result can be useful as an *input-prefix* for attempting a second match on *input* starting from the end of the first match. In that case, use regexp-max-lookbehind to determine an appropriate value for *count*.

```
(regexp-match-positions/end  pattern
                             input
                            [start-pos
                             end-pos
                             input-prefix
                             count])
   (listof (cons/c exact-nonnegative-integer?
→                  exact-nonnegative-integer?))
   (or/c #f bytes?)
 pattern : (or/c string? bytes? regexp? byte-regexp?)
 input : (or/c string? bytes? path? input-port?)
 start-pos : exact-nonnegative-integer? = 0
 end-pos : (or/c exact-nonnegative-integer? #f) = #f
 input-prefix : bytes? = #""
 count : exact-nonnegative-integer? = 1
```

```
(regexp-match-peek-positions/end  pattern
                                  input
                                 [start-pos
                                  end-pos
                                  progress
                                  input-prefix
                                  count])
   (or/c (cons/c (cons/c exact-nonnegative-integer?
                         exact-nonnegative-integer?)
                 (listof (or/c (cons/c exact-nonnegative-integer?
 →                                     exact-nonnegative-integer?)
                               #f)))
         #f)
   (or/c #f bytes?)
  pattern : (or/c string? bytes? regexp? byte-regexp?)
  input : input-port?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : (or/c exact-nonnegative-integer? #f) = #f
  progress : (or/c evt #f) = #f
  input-prefix : bytes? = #""
  count : exact-nonnegative-integer? = 1
(regexp-match-peek-positions-immediate/end  pattern
                                            input
                                           [start-pos
                                            end-pos
                                            progress
                                            input-prefix
                                            count])
   (or/c (cons/c (cons/c exact-nonnegative-integer?
                         exact-nonnegative-integer?)
                 (listof (or/c (cons/c exact-nonnegative-integer?
 →                                     exact-nonnegative-integer?)
                               #f)))
         #f)
   (or/c #f bytes?)
  pattern : (or/c string? bytes? regexp? byte-regexp?)
  input : input-port?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : (or/c exact-nonnegative-integer? #f) = #f
  progress : (or/c evt #f) = #f
  input-prefix : bytes? = #""
  count : exact-nonnegative-integer? = 1
```

Like `regexp-match-positions`, etc., but with a second result like `regexp-match/end`.

3.7.5 Regexp Splitting

```
(regexp-split  pattern
               input
              [start-pos
               end-pos
               input-prefix])
   (if (and (or (string? pattern) (regexp? pattern))
            (string? input))
→       (cons/c string? (listof string?))
        (cons/c bytes? (listof bytes?)))
 pattern : (or/c string? bytes? regexp? byte-regexp?)
 input : (or/c string? bytes? input-port?)
 start-pos : exact-nonnegative-integer? = 0
 end-pos : (or/c exact-nonnegative-integer? #f) = #f
 input-prefix : bytes? = #""
```

The complement of `regexp-match*`: the result is a list of strings (if *pattern* is a string or character regexp and *input* is a string) or byte strings (otherwise) from *input* that are separated by matches to *pattern*. Adjacent matches are separated with `""` or `#""`. Zero-length matches are treated the same as for `regexp-match*`.

If *input* contains no matches (in the range *start-pos* to *end-pos*), the result is a list containing *input*'s content (from *start-pos* to *end-pos*) as a single element. If a match occurs at the beginning of *input* (at *start-pos*), the resulting list will start with an empty string or byte string, and if a match occurs at the end (at *end-pos*), the list will end with an empty string or byte string. The *end-pos* argument can be `#f`, in which case splitting goes to the end of *input* (which corresponds to an end-of-file if *input* is an input port).

Examples:

```
> (regexp-split #rx" +" "12  34")
'("12" "34")
> (regexp-split #rx"." "12  34")
'("" "" "" "" "" "" "")
> (regexp-split #rx"" "12  34")
'("" "1" "2" " " " " " " "3" "4" "")
> (regexp-split #rx" *" "12  34")
'("" "1" "2" "" "3" "4" "")
> (regexp-split #px"\\b" "12, 13 and 14.")
'("" "12" ", " "13" " " "and" " " "14" ".")
> (regexp-split #rx" +" "")
'("")
```

3.7.6 Regexp Substitution

```
(regexp-replace  pattern
                 input
                 insert
                [input-prefix])
   (if (and (or (string? pattern) (regexp? pattern))
            (string? input))
→       string?
        bytes?)
 pattern : (or/c string? bytes? regexp? byte-regexp?)
 input : (or/c string? bytes?)
          (or/c string? bytes?
 insert :       ((string?) () #:rest (listof string?) . ->* . string?)
                ((bytes?) () #:rest (listof bytes?) . ->* . bytes?))
 input-prefix : bytes? = #""
```

Performs a match using *pattern* on *input*, and then returns a string or byte string in which the matching portion of *input* is replaced with *insert*. If *pattern* matches no part of *input*, then *input* is returned unmodified.

The *insert* argument can be either a (byte) string, or a function that returns a (byte) string. In the latter case, the function is applied on the list of values that `regexp-match` would return (i.e., the first argument is the complete match, and then one argument for each parenthesized sub-expression) to obtain a replacement (byte) string.

If *pattern* is a string or character regexp and *input* is a string, then *insert* must be a string or a procedure that accept strings, and the result is a string. If *pattern* is a byte string or byte regexp, or if *input* is a byte string, then *insert* as a string is converted to a byte string, *insert* as a procedure is called with a byte string, and the result is a byte string.

If *insert* contains `&`, then `&` is replaced with the matching portion of *input* before it is substituted into the match's place. If *insert* contains `\`⟨*n*⟩ for some integer ⟨*n*⟩, then it is replaced with the ⟨*n*⟩th matching sub-expression from *input*. A `&` and `\0` are aliases. If the ⟨*n*⟩th sub-expression was not used in the match, or if ⟨*n*⟩ is greater than the number of sub-expressions in *pattern*, then `\`⟨*n*⟩ is replaced with the empty string.

To substitute a literal `&` or `\`, use `\&` and `\\`, respectively, in *insert*. A `\$` in *insert* is equivalent to an empty sequence; this can be used to terminate a number ⟨*n*⟩ following `\`. If a `\` in *insert* is followed by anything other than a digit, `&`, `\`, or `$`, then the `\` by itself is treated as `\0`.

Note that the `\` described in the previous paragraphs is a character or byte of *input*. To write such an *input* as a Racket string literal, an escaping `\` is needed before the `\`. For example, the Racket constant `"\\1"` is `\1`.

Examples:

```
> (regexp-replace "mi" "mi casa" "su")
```

```
"su casa"
> (regexp-replace "mi" "mi casa" string-upcase)
"MI casa"
> (regexp-replace "([Mm])i ([a-zA-Z]*)" "Mi Casa" "\\1y \\2")
"My Casa"
> (regexp-replace "([Mm])i ([a-zA-Z]*)" "mi cerveza Mi Mi Mi"
                  "\\1y \\2")
"my cerveza Mi Mi Mi"
> (regexp-replace #rx"x" "12x4x6" "\\\\")
"12\\4x6"
> (display (regexp-replace #rx"x" "12x4x6" "\\\\"))
12\4x6
```

```
(regexp-replace* pattern
                 input
                 insert
                 [start-pos
                  end-pos
                  input-prefix]) → (or/c string? bytes?)
  pattern : (or/c string? bytes? regexp? byte-regexp?)
  input : (or/c string? bytes?)
          (or/c string? bytes?
  insert :       ((string?) () #:rest (listof string?) . ->* . string?)
                 ((bytes?) () #:rest (listof bytes?) . ->* . bytes?))
  start-pos : exact-nonnegative-integer? = 0
  end-pos : (or/c exact-nonnegative-integer? #f) = #f
  input-prefix : bytes? = #""
```

Like `regexp-replace`, except that every instance of *pattern* in *input* is replaced with *insert*, instead of just the first match. Only non-overlapping instances of *pattern* in *input* are replaced, so instances of *pattern* within inserted strings are *not* replaced recursively. Zero-length matches are treated the same as in `regexp-match*`.

The optional *start-pos* and *end-pos* arguments select a portion of *input* for matching; the default is the entire string or the stream up to an end-of-file.

Examples:

```
> (regexp-replace* "([Mm])i ([a-zA-Z]*)" "mi cerveza Mi Mi Mi"
                   "\\1y \\2")
"my cerveza My Mi Mi"
> (regexp-replace* "([Mm])i ([a-zA-Z]*)" "mi cerveza Mi Mi Mi"
                   (lambda (all one two)
                     (string-append (string-downcase one) "y"
                                    (string-upcase two))))
```

```
"myCERVEZA myMI Mi"
> (regexp-replace* #px"\\w" "hello world" string-upcase 0 5)
"HELLO world"
> (display (regexp-replace* #rx"x" "12x4x6" "\\\\"))
12\4\6
```

```
(regexp-replaces input replacements) → (or/c string? bytes?)
  input : (or/c string? bytes?)
                (listof
                 (list/c (or/c string? bytes? regexp? byte-regexp?)
  replacements :          (or/c string? bytes?
                                ((string?) () #:rest (listof string?) . ->* . string?)
                                ((bytes?) () #:rest (listof bytes?) . ->* . bytes?))))
```

Performs a chain of `regexp-replace*` operations, where each element in *replacements* specifies a replacement as a `(list pattern replacement)`. The replacements are done in order, so later replacements can apply to previous insertions.

Examples:

```
> (regexp-replaces "zero-or-more?"
                   '([#rx"-" "_"] [#rx"(.*)\\?$" "is_\\1"]))
"is_zero_or_more"
> (regexp-replaces "zero-or-more?"
                   '(["e" "o"] ["o" "oo"]))
"zooroo-oor-mooroo?"
```

```
(regexp-replace-quote str) → string?
  str : string?
(regexp-replace-quote bstr) → bytes?
  bstr : bytes?
```

Produces a string suitable for use as the third argument to `regexp-replace` to insert the literal sequence of characters in *str* or bytes in *bstr* as a replacement. Concretely, every `\` and `&` in *str* or *bstr* is protected by a quoting `\`.

Examples:

```
> (regexp-replace "UT" "Go UT!" "A&M")
"Go AUTM!"
> (regexp-replace "UT" "Go UT!" (regexp-replace-quote "A&M"))
"Go A&M!"
```

3.8 Keywords

§3.7 “Keywords” in ***The Racket Guide*** introduces keywords.

A *keyword* is like an interned symbol, but its printed form starts with `#:`, and a keyword cannot be used as an identifier. Furthermore, a keyword by itself is not a valid expression,

though a keyword can be `quoted` to form an expression that produces the symbol.

Two keywords are `eq?` if and only if they print the same (i.e., keywords are always interned).

Like symbols, keywords are only weakly held by the internal keyword table; see §3.6 “Symbols” for more information.

See §1.3.14 “Reading Keywords” for information on `read`ing keywords and §1.4.11 “Printing Keywords” for information on `print`ing keywords.

```
(keyword? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a keyword, `#f` otherwise.

```
(keyword->string keyword) → string?
  keyword : keyword?
```

Returns a string for the `display`ed form of *keyword*, not including the leading `#:`.

```
(string->keyword str) → keyword?
  str : string?
```

Returns a keyword whose `display`ed form is the same as that of *str*, but with a leading `#:`.

```
(keyword<? a-keyword b-keyword ...+) → boolean?
  a-keyword : keyword?
  b-keyword : keyword?
```

Returns `#t` if the arguments are sorted, where the comparison for each pair of keywords is the same as using `keyword->string` and `string<?`.

3.9 Pairs and Lists

§3.8 “Pairs and Lists” in ***The Racket Guide*** introduces pairs and lists.

A *pair* combines exactly two values. The first value is accessed with the `car` procedure, and the second value is accessed with the `cdr` procedure. Pairs are not mutable (but see §3.10 “Mutable Pairs and Lists”).

A *list* is recursively defined: it is either the constant `null`, or it is a pair whose second value is a list.

A list can be used as a single-valued sequence (see §3.14.1 “Sequences”). The elements of the list serve as elements of the sequence. See also `in-list`.

Cyclic data structures can be created using only immutable pairs via `read` or `make-reader-graph`. If starting with a pair and using some number of `cdr`s returns to the starting pair, then the pair is not a list.

See §1.3.5 “Reading Pairs and Lists” for information on `read`ing pairs and lists and §1.4.4 “Printing Pairs and Lists” for information on `print`ing pairs and lists.

3.9.1 Pair Constructors and Selectors

```
(pair? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a pair, `#f` otherwise.

Examples:

```
> (pair? 1)
#f
> (pair? (cons 1 2))
#t
> (pair? (list 1 2))
#t
> (pair? '(1 2))
#t
> (pair? '())
#f
```

```
(null? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is the empty list, `#f` otherwise.

Examples:

```
> (null? 1)
#f
> (null? '(1 2))
#f
> (null? '())
#t
> (null? (cdr (list 1)))
#t
```

```
(cons a d) → pair?
  a : any/c
  d : any/c
```

Returns a newly allocated pair whose first element is `a` and second element is `d`.

Examples:

```
> (cons 1 2)
'(1 . 2)
> (cons 1 '())
'(1)
```

```
(car p) → any/c
  p : pair?
```

Returns the first element of the pair `p`.

Examples:

```
> (car '(1 2))
1
> (car (cons 2 3))
2
```

```
(cdr p) → any/c
  p : pair?
```

Returns the second element of the pair `p`.

Examples:

```
> (cdr '(1 2))
'(2)
> (cdr '(1))
'()
```

```
null : null?
```

The empty list.

Examples:

```
> null
'()
> '()
'()
> (eq? '() null)
#t
```

```
(list? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a list: either the empty list, or a pair whose second element is a list. This procedure effectively takes constant time due to internal caching (so that any necessary traversals of pairs can in principle count as an extra cost of allocating the pairs).

Examples:

```
> (list? '(1 2))
#t
> (list? (cons 1 (cons 2 '())))
#t
> (list? (cons 1 2))
#f
```

```
(list v ...) → list?
  v : any/c
```

Returns a newly allocated list containing the *v*s as its elements.

Examples:

```
> (list 1 2 3 4)
'(1 2 3 4)
> (list (list 1 2) (list 3 4))
'((1 2) (3 4))
```

```
(list* v ... tail) → any/c
  v : any/c
  tail : any/c
```

Like `list`, but the last argument is used as the tail of the result, instead of the final element. The result is a list only if the last argument is a list.

Examples:

```
> (list* 1 2)
'(1 . 2)
> (list* 1 2 (list 3 4))
'(1 2 3 4)
```

```
(build-list n proc) → list?
  n : exact-nonnegative-integer?
  proc : (exact-nonnegative-integer? . -> . any)
```

Creates a list of *n* elements by applying *proc* to the integers from `0` to `(sub1 n)` in order. If *lst* is the resulting list, then `(list-ref lst i)` is the value produced by `(proc i)`.

Examples:

```
> (build-list 10 values)
'(0 1 2 3 4 5 6 7 8 9)
> (build-list 5 (lambda (x) (* x x)))
'(0 1 4 9 16)
```

3.9.2 List Operations

```
(length lst) → exact-nonnegative-integer?
  lst : list?
```

Returns the number of elements in *lst*.

Examples:

```
> (length (list 1 2 3 4))
4
> (length '())
0
```

```
(list-ref lst pos) → any/c
  lst : pair?
  pos : exact-nonnegative-integer?
```

Returns the element of *lst* at position *pos*, where the list's first element is position 0. If the list has *pos* or fewer elements, then the exn:fail:contract exception is raised.

The *lst* argument need not actually be a list; *lst* must merely start with a chain of at least (add1 *pos*) pairs.

Examples:

```
> (list-ref (list 'a 'b 'c) 0)
'a
> (list-ref (list 'a 'b 'c) 1)
'b
> (list-ref (list 'a 'b 'c) 2)
'c
> (list-ref (cons 1 2) 0)
1
```

```
(list-tail lst pos) → any/c
  lst : any/c
  pos : exact-nonnegative-integer?
```

Returns the list after the first *pos* elements of *lst*. If the list has fewer than *pos* elements, then the exn:fail:contract exception is raised.

The *lst* argument need not actually be a list; *lst* must merely start with a chain of at least *pos* pairs.

Examples:

```
> (list-tail (list 1 2 3 4) 2)
'(3 4)
> (list-ref (cons 1 2) 1)
list-ref: index reaches a non-pair
  index: 1
  in: '(1 . 2)
> (list-ref 'not-a-pair 0)
list-ref: contract violation
  expected: pair?
  given: 'not-a-pair
  argument position: 1st
  other arguments...:
   0
```

```
(append lst ...) → list?
  lst : list?
(append lst ... v) → any/c
  lst : list?
  v : any/c
```

When given all list arguments, the result is a list that contains all of the elements of the given lists in order. The last argument is used directly in the tail of the result.

The last argument need not be a list, in which case the result is an "improper list."

Examples:

```
> (append (list 1 2) (list 3 4))
'(1 2 3 4)
> (append (list 1 2) (list 3 4) (list 5 6) (list 7 8))
'(1 2 3 4 5 6 7 8)
```

```
(reverse lst) → list?
  lst : list?
```

Returns a list that has the same elements as *lst*, but in reverse order.

Example:

```
> (reverse (list 1 2 3 4))
'(4 3 2 1)
```

3.9.3 List Iteration

```
(map proc lst ...+) → list?
  proc : procedure?
  lst : list?
```

Applies *proc* to the elements of the *lst*s from the first elements to the last. The *proc* argument must accept the same number of arguments as the number of supplied *lst*s, and all *lst*s must have the same number of elements. The result is a list containing each result of *proc* in order.

Examples:

```
> (map (lambda (number)
         (+ 1 number))
       '(1 2 3 4))
'(2 3 4 5)
> (map (lambda (number1 number2)
         (+ number1 number2))
       '(1 2 3 4)
       '(10 100 1000 10000))
'(11 102 1003 10004)
```

(andmap *proc lst ...+*) → any
  proc : procedure?
  lst : list?

Similar to map in the sense that *proc* is applied to each element of racket[lst], but

- the result is #f if any application of *proc* produces #f, in which case *proc* is not applied to later elements of the *lst*s; and
- the result is that of *proc* applied to the last elements of the *lst*s; more specifically, the application of *proc* to the last elements in the *lst*s is in tail position with respect to the andmap call.

The andmap function is actually closer to foldl than map, since andmap doesn't produce a list. Still, (andmap f (list x y z)) is equivalent to (and (f x) (f y) (f z)) in the same way that (map f (list x y z)) is equivalent to (list (f x) (f y) (f z)).

If the *lst*s are empty, then #t is returned.

Examples:

```
> (andmap positive? '(1 2 3))
#t
> (andmap positive? '(1 2 a))
positive?: contract violation
  expected: real?
  given: 'a
> (andmap positive? '(1 -2 a))
#f
> (andmap + '(1 2 3) '(4 5 6))
9
```

(ormap *proc lst ...+*) → any
  proc : procedure?
  lst : list?

Similar to map in the sense that *proc* is applied to each element of racket[lst], but

- the result is #f if every application of *proc* produces #f; and
- the result is that of the first application of *proc* producing a value other than #f, in which case *proc* is not applied to later elements of the *lst*s; the application of *proc* to the last elements of the *lst*s is in tail position with respect to the ormap call.

To continue the andmap note above, (ormap f (list x y z)) is equivalent to (or (f x) (f y) (f z)).

If the *lst*s are empty, then #f is returned.

Examples:

```
> (ormap eq? '(a b c) '(a b c))
#t
> (ormap positive? '(1 2 a))
#t
> (ormap + '(1 2 3) '(4 5 6))
5
```

(for-each *proc lst* ...+) → void?
  proc : procedure?
  lst : list?

Similar to map, but *proc* is called only for its effect, and its result (which can be any number of values) is ignored.

Example:

```
> (for-each (lambda (arg)
              (printf "Got ~a\n" arg)
              23)
            '(1 2 3 4))
Got 1
Got 2
Got 3
Got 4
```

(foldl *proc init lst* ...+) → any/c
  proc : procedure?
  init : any/c
  lst : list?

Like map, foldl applies a procedure to the elements of one or more lists. Whereas map combines the return values into a list, foldl combines the return values in an arbitrary way that is determined by *proc*.

If `foldl` is called with *n* lists, then *proc* must take *n*+1 arguments. The extra argument is the combined return values so far. The *proc* is initially invoked with the first item of each list, and the final argument is *init*. In subsequent invocations of *proc*, the last argument is the return value from the previous invocation of *proc*. The input *lst*s are traversed from left to right, and the result of the whole `foldl` application is the result of the last application of *proc*. If the *lst*s are empty, the result is *init*.

Unlike `foldr`, `foldl` processes the *lst*s in constant space (plus the space for each call to *proc*).

Examples:

```
> (foldl cons '() '(1 2 3 4))
'(4 3 2 1)
> (foldl + 0 '(1 2 3 4))
10
> (foldl (lambda (a b result)
           (* result (- a b)))
         1
         '(1 2 3)
         '(4 5 6))
-27
```

```
(foldr proc init lst ...+) → any/c
  proc : procedure?
  init : any/c
  lst : list?
```

Like `foldl`, but the lists are traversed from right to left. Unlike `foldl`, `foldr` processes the *lst*s in space proportional to the length of *lst*s (plus the space for each call to *proc*).

Examples:

```
> (foldr cons '() '(1 2 3 4))
'(1 2 3 4)
> (foldr (lambda (v l) (cons (add1 v) l)) '() '(1 2 3 4))
'(2 3 4 5)
```

3.9.4 List Filtering

```
(filter pred lst) → list?
  pred : procedure?
  lst : list?
```

Returns a list with the elements of *lst* for which *pred* produces a true value. The *pred* procedure is applied to each element from first to last.

Example:

```
> (filter positive? '(1 -2 3 4 -5))
'(1 3 4)
```

```
(remove v lst [proc]) → list?
  v : any/c
  lst : list?
  proc : procedure? = equal?
```

Returns a list that is like *lst*, omitting the first element of *lst* that is equal to *v* using the comparison procedure *proc* (which must accept two arguments).

Examples:

```
> (remove 2 (list 1 2 3 2 4))
'(1 3 2 4)
> (remove 2 (list 1 2 3 2 4) =)
'(1 3 2 4)
> (remove '(2) (list '(1) '(2) '(3)))
'((1) (3))
> (remove "2" (list "1" "2" "3"))
'("1" "3")
> (remove #\c (list #\a #\b #\c))
'(#\a #\b)
```

```
(remq v lst) → list?
  v : any/c
  lst : list?
```

Returns (remove *v* *lst* eq?).

Examples:

```
> (remq 2 (list 1 2 3 4 5))
'(1 3 4 5)
> (remq '(2) (list '(1) '(2) '(3)))
'((1) (2) (3))
> (remq "2" (list "1" "2" "3"))
'("1" "3")
> (remq #\c (list #\a #\b #\c))
'(#\a #\b)
```

```
(remv v lst) → list?
  v : any/c
  lst : list?
```

Returns (remove v lst eqv?).

Examples:

```
> (remv 2 (list 1 2 3 4 5))
'(1 3 4 5)
> (remv '(2) (list '(1) '(2) '(3)))
'((1) (2) (3))
> (remv "2" (list "1" "2" "3"))
'("1" "3")
> (remv #\c (list #\a #\b #\c))
'(#\a #\b)
```

```
(remove* v-lst lst [proc]) → list?
  v-lst : list?
  lst : list?
  proc : procedure? = equal?
```

Like remove, but removes from lst every instance of every element of v-lst.

Example:

```
> (remove* (list 1 2) (list 1 2 3 2 4 5 2))
'(3 4 5)
```

```
(remq* v-lst lst) → list?
  v-lst : list?
  lst : list?
```

Returns (remove* v-lst lst eq?).

Example:

```
> (remq* (list 1 2) (list 1 2 3 2 4 5 2))
'(3 4 5)
```

```
(remv* v-lst lst) → list?
  v-lst : list?
  lst : list?
```

Returns (remove* v-lst lst eqv?).

Example:

```
> (remv* (list 1 2) (list 1 2 3 2 4 5 2))
'(3 4 5)
```

```
(sort  lst
       less-than?
      [#:key extract-key
       #:cache-keys? cache-keys?]) → list?
  lst : list?
  less-than? : (any/c any/c . -> . any/c)
  extract-key : (any/c . -> . any/c) = (lambda (x) x)
  cache-keys? : boolean? = #f
```

Returns a list sorted according to the *less-than?* procedure, which takes two elements of *lst* and returns a true value if the first is less (i.e., should be sorted earlier) than the second.

The sort is stable; if two elements of *lst* are "equal" (i.e., `proc` does not return a true value when given the pair in either order), then the elements preserve their relative order from *lst* in the output list. To preserve this guarantee, use `sort` with a strict comparison functions (e.g., `<` or `string<?`; not `<=` or `string<=?`).

The `#:key` argument *extract-key* is used to extract a key value for comparison from each list element. That is, the full comparison procedure is essentially

```
(lambda (x y)
  (less-than? (extract-key x) (extract-key y)))
```

By default, *extract-key* is applied to two list elements for every comparison, but if *cache-keys?* is true, then the *extract-key* function is used exactly once for each list item. Supply a true value for *cache-keys?* when *extract-key* is an expensive operation; for example, if `file-or-directory-modify-seconds` is used to extract a timestamp for every file in a list, then *cache-keys?* should be `#t` to minimize file-system calls, but if *extract-key* is `car`, then *cache-keys?* should be `#f`. As another example, providing *extract-key* as `(lambda (x) (random))` and `#t` for *cache-keys?* effectively shuffles the list.

Examples:

```
> (sort '(1 3 4 2) <)
'(1 2 3 4)
> (sort '("aardvark" "dingo" "cow" "bear") string<?)
'("aardvark" "bear" "cow" "dingo")
> (sort '(("aardvark") ("dingo") ("cow") ("bear"))
        #:key car string<?)
'(("aardvark") ("bear") ("cow") ("dingo"))
```

3.9.5 List Searching

```
(member v lst) → (or/c list? #f)
  v : any/c
  lst : list?
```

Locates the first element of *lst* that is equal? to *v*. If such an element exists, the tail of *lst* starting with that element is returned. Otherwise, the result is #f.

Examples:

```
> (member 2 (list 1 2 3 4))
'(2 3 4)
> (member 9 (list 1 2 3 4))
#f
```

```
(memv v lst) → (or/c list? #f)
  v : any/c
  lst : list?
```

Like member, but finds an element using eqv?.

Examples:

```
> (memv 2 (list 1 2 3 4))
'(2 3 4)
> (memv 9 (list 1 2 3 4))
#f
```

```
(memq v lst) → (or/c list? #f)
  v : any/c
  lst : list?
```

Like member, but finds an element using eq?.

Examples:

```
> (memq 2 (list 1 2 3 4))
'(2 3 4)
> (memq 9 (list 1 2 3 4))
#f
```

```
(memf proc lst) → (or/c list? #f)
  proc : procedure?
  lst : list?
```

Like member, but finds an element using the predicate *proc*; an element is found when *proc* applied to the element returns a true value.

Example:

```
> (memf (lambda (arg)
          (> arg 9))
        '(7 8 9 10 11))
'(10 11)
```

```
(findf proc lst) → any/c
  proc : procedure?
  lst : list?
```

Like memf, but returns the element or #f instead of a tail of *lst* or #f.

Example:

```
> (findf (lambda (arg)
          (> arg 9))
        '(7 8 9 10 11))
10
```

```
(assoc v lst [is-equal?]) → (or/c pair? #f)
  v : any/c
  lst : (listof pair?)
  is-equal? : (any/c any/c -> any/c) = equal?
```

Locates the first element of *lst* whose car is equal to *v* according to *is-equal?*. If such an element exists, the pair (i.e., an element of *lst*) is returned. Otherwise, the result is #f.

Examples:

```
> (assoc 3 (list (list 1 2) (list 3 4) (list 5 6)))
'(3 4)
> (assoc 9 (list (list 1 2) (list 3 4) (list 5 6)))
#f
> (assoc 3.5
         (list (list 1 2) (list 3 4) (list 5 6))
         (lambda (a b) (< (abs (- a b)) 1)))
'(3 4)
```

```
(assv v lst) → (or/c pair? #f)
  v : any/c
  lst : (listof pair?)
```

Like assoc, but finds an element using eqv?.

Example:

```
> (assv 3 (list (list 1 2) (list 3 4) (list 5 6)))
'(3 4)
```

```
(assq v lst) → (or/c pair? #f)
  v : any/c
  lst : (listof pair?)
```

Like assoc, but finds an element using eq?.

Example:

```
> (assq 3 (list (list 1 2) (list 3 4) (list 5 6)))
'(3 4)
```

```
(assf proc lst) → (or/c list? #f)
  proc : procedure?
  lst : list?
```

Like assoc, but finds an element using the predicate *proc*; an element is found when *proc* applied to the car of an *lst* element returns a true value.

Example:

```
> (assf (lambda (arg)
          (> arg 2))
        (list (list 1 2) (list 3 4) (list 5 6)))
'(3 4)
```

3.9.6 Pair Accessor Shorthands

```
(caar v) → any/c
  v : (cons/c pair? any/c)
```

Returns (car (car p))

Example:

```
> (caar '((1 2) 3 4))
1
```

```
(cadr v) → any/c
  v : (cons/c any/c pair?)
```

Returns (car (cdr p))

Example:

```
> (cadr '((1 2) 3 4))
3
```

```
(cdar v) → any/c
  v : (cons/c pair? any/c)
```

Returns `(cdr (car p))`

Example:

```
> (cdar '((7 6 5 4 3 2 1) 8 9))
'(6 5 4 3 2 1)
```

```
(cddr v) → any/c
  v : (cons/c any/c pair?)
```

Returns `(cdr (cdr p))`

Example:

```
> (cddr '(2 1))
'()
```

```
(caaar v) → any/c
  v : (cons/c (cons/c pair? any/c) any/c)
```

Returns `(car (car (car p)))`

Example:

```
> (caaar '(((6 5 4 3 2 1) 7) 8 9))
6
```

```
(caadr v) → any/c
  v : (cons/c any/c (cons/c pair? any/c))
```

Returns `(car (car (cdr p)))`

Example:

```
> (caadr '(9 (7 6 5 4 3 2 1) 8))
7
```

```
(cadar v) → any/c
  v : (cons/c (cons/c any/c pair?) any/c)
```

Returns `(car (cdr (car p)))`

Example:

```
> (cadar '((7 6 5 4 3 2 1) 8 9))
6
```

```
(caddr v) → any/c
  v : (cons/c any/c (cons/c any/c pair?))
```

Returns `(car (cdr (cdr p)))`

Example:

```
> (caddr '(3 2 1))
1
```

```
(cdaar v) → any/c
  v : (cons/c (cons/c pair? any/c) any/c)
```

Returns `(cdr (car (car p)))`

Example:

```
> (cdaar '(((6 5 4 3 2 1) 7) 8 9))
'(5 4 3 2 1)
```

```
(cdadr v) → any/c
  v : (cons/c any/c (cons/c pair? any/c))
```

Returns `(cdr (car (cdr p)))`

Example:

```
> (cdadr '(9 (7 6 5 4 3 2 1) 8))
'(6 5 4 3 2 1)
```

```
(cddar v) → any/c
  v : (cons/c (cons/c any/c pair?) any/c)
```

Returns `(cdr (cdr (car p)))`

Example:

```
> (cddar '((7 6 5 4 3 2 1) 8 9))
'(5 4 3 2 1)
```

```
(cdddr v) → any/c
  v : (cons/c any/c (cons/c any/c pair?))
```

Returns `(cdr (cdr (cdr p)))`

Example:

```
> (cdddr '(3 2 1))
'()
```

```
(caaaar v) → any/c
  v : (cons/c (cons/c (cons/c pair? any/c) any/c) any/c)
```

Returns `(car (car (car (car p))))`

Example:

```
> (caaaar '((((5 4 3 2 1) 6) 7) 8 9))
5
```

```
(caaadr v) → any/c
  v : (cons/c any/c (cons/c (cons/c pair? any/c) any/c))
```

Returns `(car (car (car (cdr p))))`

Example:

```
> (caaadr '(9 ((6 5 4 3 2 1) 7) 8))
6
```

```
(caadar v) → any/c
  v : (cons/c (cons/c any/c (cons/c pair? any/c)) any/c)
```

Returns `(car (car (cdr (car p))))`

Example:

```
> (caadar '((7 (5 4 3 2 1) 6) 8 9))
5
```

```
(caaddr v) → any/c
  v : (cons/c any/c (cons/c any/c (cons/c pair? any/c)))
```

Returns `(car (car (cdr (cdr p))))`

Example:

```
> (caaddr '(9 8 (6 5 4 3 2 1) 7))
6
```

```
(cadaar v) → any/c
  v : (cons/c (cons/c (cons/c any/c pair?) any/c) any/c)
```

Returns `(car (cdr (car (car p))))`

Example:

```
> (cadaar '(((6 5 4 3 2 1) 7) 8 9))
5
```

```
(cadadr v) → any/c
  v : (cons/c any/c (cons/c (cons/c any/c pair?) any/c))
```

Returns `(car (cdr (car (cdr p))))`

Example:

```
> (cadadr '(9 (7 6 5 4 3 2 1) 8))
6
```

```
(caddar v) → any/c
  v : (cons/c (cons/c any/c (cons/c any/c pair?)) any/c)
```

Returns `(car (cdr (cdr (car p))))`

Example:

```
> (caddar '((7 6 5 4 3 2 1) 8 9))
5
```

```
(cadddr v) → any/c
  v : (cons/c any/c (cons/c any/c (cons/c any/c pair?)))
```

Returns `(car (cdr (cdr (cdr p))))`

Example:

```
> (cadddr '(4 3 2 1))
1
```

```
(cdaaar v) → any/c
  v : (cons/c (cons/c (cons/c pair? any/c) any/c) any/c)
```

Returns `(cdr (car (car (car p))))`

Example:

```
> (cdaaar '((((5 4 3 2 1) 6) 7) 8 9))
'(4 3 2 1)
```

```
(cdaadr v) → any/c
  v : (cons/c any/c (cons/c (cons/c pair? any/c) any/c))
```

Returns `(cdr (car (car (cdr p))))`

Example:

```
> (cdaadr '(9 ((6 5 4 3 2 1) 7) 8))
'(5 4 3 2 1)
```

```
(cdadar v) → any/c
  v : (cons/c (cons/c any/c (cons/c pair? any/c)) any/c)
```

Returns `(cdr (car (cdr (car p))))`

Example:

```
> (cdadar '((7 (5 4 3 2 1) 6) 8 9))
'(4 3 2 1)
```

```
(cdaddr v) → any/c
  v : (cons/c any/c (cons/c any/c (cons/c pair? any/c)))
```

Returns `(cdr (car (cdr (cdr p))))`

Example:

```
> (cdaddr '(9 8 (6 5 4 3 2 1) 7))
'(5 4 3 2 1)
```

```
(cddaar v) → any/c
  v : (cons/c (cons/c (cons/c any/c pair?) any/c) any/c)
```

Returns `(cdr (cdr (car (car p))))`

Example:

```
> (cddaar '(((6 5 4 3 2 1) 7) 8 9))
'(4 3 2 1)
```

```
(cddadr v) → any/c
  v : (cons/c any/c (cons/c (cons/c any/c pair?) any/c))
```

Returns `(cdr (cdr (car (cdr p))))`

Example:

```
> (cddadr '(9 (7 6 5 4 3 2 1) 8))
'(5 4 3 2 1)
```

```
(cdddar v) → any/c
  v : (cons/c (cons/c any/c (cons/c any/c pair?)) any/c)
```

Returns `(cdr (cdr (cdr (car p))))`

Example:

```
> (cdddar '((7 6 5 4 3 2 1) 8 9))
'(4 3 2 1)
```

```
(cddddr v) → any/c
  v : (cons/c any/c (cons/c any/c (cons/c any/c pair?)))
```

Returns `(cdr (cdr (cdr (cdr p))))`

Example:

```
> (cddddr '(4 3 2 1))
'()
```

3.9.7 Additional List Functions and Synonyms

```
(require racket/list)
```

The bindings documented in this section are provided by the `racket/list` and `racket` libraries, but not `racket/base`.

```
empty : null?
```

The empty list.

Examples:

```
> empty
'()
> (eq? empty null)
#t
```

```
(cons? v) → boolean?
  v : any/c
```

The same as `(pair? v)`.

Example:

```
> (cons? '(1 2))
#t
```

```
(empty? v) → boolean?
  v : any/c
```

The same as `(null? v)`.

Examples:

```
> (empty? '(1 2))
#f
> (empty? '())
#t
```

```
(first lst) → any/c
  lst : list?
```

The same as `(car lst)`, but only for lists (that are not empty).

Example:

```
> (first '(1 2 3 4 5 6 7 8 9 10))
1
```

```
(rest lst) → list?
  lst : list?
```

The same as `(cdr lst)`, but only for lists (that are not empty).

Example:

```
> (rest '(1 2 3 4 5 6 7 8 9 10))
'(2 3 4 5 6 7 8 9 10)
```

```
(second lst) → any
  lst : list?
```

Returns the second element of the list.

Example:

```
> (second '(1 2 3 4 5 6 7 8 9 10))
2
```

```
(third lst) → any
  lst : list?
```

Returns the third element of the list.

Example:

```
> (third '(1 2 3 4 5 6 7 8 9 10))
3
```

```
(fourth lst) → any
  lst : list?
```

Returns the fourth element of the list.

Example:

```
> (fourth '(1 2 3 4 5 6 7 8 9 10))
4
```

```
(fifth lst) → any
  lst : list?
```

Returns the fifth element of the list.

Example:

```
> (fifth '(1 2 3 4 5 6 7 8 9 10))
5
```

```
(sixth lst) → any
  lst : list?
```

Returns the sixth element of the list.

Example:

```
> (sixth '(1 2 3 4 5 6 7 8 9 10))
6
```

```
(seventh lst) → any
  lst : list?
```

Returns the seventh element of the list.

Example:

```
> (seventh '(1 2 3 4 5 6 7 8 9 10))
7
```

```
(eighth lst) → any
  lst : list?
```

Returns the eighth element of the list.

Example:

```
> (eighth '(1 2 3 4 5 6 7 8 9 10))
8
```

```
(ninth lst) → any
  lst : list?
```

Returns the ninth element of the list.

Example:

```
> (ninth '(1 2 3 4 5 6 7 8 9 10))
9
```

```
(tenth lst) → any
  lst : list?
```

Returns the tenth element of the list.

Example:

```
> (tenth '(1 2 3 4 5 6 7 8 9 10))
10
```

```
(last lst) → any
  lst : list?
```

Returns the last element of the list.

Example:

```
> (last '(1 2 3 4 5 6 7 8 9 10))
10
```

```
(last-pair p) → pair?
  p : pair?
```

Returns the last pair of a (possibly improper) list.

Example:

```
> (last-pair '(1 2 3 4))
'(4)
```

```
(make-list k v) → list?
  k : exact-nonnegative-integer?
  v : any?
```

Returns a newly constructed list of length *k*, holding *v* in all positions.

Example:

```
> (make-list 7 'foo)
'(foo foo foo foo foo foo foo)
```

```
(take lst pos) → list?
  lst : any/c
  pos : exact-nonnegative-integer?
```

Returns a fresh list whose elements are the first *pos* elements of *lst*. If *lst* has fewer than *pos* elements, the `exn:fail:contract` exception is raised.

The *lst* argument need not actually be a list; *lst* must merely start with a chain of at least *pos* pairs.

Examples:

```
> (take '(1 2 3 4) 2)
'(1 2)
> (take 'non-list 0)
'()
```

```
(drop lst pos) → any/c
  lst : any/c
  pos : exact-nonnegative-integer?
```

Just like `list-tail`.

```
(split-at lst pos) → list? any/c
  lst : any/c
  pos : exact-nonnegative-integer?
```

Returns the same result as

```
(values (take lst pos) (drop lst pos))
```

except that it can be faster.

```
(take-right lst pos) → any/c
  lst : any/c
  pos : exact-nonnegative-integer?
```

Returns the `list`'s *pos*-length tail. If *lst* has fewer than *pos* elements, then the `exn:fail:contract` exception is raised.

The *lst* argument need not actually be a list; *lst* must merely end with a chain of at least *pos* pairs.

Examples:

```
> (take-right '(1 2 3 4) 2)
'(3 4)
> (take-right 'non-list 0)
'non-list
```

```
(drop-right lst pos) → list?
  lst : any/c
  pos : exact-nonnegative-integer?
```

Returns a fresh list whose elements are the prefix of *lst*, dropping its *pos*-length tail. If *lst* has fewer than *pos* elements, then the `exn:fail:contract` exception is raised.

The *lst* argument need not actually be a list; *lst* must merely end with a chain of at least *pos* pairs.

Examples:

```
> (drop-right '(1 2 3 4) 2)
'(1 2)
> (drop-right 'non-list 0)
'()
```

```
(split-at-right lst pos) → list? any/c
  lst : any/c
  pos : exact-nonnegative-integer?
```

Returns the same result as

```
(values (drop-right lst pos) (take-right lst pos))
```

except that it can be faster.

Examples:

```
> (split-at-right '(1 2 3 4 5 6) 3)
'(1 2 3)
'(4 5 6)
> (split-at-right '(1 2 3 4 5 6) 4)
'(1 2)
'(3 4 5 6)
```

```
(add-between lst
             v
            [#:before-first before-first
             #:before-last before-last
             #:after-last after-last
             #:splice? splice?])          → list?
  lst : list?
  v : any/c
  before-first : list? = '()
  before-last : any/c = v
  after-last : list? = '()
  splice? : any/c = #f
```

Returns a list with the same elements as *lst*, but with *v* between each pair of elements in *lst*; the last pair of elements will have *before-last* between them, instead of *v* (but *before-last* defaults to *v*).

If *splice?* is true, then *v* and *before-last* should be lists, and the list elements are spliced into the result. In addition, when *splice?* is true, *before-first* and *after-last* are inserted before the first element and after the last element respectively.

Examples:

```
> (add-between '(x y z) 'and)
'(x and y and z)
> (add-between '(x) 'and)
'(x)
> (add-between '("a" "b" "c" "d") "," #:before-last "and")
'("a" "," "b" "," "c" "and" "d")
> (add-between '(x y z) '(-) #:before-last '(- -)
               #:before-first '(begin) #:after-last '(end LF)
               #:splice? #t)
'(begin x - y - - z end LF)
```

```
(append* lst ... lsts) → list?
  lst : list?
  lsts : (listof list?)
(append* lst ... lsts) → any/c
  lst : list?
  lsts : list?
```

Like append, but the last argument is used as a list of arguments for append, so (append* *lst* ... *lsts*) is the same as (apply append *lst* ... *lsts*). In other words, the relationship between append and append* is similar to the one between list and list*.

Examples:

```
> (append* '(a) '(b) '((c) (d)))
'(a b c d)
> (cdr (append* (map (lambda (x) (list ", " x))
                     '("Alpha" "Beta" "Gamma"))))
'("Alpha" ", " "Beta" ", " "Gamma")
```

```
(flatten v) → list?
  v : any/c
```

Flattens an arbitrary S-expression structure of pairs into a list. More precisely, *v* is treated as a binary tree where pairs are interior nodes, and the resulting list contains all of the non-null leaves of the tree in the same order as an inorder traversal.

Examples:

```
> (flatten '((a) b (c (d) . e) ()))
'(a b c d e)
> (flatten 'a)
'(a)
```

```
(remove-duplicates lst
                   [same?
                    #:key extract-key]) → list?
  lst : list?
  same? : (any/c any/c . -> . any/c) = equal?
  extract-key : (any/c . -> . any/c) = (lambda (x) x)
```

Returns a list that has all items in *lst*, but without duplicate items, where *same?* determines whether two elements of the list are equivalent. The resulting list is in the same order as *lst*, and for any item that occurs multiple times, the first one is kept.

The `#:key` argument *extract-key* is used to extract a key value from each list element, so two items are considered equal if `(same? (extract-key x) (extract-key y))` is true.

Examples:

```
> (remove-duplicates '(a b b a))
'(a b)
> (remove-duplicates '(1 2 1.0 0))
'(1 2 1.0 0)
> (remove-duplicates '(1 2 1.0 0) =)
'(1 2 0)
```

```
(filter-map proc lst ...+) → list?
  proc : procedure?
  lst : list?
```

Returns `(filter (lambda (x) x) (map proc lst ...))`, but without building the intermediate list.

Example:

```
> (filter-map (lambda (x) (and (positive? x) x)) '(1 2 3 -2 8))
'(1 2 3 8)
```

```
(count proc lst ...+) → exact-nonnegative-integer?
  proc : procedure?
  lst : list?
```

Returns `(length (filter proc lst ...))`, but without building the intermediate list.

Example:

```
> (count positive? '(1 -1 2 3 -2 5))
4
```

```
(partition pred lst) → list? list?
  pred : procedure?
  lst : list?
```

Similar to `filter`, except that two values are returned: the items for which *pred* returns a true value, and the items for which *pred* returns `#f`.

The result is the same as

```
(values (filter pred lst) (filter (negate pred) lst))
```

but *pred* is applied to each item in *lst* only once.

Example:

```
> (partition even? '(1 2 3 4 5 6))
'(2 4 6)
'(1 3 5)
```

```
(append-map proc lst ...+) → list?
  proc : procedure?
  lst : list?
```

Returns `(append* (map proc lst ...))`.

Example:

```
> (append-map vector->list '(#(1) #(2 3) #(4)))
'(1 2 3 4)
```

```
(filter-not pred lst) → list?
  pred : (any/c . -> . any/c)
  lst : list?
```

Like `filter`, but the meaning of the *pred* predicate is reversed: the result is a list of all items for which *pred* returns `#f`.

Example:

```
> (filter-not even? '(1 2 3 4 5 6))
'(1 3 5)
```

```
(shuffle lst) → list?
  lst : list?
```

Returns a list with all elements from *lst*, randomly shuffled.

Example:

```
> (shuffle '(1 2 3 4 5 6))
'(4 1 2 6 5 3)
```

```
(argmin proc lst) → any/c
  proc : (-> any/c real?)
  lst : (and/c pair? list?)
```

Returns the first element in the list *lst* that minimizes the result of *proc*. Signals an error on an empty list.

Examples:

```
> (argmin car '((3 pears) (1 banana) (2 apples)))
'(1 banana)
> (argmin car '((1 banana) (1 orange)))
'(1 banana)
```

```
(argmax proc lst) → any/c
  proc : (-> any/c real?)
  lst : (and/c pair? list?)
```

Returns the first element in the list *lst* that maximizes the result of *proc*. Signals an error on an empty list.

Examples:

```
> (argmax car '((3 pears) (1 banana) (2 apples)))
'(3 pears)
> (argmax car '((3 pears) (3 oranges)))
'(3 pears)
```

```
(range end) → list?
  end : real?
(range start end [step]) → list?
  start : real?
  end : real?
  step : real? = 1
```

Similar to in-range, but returns lists. Returns a list of numbers starting at *start* and whose successive elements are computed by adding *step* to their predecessor until *end* (excluded) is reached. If no starting point is provided, 0 is used. If no *step* argument is provided, 1 is used.

Examples:

```
> (range 10)
'(0 1 2 3 4 5 6 7 8 9)
> (range 10 20)
'(10 11 12 13 14 15 16 17 18 19)
> (range 20 40 2)
'(20 22 24 26 28 30 32 34 36 38)
> (range 20 10 -1)
'(20 19 18 17 16 15 14 13 12 11)
> (range 10 15 1.5)
'(10 11.5 13.0 14.5)
```

3.9.8 Immutable Cyclic Data

```
(make-reader-graph v) → any/c
  v : any/c
```

Returns a value like `v`, with placeholders created by `make-placeholder` replaced with the values that they contain, and with placeholders created by `make-hash-placeholder` with an immutable hash table. No part of `v` is mutated; instead, parts of `v` are copied as necessary to construct the resulting graph, where at most one copy is created for any given value.

Since the copied values can be immutable, and since the copy is also immutable, `make-reader-graph` can create cycles involving only immutable pairs, vectors, boxes, and hash tables.

Only the following kinds of values are copied and traversed to detect placeholders:

- pairs
- vectors, both mutable and immutable
- boxes, both mutable and immutable
- hash tables, both mutable and immutable
- instances of a prefab structure type
- placeholders created by `make-placeholder` and `make-hash-placeholder`

Due to these restrictions, `make-reader-graph` creates exactly the same sort of cyclic values as `read`.

Example:

```
> (let* ([ph (make-placeholder #f)]
         [x (cons 1 ph)])
    (placeholder-set! ph x)
    (make-reader-graph x))
#0='(1 . #0#)
```

(placeholder? *v*) → boolean?
  v : any/c

Returns #t if *v* is a placeholder created by make-placeholder, #f otherwise.

(make-placeholder *v*) → placeholder?
  v : any/c

Returns a placeholder for use with placeholder-set! and make-reader-graph. The *v* argument supplies the initial value for the placeholder.

(placeholder-set! *ph datum*) → void?
  ph : placeholder?
  datum : any/c

Changes the value of *ph* to v.

(placeholder-get *ph*) → any/c
  ph : placeholder?

Returns the value of *ph*.

(hash-placeholder? *v*) → boolean?
  v : any/c

Returns #t if *v* is a placeholder created by make-hash-placeholder, #f otherwise.

(make-hash-placeholder *assocs*) → hash-placeholder?
  assocs : (listof pair?)

Like make-immutable-hash, but produces a table placeholder for use with make-reader-graph.

(make-hasheq-placeholder *assocs*) → hash-placeholder?
  assocs : (listof pair?)

Like make-immutable-hasheq, but produces a table placeholder for use with make-reader-graph.

(make-hasheqv-placeholder *assocs*) → hash-placeholder?
  assocs : (listof pair?)

Like make-immutable-hasheqv, but produces a table placeholder for use with make-reader-graph.

3.10 Mutable Pairs and Lists

A *mutable pair* is like a pair created by `cons`, but it supports `set-mcar!` and `set-mcdr!` mutation operations to change the parts of the mutable pair (like traditional Lisp and Scheme pairs).

A *mutable list* is analogous to a list created with pairs, but instead created with mutable pairs.

A mutable pair is not a pair; they are completely separate datatypes. Similarly, a mutable list is not a list, except that the empty list is also the empty mutable list. Instead of programming with mutable pairs and mutable lists, data structures such as pairs, lists, and hash tables are practically always better choices.

A mutable list can be used as a single-valued sequence (see §3.14.1 "Sequences"). The elements of the mutable list serve as elements of the sequence. See also `in-mlist`.

3.10.1 Mutable Pair Constructors and Selectors

```
(mpair? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a mutable pair, `#f` otherwise.

```
(mcons a d) → pair?
  a : any/c
  d : any/c
```

Returns a newly allocated mutable pair whose first element is `a` and second element is `d`.

```
(mcar p) → any/c
  p : mpair?
```

Returns the first element of the mutable pair `p`.

```
(mcdr p) → any/c
  p : mpair?
```

Returns the second element of the mutable pair `p`.

```
(set-mcar! p v) → void?
  p : mpair?
  v : any/v
```

Changes the mutable pair `p` so that its first element is `v`.

```
(set-mcdr! p v) → void?
  p : mpair?
  v : any/v
```

Changes the mutable pair *p* so that its second element is *v*.

3.10.2 Mutable List Functions

```
(require racket/mpair)
```

The bindings documented in this section are provided by the racket/mpair library, not racket/base or racket.

For functions described in this section, contracts are not directly enforced. In particular, when a mutable list is expected, supplying any other kind of value (or mutating a value that starts as a mutable list) tends to produce an exception from mcar or mcdr.

```
(mlist? v) → boolean?
  v : any/c
```

Returns #t if *v* is a mutable list: either the empty list, or a mutable pair whose second element is a mutable list.

```
(mlist v ...) → mlist?
  v : any/c
```

Returns a newly allocated mutable list containing the *v*s as its elements.

```
(list->mlist lst) → mlist?
  lst : list?
```

Returns a newly allocated mutable list with the same elements as *lst*.

```
(mlist->list mlst) → list?
  mlst : mlist?
```

Returns a newly allocated list with the same elements as *mlst*.

```
(mlength mlst) → exact-nonnegative-integer?
  mlst : mlist?
```

Returns the number of elements in *mlst*.

```
(mlist-ref mlst pos) → any/c
  mlst : mlist?
  pos : exact-nonnegative-integer?
```

Like `list-ref`, but for mutable lists.

```
(mlist-tail mlst pos) → any/c
  mlst : mlist?
  pos : exact-nonnegative-integer?
```

Like `list-tail`, but for mutable lists.

```
(mappend mlst ...) → mlist?
  mlst : mlist?
(mappend mlst ... v) → any/c
  mlst : mlist?
  v : any/c
```

Like `append`, but for mutable lists.

```
(mappend! mlst ...) → mlist?
  mlst : mlist?
(mappend! mlst ... v) → any/c
  mlst : mlist?
  v : any/c
```

The `mappend!` procedure appends the given mutable lists by mutating the tail of each to refer to the next, using `set-mcdr!`. Empty lists are dropped; in particular, the result of calling `mappend!` with one or more empty lists is the same as the result of the call with the empty lists removed from the set of arguments.

```
(mreverse mlst) → mlist?
  mlst : mlist?
```

Like `reverse`, but for mutable lists.

```
(mreverse! mlst) → mlist?
  mlst : mlist?
```

Like `mreverse`, but destructively reverses the mutable list by using all of the mutable pairs in *mlst* and changing them with `set-mcdr!`.

```
(mmap proc mlst ...+) → mlist?
  proc : procedure?
  mlst : mlist?
```

Like map, but for mutable lists.

```
(mfor-each proc mlst ...+) → void?
  proc : procedure?
  mlst : mlist?
```

Like for-each, but for mutable lists.

```
(mmember v mlst) → (or/c mlist? #f)
  v : any/c
  mlst : mlist?
```

Like member, but for mutable lists.

```
(mmemv v mlst) → (or/c mlist? #f)
  v : any/c
  mlst : mlist?
```

Like memv, but for mutable lists.

```
(mmemq v mlst) → (or/c list? #f)
  v : any/c
  mlst : mlist?
```

Like memq, but for mutable lists.

```
(massoc v mlst) → (or/c mpair? #f)
  v : any/c
  mlst : (mlistof mpair?)
```

Like assoc, but for mutable lists of mutable pairs.

```
(massv v mlst) → (or/c mpair? #f)
  v : any/c
  mlst : (mlistof mpair?)
```

Like assv, but for mutable lists of mutable pairs.

```
(massq v mlst) → (or/c mpair? #f)
  v : any/c
  mlst : (mlistof mpair?)
```

Like assq, but for mutable lists of mutable pairs.

```
(mlistof pred) → (any/c . -> . boolean?)
  pred : (any/c . -> . any/c)
```

Returns a procedure that returns #t when given a mutable list for which *pred* returns a true value for all elements.

3.11 Vectors

§3.9 “Vectors” in ***The Racket Guide*** introduces vectors.

A *vector* is a fixed-length array with constant-time access and update of the vector slots, which are numbered from `0` to one less than the number of slots in the vector.

Two vectors are `equal?` if they have the same length, and if the values in corresponding slots of the vectors are `equal?`.

A vector can be *mutable* or *immutable*. When an immutable vector is provided to a procedure like `vector-set!`, the `exn:fail:contract` exception is raised. Vectors generated by the default reader (see §1.3.6 “Reading Strings”) are immutable.

A vector can be used as a single-valued sequence (see §3.14.1 “Sequences”). The elements of the vector serve as elements of the sequence. See also `in-vector`.

A literal or printed vector starts with `#(`, optionally with a number between the `#` and `(`. See §1.3.9 “Reading Vectors” for information on `read`ing vectors and §1.4.6 “Printing Vectors” for information on `print`ing vectors.

```
(vector? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a vector, `#f` otherwise.

```
(make-vector size [v]) → vector?
  size : exact-nonnegative-integer?
  v : any/c = 0
```

Returns a mutable vector with `size` slots, where all slots are initialized to contain `v`.

```
(vector v ...) → vector?
  v : any/c
```

Returns a newly allocated mutable vector with as many slots as provided `v`s, where the slots are initialized to contain the given `v`s in order.

```
(vector-immutable v ...) → (and/c vector?
                                  immutable?)
  v : any/c
```

Returns a newly allocated immutable vector with as many slots as provided `v`s, where the slots are contain the given `v`s in order.

```
(vector-length vec) → exact-nonnegative-integer?
  vec : vector?
```

Returns the length of *vec* (i.e., the number of slots in the vector).

```
(vector-ref vec pos) → any/c
  vec : vector?
  pos : exact-nonnegative-integer?
```

Returns the element in slot *pos* of *vec*. The first slot is position 0, and the last slot is one less than (vector-length *vec*).

```
(vector-set! vec pos v) → void?
  vec : (and/c vector? (not/c immutable?))
  pos : exact-nonnegative-integer?
  v : any/c
```

Updates the slot *pos* of *vec* to contain *v*.

```
(vector->list vec) → list?
  vec : vector?
```

Returns a list with the same length and elements as *vec*.

```
(list->vector lst) → vector?
  lst : list?
```

Returns a mutable vector with the same length and elements as *lst*.

```
(vector->immutable-vector vec) → (and/c vector? immutable?)
  vec : vector?
```

Returns an immutable vector with the same length and elements as *vec*. If *vec* is itself immutable, then it is returned as the result.

```
(vector-fill! vec v) → void?
  vec : (and/c vector? (not/c immutable?))
  v : any/c
```

Changes all slots of *vec* to contain *v*.

```
(vector-copy!  dest
               dest-start
               src
              [src-start
               src-end])   → void?
  dest : (and/c vector? (not/c immutable?))
  dest-start : exact-nonnegative-integer?
  src : vector?
  src-start : exact-nonnegative-integer? = 0
  src-end : exact-nonnegative-integer? = (vector-length src)
```

Changes the elements of *dest* starting at position *dest-start* to match the elements in *src* from *src-start* (inclusive) to *src-end* (exclusive). The vectors *dest* and *src* can be the same vector, and in that case the destination region can overlap with the source region; the destination elements after the copy match the source elements from before the copy. If any of *dest-start*, *src-start*, or *src-end* are out of range (taking into account the sizes of the vectors and the source and destination regions), the `exn:fail:contract` exception is raised.

Examples:

```
> (define v (vector 'A 'p 'p 'l 'e))

> (vector-copy! v 4 #(y))

> (vector-copy! v 0 v 3 4)

> v
'#(l p p l y)
```

```
(vector->values vec [start-pos end-pos]) → any
  vec : vector?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (vector-length vec)
```

Returns *end-pos* - *start-pos* values, which are the elements of *vec* from *start-pos* (inclusive) to *end-pos* (exclusive). If *start-pos* or *end-pos* are greater than `(vector-length vec)`, or if *end-pos* is less than *start-pos*, the `exn:fail:contract` exception is raised.

```
(build-vector n proc) → vector?
  n : exact-nonnegative-integer?
  proc : (exact-nonnegative-integer? . -> . any/c)
```

Creates a vector of *n* elements by applying *proc* to the integers from `0` to `(sub1 n)` in order. If *vec* is the resulting vector, then `(vector-ref vec i)` is the value produced by `(proc i)`.

Example:

```
> (build-vector 5 add1)
'#(1 2 3 4 5)
```

3.11.1 Additional Vector Functions

```
(require racket/vector)
```

The bindings documented in this section are provided by the racket/vector and racket libraries, but not racket/base.

```
(vector-set*! vec pos v ... ...) → void?
  vec : (and/c vector? (not/c immutable?))
  pos : exact-nonnegative-integer?
  v : any/c
```

Updates each slot *pos* of *vec* to contain each *v*. The update takes place from the left so later updates overwrite earlier updates.

```
(vector-map proc vec ...+) → vector?
  proc : procedure?
  vec : vector?
```

Applies *proc* to the elements of the *vec*s from the first elements to the last. The *proc* argument must accept the same number of arguments as the number of supplied *vec*s, and all *vec*s must have the same number of elements. The result is a fresh vector containing each result of *proc* in order.

Example:

```
> (vector-map + #(1 2) #(3 4))
'#(4 6)
```

```
(vector-map! proc vec ...+) → vector?
  proc : procedure?
  vec : vector?
```

Applies *proc* to the elements of the *vec*s from the first elements to the last. The *proc* argument must accept the same number of arguments as the number of supplied *vec*s, and all *vec*s must have the same number of elements. The each result of *proc* is inserted into the first *vec* at the index that the arguments to *proc* was taken from. The result is the first *vec*.

Examples:

```
> (define v #(1 2 3 4))

> (vector-map! add1 v)
'#(2 3 4 5)
> v
'#(2 3 4 5)
```

```
(vector-append vec ...) → vector?
  vec : vector?
```

Creates a fresh vector that contains all of the elements of the given vectors in order.

Example:

```
> (vector-append #(1 2) #(3 4))
'#(1 2 3 4)
```

```
(vector-take vec pos) → vector?
  vec : vector?
  pos : exact-nonnegative-integer?
```

Returns a fresh vector whose elements are the first *pos* elements of *vec*. If *vec* has fewer than *pos* elements, then the `exn:fail:contract` exception is raised.

Example:

```
> (vector-take #(1 2 3 4) 2)
'#(1 2)
```

```
(vector-take-right vec pos) → vector?
  vec : vector?
  pos : exact-nonnegative-integer?
```

Returns a fresh vector whose elements are the last *pos* elements of *vec*. If *vec* has fewer than *pos* elements, then the `exn:fail:contract` exception is raised.

Example:

```
> (vector-take-right #(1 2 3 4) 2)
'#(3 4)
```

```
(vector-drop vec pos) → vector?
  vec : vector?
  pos : exact-nonnegative-integer?
```

Returns a fresh vector whose elements are the elements of *vec* after the first *pos* elements. If *vec* has fewer than *pos* elements, then the `exn:fail:contract` exception is raised.

Example:

```
> (vector-drop #(1 2 3 4) 2)
'#(3 4)
```

```
(vector-drop-right vec pos) → vector?
  vec : vector?
  pos : exact-nonnegative-integer?
```

Returns a fresh vector whose elements are the elements of *vec* before the first *pos* elements. If *vec* has fewer than *pos* elements, then the `exn:fail:contract` exception is raised.

Example:

```
> (vector-drop-right #(1 2 3 4) 2)
'#(1 2)
```

```
(vector-split-at vec pos) → vector? vector?
  vec : vector?
  pos : exact-nonnegative-integer?
```

Returns the same result as

```
(values (vector-take vec pos) (vector-drop vec pos))
```

except that it can be faster.

Example:

```
> (vector-split-at #(1 2 3 4 5) 2)
'#(1 2)
'#(3 4 5)
```

```
(vector-split-at-right vec pos) → vector? vector?
  vec : vector?
  pos : exact-nonnegative-integer?
```

Returns the same result as

```
(values (vector-take-right vec pos) (vector-drop-right vec pos))
```

except that it can be faster.

Example:

```
> (vector-split-at-right #(1 2 3 4 5) 2)
'#(1 2 3)
'#(4 5)
```

```
(vector-copy vec [start end]) → vector?
  vec : vector?
  start : exact-nonnegative-integer? = 0
  end : exact-nonnegative-integer? = (vector-length v)
```

Creates a fresh vector of size `(- end start)`, with all of the elements of *vec* from *start* (inclusive) to *end* (exclusive).

Examples:

```
> (vector-copy #(1 2 3 4))
'#(1 2 3 4)
> (vector-copy #(1 2 3 4) 3)
'#(4)
> (vector-copy #(1 2 3 4) 2 3)
'#(3)
```

(vector-filter *pred* *vec*) → vector?
  pred : procedure?
  vec : vector?

Returns a fresh vector with the elements of *vec* for which *pred* produces a true value. The *pred* procedure is applied to each element from first to last.

Example:

```
> (vector-filter even? #(1 2 3 4 5 6))
'#(2 4 6)
```

(vector-filter-not *pred* *vec*) → vector?
  pred : procedure?
  vec : vector?

Like vector-filter, but the meaning of the *pred* predicate is reversed: the result is a vector of all items for which *pred* returns #f.

Example:

```
> (vector-filter-not even? #(1 2 3 4 5 6))
'#(1 3 5)
```

(vector-count *proc* *vec* ...+) → exact-nonnegative-integer?
  proc : procedure?
  vec : vector?

Returns the number of elements of the *vec* ... (taken in parallel) on which *proc* does not evaluate to #f.

Examples:

```
> (vector-count even? #(1 2 3 4 5))
2
> (vector-count = #(1 2 3 4 5) #(5 4 3 2 1))
1
```

```
(vector-argmin proc vec) → any/c
  proc : (-> any/c real?)
  vec : vector?
```

This returns the first element in the non-empty vector *vec* that minimizes the result of *proc*.

Examples:

```
> (vector-argmin car #((3 pears) (1 banana) (2 apples)))
'(1 banana)
> (vector-argmin car #((1 banana) (1 orange)))
'(1 banana)
```

```
(vector-argmax proc vec) → any/c
  proc : (-> any/c real?)
  vec : vector?
```

This returns the first element in the non-empty vector *vec* that maximizes the result of *proc*.

Examples:

```
> (vector-argmax car #((3 pears) (1 banana) (2 apples)))
'(3 pears)
> (vector-argmax car #((3 pears) (3 oranges)))
'(3 pears)
```

```
(vector-member v lst) → (or/c natural-number/c #f)
  v : any/c
  lst : vector?
```

Locates the first element of `vec` that is `equal?` to `v`. If such an element exists, the index of that element in `vec` is returned. Otherwise, the result is `#f`.

Examples:

```
> (vector-member 2 (vector 1 2 3 4))
1
> (vector-member 9 (vector 1 2 3 4))
#f
```

```
(vector-memv v vec) → (or/c natural-number/c #f)
  v : any/c
  vec : vector?
```

Like `vector-member`, but finds an element using `eqv?`.

Examples:

```
> (vector-memv 2 (vector 1 2 3 4))
1
> (vector-memv 9 (vector 1 2 3 4))
#f
```

```
(vector-memq v vec) → (or/c natural-number/c #f)
  v : any/c
  vec : vector?
```

Like `vector-member`, but finds an element using `eq?`.

Examples:

```
> (vector-memq 2 (vector 1 2 3 4))
1
> (vector-memq 9 (vector 1 2 3 4))
#f
```

3.12 Boxes

§3.11 “Boxes” in ***The Racket Guide*** introduces boxes.

A *box* is like a single-element vector, normally used as minimal mutable storage.

A literal or printed box starts with `#&`. See §1.3.12 “Reading Boxes” for information on `read`ing boxes and §1.4.9 “Printing Boxes” for information on `print`ing boxes.

```
(box? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a box, `#f` otherwise.

```
(box v) → box?
  v : any/c
```

Returns a new mutable box that contains `v`.

```
(box-immutable v) → (and/c box? immutable?)
  v : any/c
```

Returns a new immutable box that contains `v`.

```
(unbox box) → any/c
  box : box?
```

Returns the content of *box*.

For any v, (unbox (box v)) returns v.

```
(set-box! box v) → void?
  box : (and/c box? (not/c immutable?))
  v : any/c
```

Sets the content of *box* to *v*.

```
(box-cas! box old new) → boolean?
  box : (and/c box? (not/c immutable?) (not/c impersonator?))
  old : any/c
  new : any/c
```

Atomically updates the contents of *box* to *new*, provided that *box* currently contains a value that is eq? to *old*, and returns #t in that case. If *box* does not contain *old*, then the result is #f.

If no other threads or futures attempt to access *box*, the operation is equivalent to

```
(and (eq? old (unbox loc)) (set-box! loc new) #t)
```

When Racket is compiled with support for futures, box-cas! uses a hardware *compare and set* operation. Uses of box-cas! be performed safely in a future (i.e., allowing the future thunk to continue in parallel).

3.13 Hash Tables

§3.10 “Hash Tables” in ***The Racket Guide*** introduces hash tables.

A *hash table* (or simply *hash*) maps each of its keys to a single value. For a given hash table, keys are equivalent via equal?, eqv?, or eq?, and keys are retained either strongly or weakly (see §15.1 “Weak Boxes”). A hash table is also either mutable or immutable. Immutable hash tables support effectively constant-time access and update, just like mutable hash tables; the constant on immutable operations is usually larger, but the functional nature of immutable hash tables can pay off in certain algorithms.

Immutable hash tables actually provide *O*(*log N*) access and update. Since *N* is limited by the address space so that *log N* is limited to less than 30 or 62 (depending on the platform), *log N* can be treated reasonably as a constant.

A hash table can be used as a two-valued sequence (see §3.14.1 “Sequences”). The keys and values of the hash table serve as elements of the sequence (i.e., each element is a key and its associated value). If a mapping is added to or removed from the hash table during iteration, then an iteration step may fail with exn:fail:contract, or the iteration may skip or duplicate keys and values. See also in-hash, in-hash-keys, in-hash-values, and in-hash-pairs.

Two hash tables cannot be equal? unless they use the same key-comparison procedure (equal?, eqv?, or eq?), both hold keys strongly or weakly, and have the same mutability.

Caveats concerning concurrent modification: A mutable hash table can be manipulated with `hash-ref`, `hash-set!`, and `hash-remove!` concurrently by multiple threads, and the operations are protected by a table-specific semaphore as needed. Three caveats apply, however:

- If a thread is terminated while applying `hash-ref`, `hash-set!`, `hash-remove!`, `hash-ref!`, or `hash-update!` to a hash table that uses `equal?` or `eqv?` key comparisons, all current and future operations on the hash table may block indefinitely.
- The `hash-map` and `hash-for-each` procedures do not use the table's semaphore to guard the traversal as a whole. Changes by one thread to a hash table can affect the keys and values seen by another thread part-way through its traversal of the same hash table.
- The `hash-update!` and `hash-ref!` functions use a table's semaphore independently for the `hash-ref` and `hash-set!` parts of their functionality, which means that the update as a whole is not "atomic."

Caveat concerning mutable keys: If a key in an `equal?`-based hash table is mutated (e.g., a key string is modified with `string-set!`), then the hash table's behavior for insertion and lookup operations becomes unpredictable.

A literal or printed hash table starts with `#hash`, `#hasheqv`, or `#hasheq`. See §1.3.11 "Reading Hash Tables" for information on `read`ing hash tables and §1.4.8 "Printing Hash Tables" for information on `print`ing hash tables.

```
(hash? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a hash table, `#f` otherwise.

```
(hash-equal? hash) → boolean?
  hash : hash?
```

Returns `#t` if *hash* compares keys with `equal?`, `#f` if it compares with `eq?` or `eqv?`.

```
(hash-eqv? hash) → boolean?
  hash : hash?
```

Returns `#t` if *hash* compares keys with `eqv?`, `#f` if it compares with `equal?` or `eq?`.

```
(hash-eq? hash) → boolean?
  hash : hash?
```

Returns `#t` if *hash* compares keys with `eq?`, `#f` if it compares with `equal?` or `eqv?`.

```
(hash-weak? hash) → boolean?
  hash : hash?
```

Returns `#t` if *hash* retains its keys weakly, `#f` if it retains keys strongly.

```
(hash key val ... ...) → (and/c hash? hash-equal? immutable?)
  key : any/c
  val : any/c
(hasheq key val ... ...) → (and/c hash? hash-eq? immutable?)
  key : any/c
  val : any/c
(hasheqv key val ... ...) → (and/c hash? hash-eqv? immutable?)
  key : any/c
  val : any/c
```

Creates an immutable hash table with each given *key* mapped to the following *val*; each *key* must have a *val*, so the total number of arguments to `hash` must be even.

The `hash` procedure creates a table where keys are compared with `equal?`, `hasheq` procedure creates a table where keys are compared with `eq?`, and `hasheqv` procedure creates a table where keys are compared with `eqv?`.

The *key* to *val* mappings are added to the table in the order that they appear in the argument list, so later mappings can hide earlier mappings if the *key*s are equal.

```
(make-hash [assocs]) → (and/c hash? hash-equal?)
  assocs : (listof pair?) = null
(make-hasheqv [assocs]) → (and/c hash? hash-eqv?)
  assocs : (listof pair?) = null
(make-hasheq [assocs]) → (and/c hash? hash-eq?)
  assocs : (listof pair?) = null
```

Creates a mutable hash table that holds keys strongly.

The `make-hash` procedure creates a table where keys are compared with `equal?`, `make-hasheq` procedure creates a table where keys are compared with `eq?`, and `make-hasheqv` procedure creates a table where keys are compared with `eqv?`.

The table is initialized with the content of *assocs*. In each element of *assocs*, the `car` is a key, and the `cdr` is the corresponding value. The mappings are added to the table in the order that they appear in *assocs*, so later mappings can hide earlier mappings.

See also `make-custom-hash`.

```
(make-weak-hash [assocs]) → (and/c hash? hash-equal? hash-weak?)
  assocs : (listof pair?) = null
```

```
(make-weak-hasheqv [assocs]) → (and/c hash? hash-eqv? hash-weak?)
  assocs : (listof pair?) = null
(make-weak-hasheq [assocs]) → (and/c hash? hash-eq? hash-weak?)
  assocs : (listof pair?) = null
```

Like make-hash, make-hasheq, and make-hasheqv, but creates a mutable hash table that holds keys weakly.

```
(make-immutable-hash [assocs])
 → (and/c hash? hash-equal? immutable?)
  assocs : (listof pair?) = null
(make-immutable-hasheqv [assocs])
 → (and/c hash? hash-eqv? immutable?)
  assocs : (listof pair?) = null
(make-immutable-hasheq [assocs])
 → (and/c hash? hash-eq? immutable?)
  assocs : (listof pair?) = null
```

Like hash, hasheq, and hasheqv, but accepts the key–value mapping in association-list form like make-hash, make-hasheq, and make-hasheqv.

```
(hash-set! hash key v) → void?
  hash : (and/c hash? (not/c immutable?))
  key : any/c
  v : any/c
```

Maps *key* to *v* in *hash*, overwriting any existing mapping for *key*.

See also the caveats concerning concurrent modification and the caveat concerning mutable keys above.

```
(hash-set*! hash key v ... ...) → void?
  hash : (and/c hash? (not/c immutable?))
  key : any/c
  v : any/c
```

Maps each *key* to each *v* in *hash*, overwriting any existing mapping for each *key*. Mappings are added from the left, so later mappings overwrite earlier mappings.

See also the caveats concerning concurrent modification and the caveat concerning mutable keys above.

```
(hash-set hash key v) → (and/c hash? immutable?)
  hash : (and/c hash? immutable?)
  key : any/c
  v : any/c
```

Functionally extends *hash* by mapping *key* to *v*, overwriting any existing mapping for *key*, and returning the extended hash table.

See also the caveat concerning mutable keys above.

```
(hash-set* hash key v ... ...) → (and/c hash? immutable?)
  hash : (and/c hash? immutable?)
  key : any/c
  v : any/c
```

Functionally extends *hash* by mapping each *key* to *v*, overwriting any existing mapping for each *key*, and returning the extended hash table. Mappings are added from the left, so later mappings overwrite earlier mappings.

See also the caveat concerning mutable keys above.

```
(hash-ref hash key [failure-result]) → any
  hash : hash?
  key : any/c
  failure-result : any/c
                 = (lambda ()
                     (raise (make-exn:fail:contract ....)))
```

Returns the value for *key* in *hash*. If no value is found for *key*, then *failure-result* determines the result:

- If *failure-result* is a procedure, it is called (through a tail call) with no arguments to produce the result.
- Otherwise, *failure-result* is returned as the result.

See also the caveats concerning concurrent modification and the caveat concerning mutable keys above.

```
(hash-ref! hash key to-set) → any
  hash : hash?
  key : any/c
  to-set : any/c
```

Returns the value for *key* in *hash*. If no value is found for *key*, then *to-set* determines the result as in hash-ref (i.e., it is either a thunk that computes a value or a plain value), and this result is stored in *hash* for the *key*. (Note that if *to-set* is a thunk, it is not invoked in tail position.)

See also the caveats concerning concurrent modification and the caveat concerning mutable keys above.

```
(hash-has-key? hash key) → boolean?
  hash : hash?
  key : any/c
```

Returns `#t` if *hash* contains a value for the given *key*, `#f` otherwise.

```
(hash-update! hash
              key
              updater
             [failure-result]) → void?
  hash : (and/c hash? (not/c immutable?))
  key : any/c
  updater : (any/c . -> . any/c)
  failure-result : any/c
                 = (lambda ()
                     (raise (make-exn:fail:contract ....)))
```

Composes `hash-ref` and `hash-set!` to update an existing mapping in *hash*, where the optional *failure-result* argument is used as in `hash-ref` when no mapping exists for *key* already. See the caveat above about concurrent updates.

See also the caveats concerning concurrent modification and the caveat concerning mutable keys above.

```
(hash-update hash key updater [failure-result])
 → (and/c hash? immutable?)
  hash : (and/c hash? immutable?)
  key : any/c
  updater : (any/c . -> . any/c)
  failure-result : any/c
                 = (lambda ()
                     (raise (make-exn:fail:contract ....)))
```

Composes `hash-ref` and `hash-set` to functionally update an existing mapping in *hash*, where the optional *failure-result* argument is used as in `hash-ref` when no mapping exists for *key* already.

See also the caveat concerning mutable keys above.

```
(hash-remove! hash key) → void?
  hash : (and/c hash? (not/c immutable?))
  key : any/c
```

Removes any existing mapping for *key* in *hash*.

See also the caveats concerning concurrent modification and the caveat concerning mutable keys above.

```
(hash-remove hash key) → (and/c hash? immutable?)
  hash : (and/c hash? immutable?)
  key : any/c
```

Functionally removes any existing mapping for *key* in *hash*, returning the fresh hash table.

See also the caveat concerning mutable keys above.

```
(hash-map hash proc) → (listof any/c)
  hash : hash?
  proc : (any/c any/c . -> . any/c)
```

Applies the procedure *proc* to each element in *hash* in an unspecified order, accumulating the results into a list. The procedure *proc* is called each time with a key and its value.

If a hash table is extended with new keys (either through *proc* or by another thread) while a hash-map or hash-for-each traversal is in process, arbitrary key–value pairs can be dropped or duplicated in the traversal. Key mappings can be deleted or remapped (by any thread) with no adverse affects; the change does not affect a traversal if the key has been seen already, otherwise the traversal skips a deleted key or uses the remapped key's new value.

See also the caveats concerning concurrent modification above.

```
(hash-keys hash) → (listof any/c)
  hash : hash?
```

Returns a list of the keys of *hash* in an unspecified order.

See hash-map for information about modifying *hash* during hash-keys.

See also the caveats concerning concurrent modification above.

```
(hash-values hash) → (listof any/c)
  hash : hash?
```

Returns a list of the values of *hash* in an unspecified order.

See hash-map for information about modifying *hash* during hash-values.

See also the caveats concerning concurrent modification above.

```
(hash->list hash) → (listof (cons/c any/c any/c))
  hash : hash?
```

Returns a list of the key–value pairs of *hash* in an unspecified order.

See hash-map for information about modifying *hash* during hash->list.

See also the caveats concerning concurrent modification above.

```
(hash-for-each hash proc) → void?
  hash : hash?
  proc : (any/c any/c . -> . any)
```

Applies *proc* to each element in *hash* (for the side-effects of *proc*) in an unspecified order. The procedure *proc* is called each time with a key and its value.

See hash-map for information about modifying *hash* within *proc*.

See also the caveats concerning concurrent modification above.

```
(hash-count hash) → exact-nonnegative-integer?
  hash : hash?
```

Returns the number of keys mapped by *hash*. Unless *hash* retains keys weakly, the result is computed in constant time and atomically. If *hash* retains it keys weakly, a traversal is required to count the keys.

```
(hash-iterate-first hash)
 → (or/c #f exact-nonnegative-integer?)
  hash : hash?
```

Returns #f if *hash* contains no elements, otherwise it returns an integer that is an index to the first element in the hash table; "first" refers to an unspecified ordering of the table elements, and the index values are not necessarily consecutive integers. For a mutable *hash*, this index is guaranteed to refer to the first item only as long as no items are added to or removed from *hash*.

```
(hash-iterate-next hash pos)
 → (or/c #f exact-nonnegative-integer?)
  hash : hash?
  pos : exact-nonnegative-integer?
```

Returns either an integer that is an index to the element in *hash* after the element indexed by *pos* (which is not necessarily one more than *pos*) or #f if *pos* refers to the last element in *hash*. If *pos* is not a valid index, then the exn:fail:contract exception is raised. For a mutable *hash*, the result index is guaranteed to refer to its item only as long as no items are added to or removed from *hash*.

```
(hash-iterate-key hash pos) → any
  hash : hash?
  pos : exact-nonnegative-integer?
```

Returns the key for the element in *hash* at index *pos*. If *pos* is not a valid index for *hash*, the exn:fail:contract exception is raised.

```
(hash-iterate-value hash pos) → any
  hash : hash?
  pos : exact-nonnegative-integer?
```

Returns the value for the element in *hash* at index *pos*. If *pos* is not a valid index for *hash*, the exn:fail:contract exception is raised.

```
(hash-copy hash) → (and/c hash? (not/c immutable?))
  hash : hash?
```

Returns a mutable hash table with the same mappings, same key-comparison mode, and same key-holding strength as *hash*.

```
(eq-hash-code v) → fixnum?
  v : any/c
```

Returns a fixnum; for any two calls with eq? values, the returned number is the same.

Equal fixnums are always eq?.

```
(eqv-hash-code v) → fixnum?
  v : any/c
```

Returns a fixnum; for any two calls with eqv? values, the returned number is the same.

```
(equal-hash-code v) → fixnum?
  v : any/c
```

Returns a fixnum; for any two calls with equal? values, the returned number is the same. A hash code is computed even when *v* contains a cycle through pairs, vectors, boxes, and/or inspectable structure fields. See also gen:equal+hash.

```
(equal-secondary-hash-code v) → fixnum?
  v : any/c
```

Like equal-hash-code, but computes a secondary value suitable for use in double hashing.

3.14 Sequences and Streams

Sequences and streams abstract over iteration of elements in a collection. Sequences allow iteration with `for` macros or with sequence operations such as sequence-map. Streams are functional sequences that can be used either in a generic way or a stream-specific way. Generators are closely related stateful objects that can be converted to a sequence and vice-versa.

3.14.1 Sequences

§11.1 “Sequence Constructors” in ***The Racket Guide*** introduces sequences.

A *sequence* encapsulates an ordered collection of values. The elements of a sequence can be extracted with one of the `for` syntactic forms, with the procedures returned by `sequence-generate`, or by converting the sequence into a stream.

The sequence datatype overlaps with many other datatypes. Among built-in datatypes, the sequence datatype includes the following:

- exact nonnegative integers (see below)
- strings (see §3.3 “Strings”)
- byte strings (see §3.4 “Byte Strings”)
- lists (see §3.9 “Pairs and Lists”)
- mutable lists (see §3.10 “Mutable Pairs and Lists”)
- vectors (see §3.11 “Vectors”)
- hash tables (see §3.13 “Hash Tables”)
- dictionaries (see §3.15 “Dictionaries”)
- sets (see §3.16 “Sets”)
- input ports (see §12.1 “Ports”)
- streams (see §3.14.2 “Streams”)

An exact number *k* that is a non-negative integer acts as a sequence similar to `(in-range k)`, except that *k* by itself is not a stream.

Custom sequences can be defined using structure type properties. The easiest method to define a custom sequence is to use the `gen:stream` generic interface. Streams are a suitable abstraction for data structures that are directly iterable. For example, a list is directly iterable with `first` and `rest`. On the other hand, vectors are not directly iterable: iteration has to go through an index. For data structures that are not directly iterable, the *iterator* for the data structure can be defined to be a stream (e.g., a structure containing the index of a vector).

For example, unrolled linked lists (represented as a list of vectors) themeselves do not fit the stream abstraction, but have index-based iterators that can be represented as streams:

Examples:

```
> (struct unrolled-list-iterator (idx lst)
    #:methods gen:stream
    [(define (stream-empty? iter)
```

```
    (define lst (unrolled-list-iterator-lst iter))
    (or (null? lst)
        (and (>= (unrolled-list-iterator-idx iter)
                 (vector-length (first lst)))
             (null? (rest lst)))))
  (define (stream-first iter)
    (vector-ref (first (unrolled-list-iterator-lst iter))
                (unrolled-list-iterator-idx iter)))
  (define (stream-rest iter)
    (define idx (unrolled-list-iterator-idx iter))
    (define lst (unrolled-list-iterator-lst iter))
    (if (>= idx (sub1 (vector-length (first lst))))
        (unrolled-list-iterator 0 (rest lst))
        (unrolled-list-iterator (add1 idx) lst)))])

> (define (make-unrolled-list-iterator ul)
    (unrolled-list-iterator 0 (unrolled-list-lov ul)))

> (struct unrolled-list (lov)
    #:property prop:sequence
    make-unrolled-list-iterator)

> (define ul1 (unrolled-list '(#(cracker biscuit) #(cookie scone))))

> (for/list ([x ul1]) x)
'(cracker biscuit cookie scone)
```

The `prop:sequence` property provides more flexibility in specifying iteration, such as when a pre-processing step is needed to prepare the data for iteration. The `make-do-sequence` function creates a sequence given a thunk that returns procedures to implement a sequence, and the `prop:sequence` property can be associated with a structure type to implement its implicit conversion to a sequence.

For most sequence types, extracting elements from a sequence has no side-effect on the original sequence value; for example, extracting the sequence of elements from a list does not change the list. For other sequence types, each extraction implies a side effect; for example, extracting the sequence of bytes from a port causes the bytes to be read from the port. A sequence's state may either span all uses of the sequence, as for a port, or it may be confined to each distinct time that a sequence is *initiate*d by a `for` form, `sequence->stream`, `sequence-generate`, or `sequence-generate*`. Concretely, the thunk passed to `make-do-sequence` is called to initiate the sequence each time the sequence is used.

Individual elements of a sequence typically correspond to single values, but an element may also correspond to multiple values. For example, a hash table generates two values—a key and its value—for each element in the sequence.

Sequence Predicate and Constructors

```
(sequence? v) → boolean?
  v : any/c
```

Returns `#t` if *v* can be used as a sequence, `#f` otherwise.

```
(in-range end) → stream?
  end : number?
(in-range start end [step]) → stream?
  start : number?
  end : number?
  step : number? = 1
```

Returns a sequence (that is also a stream) whose elements are numbers. The single-argument case `(in-range end)` is equivalent to `(in-range 0 end 1)`. The first number in the sequence is *start*, and each successive element is generated by adding *step* to the previous element. The sequence stops before an element that would be greater or equal to *end* if *step* is non-negative, or less or equal to *end* if *step* is negative.

An `in-range` application can provide better performance for number iteration when it appears directly in a `for` clause.

```
(in-naturals [start]) → stream?
  start : exact-nonnegative-integer? = 0
```

Returns an infinite sequence (that is also a stream) of exact integers starting with *start*, where each element is one more than the preceding element.

An `in-naturals` application can provide better performance for integer iteration when it appears directly in a `for` clause.

```
(in-list lst) → stream?
  lst : list?
```

Returns a sequence (that is also a stream) that is equivalent to using *lst* directly as a sequence.

See §3.9 “Pairs and Lists” for information on using lists as sequences.

An `in-list` application can provide better performance for list iteration when it appears directly in a `for` clause.

```
(in-mlist mlst) → sequence?
  mlst : mlist?
```

Returns a sequence equivalent to *mlst*.

See §3.10 “Mutable Pairs and Lists” for information on using mutable lists as sequences.

An in-mlist application can provide better performance for mutable list iteration when it appears directly in a for clause.

```
(in-vector vec [start stop step]) → sequence?
  vec : vector?
  start : exact-nonnegative-integer? = 0
  stop : (or/c exact-integer? #f) = #f
  step : (and/c exact-integer? (not/c zero?)) = 1
```

Returns a sequence equivalent to *vec* when no optional arguments are supplied.

See §3.11 “Vectors” for information on using vectors as sequences.

The optional arguments *start*, *stop*, and *step* are analogous to in-range, except that a #f value for *stop* is equivalent to (vector-length *vec*). That is, the first element in the sequence is (vector-ref *vec* *start*), and each successive element is generated by adding *step* to index of the previous element. The sequence stops before an index that would be greater or equal to end if *step* is non-negative, or less or equal to end if *step* is negative.

If *start* is not a valid index, or *stop* is not in [-1, (vector-length *vec*)] then the exn:fail:contract exception is raised. If *start* is less than *stop* and *step* is negative, then the exn:fail:contract:mismatch exception is raised. Similarly, if *start* is more than *stop* and *step* is positive, then the exn:fail:contract:mismatch exception is raised.

An in-vector application can provide better performance for vector iteration when it appears directly in a for clause.

```
(in-string str [start stop step]) → sequence?
  str : string?
  start : exact-nonnegative-integer? = 0
  stop : (or/c exact-integer? #f) = #f
  step : (and/c exact-integer? (not/c zero?)) = 1
```

Returns a sequence equivalent to *str* when no optional arguments are supplied.

See §3.3 “Strings” for information on using strings as sequences.

The optional arguments *start*, *stop*, and *step* are as in in-vector.

An in-string application can provide better performance for string iteration when it appears directly in a for clause.

```
(in-bytes bstr [start stop step]) → sequence?
  bstr : bytes?
  start : exact-nonnegative-integer? = 0
  stop : (or/c exact-integer? #f) = #f
  step : (and/c exact-integer? (not/c zero?)) = 1
```

Returns a sequence equivalent to *bstr* when no optional arguments are supplied.

See §3.4 “Byte Strings” for information on using byte strings as sequences.

The optional arguments *start*, *stop*, and *step* are as in in-vector.

An in-bytes application can provide better performance for byte string iteration when it appears directly in a for clause.

```
(in-port [r in]) → sequence?
  r : (input-port? . -> . any/c) = read
  in : input-port? = (current-input-port)
```

Returns a sequence whose elements are produced by calling *r* on *in* until it produces eof.

```
(in-input-port-bytes in) → sequence?
  in : input-port?
```

Returns a sequence equivalent to (in-port read-byte *in*).

```
(in-input-port-chars in) → sequence?
  in : input-port?
```

Returns a sequence whose elements are read as characters from *in* (equivalent to (in-port read-char *in*)).

```
(in-lines [in mode]) → sequence?
  in : input-port? = (current-input-port)
  mode : (or/c 'linefeed 'return 'return-linefeed 'any 'any-one)
       = 'any
```

Returns a sequence equivalent to (in-port (lambda (p) (read-line p *mode*)) *in*). Note that the default mode is 'any, whereas the default mode of read-line is 'linefeed.

```
(in-bytes-lines [in mode]) → sequence?
  in : input-port? = (current-input-port)
  mode : (or/c 'linefeed 'return 'return-linefeed 'any 'any-one)
       = 'any
```

Returns a sequence equivalent to (in-port (lambda (p) (read-bytes-line p *mode*)) *in*). Note that the default mode is 'any, whereas the default mode of read-bytes-line is 'linefeed.

```
(in-hash hash) → sequence?
  hash : hash?
```

Returns a sequence equivalent to *hash*.

Examples:

```
> (define table (hash 'a 1 'b 2))

> (for ([(key value) (in-hash table)])
    (printf "key: ~a value: ~a\n" key value))
key: b value: 2
key: a value: 1
```

See §3.13 “Hash Tables” for information on using hash tables as sequences.

```
(in-hash-keys hash) → sequence?
  hash : hash?
```

Returns a sequence whose elements are the keys of *hash*.

Examples:

```
> (define table (hash 'a 1 'b 2))

> (for ([key (in-hash-keys table)])
    (printf "key: ~a\n" key))
key: b
key: a
```

```
(in-hash-values hash) → sequence?
  hash : hash?
```

Returns a sequence whose elements are the values of *hash*.

Examples:

```
> (define table (hash 'a 1 'b 2))

> (for ([value (in-hash-values table)])
    (printf "value: ~a\n" value))
value: 2
value: 1
```

```
(in-hash-pairs hash) → sequence?
  hash : hash?
```

Returns a sequence whose elements are pairs, each containing a key and its value from *hash* (as opposed to using *hash* directly as a sequence to get the key and value as separate values for each element).

Examples:

```
> (define table (hash 'a 1 'b 2))

> (for ([key+value (in-hash-pairs table)])
    (printf "key and value: ~a\n" key+value))
key and value: (b . 2)
key and value: (a . 1)
```

```
(in-directory [dir]) → sequence?
  dir : (or/c #f path-string?) = #f
```

Returns a sequence that produces all of the paths for files, directories, and links within *dir*. If *dir* is not `#f`, then every produced path starts with *dir* as its prefix. If *dir* is `#f`, then paths in and relative to the current directory are produced.

An `in-directory` sequence traverses nested subdirectories recursively. To generate a sequence that includes only the immediate content of a directory, use the result of `directory-list` as a sequence.

```
(in-producer producer stop args ...) → sequence?
  producer : procedure?
  stop : any/c
  args : any/c
```

Returns a sequence that contains values from sequential calls to *producer*. A *stop* value returned by *producer* marks the end of the sequence (and the *stop* value is not included in the sequence); *stop* can be a predicate that is applied to the results of *producer*, or it can be a value that is tested against the result of with `eq?`. (The *stop* argument must be a predicate if the stop value is itself a function or if *producer* returns multiple values.)

```
(in-value v) → sequence?
  v : any/c
```

Returns a sequence that produces a single value: *v*. This form is mostly useful for `let`-like bindings in forms such as `for*/list`.

```
(in-indexed seq) → sequence?
  seq : sequence?
```

Returns a sequence where each element has two values: the value produced by *seq*, and a non-negative exact integer starting with `0`. The elements of *seq* must be single-valued.

```
(in-sequences seq ...) → sequence?
  seq : sequence?
```

Returns a sequence that is made of all input sequences, one after the other. Each *seq* is initiated only after the preceding *seq* is exhausted. If a single *seq* is provided, then *seq* is returned; otherwise, the elements of each *seq* must all have the same number of values.

```
(in-cycle seq ...) → sequence?
  seq : sequence?
```

Similar to `in-sequences`, but the sequences are repeated in an infinite cycle, where each *seq* is initiated afresh in each iteration. Beware that if no *seq*s are provided or if all *seq*s become empty, then the sequence produced by `in-cycle` never returns when an element is demanded—or even when the sequence is initiated, if all *seq*s are initially empty.

```
(in-parallel seq ...) → sequence?
  seq : sequence?
```

Returns a sequence where each element has as many values as the number of supplied *seq*s; the values, in order, are the values of each *seq*. The elements of each *seq* must be single-valued.

```
(in-values-sequence seq) → sequence?
  seq : sequence?
```

Returns a sequence that is like *seq*, but it combines multiple values for each element from *seq* as a list of elements.

```
(in-values*-sequence seq) → sequence?
  seq : sequence?
```

Returns a sequence that is like *seq*, but when an element of *seq* has multiple values or a single list value, then the values are combined in a list. In other words, `in-values*-sequence` is like `in-values-sequence`, except that non-list, single-valued elements are not wrapped in a list.

```
(stop-before seq pred) → sequence?
  seq : sequence?
  pred : (any/c . -> . any)
```

Returns a sequence that contains the elements of *seq* (which must be single-valued), but only until the last element for which applying *pred* to the element produces `#t`, after which the sequence ends.

```
(stop-after seq pred) → sequence?
  seq : sequence?
  pred : (any/c . -> . any)
```

Returns a sequence that contains the elements of *seq* (which must be single-valued), but only until the element (inclusive) for which applying *pred* to the element produces `#t`, after which the sequence ends.

```
(make-do-sequence thunk) → sequence?
          (-> (values (any/c . -> . any)
                      (any/c . -> . any/c)
  thunk :             any/c
                      (or/c (any/c . -> . any/c) #f)
                      (or/c (() () #:rest list? . ->* . any/c) #f)
                      (or/c ((any/c) () #:rest list? . ->* . any/c) #f)))
```

Returns a sequence whose elements are generated by the procedures and initial value returned by the thunk, which is called to initiate the sequence. The initiated sequence is defined in terms of a *position*, which is initialized to the third result of the thunk, and the *element*, which may consist of multiple values.

The *thunk* results define the generated elements as follows:

- The first result is a *pos->element* procedure that takes the current position and returns the value(s) for the current element.
- The second result is a *next-pos* procedure that takes the current position and returns the next position.
- The third result is the initial position.
- The fourth result is a *continue-with-pos?* function that takes the current position and returns a true result if the sequence includes the value(s) for the current position, and false if the sequence should end instead of including the value(s). Alternatively, the fourth result can be `#f` to indicate that the sequence should always include the current value(s).
- The fifth result is a *continue-with-val?* function that is like the fourth result, but it takes the current element value(s) instead of the current position. Alternatively, the fifth result can be `#f` to indicate that the sequence should always include the value(s) at the current position.
- The sixth result is a *continue-after-pos+val?* procedure that takes both the current position and the current element value(s) and determines whether the sequence ends after the current element is already included in the sequence. Alternatively, the sixth result can be `#f` to indicate that the sequence can always continue after the current value(s).

Each of the procedures listed above is called only once per position. Among the last three procedures, as soon as one of the procedures returns `#f`, the sequence ends, and none are called again. Typically, one of the functions determines the end condition, and `#f` is used in place of the other two functions.

prop:sequence : struct-type-property?

Associates a procedure to a structure type that takes an instance of the structure and returns a sequence. If `v` is an instance of a structure type with this property, then `(sequence? v)` produces `#t`.

Using a pre-existing sequence:

Examples:

```
> (struct my-set (table)
    #:property prop:sequence
    (lambda (s)
      (in-hash-keys (my-set-table s))))

> (define (make-set . xs)
    (my-set (for/hash ([x (in-list xs)])
              (values x #t))))

> (for/list ([c (make-set 'celeriac 'carrot 'potato)])
    c)
'(potato carrot celeriac)
```

Using `make-do-sequence`:

Examples:

```
> (define-struct train (car next)
    #:property prop:sequence (lambda (t)
                               (make-do-sequence
                                (lambda ()
                                  (values train-car
                                          train-next
                                          t
                                          (lambda (t) t)
                                          (lambda (v) #t)
                                          (lambda (t v) #t))))))

> (for/list ([c (make-train 'engine
                            (make-train 'boxcar
                                        (make-train 'caboose
                                                    #f)))])
    c)
'(engine boxcar caboose)
```

Sequence Conversion

```
(sequence->stream seq) → stream?
  seq : sequence?
```

Coverts a sequence to a stream, which supports the `stream-first` and `stream-rest` operations. Creation of the stream eagerly initiates the sequence, but the stream lazily draws elements from the sequence, caching each element so that `stream-first` produces the same result each time is applied to a stream.

If extracting an element from *seq* involves a side-effect, then the effect is performed each time that either `stream-first` or `stream-rest` is first used to access or skip an element.

```
(sequence-generate seq) → (-> boolean?) (-> any)
  seq : sequence?
```

Initiates a sequence and returns two thunks to extract elements from the sequence. The first returns `#t` if more values are available for the sequence. The second returns the next element (which may be multiple values) from the sequence; if no more elements are available, the `exn:fail:contract` exception is raised.

```
(sequence-generate* seq)
  →  (or/c list? #f)
     (-> (values (or/c list? #f) procedure?))
  seq : sequence?
```

Like `sequence-generate`, but avoids state (aside from any inherent in the sequence) by returning a list of values for the sequence's first element—or `#f` if the sequence is empty—and a thunk to continue with the sequence; the result of the thunk is the same as the result of `sequence-generate*`, but for the second element of the sequence, and so on. If the thunk is called when the element result is `#f` (indicating no further values in the sequence), the `exn:fail:contract` exception is raised.

Sequence Combinations

```
(require racket/sequence)
```

The bindings documented in this section are provided by the `racket/sequence` and `racket` libraries, but not `racket/base`.

```
empty-sequence : sequence?
```

A sequence with no elements.

```
(sequence->list s) → list?
  s : sequence?
```

Returns a list whose elements are the elements of `s`, each of which must be a single value. If `s` is infinite, this function does not terminate.

```
(sequence-length s) → exact-nonnegative-integer?
  s : sequence?
```

Returns the number of elements of `s` by extracting and discarding all of them. If `s` is infinite, this function does not terminate.

```
(sequence-ref s i) → any
  s : sequence?
  i : exact-nonnegative-integer?
```

Returns the `i`th element of `s` (which may be multiple values).

```
(sequence-tail s i) → sequence?
  s : sequence?
  i : exact-nonnegative-integer?
```

Returns a sequence equivalent to `s`, except that the first `i` elements are omitted.

In case initiating `s` involves a side effect, the sequence `s` is not initiated until the resulting sequence is initiated, at which point the first `i` elements are extracted from the sequence.

```
(sequence-append s ...) → sequence?
  s : sequence?
```

Returns a sequence that contains all elements of each sequence in the order they appear in the original sequences. The new sequence is constructed lazily.

If all given `s`s are streams, the result is also a stream.

```
(sequence-map f s) → sequence?
  f : procedure?
  s : sequence?
```

Returns a sequence that contains `f` applied to each element of `s`. The new sequence is constructed lazily.

If `s` is a stream, then the result is also a stream.

```
(sequence-andmap f s) → boolean?
  f : (-> any/c ... boolean?)
  s : sequence?
```

Returns `#t` if `f` returns a true result on every element of `s`. If `s` is infinite and `f` never returns a false result, this function does not terminate.

```
(sequence-ormap f s) → boolean?
  f : (-> any/c ... boolean?)
  s : sequence?
```

Returns `#t` if `f` returns a true result on some element of `s`. If `s` is infinite and `f` never returns a true result, this function does not terminate.

```
(sequence-for-each f s) → void?
  f : (-> any/c ... any)
  s : sequence?
```

Applies `f` to each element of `s`. If `s` is infinite, this function does not terminate.

```
(sequence-fold f i s) → any/c
  f : (-> any/c any/c ... any/c)
  i : any/c
  s : sequence?
```

Folds `f` over each element of `s` with `i` as the initial accumulator. If `s` is infinite, this function does not terminate. The `f` function takes the accumulator as its first argument and the next sequence element as its second.

```
(sequence-count f s) → exact-nonnegative-integer?
  f : procedure?
  s : sequence?
```

Returns the number of elements in `s` for which `f` returns a true result. If `s` is infinite, this function does not terminate.

```
(sequence-filter f s) → sequence?
  f : (-> any/c ... boolean?)
  s : sequence?
```

Returns a sequence whose elements are the elements of `s` for which `f` returns a true result. Although the new sequence is constructed lazily, if `s` has an infinite number of elements where `f` returns a false result in between two elements where `f` returns a true result, then operations on this sequence will not terminate during the infinite sub-sequence.

If `s` is a stream, then the result is also a stream.

```
(sequence-add-between s e) → sequence?
  s : sequence?
  e : any/c
```

Returns a sequence whose elements are the elements of `s`, but with `e` between each pair of elements in `s`. The new sequence is constructed lazily.

If `s` is a stream, then the result is also a stream.

3.14.2 Streams

A *stream* is a kind of sequence that supports functional iteration via `stream-first` and `stream-rest`. The `stream-cons` form constructs a lazy stream, but plain lists can be used as stream, and functions such as `in-range` and `in-naturals` also create streams.

```
(require racket/stream)
```

The bindings documented in this section are provided by the `racket/stream` and `racket` libraries, but not `racket/base`.

```
(stream? v) → boolean?
  v : any/c
```

Returns `#t` if `v` can be used as a stream, `#f` otherwise.

```
(stream-empty? s) → boolean?
  s : stream?
```

Returns `#f` if `s` has no elements, `#f` otherwise.

```
(stream-first s) → any
  s : (and/c stream? (not/c stream-empty?))
```

Returns the value(s) of the first element in `s`.

```
(stream-rest s) → stream?
  s : (and/c stream? (not/c stream-empty?))
```

Returns a stream that is equivalent to `s` without its first element.

```
(stream-cons first-expr rest-expr)
```

Produces a lazy stream for which `stream-first` forces the evaluation of *first-expr* to produce the first element of the stream, and `stream-rest` forces the evaluation of *rest-expr* to produce a stream for the rest of the returned stream.

The first element of the stream as produced by *first-expr* must be a single value. The *rest-expr* must produce a stream when it is evaluated, otherwise the `exn:fail:contract?` exception is raised.

```
(stream expr ...)
```

A shorthand for nested `stream-cons`es ending with `empty-stream`.

```
(in-stream s) → sequence?
  s : stream?
```

Returns a sequence that is equivalent to *s*.

An `in-stream` application can provide better performance for streams iteration when it appears directly in a `for` clause.

```
empty-stream : stream?
```

A stream with no elements.

```
(stream->list s) → list?
  s : stream?
```

Returns a list whose elements are the elements of *s*, each of which must be a single value. If *s* is infinite, this function does not terminate.

```
(stream-length s) → exact-nonnegative-integer?
  s : stream?
```

Returns the number of elements of *s*. If *s* is infinite, this function does not terminate.

In the case of lazy streams, this function forces evaluation only of the sub-streams, and not the stream's elements.

```
(stream-ref s i) → any
  s : stream?
  i : exact-nonnegative-integer?
```

Returns the *i* th element of *s* (which may be multiple values).

```
(stream-tail s i) → stream?
  s : stream?
  i : exact-nonnegative-integer?
```

Returns a stream equivalent to `s`, except that the first `i` elements are omitted.

In case extracting elements from `s` involves a side effect, they will not be extracted until the first element is extracted from the resulting stream.

```
(stream-append s ...) → stream?
  s : stream?
```

Returns a stream that contains all elements of each stream in the order they appear in the original streams. The new stream is constructed lazily.

```
(stream-map f s) → stream?
  f : procedure?
  s : stream?
```

Returns a stream that contains `f` applied to each element of `s`. The new stream is constructed lazily.

```
(stream-andmap f s) → boolean?
  f : (-> any/c ... boolean?)
  s : stream?
```

Returns `#t` if `f` returns a true result on every element of `s`. If `s` is infinite and `f` never returns a false result, this function does not terminate.

```
(stream-ormap f s) → boolean?
  f : (-> any/c ... boolean?)
  s : stream?
```

Returns `#t` if `f` returns a true result on some element of `s`. If `s` is infinite and `f` never returns a true result, this function does not terminate.

```
(stream-for-each f s) → void?
  f : (-> any/c ... any)
  s : stream?
```

Applies `f` to each element of `s`. If `s` is infinite, this function does not terminate.

```
(stream-fold f i s) → any/c
  f : (-> any/c any/c ... any/c)
  i : any/c
  s : stream?
```

Folds `f` over each element of `s` with `i` as the initial accumulator. If `s` is infinite, this function does not terminate.

```
(stream-count f s) → exact-nonnegative-integer?
  f : procedure?
  s : stream?
```

Returns the number of elements in `s` for which `f` returns a true result. If `s` is infinite, this function does not terminate.

```
(stream-filter f s) → stream?
  f : (-> any/c ... boolean?)
  s : stream?
```

Returns a stream whose elements are the elements of `s` for which `f` returns a true result. Although the new stream is constructed lazily, if `s` has an infinite number of elements where `f` returns a false result in between two elements where `f` returns a true result, then operations on this stream will not terminate during the infinite sub-stream.

```
(stream-add-between s e) → stream?
  s : stream?
  e : any/c
```

Returns a stream whose elements are the elements of `s`, but with `e` between each pair of elements in `s`. The new stream is constructed lazily.

```
gen:stream : any/c
```

Associates three methods to a structure type to implement the generic interface (see §4.4 “Generic Interfaces”) for streams.

To supply method implementations, the `#:methods` keyword should be used in a structure type definition. The following three methods should be implemented:

- `stream-empty?` : accepts one argument
- `stream-first` : accepts one argument
- `stream-rest` : accepts one argument

Examples:

```
> (define-struct list-stream (v)
    #:methods gen:stream
    [(define (stream-empty? stream)
       (empty? (list-stream-v stream)))
     (define (stream-first stream)
       (first (list-stream-v stream)))
     (define (stream-rest stream)
       (rest (list-stream-v stream)))])
```

```
> (define l1 (list-stream '(1 2)))

> (stream? l1)
#t
> (stream-first l1)
1
```

```
prop:stream : struct-type-property?
```

A deprecated structure type property used to define custom extensions to the stream API. Use `gen:stream` instead. Accepts a vector of three procedures taking the same arguments as the methods in `gen:stream`.

3.14.3 Generators

A *generator* is a procedure that returns a sequence of values, incrementing the sequence each time that the generator is called. In particular, the `generator` form implements a generator by evaluating a body that calls `yield` to return values from the generator.

```
(require racket/generator)
```

```
(generator? v) → boolean?
  v : any/c
```

Return `#t` if `v` is a generator, `#f` otherwise.

```
(generator formals body ...+)
```

Creates a generator, where *formals* is like the *formals* of `case-lambda` (i.e., the *kw-formals* of `lambda` restricted to non-optional and non-keyword arguments).

For the first call to a generator, the arguments are bound to the *formals* and evaluation of *body* starts. During the dynamic extent of *body*, the generator can return immediately using the `yield` function. A second call to the generator resumes at the `yield` call, producing the arguments of the second call as the results of the `yield`, and so on. The eventual results of *body* are supplied to an implicit final `yield`; after that final `yield`, calling the generator again returns the same values, but all such calls must provide 0 arguments to the generator.

Examples:

```
> (define g (generator ()
              (let loop ([x '(a b c)])
                (if (null? x)
```

```
                    0
                    (begin
                      (yield (car x))
                      (loop (cdr x)))))))

> (g)
'a
> (g)
'b
> (g)
'c
> (g)
0
> (g)
0
```

(yield v ...) → any
  v : any/c

Returns vs from a generator, saving the point of execution inside a generator (i.e., within the dynamic extent of a generator body) to be resumed by the next call to the generator. The results of yield are the arguments that are provided to the next call of the generator.

When not in the dynamic extent of a generator, infinite-generator, or in-generator body, yield raises exn:fail after evaluating its exprs.

Examples:

```
> (define my-generator (generator () (yield 1) (yield 2 3 4)))

> (my-generator)
1
> (my-generator)
2
3
4
```

Examples:

```
> (define pass-values-generator
    (generator ()
      (let* ([from-user (yield 2)]
             [from-user-again (yield (add1 from-user))])
        (yield from-user-again))))

> (pass-values-generator)
2
```

```
> (pass-values-generator 5)
6
> (pass-values-generator 12)
12
```

(infinite-generator *body* ...+)

Like generator, but repeats evaluation of the *body*s when the last *body* completes without implicitly yielding.

Examples:

```
> (define welcome
    (infinite-generator
      (yield 'hello)
      (yield 'goodbye)))

> (welcome)
'hello
> (welcome)
'goodbye
> (welcome)
'hello
> (welcome)
'goodbye
```

(in-generator *body* ...+)

Produces a sequence that encapsulates the generator formed by (generator () *body* ...+). The values produced by the generator form the elements of the sequence.

Example:

```
> (for/list ([i (in-generator
                 (let loop ([x '(a b c)])
                   (when (not (null? x))
                     (yield (car x))
                     (loop (cdr x)))))])
    i)
'(a b c)
```

To use an existing generator as a sequence, use in-producer with a stop-value known for the generator.

Examples:

```
> (define my-stop-value (gensym))
```

```
> (define my-generator (generator ()
                         (let loop ([x '(a b c)])
                           (if (null? x)
                               my-stop-value
                               (begin
                                 (yield (car x))
                                 (loop (cdr x)))))))

> (for/list ([i (in-producer my-generator my-stop-value)])
    i)
'(a b c)
```

(generator-state *g*) → symbol?
g : generator?

Returns a symbol that describes the state of the generator.

- `'fresh` — The generator has been freshly created and has not been called yet.
- `'suspended` — Control within the generator has been suspended due to a call to `yield`. The generator can be called.
- `'running` — The generator is currently executing.
- `'done` — The generator has executed its entire body and will continue to produce the same result as from the last call.

Examples:

```
> (define my-generator (generator () (yield 1) (yield 2)))

> (generator-state my-generator)
'fresh
> (my-generator)
1
> (generator-state my-generator)
'suspended
> (my-generator)
2
> (generator-state my-generator)
'suspended
> (my-generator)

> (generator-state my-generator)
'done
> (define introspective-generator (generator () ((yield 1))))
```

```
> (introspective-generator)
1
> (introspective-generator
   (lambda () (generator-state introspective-generator)))
'running
> (generator-state introspective-generator)
'done
> (introspective-generator)
'running
```

```
(sequence->generator s) → (-> any)
  s : sequence?
```

Converts a sequence to a generator. The generator returns the next element of the sequence each time the generator is invoked, where each element of the sequence must be a single value. When the sequence ends, the generator returns `#<void>` as its final result.

```
(sequence->repeated-generator s) → (-> any)
  s : sequence?
```

Like `sequence->generator`, but when `s` has no further values, the generator starts the sequence again (so that the generator never stops producing values).

3.15 Dictionaries

A *dictionary* is an instance of a datatype that maps keys to values. The following datatypes are all dictionaries:

- hash tables;
- vectors (using only exact integers as keys);
- lists of pairs (an *association list* using `equal?` to compare keys); and
- structures whose types implement the `gen:dict` generic interface.

A dictionary can be used as a two-valued sequence (see §3.14.1 “Sequences”). The associations of the dictionary serve as elements of the sequence. See also `in-dict`, `in-dict-keys`, and `in-dict-values`.

```
(require racket/dict)
```

The bindings documented in this section are provided by the `racket/dict` and `racket` libraries, but not `racket/base`.

```
(dict? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a dictionary, `#f` otherwise.

Beware that `dict?` is not a constant-time test on pairs, since checking that `v` is an association list may require traversing the list.

Examples:

```
> (dict? #hash((a . "apple")))
#t
> (dict? '#("apple" "banana"))
#t
> (dict? '("apple" "banana"))
#f
> (dict? '((a . "apple") (b . "banana")))
#t
```

```
(dict-mutable? d) → boolean?
  d : dict?
```

Returns `#t` if `d` is mutable via `dict-set!` and maybe `dict-remove!`, `#f` otherwise.

Examples:

```
> (dict-mutable? #hash((a . "apple")))
#f
> (dict-mutable? (make-hash))
#t
> (dict-mutable? '#("apple" "banana"))
#f
> (dict-mutable? (vector "apple" "banana"))
#t
> (dict-mutable? '((a . "apple") (b . "banana")))
#f
```

```
(dict-can-remove-keys? d) → boolean?
  d : dict?
```

Returns `#t` if `d` supports removing mappings via `dict-remove!` and/or `dict-remove`, `#f` otherwise.

Examples:

```
> (dict-can-remove-keys? #hash((a . "apple")))
#t
> (dict-can-remove-keys? '#("apple" "banana"))
#f
> (dict-can-remove-keys? '((a . "apple") (b . "banana")))
#t
```

```
(dict-can-functional-set? d) → boolean?
  d : dict?
```

Returns `#t` if *d* supports functional update via `dict-set` and maybe `dict-remove`, `#f` otherwise.

Examples:

```
> (dict-can-functional-set? #hash((a . "apple")))
#t
> (dict-can-functional-set? (make-hash))
#f
> (dict-can-functional-set? '#("apple" "banana"))
#f
> (dict-can-functional-set? '((a . "apple") (b . "banana")))
#t
```

```
(dict-set! dict key v) → void?
  dict : (and/c dict? (not/c immutable?))
  key : any/c
  v : any/c
```

Maps *key* to *v* in *dict*, overwriting any existing mapping for *key*. The update can fail with a `exn:fail:contract` exception if *dict* is not mutable or if *key* is not an allowed key for the dictionary (e.g., not an exact integer in the appropriate range when *dict* is a vector).

Examples:

```
> (define h (make-hash))

> (dict-set! h 'a "apple")

> h
'#hash((a . "apple"))
> (define v (vector #f #f #f))

> (dict-set! v 0 "apple")

> v
'#("apple" #f #f)
```

```
(dict-set*! dict key v ... ...) → void?
  dict : (and/c dict? (not/c immutable?))
  key : any/c
  v : any/c
```

Maps each *key* to each *v* in *dict*, overwriting any existing mapping for each *key*. The update can fail with a `exn:fail:contract` exception if *dict* is not mutable or if any *key* is not an allowed key for the dictionary (e.g., not an exact integer in the appropriate range when *dict* is a vector). The update takes place from the left, so later mappings overwrite earlier mappings.

Examples:

```
> (define h (make-hash))

> (dict-set*! h 'a "apple" 'b "banana")

> h
'#hash((b . "banana") (a . "apple"))
> (define v1 (vector #f #f #f))

> (dict-set*! v1 0 "apple" 1 "banana")

> v1
'#("apple" "banana" #f)
> (define v2 (vector #f #f #f))

> (dict-set*! v2 0 "apple" 0 "banana")

> v2
'#("banana" #f #f)
```

```
(dict-set dict key v) → (and/c dict? immutable?)
  dict : (and/c dict? immutable?)
  key : any/c
  v : any/c
```

Functionally extends *dict* by mapping *key* to *v*, overwriting any existing mapping for *key*, and returning an extended dictionary. The update can fail with a `exn:fail:contract` exception if *dict* does not support functional extension or if *key* is not an allowed key for the dictionary.

Examples:

```
> (dict-set #hash() 'a "apple")
```

```
'#hash((a . "apple"))
> (dict-set #hash((a . "apple") (b . "beer")) 'b "banana")
'#hash((b . "banana") (a . "apple"))
> (dict-set '() 'a "apple")
'((a . "apple"))
> (dict-set '((a . "apple") (b . "beer")) 'b "banana")
'((a . "apple") (b . "banana"))
```

```
(dict-set* dict key v ... ...) → (and/c dict? immutable?)
  dict : (and/c dict? immutable?)
  key : any/c
  v : any/c
```

Functionally extends *dict* by mapping each *key* to each *v*, overwriting any existing mapping for each *key*, and returning an extended dictionary. The update can fail with a exn:fail:contract exception if *dict* does not support functional extension or if any *key* is not an allowed key for the dictionary. The update takes place from the left, so later mappings overwrite earlier mappings.

Examples:

```
> (dict-set* #hash() 'a "apple" 'b "beer")
'#hash((b . "beer") (a . "apple"))
> (dict-set* #hash((a . "apple") (b . "beer")) 'b "banana" 'a "anchor")
'#hash((b . "banana") (a . "anchor"))
> (dict-set* '() 'a "apple" 'b "beer")
'((a . "apple") (b . "beer"))
> (dict-set* '((a . "apple") (b . "beer")) 'b "banana" 'a "anchor")
'((a . "anchor") (b . "banana"))
> (dict-set* '((a . "apple") (b . "beer")) 'b "banana" 'b "ballistic")
'((a . "apple") (b . "ballistic"))
```

```
(dict-has-key? dict key) → boolean?
  dict : dict?
  key : any/c
```

Returns #t if *dict* contains a value for the given *key*, #f otherwise.

Examples:

```
> (dict-has-key? #hash((a . "apple") (b . "beer")) 'a)
#t
> (dict-has-key? #hash((a . "apple") (b . "beer")) 'c)
#f
> (dict-has-key? '((a . "apple") (b . "banana")) 'b)
#t
```

```
> (dict-has-key? #("apple" "banana") 1)
#t
> (dict-has-key? #("apple" "banana") 3)
#f
> (dict-has-key? #("apple" "banana") -3)
#f
```

```
(dict-ref dict key [failure-result]) → any
  dict : dict?
  key : any/c
  failure-result : any/c
                 = (lambda () (raise (make-exn:fail ....)))
```

Returns the value for *key* in *dict*. If no value is found for *key*, then *failure-result* determines the result:

- If *failure-result* is a procedure, it is called (through a tail call) with no arguments to produce the result.
- Otherwise, *failure-result* is returned as the result.

Examples:

```
> (dict-ref #hash((a . "apple") (b . "beer")) 'a)
"apple"
> (dict-ref #hash((a . "apple") (b . "beer")) 'c)
hash-ref: no value found for key
  key: 'c
> (dict-ref #hash((a . "apple") (b . "beer")) 'c #f)
#f
> (dict-ref '((a . "apple") (b . "banana")) 'b)
"banana"
> (dict-ref #("apple" "banana") 1)
"banana"
> (dict-ref #("apple" "banana") 3 #f)
#f
> (dict-ref #("apple" "banana") -3 #f)
dict-ref: contract violation
  expected: exact-nonnegative-integer?, given: -3
  in: the k argument of
      (->i
       ((d d:dict?) (k (d) ...))
       ((default any/c))
       any)
  contract from: <collects>/racket/dict.rkt
  blaming: top-level
  at: <collects>/racket/dict.rkt:143.2
```

```
(dict-ref! dict key to-set) → any
  dict : dict?
  key : any/c
  to-set : any/c
```

Returns the value for *key* in *dict*. If no value is found for *key*, then *to-set* determines the result as in dict-ref (i.e., it is either a thunk that computes a value or a plain value), and this result is stored in *dict* for the *key*. (Note that if *to-set* is a thunk, it is not invoked in tail position.)

Examples:

```
> (dict-ref! (make-hasheq '((a . "apple") (b . "beer"))) 'a)
dict-ref!: arity mismatch;
 the expected number of arguments does not match the given
number
  expected: 3
  given: 2
  arguments...:
   '#hasheq((a . "apple") (b . "beer"))
   'a
> (dict-ref! (make-hasheq '((a . "apple") (b .
"beer"))) 'c 'cabbage)
'cabbage
> (define h (make-hasheq '((a . "apple") (b . "beer"))))

> (dict-ref h 'c)
hash-ref: no value found for key
  key: 'c
> (dict-ref! h 'c (λ () 'cabbage))
'cabbage
> (dict-ref h 'c)
'cabbage
```

```
(dict-update! dict
              key
              updater
             [failure-result]) → void?
  dict : (and/c dict? (not/c immutable?))
  key : any/c
  updater : (any/c . -> . any/c)
  failure-result : any/c
                 = (lambda () (raise (make-exn:fail ....)))
```

Composes dict-ref and dict-set! to update an existing mapping in *dict*, where the optional *failure-result* argument is used as in dict-ref when no mapping exists for

key already.

Examples:

```
> (define h (make-hash))

> (dict-update! h 'a add1)
hash-ref: no value found for key
  key: 'a
> (dict-update! h 'a add1 0)

> h
'#hash((a . 1))
> (define v (vector #f #f #f))

> (dict-update! v 0 not)

> v
'#(#t #f #f)
```

```
(dict-update dict key updater [failure-result])
 → (and/c dict? immutable?)
  dict : dict?
  key : any/c
  updater : (any/c . -> . any/c)
  failure-result : any/c
                 = (lambda () (raise (make-exn:fail ....)))
```

Composes `dict-ref` and `dict-set` to functionally update an existing mapping in *dict*, where the optional *failure-result* argument is used as in `dict-ref` when no mapping exists for *key* already.

Examples:

```
> (dict-update #hash() 'a add1)
hash-ref: no value found for key
  key: 'a
> (dict-update #hash() 'a add1 0)
'#hash((a . 1))
> (dict-update #hash((a . "apple") (b . "beer")) 'b string-length)
'#hash((b . 4) (a . "apple"))
```

```
(dict-remove! dict key) → void?
  dict : (and/c dict? (not/c immutable?))
  key : any/c
```

Removes any existing mapping for *key* in *dict*. The update can fail if *dict* is not mutable or does not support removing keys (as is the case for vectors, for example).

Examples:

```
> (define h (make-hash))

> (dict-set! h 'a "apple")

> h
'#hash((a . "apple"))
> (dict-remove! h 'a)

> h
'#hash()
```

```
(dict-remove dict key) → (and/c dict? immutable?)
  dict : (and/c dict? immutable?)
  key : any/c
```

Functionally removes any existing mapping for *key* in *dict*, returning the fresh dictionary. The update can fail if *dict* does not support functional update or does not support removing keys.

Examples:

```
> (define h #hash())

> (define h (dict-set h 'a "apple"))

> h
'#hash((a . "apple"))
> (dict-remove h 'a)
'#hash()
> h
'#hash((a . "apple"))
> (dict-remove h 'z)
'#hash((a . "apple"))
> (dict-remove '((a . "apple") (b . "banana")) 'a)
'((b . "banana"))
```

```
(dict-map dict proc) → (listof any/c)
  dict : dict?
  proc : (any/c any/c . -> . any/c)
```

Applies the procedure *proc* to each element in *dict* in an unspecified order, accumulating the results into a list. The procedure *proc* is called each time with a key and its value.

Example:

```
> (dict-map #hash((a . "apple") (b . "banana")) vector)
'(#(b "banana") #(a "apple"))
```

```
(dict-for-each dict proc) → void?
  dict : dict?
  proc : (any/c any/c . -> . any)
```

Applies *proc* to each element in *dict* (for the side-effects of *proc*) in an unspecified order. The procedure *proc* is called each time with a key and its value.

Example:

```
> (dict-for-each #hash((a . "apple") (b . "banana"))
                 (lambda (k v)
                   (printf "~a = ~s\n" k v)))
b = "banana"
a = "apple"
```

```
(dict-count dict) → exact-nonnegative-integer?
  dict : dict?
```

Returns the number of keys mapped by *dict*, usually in constant time.

Examples:

```
> (dict-count #hash((a . "apple") (b . "banana")))
2
> (dict-count #("apple" "banana"))
2
```

```
(dict-iterate-first dict) → any/c
  dict : dict?
```

Returns `#f` if *dict* contains no elements, otherwise it returns a non-`#f` value that is an index to the first element in the dict table; "first" refers to an unspecified ordering of the dictionary elements. For a mutable *dict*, this index is guaranteed to refer to the first item only as long as no mappings are added to or removed from *dict*.

Examples:

```
> (dict-iterate-first #hash((a . "apple") (b . "banana")))
0
> (dict-iterate-first #hash())
#f
> (dict-iterate-first #("apple" "banana"))
0
> (dict-iterate-first '((a . "apple") (b . "banana")))
#<assoc-iter>
```

```
(dict-iterate-next dict pos) → any/c
  dict : dict?
  pos : any/c
```

Returns either a non-`#f` that is an index to the element in *dict* after the element indexed by *pos* or `#f` if *pos* refers to the last element in *dict*. If *pos* is not a valid index, then the `exn:fail:contract` exception is raised. For a mutable *dict*, the result index is guaranteed to refer to its item only as long as no items are added to or removed from *dict*. The `dict-iterate-next` operation should take constant time.

Examples:

```
> (define h #hash((a . "apple") (b . "banana")))

> (define i (dict-iterate-first h))

> i
0
> (dict-iterate-next h i)
1
> (dict-iterate-next h (dict-iterate-next h i))
#f
```

```
(dict-iterate-key dict pos) → any
  dict : dict?
  pos : any/c
```

Returns the key for the element in *dict* at index *pos*. If *pos* is not a valid index for *dict*, the `exn:fail:contract` exception is raised. The `dict-iterate-key` operation should take constant time.

Examples:

```
> (define h '((a . "apple") (b . "banana")))

> (define i (dict-iterate-first h))

> (dict-iterate-key h i)
'a
> (dict-iterate-key h (dict-iterate-next h i))
'b
```

```
(dict-iterate-value dict pos) → any
  dict : dict?
  pos : any/c
```

Returns the value for the element in *dict* at index *pos*. If *pos* is not a valid index for *dict*, the exn:fail:contract exception is raised. The dict-iterate-key operation should take constant time.

Examples:

```
> (define h '((a . "apple") (b . "banana")))

> (define i (dict-iterate-first h))

> (dict-iterate-value h i)
"apple"
> (dict-iterate-value h (dict-iterate-next h i))
"banana"
```

(in-dict *dict*) → sequence?
  dict : dict?

Returns a sequence whose each element is two values: a key and corresponding value from *dict*.

Examples:

```
> (define h #hash((a . "apple") (b . "banana")))

> (for/list ([(k v) (in-dict h)])
    (format "~a = ~s" k v))
'("b = \"banana\"" "a = \"apple\"")
```

(in-dict-keys *dict*) → sequence?
  dict : dict?

Returns a sequence whose elements are the keys of *dict*.

Examples:

```
> (define h #hash((a . "apple") (b . "banana")))

> (for/list ([k (in-dict-keys h)])
    k)
'(b a)
```

(in-dict-values *dict*) → sequence?
  dict : dict?

Returns a sequence whose elements are the values of *dict*.

Examples:

```
> (define h #hash((a . "apple") (b . "banana")))

> (for/list ([v (in-dict-values h)])
    v)
'("banana" "apple")
```

```
(in-dict-pairs dict) → sequence?
  dict : dict?
```

Returns a sequence whose elements are pairs, each containing a key and its value from *dict* (as opposed to using in-dict, which gets the key and value as separate values for each element).

Examples:

```
> (define h #hash((a . "apple") (b . "banana")))

> (for/list ([p (in-dict-pairs h)])
    p)
'((b . "banana") (a . "apple"))
```

```
(dict-keys dict) → list?
  dict : dict?
```

Returns a list of the keys from *dict* in an unspecified order.

Examples:

```
> (define h #hash((a . "apple") (b . "banana")))

> (dict-keys h)
'(b a)
```

```
(dict-values dict) → list?
  dict : dict?
```

Returns a list of the values from *dict* in an unspecified order.

Examples:

```
> (define h #hash((a . "apple") (b . "banana")))

> (dict-values h)
'("banana" "apple")
```

```
(dict->list dict) → list?
  dict : dict?
```

Returns a list of the associations from *dict* in an unspecified order.

Examples:

```
> (define h #hash((a . "apple") (b . "banana")))

> (dict->list h)
'((b . "banana") (a . "apple"))
```

```
gen:dict : any/c
```

A generic interface (see §4.4 “Generic Interfaces”) that supplies dictionary method implementations for a structure type. To supply method implementations, the `#:methods` form should be used. The provided implementations are applied only to instances of the structure type. The following methods can be implemented:

- `dict-ref` : accepts either two or three arguments
- `dict-set!` : accepts three arguments, left unimplemented if mutation is not supported
- `dict-set` : accepts three arguments and returns an updated dictionary, left unimplemented if functional update is not supported
- `dict-remove!` : accepts two arguments, left unimplemented if mutation is not supported or if key removal is not supported
- `dict-remove` : accepts two arguments and returns an updated dictionary, left unimplemented if functional update or key removal is not supported
- `dict-count` : accepts one argument
- `dict-iterate-first` : accepts one argument
- `dict-iterate-next` : accepts two arguments; the procedure is responsible for checking that the second argument is a valid position for the first argument
- `dict-iterate-key` : accepts two arguments; the procedure is responsible for checking that the second argument is a valid position for the first argument
- `dict-iterate-value` : accepts two arguments; the procedure is responsible for checking that the second argument is a valid position for the first argument

Examples:

```
> (struct alist (v)
    #:methods gen:dict
    [(define (dict-ref dict key
                       [default (lambda () (error "key not
found" key))])
       (cond [(assoc key (alist-v dict)) => cdr]
             [else (if (procedure? default) (default) default)]))
     (define (dict-set dict key val)
       (alist (cons (cons key val) (alist-v dict))))
     (define (dict-remove dict key)
       (define al (alist-v dict))
       (remove* (assoc key al) al))
     (define (dict-count dict #:default [x #f])
       (or x
           (length (remove-duplicates (alist-v dict) #:key car))))
     ; etc. other methods])

> (define d1 '((1 . a) (2 . b)))

> (dict? d1)
#t
> (dict-ref d1 1)
'a
```

prop:dict : struct-type-property?

A deprecated structure type property used to define custom extensions to the dictionary API. Use gen:dict instead. Accepts a vector of 10 procedures with the same arguments as the methods of gen:dict.

prop:dict/contract : struct-type-property?

A structure type property for defining dictionaries with contracts. The value associated with prop:dict/contract must be a list of two immutable vectors:

```
(list dict-vector
      (vector type-key-contract
              type-value-contract
              type-iter-contract
              instance-key-contract
              instance-value-contract
              instance-iter-contract))
```

The first vector must be a vector of 10 procedures which match the gen:dict generic interface (in addition, it must be an immutable vector). The second vector must contain six

elements; each of the first three is a contract for the dictionary type's keys, values, and positions, respectively. Each of the second three is either `#f` or a procedure used to extract the contract from a dictionary instance.

```
(dict-key-contract d) → contract?
  d : dict?
(dict-value-contract d) → contract?
  d : dict?
(dict-iter-contract d) → contract?
  d : dict?
```

Returns the contract that *d* imposes on its keys, values, or iterators, respectively, if *d* implements the `prop:dict/contract` interface.

```
(make-custom-hash eql? hash-proc [hash2-proc]) → dict?
  eql? : (any/c any/c . -> . any/c)
  hash-proc : (any/c . -> . exact-integer?)
  hash2-proc : (any/c . -> . exact-integer?)
             = (lambda (v) 10001)
(make-immutable-custom-hash  eql?
                             hash-proc
                            [hash2-proc]) → dict?
  eql? : (any/c any/c . -> . any/c)
  hash-proc : (any/c . -> . exact-integer?)
  hash2-proc : (any/c . -> . exact-integer?)
             = (lambda (v) 10001)
(make-weak-custom-hash  eql?
                        hash-proc
                       [hash2-proc]) → dict?
  eql? : (any/c any/c . -> . any/c)
  hash-proc : (any/c . -> . exact-integer?)
  hash2-proc : (any/c . -> . exact-integer?)
             = (lambda (v) 10001)
```

Creates a dictionary that is implemented in terms of a hash table where keys are compared with *eql?* and hashed with *hash-proc* and *hash2-proc*. See `gen:equal+hash` for information on suitable equality and hashing functions.

The `make-custom-hash` and `make-weak-custom-hash` functions create a mutable dictionary that does not support functional update, while `make-immutable-custom-hash` creates an immutable dictionary that supports functional update. The dictionary created by `make-weak-custom-hash` retains its keys weakly, like the result of `make-weak-hash`.

Dictionaries created by `make-custom-hash` and company are `equal?` when they have the same mutability and key strength, the associated procedures are `equal?`, and the key–value mappings are the same when keys and values are compared with `equal?`.

Examples:

```
> (define h (make-custom-hash (lambda (a b)
                                (string=? (format "~a" a)
                                          (format "~a" b)))
                              (lambda (a)
                                (equal-hash-code
                                 (format "~a" a)))))

> (dict-set! h 1 'one)

> (dict-ref h "1")
'one
```

3.16 Sets

A *set* represents a set of distinct elements. For a given set, elements are equivalent via `equal?`, `eqv?`, or `eq?`. Two sets are `equal?` when they use the same element-comparison procedure (`equal?`, `eqv?`, or `eq?`) and have equivalent elements.

A set can be used as a single-valued sequence (see §3.14.1 "Sequences"). The elements of the set serve as elements of the sequence. See also `in-set`.

Operations on sets that contain elements that are mutated are unpredictable in much the same way that hash table operations are unpredictable when keys are mutated.

```
(require racket/set)
```

The bindings documented in this section are provided by the `racket/set` and `racket` libraries, but not `racket/base`.

```
(set v ...) → set?
  v : any/c
(seteqv v ...) → set?
  v : any/c
(seteq v ...) → set?
  v : any/c
```

Creates a set that uses `equal?`, `eqv?`, or `eq?`, respectively, to compare elements. The given `v`s are added to the set. The elements are added in the order that they appear as `v`s, so in the first two cases, an earlier element that is `equal?` or `eqv?` but not `eq?` to a later element takes precedence over the later element.

```
(set-empty? st) → boolean?
  st : set?
```

Returns `#t` if *st* has no members, `#f` otherwise.

```
(set-count st) → exact-nonnegative-integer?
  st : set?
```

Returns the number of elements in *st*.

```
(set-member? st v) → boolean?
  st : set?
  v : any/c
```

Returns `#t` if *v* is in *st*, `#f` otherwise.

```
(set-add st v) → set?
  st : set?
  v : any/c
```

Like operations on immutable hash tables, "constant time" set operations actually require $O(\log N)$ time for a set of size N.

Produces a set that includes *v* plus all elements of *st*. This operation runs in constant time.

```
(set-remove st v) → set?
  st : set?
  v : any/c
```

Produces a set that includes all elements of *st* except *v*. This operation runs in constant time.

```
(set-union st ...+) → set?
  st : set?
```

Produces a set that includes all elements of all given *st*s, which must all use the same equivalence predicate (`equal?`, `eq?`, or `eqv?`). This operation runs in time proportional to the total size of all given *st*s except for the largest.

At least one set must be provided to `set-union` even though mathematically `set-union` could accept zero arguments. Since there are multiple types of sets (`eq?`, `eqv?`, and `equal?`) there is no obvious choice for a default empty set to be returned. If there is a case where `set-union` may be applied to zero arguments, instead pass an empty set of the type you desire.

Examples:

```
> (set-union (set))
(set)
> (set-union (seteq))
(seteq)
```

```
> (set-union (set 1) (set 2))
(set 1 2)
> (set-union (set 1) (seteq 2))
set-union: second set's equivalence predicate is not the
same as the first set's
  first set: (set 1)
  second set: (seteq 2)
; Sets of different types cannot be unioned
```

(set-intersect st ...+) → set?
  st : set?

Produces a set that includes only the elements in all of the given sts, which must all use the same equivalence predicate (equal?, eq?, or eqv?). This operation runs in time proportional to the total size of all given sts except for the largest.

(set-subtract st ...+) → set?
  st : set?

Produces a set that includes all elements the first sts that are not present in any of the other given stss. All of the given sts must use the same equivalence predicate (equal?, eq?, or eqv?). This operation runs in time proportional to the total size of all given sts except the first one.

(set-symmetric-difference st ...+) → set?
  st : set?

Produces a set containing only those elements found in each st an odd number of times. All of the given sts must use the same equivalence predicate (equal?, eq?, or eqv?). This operation runs in time proportional to the total size of all given sts except the first one.

Example:

```
> (set-symmetric-difference (set 1) (set 1 2) (set 1 2 3))
(set 1 3)
```

(set=? st st2) → boolean?
  st : set?
  st2 : set?

Returns #t if st and st2 contain the same members, #f otherwise. The st and st2 must use the same equivalence predicate (equal?, eq?, or eqv?). This operation runs in time proportional to the size of st.

Equivalent to (equal? st st2).

Examples:

```
> (set=? (set 1) (set 1 2 3))
#f
> (set=? (set 1 2 3) (set 1))
#f
> (set=? (set 1 2 3) (set 1 2 3))
#t
```

(subset? *st st2*) → boolean?
  st : set?
  st2 : set?

Returns #t if every member of *st* is in *st2*, #f otherwise. The *st* and *st2* must use the same equivalence predicate (equal?, eq?, or eqv?). This operation runs in time proportional to the size of *st*.

Examples:

```
> (subset? (set 1) (set 1 2 3))
#t
> (subset? (set 1 2 3) (set 1))
#f
> (subset? (set 1 2 3) (set 1 2 3))
#t
```

(proper-subset? *st st2*) → boolean?
  st : set?
  st2 : set?

Returns #t if every member of *st* is in *st2* and there is some member of *st2* that is not a member of *st*, #f otherwise. The *st* and *st2* must use the same equivalence predicate (equal?, eq?, or eqv?). This operation runs in time proportional to the size of *st*.

Examples:

```
> (proper-subset? (set 1) (set 1 2 3))
#t
> (proper-subset? (set 1 2 3) (set 1))
#f
> (proper-subset? (set 1 2 3) (set 1 2 3))
#f
```

(set-map *st proc*) → (listof any/c)
  st : set?
  proc : (any/c . -> . any/c)

Applies the procedure *proc* to each element in *st* in an unspecified order, accumulating the results into a list.

```
(set-for-each st proc) → void?
  st : set?
  proc : (any/c . -> . any)
```

Applies *proc* to each element in *st* (for the side-effects of *proc*) in an unspecified order.

```
(set? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a set, `#f` otherwise.

```
(set-equal? st) → boolean?
  st : set?
```

Returns `#t` if *st* compares elements with `equal?`, `#f` if it compares with `eqv?` or `eq?`.

```
(set-eqv? st) → boolean?
  st : set?
```

Returns `#t` if *st* compares elements with `eqv?`, `#f` if it compares with `equal?` or `eq?`.

```
(set-eq? st) → boolean?
  st : set?
```

Returns `#t` if *st* compares elements with `eq?`, `#f` if it compares with `equal?` or `eqv?`.

```
(set/c contract [#:cmp cmp]) → contract?
  contract : chaperone-contract?
  cmp : (or/c 'dont-care 'equal 'eqv 'eq) = 'dont-care
```

Constructs a contract that recognizes sets whose elements match *contract*.

If *cmp* is `'dont-care`, then the equality notion of the set is not considered when checking the contract. Otherwise, the contract accepts only sets with the corresponding notion of equality.

If *cmp* is `'eq` or `'eqv`, then *contract* must be a flat contract. If *contract* is not a flat contract, then *cmp* cannot be `'eq` or `'eqv` and the resulting contract can only be applied to sets that use `equal?` for equality.

```
(in-set st) → sequence?
  st : set?
```

Explicitly converts a set to a sequence for use with `for` and other forms.

```
(for/set (for-clause ...) body ...+)
(for/seteq (for-clause ...) body ...+)
(for/seteqv (for-clause ...) body ...+)
(for*/set (for-clause ...) body ...+)
(for*/seteq (for-clause ...) body ...+)
(for*/seteqv (for-clause ...) body ...+)
```

Analogous to `for/list` and `for*/list`, but to construct a set instead of a list.

```
(list->set lst) → set?
  lst : list?
(list->seteq lst) → set?
  lst : list?
(list->seteqv lst) → set?
  lst : list?
```

Produces the appropriate type of set containing the elements of the given list. Equivalent to `(apply set lst)`, `(apply seteq lst)`, and `(apply seteqv lst)`, respectively.

```
(set->list st) → list?
  st : set?
```

Produces a list containing the elements of *st*.

3.17 Procedures

```
(procedure? v) → boolean
  v : any/c
```

Returns `#t` if *v* is a procedure, `#f` otherwise.

```
(apply proc v ... lst #:<kw> kw-arg ...) → any
  proc : procedure?
  v : any/c
  lst : list?
  kw-arg : any/c
```

Applies *proc* using the content of `(list* v ... lst)` as the (by-position) arguments. The `#:<kw> kw-arg` sequence is also supplied as keyword arguments to *proc*, where `#:<kw>` stands for any keyword.

§4.3.3 "The `apply` Function" in ***The Racket Guide*** introduces `apply`.

The given *proc* must accept as many arguments as the number of *v*s plus length of *lst*, it must accept the supplied keyword arguments, and it must not require any other keyword arguments; otherwise, the exn:fail:contract exception is raised. The given *proc* is called in tail position with respect to the apply call.

Examples:

```
> (apply + '(1 2 3))
6
> (apply + 1 2 '(3))
6
> (apply + '())
0
> (apply sort (list (list '(2) '(1)) <) #:key car)
'((1) (2))
```

(compose *proc* ...) → procedure?
  proc : procedure?
(compose1 *proc* ...) → procedure?
  proc : procedure?

Returns a procedure that composes the given functions, applying the last *proc* first and the first *proc* last. The compose function allows the given functions to consume and produce any number of values, as long as each function produces as many values as the preceding function consumes, while compose1 restricts the internal value passing to a single value. In both cases, the input arity of the last function and the output arity of the first are unrestricted, and they become the corresponding arity of the resulting composition (including keyword arguments for the input side).

When no *proc* arguments are given, the result is values. When exactly one is given, it is returned.

Examples:

```
> ((compose1 - sqrt) 10)
-3.1622776601683795
> ((compose1 sqrt -) 10)
0+3.1622776601683795i
> ((compose list split-path) (bytes->path #"/a" 'unix))
'(#<path:/> #<path:a> #f)
```

Note that in many cases, compose1 is preferred. For example, using compose with two library functions may lead to problems when one function is extended to return two values, and the preceding one has an optional input with different semantics. In addition, compose1 may create faster compositions.

(procedure-rename *proc* *name*) → procedure?
  proc : procedure?
  name : symbol?

Returns a procedure that is like *proc*, except that its name as returned by `object-name` (and as printed for debugging) is *name*.

The given *name* is used for printing an error message if the resulting procedure is applied to the wrong number of arguments. In addition, if *proc* is an accessor or mutator produced by `struct`, `make-struct-field-accessor`, or `make-struct-field-mutator`, the resulting procedure also uses *name* when its (first) argument has the wrong type. More typically, however, *name* is not used for reporting errors, since the procedure name is typically hard-wired into an internal check.

```
(procedure->method proc) → procedure?
  proc : procedure?
```

Returns a procedure that is like *proc* except that, when applied to the wrong number of arguments, the resulting error hides the first argument as if the procedure had been compiled with the `'method-arity-error` syntax property.

3.17.1 Keywords and Arity

```
(keyword-apply proc
               kw-lst
               kw-val-lst
               v ...
               lst
               #:<kw> kw-arg ...) → any
  proc : procedure?
  kw-lst : (listof keyword?)
  kw-val-lst : list?
  v : any/c
  lst : list?
  kw-arg : any/c
```

§4.3.3 "The `apply` Function" in ***The Racket Guide*** introduces `keyword-apply`.

Like `apply`, but *kw-lst* and *kw-val-lst* supply by-keyword arguments in addition to the by-position arguments of the *v*s and *lst*, and in addition to the directly supplied keyword arguments in the `#:<kw>` *kw-arg* sequence, where `#:<kw>` stands for any keyword.

The given *kw-lst* must be sorted using `keyword<?`. No keyword can appear twice in *kw-lst* or in both `kw-list` and as a `#:<kw>`, otherwise, the `exn:fail:contract` exception is raised. The given *kw-val-lst* must have the same length as *kw-lst*, otherwise, the `exn:fail:contract` exception is raised. The given *proc* must accept all of the keywords in *kw-lst* plus the `#:<kw>`s, it must not require any other keywords, and it must accept as many by-position arguments as supplied via the *v*s and *lst*; otherwise, the `exn:fail:contract` exception is raised.

Examples:

```
(define (f x #:y y #:z [z 10])
  (list x y z))

> (keyword-apply f '(#:y) '(2) '(1))
'(1 2 10)
> (keyword-apply f '(#:y #:z) '(2 3) '(1))
'(1 2 3)
> (keyword-apply f #:z 7 '(#:y) '(2) '(1))
'(1 2 7)
```

(procedure-arity *proc*) → procedure-arity?
  proc : procedure?

Returns information about the number of by-position arguments accepted by *proc*. See also procedure-arity?.

(procedure-arity? *v*) → boolean?
  v : any/c

A valid arity a is one of the following:

- An exact non-negative integer, which means that the procedure accepts a arguments, only.
- A arity-at-least instance, which means that the procedure accepts (arity-at-least-value a) or more arguments.
- A list containing integers and arity-at-least instances, which means that the procedure accepts any number of arguments that can match one of the elements of a.

Generally, procedure-arity always produces an arity that is normalized. Specifically, it is either the empty list (corresponding to the procedure (case-lambda)), one of the first two cases above, or a list that contains at least two elements. If it is a list, there is at most one arity-at-least instance that appears as the last element of the list, all of the other elements are sorted in ascending order, and there are no duplicate elements.

Examples:

```
> (procedure-arity cons)
2
> (procedure-arity list)
(arity-at-least 0)
> (arity-at-least? (procedure-arity list))
#t
> (arity-at-least-value (procedure-arity list))
0
```

```
> (arity-at-least-value (procedure-arity (lambda (x . y) x)))
1
> (procedure-arity (case-lambda [(x) 0] [(x y) 1]))
'(1 2)
```

```
(procedure-arity-includes? proc k [kws-ok?]) → boolean?
  proc : procedure?
  k : exact-nonnegative-integer?
  kws-ok? : any/c = #f
```

Returns `#t` if the procedure can accept *k* by-position arguments, `#f` otherwise. If *kws-ok?* is `#f`, the result is `#t` only if *proc* has no required keyword arguments.

Examples:

```
> (procedure-arity-includes? cons 2)
#t
> (procedure-arity-includes? display 3)
#f
> (procedure-arity-includes? (lambda (x #:y y) x) 1)
#f
> (procedure-arity-includes? (lambda (x #:y y) x) 1 #t)
#t
```

```
(procedure-reduce-arity proc arity) → procedure?
  proc : procedure?
  arity : procedure-arity?
```

Returns a procedure that is the same as *proc* (including the same name returned by `object-name`), but that accepts only arguments consistent with *arity*. In particular, when `procedure-arity` is applied to the generated procedure, it returns a value that is `equal?` to *arity*.

If the *arity* specification allows arguments that are not in `(procedure-arity proc)`, the `exn:fail:contract` exception is raised. If *proc* accepts keyword argument, either the keyword arguments must be all optional (and they are not accepted in by the arity-reduced procedure) or *arity* must be the empty list (which makes a procedure that cannot be called); otherwise, the `exn:fail:contract` exception is raised.

Examples:

```
> (define my+ (procedure-reduce-arity + 2))

> (my+ 1 2)
3
> (my+ 1 2 3)
+: arity mismatch;
```

the expected number of arguments does not match the given number
expected: 2
given: 3
arguments...:
1
2
3

```
(procedure-keywords proc) → (listof keyword?)
                            (or/c (listof keyword?) #f)
  proc : procedure?
```

Returns information about the keyword arguments required and accepted by a procedure. The first result is a list of distinct keywords (sorted by `keyword<?`) that are required when applying *proc*. The second result is a list of distinct accepted keywords (sorted by `keyword<?`), or `#f` to mean that any keyword is accepted. When the second result is a list, every element in the first list is also in the second list.

Examples:

```
> (procedure-keywords +)
'()
'()
> (procedure-keywords (lambda (#:tag t #:mode m) t))
'(#:mode #:tag)
'(#:mode #:tag)
> (procedure-keywords (lambda (#:tag t #:mode [m #f]) t))
'(#:tag)
'(#:mode #:tag)
```

```
(make-keyword-procedure proc [plain-proc]) → procedure?
  proc : (((listof keyword?) list?) () #:rest list? . ->* . any)
  plain-proc : procedure?
             = (lambda args (apply proc null null args))
```

Returns a procedure that accepts all keyword arguments (without requiring any keyword arguments).

When the procedure returned by `make-keyword-procedure` is called with keyword arguments, then *proc* is called; the first argument is a list of distinct keywords sorted by `keyword<?`, the second argument is a parallel list containing a value for each keyword, and the remaining arguments are the by-position arguments.

When the procedure returned by `make-keyword-procedure` is called without keyword arguments, then *plain-proc* is called—possibly more efficiently than dispatching through

proc. Normally, *plain-proc* should have the same behavior as calling *proc* with empty lists as the first two arguments, but that correspondence is in no way enforced.

The result of procedure-arity and object-name on the new procedure is the same as for *plain-proc*. See also procedure-reduce-keyword-arity and procedure-rename.

Examples:

```
(define show
  (make-keyword-procedure (lambda (kws kw-args . rest)
                            (list kws kw-args rest))))

> (show 1)
'(() () (1))
> (show #:init 0 1 2 3 #:extra 4)
'((#:extra #:init) (4 0) (1 2 3))
(define show2
  (make-keyword-procedure (lambda (kws kw-args . rest)
                            (list kws kw-args rest))
                          (lambda args
                            (list->vector args))))

> (show2 1)
'#(1)
> (show2 #:init 0 1 2 3 #:extra 4)
'((#:extra #:init) (4 0) (1 2 3))
```

```
(procedure-reduce-keyword-arity proc
                                arity
                                required-kws
                                allowed-kws) → procedure?
  proc : procedure?
  arity : procedure-arity?
  required-kws : (listof keyword?)
  allowed-kws : (or/c (listof keyword?)
                      #f)
```

Like procedure-reduce-arity, but constrains the keyword arguments according to *required-kws* and *allowed-kws*, which must be sorted using keyword<? and contain no duplicates. If *allowed-kws* is #f, then the resulting procedure still accepts any keyword, otherwise the keywords in *required-kws* must be a subset of those in *allowed-kws*. The original *proc* must require no more keywords than the ones listed in *required-kws*, and it must allow at least the keywords in *allowed-kws* (or it must allow all keywords if *allowed-kws* is #f).

Examples:

```
(define orig-show
  (make-keyword-procedure (lambda (kws kw-args . rest)
                            (list kws kw-args rest))))

(define show (procedure-reduce-keyword-arity
              orig-show 3 '(#:init) '(#:extra #:init)))

> (show #:init 0 1 2 3 #:extra 4)
'((#:extra #:init) (4 0) (1 2 3))
> (show 1)
...t/private/kw.rkt:191:14: arity mismatch;
 the expected number of arguments does not match the given
number
  expected: 3 plus an argument with keyword #:init plus an
optional argument with keyword #:extra
  given: 1
  arguments...:
   1
> (show #:init 0 1 2 3 #:extra 4 #:more 7)
application: procedure does not expect an argument with
given keyword
  procedure: ...t/private/kw.rkt:191:14
  given keyword: #:more
  given arguments:
   1
   2
   3
   #:extra 4
   #:init 0
   #:more 7
```

```
(struct arity-at-least (value)
        #:extra-constructor-name make-arity-at-least)
  value : exact-nonnegative-integer?
```

A structure type used for the result of `procedure-arity`. See also `procedure-arity?`.

```
prop:procedure : struct-type-property?
```

A structure type property to identify structure types whose instances can be applied as procedures. In particular, when `procedure?` is applied to the instance, the result will be `#t`, and when an instance is used in the function position of an application expression, a procedure is extracted from the instance and used to complete the procedure call.

If the `prop:procedure` property value is an exact non-negative integer, it designates a field within the structure that should contain a procedure. The integer must be between `0` (inclusive) and the number of non-automatic fields in the structure type (exclusive, not counting

supertype fields). The designated field must also be specified as immutable, so that after an instance of the structure is created, its procedure cannot be changed. (Otherwise, the arity and name of the instance could change, and such mutations are generally not allowed for procedures.) When the instance is used as the procedure in an application expression, the value of the designated field in the instance is used to complete the procedure call. (This procedure can be another structure that acts as a procedure; the immutability of procedure fields disallows cycles in the procedure graph, so that the procedure call will eventually continue with a non-structure procedure.) That procedure receives all of the arguments from the application expression. The procedure's name (see `object-name`), arity (see `procedure-arity`), and keyword protocol (see `procedure-keywords`) are also used for the name, arity, and keyword protocol of the structure. If the value in the designated field is not a procedure, then the instance behaves like `(case-lambda)` (i.e., a procedure which does not accept any number of arguments). See also `procedure-extract-target`.

Providing an integer `proc-spec` argument to `make-struct-type` is the same as both supplying the value with the `prop:procedure` property and designating the field as immutable (so that a property binding or immutable designation is redundant and disallowed).

Examples:

```
> (struct annotated-proc (base note)
    #:property prop:procedure
               (struct-field-index base))

> (define plus1 (annotated-proc
                 (lambda (x) (+ x 1))
                 "adds 1 to its argument"))

> (procedure? plus1)
#t
> (annotated-proc? plus1)
#t
> (plus1 10)
11
> (annotated-proc-note plus1)
"adds 1 to its argument"
```

When the `prop:procedure` value is a procedure, it should accept at least one non-keyword argument. When an instance of the structure is used in an application expression, the property-value procedure is called with the instance as the first argument. The remaining arguments to the property-value procedure are the arguments from the application expression (including keyword arguments). Thus, if the application expression provides five non-keyword arguments, the property-value procedure is called with six non-keyword arguments. The name of the instance (see `object-name`) and its keyword protocol (see `procedure-keywords`) are unaffected by the property-value procedure, but the instance's arity is determined by subtracting one from every possible non-keyword argument count of the property-value procedure. If the property-value procedure cannot accept at least one argument, then

the instance behaves like (case-lambda).

Providing a procedure proc-spec argument to make-struct-type is the same as supplying the value with the prop:procedure property (so that a specific property binding is disallowed).

Examples:

```
> (struct fish (weight color)
    #:mutable
    #:property
    prop:procedure
    (lambda (f n)
      (let ([w (fish-weight f)])
        (set-fish-weight! f (+ n w)))))

> (define wanda (fish 12 'red))

> (fish? wanda)
#t
> (procedure? wanda)
#t
> (fish-weight wanda)
12
> (for-each wanda '(1 2 3))

> (fish-weight wanda)
18
```

If the value supplied for the prop:procedure property is not an exact non-negative integer or a procedure, the exn:fail:contract exception is raised.

```
(procedure-struct-type? type) → boolean?
  type : struct-type?
```

Returns #t if instances of the structure type represented by *type* are procedures (according to procedure?), #f otherwise.

```
(procedure-extract-target proc) → (or/c #f procedure?)
  proc : procedure?
```

If *proc* is an instance of a structure type with property prop:procedure, and if the property value indicates a field of the structure, and if the field value is a procedure, then procedure-extract-target returns the field value. Otherwise, the result is #f.

When a prop:procedure property value is a procedure, the procedure is *not* returned by procedure-extract-target. Such a procedure is different from one accessed through a

structure field, because it consumes an extra argument, which is always the structure that was applied as a procedure. Keeping the procedure private ensures that is it always called with a suitable first argument.

prop:arity-string : struct-type-property?

A structure type property that is used for reporting arity-mismatch errors when a structure type with the prop:procedure property is applied to the wrong number of arguments. The value of the prop:arity-string property must be a procedure that takes a single argument, which is the misapplied structure, and returns a string. The result string is used after the word "expects," and it is followed in the error message by the number of actual arguments.

Arity-mismatch reporting automatically uses procedure-extract-target when the prop:arity-string property is not associated with a procedure structure type.

Examples:

```
> (struct evens (proc)
    #:property prop:procedure (struct-field-index proc)
    #:property prop:arity-string
    (lambda (p)
      "an even number of arguments"))

> (define pairs
    (evens
     (case-lambda
      [() null]
      [(a b . more)
       (cons (cons a b)
             (apply pairs more))])))

> (pairs 1 2 3 4)
'((1 . 2) (3 . 4))
> (pairs 5)
#<procedure>: arity mismatch;
 the expected number of arguments does not match the given
number
  expected: an even number of arguments
  given: 1
  arguments...:
   5
```

prop:checked-procedure : struct-type-property?

A structure type property that is used with checked-procedure-check-and-extract, which is a hook to allow the compiler to improve the performance of keyword arguments.

The property can only be attached to a structure type without a supertype and with at least two fields.

```
(checked-procedure-check-and-extract type
                                     v
                                     proc
                                     v1
                                     v2)  → any/c
  type : struct-type?
  v : any/c
  proc : (any/c any/c any/c . -> . any/c)
  v1 : any/c
  v2 : any/c
```

Extracts a value from *v* if it is an instance of *type*, which must have the property prop:checked-procedure. If *v* is such an instance, then the first field of *v* is extracted and applied to *v1* and *v2*; if the result is a true value, the result is the value of the second field of *v*.

If *v* is not an instance of *type*, or if the first field of *v* applied to *v1* and *v2* produces #f, then *proc* is applied to *v*, *v1*, and *v2*, and its result is returned by checked-procedure-check-and-extract.

3.17.2 Reflecting on Primitives

A *primitive procedure* is a built-in procedure that is implemented in low-level language. Not all procedures of racket/base are primitives, but many are. The distinction is mainly useful to other low-level code.

```
(primitive? v) → boolean?
  v : any/c
```

Returns #t if *v* is a primitive procedure, #f otherwise.

```
(primitive-closure? v) → boolean
  v : any/c
```

Returns #t if *v* is internally implemented as a primitive closure rather than a simple primitive procedure, #f otherwise.

```
(primitive-result-arity prim) → procedure-arity?
  prim : primitive?
```

Returns the arity of the result of the primitive procedure *prim* (as opposed to the procedure's input arity as returned by procedure-arity). For most primitives, this procedure returns 1, since most primitives return a single value when applied.

3.17.3 Additional Higher-Order Functions

```
(require racket/function)
```

The bindings documented in this section are provided by the `racket/function` and `racket` libraries, but not `racket/base`.

```
(identity v) → any/c
  v : any/c
```

Returns `v`.

```
(const v) → procedure?
  v : any
```

Returns a procedure that accepts any arguments (including keyword arguments) and returns `v`.

Examples:

```
> ((const 'foo) 1 2 3)
'foo
> ((const 'foo))
'foo
```

```
(thunk  body ...+)
(thunk* body ...+)
```

The `thunk` form creates a nullary function that evaluates the given body. The `thunk*` form is similar, except that the resulting function accepts any arguments (including keyword arguments).

Examples:

```
(define th1 (thunk (define x 1) (printf "~a\n" x)))

> (th1)
1

> (th1 'x)
th1: arity mismatch;
 the expected number of arguments does not match the given
number
  expected: 0
  given: 1
  arguments...:
```

```
    'x
> (th1 #:y 'z)
application: procedure does not accept keyword arguments
  procedure: th1
  given arguments:
   #:y 'z
(define th2 (thunk* (define x 1) (printf "~a\n" x)))

> (th2)
1

> (th2 'x)
1

> (th2 #:y 'z)
1
```

```
(negate proc) → procedure?
  proc : procedure?
```

Returns a procedure that is just like *proc*, except that it returns the `not` of *proc*'s result.

Examples:

```
> (filter (negate symbol?) '(1 a 2 b 3 c))
'(1 2 3)
> (map (negate =) '(1 2 3) '(1 1 1))
'(#f #t #t)
```

```
(curry proc) → procedure?
  proc : procedure?
(curry proc v ...+) → any/c
  proc : procedure?
  v : any/c
```

Returns a procedure that is a curried version of *proc*. When the resulting procedure is first applied, unless it is given the maximum number of arguments that it can accept, the result is a procedure to accept additional arguments.

Examples:

```
> ((curry list) 1 2)
#<procedure:curried>
> ((curry cons) 1)
#<procedure:curried>
> ((curry cons) 1 2)
'(1 . 2)
```

After the first application of the result of curry, each further application accumulates arguments until an acceptable number of arguments have been accumulated, at which point the original *proc* is called.

Examples:

```
> (((curry list) 1 2) 3)
'(1 2 3)
> (((curry list) 1) 3)
'(1 3)
> ((((curry foldl) +) 0) '(1 2 3))
6
```

A function call (curry *proc v* ...) is equivalent to ((curry *proc*) *v* ...). In other words, curry itself is curried.

The curry function provides limited support for keyworded functions: only the curry call itself can receive keyworded arguments to be propagated eventually to *proc*.

Examples:

```
> (map ((curry +) 10) '(1 2 3))
'(11 12 13)
> (map (curry + 10) '(1 2 3))
'(11 12 13)
> (map (compose (curry * 2) (curry + 10)) '(1 2 3))
'(22 24 26)
> (define foo (curry (lambda (x y z) (list x y z))))

> (foo 1 2 3)
'(1 2 3)
> (((((foo) 1) 2)) 3)
'(1 2 3)
```

```
(curryr proc) → procedure?
  proc : procedure?
(curryr proc v ...+) → any/c
  proc : procedure?
  v : any/c
```

Like curry, except that the arguments are collected in the opposite direction: the first step collects the rightmost group of arguments, and following steps add arguments to the left of these.

Example:

```
> (map (curryr list 'foo) '(1 2 3))
'((1 foo) (2 foo) (3 foo))
```

3.18 Void and Undefined

The constant `#<void>` is returned by most forms and procedures that have a side-effect and no useful result. The constant `#<undefined>` is used as the initial value for `letrec` bindings.

The `#<void>` value is always `eq?` to itself, and the `#<undefined>` value is also `eq?` to itself.

```
(void? v) → void?
  v : any/c
```

Returns `#t` if `v` is the constant `#<void>`, `#f` otherwise.

```
(void v ...) → void?
  v : any/c
```

Returns the constant `#<void>`. Each `v` argument is ignored.

4 Structures

§5 “Programmer-Defined Datatypes” in ***The Racket Guide*** introduces structure types via `struct`.

A *structure type* is a record datatype composing a number of *fields*. A *structure*, an instance of a structure type, is a first-class value that contains a value for each field of the structure type. A structure instance is created with a type-specific constructor procedure, and its field values are accessed and changed with type-specific accessor and mutator procedures. In addition, each structure type has a predicate procedure that answers `#t` for instances of the structure type and `#f` for any other value.

A structure type's fields are essentially unnamed, though names are supported for error-reporting purposes. The constructor procedure takes one value for each field of the structure type, except that some of the fields of a structure type can be *automatic fields*; the automatic fields are initialized to a constant that is associated with the structure type, and the corresponding arguments are omitted from the constructor procedure. All automatic fields in a structure type follow the non-automatic fields.

A structure type can be created as a *structure subtype* of an existing base structure type. An instance of a structure subtype can always be used as an instance of the base structure type, but the subtype gets its own predicate procedure, and it may have its own fields in addition to the fields of the base type.

A structure subtype “inherits” the fields of its base type. If the base type has m fields, and if n fields are specified for the new structure subtype, then the resulting structure type has $m+n$ fields. The value for automatic fields can be different in a subtype than in its base type.

If m' of the original m fields are non-automatic (where $m'<m$), and n' of the new fields are automatic (where $n'<n$), then $m'+n'$ field values must be provided to the subtype's constructor procedure. Values for the first m fields of a subtype instance are accessed with selector procedures for the original base type (or its supertypes), and the last n are accessed with subtype-specific selectors. Subtype-specific accessors and mutators for the first m fields do not exist.

The `struct` form and `make-struct-type` procedure typically create a new structure type, but they can also access *prefab* (i.e., previously fabricated) structure types that are globally shared, and whose instances can be parsed and written by the default reader (see §1.3 “The Reader”) and printer (see §1.4 “The Printer”). Prefab structure types can inherit only from other prefab structure types, and they cannot have guards (see §4.2 “Creating Structure Types”) or properties (see §4.3 “Structure Type Properties”). Exactly one prefab structure type exists for each combination of name, supertype, field count, automatic field count, automatic field value (when there is at least one automatic field), and field mutability.

§12.9 “Serialization” also provides information on reading and writing structures.

Two structure values are `eqv?` if and only if they are `eq?`. Two structure values are `equal?` if they are `eq?`. By default, two structure values are also `equal?` if they are instances of the same structure type, no fields are opaque, and the results of applying `struct->vector` to the structs are `equal?`. (Consequently, `equal?` testing for structures may depend on the current inspector.) A structure type can override the default `equal?` definition through the

gen:equal+hash generic interface.

4.1 Defining Structure Types: struct

§5 “Programmer-Defined Datatypes” in ***The Racket Guide*** introduces struct.

```
(struct id maybe-super (field ...)
        struct-option ...)

  maybe-super =
              | super-id

        field = field-id
              | [field-id field-option ...]

struct-option = #:mutable
              | #:super super-expr
              | #:inspector inspector-expr
              | #:auto-value auto-expr
              | #:guard guard-expr
              | #:property prop-expr val-expr
              | #:transparent
              | #:prefab
              | #:constructor-name constructor-id
              | #:extra-constructor-name constructor-id
              | #:reflection-name symbol-expr
              | #:methods gen:name method-defs
              | #:omit-define-syntaxes
              | #:omit-define-values

 field-option = #:mutable
              | #:auto

  method-defs = (definition ...)

  gen:name : identifier?
```

Creates a new structure type (or uses a pre-existing structure type if #:prefab is specified), and binds transformers and variables related to the structure type.

A struct form with *n* field s defines up to 4+2*n* names:

- struct:id, a *structure type descriptor* value that represents the structure type.
- *constructor-id* (which defaults to *id*), a *constructor* procedure that takes *m* arguments and returns a new instance of the structure type, where *m* is the number of field s that do not include an #:auto option.

- *id*, a transformer binding that encapsulates information about the structure type declaration. This binding is used to define subtypes, and it also works with the `shared` and `match` forms. For detailed information about the binding of *id*, see §4.7 “Structure Type Transformer Binding”.

  The *constructor-id* and *id* can be the same, in which case *id* performs both roles. In that case, the expansion of *id* as an expression produces an otherwise inaccessible identifier that is bound to the constructor procedure; the expanded identifier has a `'constructor-for` property whose value is an identifier that is `free-identifier=?` to *id*.

- *id*`?`, a *predicate* procedure that returns `#t` for instances of the structure type (constructed by *constructor-id* or the constructor for a subtype) and `#f` for any other value.

- *id-field-id*, for each *field*; an *accessor* procedure that takes an instance of the structure type and extracts the value for the corresponding field.

- `set-`*id-field-id*`!`, for each *field* that includes a `#:mutable` option, or when the `#:mutable` option is specified as a *struct-option*; a *mutator* procedure that takes an instance of the structure type and a new field value. The structure is destructively updated with the new value, and `#<void>` is returned.

If *super-id* is provided, it must have a transformer binding of the same sort bound to *id* (see §4.7 “Structure Type Transformer Binding”), and it specifies a supertype for the structure type. Alternately, the `#:super` option can be used to specify an expression that must produce a structure type descriptor. See §4 “Structures” for more information on structure subtypes and supertypes. If both *super-id* and `#:super` are provided, a syntax error is reported.

If the `#:mutable` option is specified for an individual field, then the field can be mutated in instances of the structure type, and a mutator procedure is bound. Supplying `#:mutable` as a *struct-option* is the same as supplying it for all *field*s. If `#:mutable` is specified as both a *field-option* and *struct-option*, a syntax error is reported.

The `#:inspector`, `#:auto-value`, and `#:guard` options specify an inspector, value for automatic fields, and guard procedure, respectively. See `make-struct-type` for more information on these attributes of a structure type. The `#:property` option, which is the only one that can be supplied multiple times, attaches a property value to the structure type; see §4.3 “Structure Type Properties” for more information on properties. The `#:transparent` option is a shorthand for `#:inspector #f`.

Use the `prop:procedure` property to implement an applicable structure, use `prop:evt` to create a structure type whose instances are synchronizable events, and so on. By convention, property names start with `prop:`.

The `#:prefab` option obtains a prefab (pre-defined, globally shared) structure type, as opposed to creating a new structure type. Such a structure type is inherently transparent and cannot have a guard or properties, so using `#:prefab` with `#:transparent`, `#:inspector`, `#:guard`, or `#:property` is a syntax error. If a supertype is specified, it must also be a prefab structure type.

If *constructor-id* is supplied, then the transformer binding of *id* records *constructor-id* as the constructor binding; as a result, for example, `struct-out` includes *constructor-id* as an export. If *constructor-id* is supplied via `#:extra-constructor-name` and it is not *id*, applying `object-name` on the constructor produces the symbolic form of *id* rather than *constructor-id*. If *constructor-id* is supplied via `#:constructor-name` and it is not the same as *id*, then *id* does not serve as a constructor, and `object-name` on the constructor produces the symbolic form of *constructor-id*. Only one of `#:extra-constructor-name` and `#:constructor-name` can be provided within a `struct` form.

If `#:reflection-name` *symbol-expr* is provided, then *symbol-expr* must produce a symbol that is used to identify the structure type in reflective operations such as `struct-type-info`. It corresponds to the first argument of `make-struct-type`. Structure printing uses the reflective name, as do the various procedures that are bound by `struct`.

If `#:methods` *gen:name method-defs* is provided, then *gen:name* must be a transformer binding for the static information about a generic group produced by `define-generics`. The *method-defs* define the methods of *gen:name*. If any method of *gen:name* is not defined, then `#f` is used to signify that the structure type does not implement the particular method. At least one method definition must be provided if this keyword is used. A `define/generic` form may appear in *method-defs*. Auxiliary definitions and expressions may also appear in *method-defs*.

If the `#:omit-define-syntaxes` option is supplied, then *id* is not bound as a transformer. If the `#:omit-define-values` option is supplied, then none of the usual variables are bound, but *id* is bound. If both are supplied, then the `struct` form is equivalent to `(begin)`.

If `#:auto` is supplied as a *field-option*, then the constructor procedure for the structure type does not accept an argument corresponding to the field. Instead, the structure type's automatic value is used for the field, as specified by the `#:auto-value` option, or as defaults to `#f` when `#:auto-value` is not supplied. The field is mutable (e.g., through reflective operations), but a mutator procedure is bound only if `#:mutable` is specified.

If a *field* includes the `#:auto` option, then all fields after it must also include `#:auto`, otherwise a syntax error is reported. If any *field-option* or *struct-option* keyword is repeated, other than `#:property`, a syntax error is reported.

For serialization, see `define-serializable-struct`.

Examples:

```
> (struct posn (x y [z #:auto])
    #:auto-value 0
    #:transparent)

> (posn 1 2)
(posn 1 2 0)
```

```
> (posn? (posn 1 2))
#t
> (posn-y (posn 1 2))
2
```

```
(struct color-posn posn (hue) #:mutable)
(define cp (color-posn 1 2 "blue"))

> (color-posn-hue cp)
"blue"
> cp
(color-posn 1 2 0 ...)
> (set-posn-z! cp 3)
set-posn-z!: undefined;
 cannot reference undefined identifier
```

```
(struct-field-index field-id)
```

This form can only appear as an expression within a `struct` form; normally, it is used with `#:property`, especially for a property like `prop:procedure`. The result of a `struct-field-index` expression is an exact, non-negative integer that corresponds to the position within the structure declaration of the field named by *field-id*.

Examples:

```
> (struct mood-procedure (base rating)
    #:property prop:procedure (struct-field-index base))

(define happy+ (mood-procedure add1 10))

> (happy+ 2)
3
> (mood-procedure-rating happy+)
10
```

```
(define-struct id-maybe-super (field ...)
               struct-option ...)

id-maybe-super = id
               | (id super-id)
```

Like `struct`, except that the syntax for supplying a *super-id* is different, and a *constructor-id* that has a `make-` prefix on *id* is implicitly supplied via `#:extra-constructor-name` if neither `#:extra-constructor-name` nor `#:constructor-name` is provided.

This form is provided for backwards compatibility; `struct` is preferred.

Examples:

```
(define-struct posn (x y [z #:auto])
   #:auto-value 0
   #:transparent)

> (make-posn 1 2)
(posn 1 2 0)
> (posn? (make-posn 1 2))
#t
> (posn-y (make-posn 1 2))
2
```

```
(define-struct/derived (id . rest-form)
  id-maybe-super (field ...) struct-option ...)
```

The same as `define-struct`, but with an extra `(id . rest-form)` sub-form that is treated as the overall form for syntax-error reporting and otherwise ignored. The only constraint on the sub-form for error reporting is that it starts with *id*. The `define-struct/derived` form is intended for use by macros that expand to `define-struct`.

Examples:

```
> (define-syntax (define-xy-struct stx)
    (syntax-case stx ()
     [(ds name . rest)
      (with-syntax ([orig stx])
        #'(define-struct/derived orig name (x y) . rest))]))

> (define-xy-struct posn)

> (posn-x (make-posn 1 2))
1
> (define-xy-struct posn #:mutable)

> (set-posn-x! (make-posn 1 2) 0)

; this next line will cause an error due to a bad keyword
> (define-xy-struct posn #:bad-option)
eval:24:0: define-xy-struct: unrecognized
struct-specification keyword
  at: #:bad-option
  in: (define-xy-struct posn #:bad-option)
```

4.2 Creating Structure Types

```
(make-struct-type  name
                   super-type
                   init-field-cnt
                   auto-field-cnt
                  [auto-v
                   props
                   inspector
                   proc-spec
                   immutables
                   guard
                   constructor-name]
                   generate
                   exercise)
   struct-type?
   struct-constructor-procedure?
 → struct-predicate-procedure?
   struct-accessor-procedure?
   struct-mutator-procedure?
  name : symbol?
  super-type : (or/c struct-type? #f)
  init-field-cnt : exact-nonnegative-integer?
  auto-field-cnt : exact-nonnegative-integer?
  auto-v : any/c = #f
  props :        (listof (cons/c struct-type-property?
                                 any/c))                  = null
  inspector : (or/c inspector? #f 'prefab) = (current-inspector)
              (or/c procedure?
  proc-spec :       exact-nonnegative-integer? = #f
                    #f)
  immutables : (listof exact-nonnegative-integer?) = null
  guard : (or/c procedure? #f) = #f
  constructor-name : (or/c symbol? #f) = #f
  generate : (-> contract? (-> int? any/c))
  exercise : (-> contract? (-> int? any/c any/c))
```

Creates a new structure type, unless *inspector* is 'prefab, in which case make-struct-type accesses a prefab structure type. The *name* argument is used as the type name. If *super-type* is not #f, the resulting type is a subtype of the corresponding structure type.

The resulting structure type has *init-field-cnt*+*auto-field-cnt* fields (in addition to any fields from *super-type*), but only *init-field-cnt* constructor arguments (in addition to any constructor arguments from *super-type*). The remaining fields are initialized with *auto-v*. The total field count (including *super-type* fields) must be no more than 32768.

The *props* argument is a list of pairs, where the car of each pair is a structure type property

descriptor, and the cdr is an arbitrary value. A property can be specified multiple times in *props* (including properties that are automatically added by properties that are directly included in *props*) only if the associated values are eq?, otherwise the exn:fail:contract exception is raised. See §4.3 "Structure Type Properties" for more information about properties. When *inspector* is 'prefab, then *props* must be null.

The *inspector* argument normally controls access to reflective information about the structure type and its instances; see §13.9 "Structure Inspectors" for more information. If *inspector* is 'prefab, then the resulting prefab structure type and its instances are always transparent.

If *proc-spec* is an integer or procedure, instances of the structure type act as procedures. See prop:procedure for further information. Providing a non-#f value for *proc-spec* is the same as pairing the value with prop:procedure at the end of *props*, plus including *proc-spec* in *immutables* when *proc-spec* is an integer.

The *immutables* argument provides a list of field positions. Each element in the list must be unique, otherwise exn:fail:contract exception is raised. Each element must also fall in the range 0 (inclusive) to *init-field-cnt* (exclusive), otherwise exn:fail:contract exception is raised.

The *guard* argument is either a procedure of $n+1$ arguments or #f, where n is the number of arguments for the new structure type's constructor (i.e., *init-field-cnt* plus constructor arguments implied by *super-type*, if any). If *guard* is a procedure, then the procedure is called whenever an instance of the type is constructed, or whenever an instance of a subtype is created. The arguments to *guard* are the values provided for the structure's first n fields, followed by the name of the instantiated structure type (which is *name*, unless a subtype is instantiated). The *guard* result must be n values, which become the actual values for the structure's fields. The *guard* can raise an exception to prevent creation of a structure with the given field values. If a structure subtype has its own guard, the subtype guard is applied first, and the first n values produced by the subtype's guard procedure become the first n arguments to *guard*. When *inspector* is 'prefab, then *guard* must be #f.

If *constructor-name* is not #f, it is used as the name of the generated constructor procedure as returned by object-name or in the printed form of the constructor value.

The *generate* argument is used to define a new generator for this structure type, which can be used to create random instances of the structure type. For more information see contract-generate.

The *exercise* argument allows you to define a function to verify that a given value is an instance of your contract. This will also be used for random generation.

The result of make-struct-type is five values:

- a structure type descriptor,

- a constructor procedure,
- a predicate procedure,
- an accessor procedure, which consumes a structure and a field index between 0 (inclusive) and *init-field-cnt*+*auto-field-cnt* (exclusive), and
- a mutator procedure, which consumes a structure, a field index, and a field value.

Examples:

```
> (define-values (struct:a make-a a? a-ref a-set!)
    (make-struct-type 'a #f 2 1 'uninitialized))

> (define an-a (make-a 'x 'y))

> (a-ref an-a 1)
'y
> (a-ref an-a 2)
'uninitialized
> (define a-first (make-struct-field-accessor a-ref 0))

> (a-first an-a)
'x

> (define-values (struct:b make-b b? b-ref b-set!)
    (make-struct-type 'b struct:a 1 2 'b-uninitialized))

> (define a-b (make-b 'x 'y 'z))

> (a-ref a-b 1)
'y
> (a-ref a-b 2)
'uninitialized
> (b-ref a-b 0)
'z
> (b-ref a-b 1)
'b-uninitialized
> (b-ref a-b 2)
'b-uninitialized


> (define-values (struct:c make-c c? c-ref c-set!)
    (make-struct-type
     'c struct:b 0 0 #f null (make-inspector) #f null
     ; guard checks for a number, and makes it inexact
     (lambda (a1 a2 b1 name)
       (unless (number? a2)
```

```
          (error (string->symbol (format "make-~a" name))
                 "second field must be a number"))
        (values a1 (exact->inexact a2) b1))))

> (make-c 'x 'y 'z)
make-c: second field must be a number
> (define a-c (make-c 'x 2 'z))

> (a-ref a-c 1)
2.0


> (define p1 #s(p a b c))

> (define-values (struct:p make-p p? p-ref p-set!)
    (make-struct-type 'p #f 3 0 #f null 'prefab #f '(0 1 2)))

> (p? p1)
#t
> (p-ref p1 0)
'a
> (make-p 'x 'y 'z)
'#s(p x y z)
```

```
(make-struct-field-accessor accessor-proc
                            field-pos
                            [field-name]) → procedure?
  accessor-proc : struct-accessor-procedure?
  field-pos : exact-nonnegative-integer?
  field-name : (or/c symbol? #f)
             = (symbol->string (format "field~a" field-pos))
```

Returns a field accessor that is equivalent to `(lambda (s) (accessor-proc s field-pos))`. The *accessor-proc* must be an accessor returned by `make-struct-type`. The name of the resulting procedure for debugging purposes is derived from *field-name* and the name of *accessor-proc*'s structure type if *field-name* is a symbol.

For examples, see `make-struct-type`.

```
(make-struct-field-mutator mutator-proc
                           field-pos
                           [field-name]) → procedure?
  mutator-proc : struct-mutator-procedure?
  field-pos : exact-nonnegative-integer?
  field-name : (or/c symbol? #f)
             = (symbol->string (format "field~a" field-pos))
```

Returns a field mutator that is equivalent to `(lambda (s v) (`*mutator-proc* `s` *field-pos* `v))`. The *mutator-proc* must be a mutator returned by `make-struct-type`. The name of the resulting procedure for debugging purposes is derived from *field-name* and the name of *mutator-proc*'s structure type if *field-name* is a symbol.

For examples, see `make-struct-type`.

4.3 Structure Type Properties

§4.4 "Generic Interfaces" provide a high-level API on top of structure type properties.

A *structure type property* allows per-type information to be associated with a structure type (as opposed to per-instance information associated with a structure value). A property value is associated with a structure type through the `make-struct-type` procedure (see §4.2 "Creating Structure Types") or through the `#:property` option of `struct`. Subtypes inherit the property values of their parent types, and subtypes can override an inherited property value with a new value.

```
(make-struct-type-property  name
                           [guard
                                       struct-type-property?
                            supers]) → procedure?
                                       procedure?
  name : symbol?
  guard : (or/c procedure? #f 'can-impersonate) = #f
  supers : (listof (cons/c struct-type-property?
                           (any/c . -> . any/c))) = null
```

Creates a new structure type property and returns three values:

- a *structure type property descriptor*, for use with `make-struct-type` and `struct`;
- a *property predicate* procedure, which takes an arbitrary value and returns `#t` if the value is a descriptor or instance of a structure type that has a value for the property, `#f` otherwise;
- a *property accessor* procedure, which returns the value associated with the structure type given its descriptor or one of its instances; if the structure type does not have a value for the property, or if any other kind of value is provided, the `exn:fail:contract` exception is raised unless a second argument, *failure-result*, is supplied to the procedure. In that case, if *failure-result* is a procedure, it is called (through a tail call) with no arguments to produce the result of the property accessor procedure; otherwise, *failure-result* is itself returned as the result.

If the optional *guard* is supplied as a procedure, it is called by `make-struct-type` before attaching the property to a new structure type. The *guard* must accept two arguments: a

value for the property supplied to make-struct-type, and a list containing information about the new structure type. The list contains the values that struct-type-info would return for the new structure type if it skipped the immediate current-inspector control check (but not the check for exposing an ancestor structure type, if any; see §13.9 “Structure Inspectors”).

The result of calling *guard* is associated with the property in the target structure type, instead of the value supplied to make-struct-type. To reject a property association (e.g., because the value supplied to make-struct-type is inappropriate for the property), the *guard* can raise an exception. Such an exception prevents make-struct-type from returning a structure type descriptor.

If *guard* is 'can-impersonate, then the property's accessor can be redirected through impersonate-struct. Otherwise, redirection of the property value through an impersonator is disallowed, since redirection is tantamount to mutation.

The optional *supers* argument is a list of properties that are automatically associated with some structure type when the newly created property is associated to the structure type. Each property in *supers* is paired with a procedure that receives the value supplied for the new property (after it is processed by *guard*) and returns a value for the associated property (which is then sent to that property's guard, of any).

Examples:

```
> (define-values (prop:p p? p-ref) (make-struct-type-property 'p))

> (define-values (struct:a make-a a? a-ref a-set!)
    (make-struct-type 'a #f 2 1 'uninitialized
                      (list (cons prop:p 8))))

> (p? struct:a)
#t
> (p? 13)
#f
> (define an-a (make-a 'x 'y))

> (p? an-a)
#t
> (p-ref an-a)
8
> (define-values (struct:b make-b b? b-ref b-set!)
    (make-struct-type 'b #f 0 0 #f))

> (p? struct:b)
#f
> (define-values (prop:q q? q-ref) (make-struct-type-property
                                    'q (lambda (v si) (add1 v))
```

```
                                    (list (cons prop:p sqrt))))

> (define-values (struct:c make-c c? c-ref c-set!)
    (make-struct-type 'c #f 0 0 'uninit
                      (list (cons prop:q 8))))

> (q-ref struct:c)
9
> (p-ref struct:c)
3
```

```
(struct-type-property? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a structure type property descriptor value, `#f` otherwise.

```
(struct-type-property-accessor-procedure? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is an accessor procedure produced by `make-struct-type-property`, `#f` otherwise.

4.4 Generic Interfaces

```
(require racket/generic)
```

A *generic interface* allows per-type methods to be associated with generic functions. Generic functions are defined using a `define-generics` form. Method implementations for a structure type are defined using the `#:methods` keyword (see §4.1 "Defining Structure Types: `struct`").

```
(define-generics id [#:defined-table defined-table-id]
  [method-id . kw-formals*]
  ...)

kw-formals* = (arg* ...)
            | (arg* ...+ . rest-id)
            | rest-id

      arg* = arg-id
           | [arg-id]
           | keyword arg-id
           | keyword [arg-id]
```

Defines

- `gen:id` as a transformer binding for the static information about a new generic interface;
- `id?` as a predicate identifying instances of structure types that implement this generic group; and
- each *method-id* as a generic procedure that calls the corresponding method on values where `id?` is true.
- `id/c` as a contract combinator that recognizes instances of structure types which implement the `gen:id` generic interface. The combinator takes pairs of *method-id*s and contracts. The contracts will be applied to each of the corresponding method implementations.

Each *method-id*'s *kw-formals** must include a required by-position argument that is `free-identifier=?` to *id*. That argument is used in the generic definition to locate the specialization.

When *defined-table-id* is provided, it is defined as a procedure that takes an instance of the generics and returns an immutable hash table that maps symbols corresponding to method names to booleans representing whether or not that method is implemented by the instance. This table is intended for use by higher-level APIs to adapt their behavior depending on method availability.

The `id/c` combinator is intended to be used to contract the range of a constructor procedure for a struct type that implements the generic interface.

```
(define/generic local-id method-id)

  local-id : identifier?
  method-id : identifier?
```

When used inside the method definitions associated with the `#:methods` keyword, binds *local-id* to the generic for *method-id*. This form is useful for method specializations to use generic methods (as opposed to the local specialization) on other values.

Using the `define/generic` form outside a `#:methods` specification in `struct` (or `define-struct`) is an syntax error.

Examples:

```
> (define-generics printable
    (gen-print printable [port])
    (gen-port-print port printable)
    (gen-print* printable [port] #:width width #:height [height]))
```

```
> (define-struct num (v)
    #:methods gen:printable
    [(define/generic super-print gen-print)
     (define (gen-print n [port (current-output-port)])
       (fprintf port "Num: ~a" (num-v n)))
     (define (gen-port-print port n)
       (super-print n port))
     (define (gen-print* n [port (current-output-port)]
                         #:width w #:height [h 0])
       (fprintf port "Num (~ax~a): ~a" w h (num-v n)))])

> (define-struct bool (v)
    #:methods gen:printable
    [(define/generic super-print gen-print)
     (define (gen-print b [port (current-output-port)])
       (fprintf port "Bool: ~a"
                (if (bool-v b) "Yes" "No")))
     (define (gen-port-print port b)
       (super-print b port))
     (define (gen-print* b [port (current-output-port)]
                         #:width w #:height [h 0])
       (fprintf port "Bool (~ax~a): ~a" w h
                (if (bool-v b) "Yes" "No")))])

> (define x (make-num 10))

> (gen-print x)
Num: 10

> (gen-port-print (current-output-port) x)
Num: 10

> (gen-print* x #:width 100 #:height 90)
Num (100x90): 10

> (define y (make-bool #t))

> (gen-print y)
Bool: Yes

> (gen-port-print (current-output-port) y)
Bool: Yes

> (gen-print* y #:width 100 #:height 90)
Bool (100x90): Yes
```

```
> (define/contract make-num-contracted
    (-> number?
        (printable/c
          [gen-print (->* (printable?) (output-port?) void?)]
          [gen-port-print (-> output-port? printable? void?)]
          [gen-print* (->* (printable? #:width exact-nonnegative-
integer?)
                           (output-port? #:height exact-
nonnegative-integer?)
                           void?)]))
     make-num)

> (define z (make-num-contracted 10))

> (gen-print* z #:width "not a number" #:height 5)
make-num-contracted: contract violation
  expected: exact-nonnegative-integer?, given: "not a
number"
  in: the #:width argument of
      the gen-print* method of
      the range of
      (->
       number?
       (printable/c
        (gen-print
         (->*
          (gen:printable?)
          (output-port?)
          void?))
        (gen-port-print
         (-> output-port? gen:printable? void?))
        (gen-print*
         (->*
          (gen:printable?
           #:width
           exact-nonnegative-integer?)
          (output-port?
           #:height
           exact-nonnegative-integer?)
          void?))))
  contract from:
      (definition make-num-contracted)
  blaming: top-level
  at: eval:13.0
```

4.5 Copying and Updating Structures

```
(struct-copy id struct-expr fld-id ...)

fld-id = [field-id expr]
       | [field-id #:parent parent-id expr]
```

Creates a new instance of the structure type *id* with the same field values as the structure produced by *struct-expr*, except that the value of each supplied *field-id* is instead determined by the corresponding *expr*. If `#:parent` is specified, the *parent-id* must be bound to a parent structure type of *id*.

The *id* must have a transformer binding that encapsulates information about a structure type (i.e., like the initial identifier bound by `struct`), and the binding must supply a constructor, a predicate, and all field accessors.

Each *field-id* is combined with *id* (or *parent-id*, if present) to form *id-field-id* (using the lexical context of *field-id*), which must be one of the accessor bindings in *id*. The accessor bindings determined by different *field-id*s must be distinct. The order of the *field-id*s need not match the order of the corresponding fields in the structure type.

The *struct-expr* is evaluated first. The result must be an instance of the *id* structure type, otherwise the `exn:fail:contract` exception is raised. Next, the field *expr*s are evaluated in order (even if the fields that correspond to the *field-id*s are in a different order). Finally, the new structure instance is created.

The result of *struct-expr* can be an instance of a sub-type of *id*, but the resulting copy is an immediate instance of *id* (not the sub-type).

Examples:

```
> (struct fish (color weight) #:transparent)

> (define marlin (fish 'orange-and-white 11))

> (define dory (struct-copy fish marlin
                            [color 'blue]))

> dory
(fish 'blue 11)
> (struct shark fish (weeks-since-eating-fish) #:transparent)

> (define bruce (shark 'grey 110 3))

> (define chum (struct-copy shark bruce
                            [weight #:parent fish 90]
                            [weeks-since-eating-fish 0]))
```

```
> chum
(shark 'grey 90 0)
; subtypes can be copied as if they were supertypes,
; but the result is an instance of the supertype
> (define not-really-chum
    (struct-copy fish bruce
                 [weight 90]))

> not-really-chum
(fish 'grey 90)
```

4.6 Structure Utilities

```
(struct->vector v [opaque-v]) → vector?
  v : any/c
  opaque-v : any/c = '...
```

Creates a vector representing `v`. The first slot of the result vector contains a symbol whose printed name has the form `struct:id`. Each remaining slot contains either the value of a field in `v`, if it is accessible via the current inspector, or `opaque-v` for a field that is not accessible. A single `opaque-v` value is used in the vector for contiguous inaccessible fields. (Consequently, the size of the vector does not match the size of the `struct` if more than one field is inaccessible.)

```
(struct? v) → any
  v : any/c
```

Returns `#t` if `struct-info` exposes any structure types of `v` with the current inspector, `#f` otherwise.

Typically, when `(struct? v)` is true, then `(struct->vector v)` exposes at least one field value. It is possible, however, for the only visible types of `v` to contribute zero fields.

```
(struct-type? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a structure type descriptor value, `#f` otherwise.

```
(struct-constructor-procedure? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a constructor procedure generated by `struct` or `make-struct-type`, `#f` otherwise.

```
(struct-predicate-procedure? v) → boolean?
  v : any/c
```

Returns #t if v is a predicate procedure generated by struct or make-struct-type, #f otherwise.

```
(struct-accessor-procedure? v) → boolean?
  v : any/c
```

Returns #t if v is an accessor procedure generated by struct, make-struct-type, or make-struct-field-accessor, #f otherwise.

```
(struct-mutator-procedure? v) → boolean?
  v : any/c
```

Returns #t if v is a mutator procedure generated by struct, make-struct-type, or make-struct-field-mutator, #f otherwise.

```
(prefab-struct-key v) → (or/c #f symbol? list?)
  v : any/c
```

Returns #f if v is not an instance of a prefab structure type. Otherwise, the result is the shorted key that could be used with make-prefab-struct to create an instance of the structure type.

Examples:

```
> (prefab-struct-key #s(cat "Garfield"))
'cat
> (struct cat (name) #:prefab)

> (struct cute-cat cat (shipping-dest) #:prefab)

> (cute-cat "Nermel" "Abu Dhabi")
'#s((cute-cat cat 1) "Nermel" "Abu Dhabi")
> (prefab-struct-key (cute-cat "Nermel" "Abu Dhabi"))
'(cute-cat cat 1)
```

```
(make-prefab-struct key v ...) → struct?
  key : prefab-key?
  v : any/c
```

Creates an instance of a prefab structure type, using the vs as field values. The *key* and the number of vs determine the prefab structure type.

A *key* identifies a structure type based on a list with the following items:

- A symbol for the structure type's name.
- An exact, nonnegative integer representing the number of non-automatic fields in the structure type, not counting fields from the supertype (if any).
- A list of two items, where the first is an exact, nonnegative integer for the number of automatic fields in the structure type that are not from the supertype (if any), and the second element is an arbitrary value that is the value for the automatic fields.
- A vector of exact, nonnegative integers that indicate mutable non-automatic fields in the structure type, counting from `0` and not including fields from the supertype (if any).
- Nothing else, if the structure type has no supertype. Otherwise, the rest of the list is the key for the supertype.

An empty vector and an auto-field list that starts with `0` can be omitted. Furthermore, the first integer (which indicates the number of non-automatic fields) can be omitted, since it can be inferred from the number of supplied *v*s. Finally, a single symbol can be used instead of a list that contains only a symbol (in the case that the structure type has no supertype, no automatic fields, and no mutable fields).

The total field count must be no more than 32768. If the number of fields indicated by *key* is inconsistent with the number of supplied *v*s, the `exn:fail:contract` exception is raised.

Examples:

```
> (make-prefab-struct 'clown "Binky" "pie")
'#s(clown "Binky" "pie")
> (make-prefab-struct '(clown 2) "Binky" "pie")
'#s(clown "Binky" "pie")
> (make-prefab-struct '(clown 2 (0 #f) #()) "Binky" "pie")
'#s(clown "Binky" "pie")
> (make-prefab-struct '(clown 1 (1 #f) #()) "Binky" "pie")
'#s((clown (1 #f)) "Binky" "pie")
> (make-prefab-struct '(clown 1 (1 #f) #(0)) "Binky" "pie")
'#s((clown (1 #f) #(0)) "Binky" "pie")
```

```
(prefab-key->struct-type key field-count) → struct-type?
  key : prefab-key?
  field-count : (integer-in 0 32768)
```

Returns a structure type descriptor for the prefab structure type specified by the combination of *key* and *field-count*.

If the number of fields indicated by *key* is inconsistent with *field-count*, the `exn:fail:contract` exception is raised.

```
(prefab-key? v) → boolean?
  v : any/c
```

Return `#t` if `v` can be a prefab structure type key, `#f` otherwise.

See `make-prefab-struct` for a description of valid key shapes.

4.7 Structure Type Transformer Binding

The `struct` form binds the name of a structure type as a transformer binding that records the other identifiers bound to the structure type, the constructor procedure, the predicate procedure, and the field accessor and mutator procedures. This information can be used during the expansion of other expressions via `syntax-local-value`.

For example, the `struct` variant for subtypes uses the base type name `t` to find the variable `struct:t` containing the base type's descriptor; it also folds the field accessor and mutator information for the base type into the information for the subtype. As another example, the `match` form uses a type name to find the predicates and field accessors for the structure type. The `struct` form in an imported signature for `unit` causes the `unit` transformer to generate information about imported structure types, so that `match` and subtyping `struct` forms work within the unit.

The expansion-time information for a structure type can be represented directly as a list of six elements (of the same sort that the encapsulated procedure must return):

- an identifier that is bound to the structure type's descriptor, or `#f` if none is known;
- an identifier that is bound to the structure type's constructor, or `#f` if none is known;
- an identifier that is bound to the structure type's predicate, or `#f` if none is known;
- a list of identifiers bound to the field accessors of the structure type, optionally with `#f` as the list's last element. A `#f` as the last element indicates that the structure type may have additional fields, otherwise the list is a reliable indicator of the number of fields in the structure type. Furthermore, the accessors are listed in reverse order for the corresponding constructor arguments. (The reverse order enables sharing in the lists for a subtype and its base type.)
- a list of identifiers bound to the field mutators of the structure type, or `#f` for each field that has no known mutator, and optionally with an extra `#f` as the list's last element (if the accessor list has such a `#f`). The list's order and the meaning of a final `#f` are the same as for the accessor identifiers, and the length of the mutator list is the same as the accessor list's length.

- an identifier that determines a super-type for the structure type, `#f` if the super-type (if any) is unknown, or `#t` if there is no super-type. If a super-type is specified, the identifier is also bound to structure-type expansion-time information.

Instead of this direct representation, the representation can be a structure created by `make-struct-info` (or an instance of a subtype of `struct:struct-info`), which encapsulates a procedure that takes no arguments and returns a list of six elements. Alternately, the representation can be a structure whose type has the `prop:struct-info` structure type property. Finally, the representation can be an instance of a structure type derived from `struct:struct-info` or with the `prop:struct-info` property that also implements `prop:procedure`, and where the instance is further is wrapped by `make-set!-transformer`. In addition, the representation may implement the `prop:struct-auto-info` property.

Use `struct-info?` to recognize all allowed forms of the information, and use `extract-struct-info` to obtain a list from any representation.

The implementor of a syntactic form can expect users of the form to know what kind of information is available about a structure type. For example, the `match` implementation works with structure information containing an incomplete set of accessor bindings, because the user is assumed to know what information is available in the context of the `match` expression. In particular, the `match` expression can appear in a `unit` form with an imported structure type, in which case the user is expected to know the set of fields that are listed in the signature for the structure type.

```
(require racket/struct-info)
```

The bindings documented in this section are provided by the `racket/struct-info` library, not `racket/base` or `racket`.

```
(struct-info? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is either a six-element list with the correct shape for representing structure-type information, a procedure encapsulated by `make-struct-info`, a structure with the `prop:struct-info` property, or a structure type derived from `struct:struct-info` or with `prop:struct-info` and wrapped with `make-set!-transformer`.

```
(checked-struct-info? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a procedure encapsulated by `make-struct-info` and produced by `struct`, but only when no parent type is specified or the parent type is also specified through a transformer binding to such a value.

```
(make-struct-info thunk) → struct-info?
  thunk : (-> (and/c struct-info? list?))
```

Encapsulates a thunk that returns structure-type information in list form.

```
(extract-struct-info v) → (and/c struct-info? list?)
  v : struct-info?
```

Extracts the list form of the structure type information represented by *v*.

```
struct:struct-info : struct-type?
```

The structure type descriptor for the structure type returned by `make-struct-info`. This structure type descriptor is mostly useful for creating structure subtypes. The structure type includes a guard that checks an instance's first field in the same way as `make-struct-info`.

```
prop:struct-info : struct-type-property?
```

The structure type property for creating new structure types like `struct:struct-info`. The property value must be a procedure of one argument that takes an instance structure and returns structure-type information in list form.

```
prop:struct-auto-info : struct-type-property?
(struct-auto-info? v) → boolean?
  v : any/c
(struct-auto-info-lists sai)
 → (list/c (listof identifier?) (listof identifier?))
  sai : struct-auto-info?
```

The `prop:struct-auto-info` property is implemented to provide static information about which of the accessor and mutator identifiers for a structure type correspond to `#:auto` fields (so that they have no corresponding argument in the constructor). The property value must be a procedure that accepts an instance structure to which the property is given, and the result must be two lists of identifiers suitable as a result from `struct-auto-info-lists`.

The `struct-auto-info?` predicate recognizes values that implement the `prop:struct-auto-info` property.

The `struct-auto-info-lists` function extracts two lists of identifiers from a value that implements the `prop:struct-auto-info` property. The first list should be a subset of the accessor identifiers for the structure type described by *sai*, and the second list should be a subset of the mutator identifiers. The two subsets correspond to `#:auto` fields.

5 Classes and Objects

§13 “Classes and Objects” in ***The Racket Guide*** introduces classes and objects.

```
(require racket/class)
```

The bindings documented in this section are provided by the `racket/class` and `racket` libraries, but not `racket/base`.

A *class* specifies

- a collection of fields;
- a collection of methods;
- initial value expressions for the fields; and
- initialization variables that are bound to initialization arguments.

In the context of the class system, an *object* is a collection of bindings for fields that are instantiated according to a class description.

The class system allows a program to define a new class (a *derived class*) in terms of an existing class (the *superclass*) using inheritance, overriding, and augmenting:

- *inheritance*: An object of a derived class supports methods and instantiates fields declared by the derived class’s superclass, as well as methods and fields declared in the derived class expression.
- *overriding*: Some methods declared in a superclass can be replaced in the derived class. References to the overridden method in the superclass use the implementation in the derived class.
- *augmenting*: Some methods declared in a superclass can be merely extended in the derived class. The superclass method specifically delegates to the augmenting method in the derived class.

An *interface* is a collection of method names to be implemented by a class, combined with a derivation requirement. A class *implements* an interface when it

- declares (or inherits) a public method for each variable in the interface;
- is derived from the class required by the interface, if any; and
- specifically declares its intention to implement the interface.

A class can implement any number of interfaces. A derived class automatically implements any interface that its superclass implements. Each class also implements an implicitly-defined interface that is associated with the class. The implicitly-defined interface contains all of the class's public method names, and it requires that all other implementations of the interface are derived from the class.

A new interface can *extend* one or more interfaces with additional method names; each class that implements the extended interface also implements the original interfaces. The derivation requirements of the original interface must be consistent, and the extended interface inherits the most specific derivation requirement from the original interfaces.

Classes, objects, and interfaces are all values. However, a class or interface is not an object (i.e., there are no "meta-classes" or "meta-interfaces").

5.1 Creating Interfaces

§13 "Classes and Objects" in ***The Racket Guide*** introduces classes, objects, and interfaces.

```
(interface (super-interface-expr ...) name-clause ...)

name-clause = id
            | (id contract-expr)
```

Produces an interface. The *id*s must be mutually distinct.

Each *super-interface-expr* is evaluated (in order) when the `interface` expression is evaluated. The result of each *super-interface-expr* must be an interface value, otherwise the `exn:fail:object` exception is raised. The interfaces returned by the *super-interface-expr*s are the new interface's superinterfaces, which are all extended by the new interface. Any class that implements the new interface also implements all of the superinterfaces.

The result of an `interface` expression is an interface that includes all of the specified *id*s, plus all identifiers from the superinterfaces. Duplicate identifier names among the superinterfaces are ignored, but if a superinterface contains one of the *id*s in the `interface` expression, the `exn:fail:object` exception is raised. A given *id* may be paired with a corresponding *contract-expr*.

If no *super-interface-expr*s are provided, then the derivation requirement of the resulting interface is trivial: any class that implements the interface must be derived from `object%`. Otherwise, the implementation requirement of the resulting interface is the most specific requirement from its superinterfaces. If the superinterfaces specify inconsistent derivation requirements, the `exn:fail:object` exception is raised.

Examples:

```
(define file-interface<%>
  (interface () open close read-byte write-byte))
```

```
(define directory-interface<%>
  (interface (file-interface<%>)
    [file-list (-> (listof (is-a?/c file-interface<%>)))]
    parent-directory))
```

```
(interface* (super-interface-expr ...)
            ([property-expr val-expr] ...)
  name-clause ...)

name-clause = id
            | (id contract-expr)
```

Like `interface`, but also associates to the interface the structure-type properties produced by the *property-expr*s with the corresponding *val-expr*s.

Whenever the resulting interface (or a sub-interface derived from it) is explicitly implemented by a class through the `class*` form, each property is attached with its value to a structure type that instantiated by instances of the class. Specifically, the property is attached to a structure type with zero immediate fields, which is extended to produce the internal structure type for instances of the class (so that no information about fields is accessible to the structure type property's guard, if any).

Example:

```
(define i<%> (interface* () ([prop:custom-write
                              (lambda (obj port mode) (void))])
               method1 method2 method3))
```

5.2 Creating Classes

§13 "Classes and Objects" in ***The Racket Guide*** introduces classes and objects.

```
object% : class?
```

A built-in class that has no methods fields, implements only its own interface (`class->interface object%`), and is transparent (i.e,. its inspector is `#f`, so all immediate instances are `equal?`). All other classes are derived from `object%`.

```
(class* superclass-expr (interface-expr ...)
  class-clause
  ...)
```

```
     class-clause = (inspect inspector-expr)
                  | (init init-decl ...)
                  | (init-field init-decl ...)
                  | (field field-decl ...)
                  | (inherit-field maybe-renamed ...)
                  | (init-rest id)
                  | (init-rest)
                  | (public maybe-renamed ...)
                  | (pubment maybe-renamed ...)
                  | (public-final maybe-renamed ...)
                  | (override maybe-renamed ...)
                  | (overment maybe-renamed ...)
                  | (override-final maybe-renamed ...)
                  | (augment maybe-renamed ...)
                  | (augride maybe-renamed ...)
                  | (augment-final maybe-renamed ...)
                  | (private id ...)
                  | (abstract id ...)
                  | (inherit maybe-renamed ...)
                  | (inherit/super maybe-renamed ...)
                  | (inherit/inner maybe-renamed ...)
                  | (rename-super renamed ...)
                  | (rename-inner renamed ...)
                  | method-definition
                  | definition
                  | expr
                  | (begin class-clause ...)

        init-decl = id
                  | (renamed)
                  | (maybe-renamed default-value-expr)

       field-decl = (maybe-renamed default-value-expr)

    maybe-renamed = id
                  | renamed

          renamed = (internal-id external-id)

method-definition = (define-values (id) method-procedure)

 method-procedure = (lambda kw-formals expr ...+)
                  | (case-lambda (formals expr ...+) ...)
                  | (#%plain-lambda formals expr ...+)
                  | (let-values ([(id) method-procedure] ...)
                      method-procedure)
                  | (letrec-values ([(id) method-procedure] ...)
                      method-procedure)
                  | (let-values ([(id) method-procedure] ...+)
                      id)
                  | (letrec-values ([(id) method-procedure] ...+)
                      id)
```

Produces a class value.

The *superclass-expr* expression is evaluated when the `class*` expression is evaluated. The result must be a class value (possibly `object%`), otherwise the `exn:fail:object` exception is raised. The result of the *superclass-expr* expression is the new class's superclass.

The *interface-expr* expressions are also evaluated when the `class*` expression is evaluated, after *superclass-expr* is evaluated. The result of each *interface-expr* must be an interface value, otherwise the `exn:fail:object` exception is raised. The interfaces returned by the *interface-expr*s are all implemented by the class. For each identifier in each interface, the class (or one of its ancestors) must declare a public method with the same name, otherwise the `exn:fail:object` exception is raised. The class's superclass must satisfy the implementation requirement of each interface, otherwise the `exn:fail:object` exception is raised.

An `inspect` *class-clause* selects an inspector (see §13.9 "Structure Inspectors") for the class extension. The *inspector-expr* must evaluate to an inspector or `#f` when the `class*` form is evaluated. Just as for structure types, an inspector controls access to the class's fields, including private fields, and also affects comparisons using `equal?`. If no `inspect` clause is provided, access to the class is controlled by the parent of the current inspector (see §13.9 "Structure Inspectors"). A syntax error is reported if more than one `inspect` clause is specified.

The other *class-clause*s define initialization arguments, public and private fields, and public and private methods. For each *id* or *maybe-renamed* in a `public`, `override`, `augment`, `pubment`, `overment`, `augride`, `public-final`, `override-final`, `augment-final`, or `private` clause, there must be one *method-definition*. All other definition *class-clause*s create private fields. All remaining *expr*s are initialization expressions to be evaluated when the class is instantiated (see §5.3 "Creating Objects").

The result of a `class*` expression is a new class, derived from the specified superclass and implementing the specified interfaces. Instances of the class are created with the `instantiate` form or `make-object` procedure, as described in §5.3 "Creating Objects".

Each *class-clause* is (partially) macro-expanded to reveal its shapes. If a *class-clause* is a `begin` expression, its sub-expressions are lifted out of the `begin` and treated as *class-clause*s, in the same way that `begin` is flattened for top-level and embedded definitions.

Within a `class*` form for instances of the new class, `this` is bound to the object itself; `this%` is bound to the class of the object; `super-instantiate`, `super-make-object`, and `super-new` are bound to forms to initialize fields in the superclass (see §5.3 "Creating Objects"); `super` is available for calling superclass methods (see §5.2.3.1 "Method Definitions"); and `inner` is available for calling subclass augmentations of methods (see §5.2.3.1 "Method Definitions").

```
(class superclass-expr class-clause ...)
```

Like `class*`, but omits the *`interface-expr`*s, for the case that none are needed.

Example:

```
(define book-class%
  (class object%
    (field (pages 5))
    (define/public (letters)
      (* pages 500))
    (super-new)))
```

this

Within a `class*` form, `this` refers to the current object (i.e., the object being initialized or whose method was called). Use outside the body of a `class*` form is a syntax error.

Examples:

```
(define (describe obj)
  (printf "Hello ~a\n" obj))

(define table%
  (class object%
    (define/public (describe-self)
      (describe this))
    (super-new)))

> (send (new table%) describe-self)
Hello #(struct:object:table% ...)
```

this%

Within a `class*` form, `this%` refers to the class of the current object (i.e., the object being initialized or whose method was called). Use outside the body of a `class*` form is a syntax error.

Examples:

```
(define account%
  (class object%
    (super-new)
    (init-field balance)
    (define/public (add n)
      (new this% [balance (+ n balance)]))))
```

```
(define savings%
  (class account%
    (super-new)
    (inherit-field balance)
    (define interest 0.04)
    (define/public (add-interest)
      (send this add (* interest balance)))))

> (let* ([acct (new savings% [balance 500])]
         [acct (send acct add 500)]
         [acct (send acct add-interest)])
    (printf "Current balance: ~a\n" (get-field balance acct)))
Current balance: 1040.0
```

(inspect *inspector-expr*)

See `class*`; use outside the body of a `class*` form is a syntax error.

(init *init-decl* ...)

See `class*` and §5.2.1 “Initialization Variables”; use outside the body of a `class*` form is a syntax error.

Example:

```
> (class object%
    (super-new)
    (init turnip
          [(internal-potato potato)]
          [carrot 'good]
          [(internal-rutabaga rutabaga) 'okay]))
#<class:eval:10:0>
```

(init-field *init-decl* ...)

See `class*`, §5.2.1 “Initialization Variables”, and §5.2.2 “Fields”; use outside the body of a `class*` form is a syntax error.

Example:

```
> (class object%
    (super-new)
    (init-field turkey
```

```
                [(internal-ostrich ostrich)]
                [chicken 7]
                [(internal-emu emu) 13]))
#<class:eval:11:0>
```

```
(field field-decl ...)
```

See `class*` and §5.2.2 “Fields”; use outside the body of a `class*` form is a syntax error.

Example:

```
> (class object%
    (super-new)
    (field [minestrone 'ready]
           [(internal-coq-au-vin coq-au-vin) 'stewing]))
#<class:eval:12:0>
```

```
(inherit-field maybe-renamed ...)
```

See `class*` and §5.2.2 “Fields”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define cookbook%
  (class object%
    (super-new)
    (field [recipes '(caldo-verde oyakodon eggs-benedict)]
           [pages 389])))

> (class cookbook%
    (super-new)
    (inherit-field recipes
                   [internal-pages pages]))
#<class:eval:14:0>
```

```
(init-rest id)
(init-rest)
```

See `class*` and §5.2.1 “Initialization Variables”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define fruit-basket%
```

```
  (class object%
    (super-new)
    (init-rest fruits)
    (displayln fruits)))

> (make-object fruit-basket% 'kiwi 'lychee 'melon)
(kiwi lychee melon)
(object:fruit-basket% ...)
```

```
(public maybe-renamed ...)
```

See `class*` and §5.2.3.1 “Method Definitions”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define jumper%
  (class object%
    (super-new)
    (define (skip) 'skip)
    (define (hop) 'hop)
    (public skip [hop jump])))

> (send (new jumper%) skip)
'skip
> (send (new jumper%) jump)
'hop
```

```
(pubment maybe-renamed ...)
```

See `class*` and §5.2.3.1 “Method Definitions”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define runner%
  (class object%
    (super-new)
    (define (run) 'run)
    (define (trot) 'trot)
    (pubment run [trot jog])))

> (send (new runner%) run)
'run
> (send (new runner%) jog)
'trot
```

(public-final *maybe-renamed* ...)

See `class*` and §5.2.3.1 “Method Definitions”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define point%
  (class object%
    (super-new)
    (init-field [x 0] [y 0])
    (define (get-x) x)
    (define (do-get-y) y)
    (public-final get-x [do-get-y get-y])))

> (send (new point% [x 1] [y 3]) get-y)
3
> (class point%
    (super-new)
    (define (get-x) 3.14)
    (override get-x))
class*: cannot override or augment final method
  method name: get-x
  class name: eval:25:0
```

(override *maybe-renamed* ...)

See `class*` and §5.2.3.1 “Method Definitions”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define sheep%
  (class object%
    (super-new)
    (define/public (bleat)
      (displayln "baaaaaaaaah"))))

(define confused-sheep%
  (class sheep%
    (super-new)
    (define (bleat)
      (super bleat)
      (displayln "???"))
    (override bleat)))
```

```
> (send (new sheep%) bleat)
baaaaaaaaah

> (send (new confused-sheep%) bleat)
baaaaaaaaah
???
```

```
(overment maybe-renamed ...)
```

See `class*` and §5.2.3.1 “Method Definitions”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define turkey%
  (class object%
    (super-new)
    (define/public (gobble)
      (displayln "gobble gobble"))))

(define extra-turkey%
  (class turkey%
    (super-new)
    (define (gobble)
      (super gobble)
      (displayln "gobble gobble gobble")
      (inner (void) gobble))
    (overment gobble)))

(define cyborg-turkey%
  (class extra-turkey%
    (super-new)
    (define/augment (gobble)
      (displayln "1100111110111111000101100010110110011001 01"))))

> (send (new extra-turkey%) gobble)
gobble gobble
gobble gobble gobble

> (send (new cyborg-turkey%) gobble)
gobble gobble
gobble gobble gobble
1100111110111111000101100010110110011001 01
```

(override-final *maybe-renamed* ...)

See `class*` and §5.2.3.1 “Method Definitions”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define meeper%
  (class object%
    (super-new)
    (define/public (meep)
      (displayln "meep"))))

(define final-meeper%
  (class meeper%
    (super-new)
    (define (meep)
      (super meep)
      (displayln "This meeping ends with me"))
    (override-final meep)))

> (send (new meeper%) meep)
meep

> (send (new final-meeper%) meep)
meep
This meeping ends with me
```

(augment *maybe-renamed* ...)

See `class*` and §5.2.3.1 “Method Definitions”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define buzzer%
  (class object%
    (super-new)
    (define/pubment (buzz)
      (displayln "bzzzt")
      (inner (void) buzz))))

(define loud-buzzer%
```

```
  (class buzzer%
    (super-new)
    (define (buzz)
      (displayln "BZZZZZZZZZT"))
    (augment buzz)))

> (send (new buzzer%) buzz)
bzzzt

> (send (new loud-buzzer%) buzz)
bzzzt
BZZZZZZZZZT
```

(augride *maybe-renamed* ...)

See `class*` and §5.2.3.1 “Method Definitions”; use outside the body of a `class*` form is a syntax error.

(augment-final *maybe-renamed* ...)

See `class*` and §5.2.3.1 “Method Definitions”; use outside the body of a `class*` form is a syntax error.

(private *id* ...)

See `class*` and §5.2.3.1 “Method Definitions”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define light%
  (class object%
    (super-new)
    (define on? #t)
    (define (toggle) (set! on? (not on?)))
    (private toggle)
    (define (flick) (toggle))
    (public flick)))

> (send (new light%) toggle)
send: no such method
  method name: toggle
  class name: light%
> (send (new light%) flick)
```

```
(abstract id ...)
```

See `class*` and §5.2.3.1 “Method Definitions”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define train%
  (class object%
    (super-new)
    (abstract get-speed)
    (init-field [position 0])
    (define/public (move)
      (new this% [position (+ position (get-speed))]))))

(define acela%
  (class train%
    (super-new)
    (define/override (get-speed) 241)))

(define talgo-350%
  (class train%
    (super-new)
    (define/override (get-speed) 330)))

> (new train%)
instantiate: cannot instantiate class with abstract methods
  class: #<class:train%>
  abstract methods:
   get-speed
> (send (new acela%) move)
(object:acela% ...)
```

```
(inherit maybe-renamed ...)
```

See `class*` and §5.2.3.2 “Inherited and Superclass Methods”; use outside the body of a `class*` form is a syntax error.

Examples:

```
(define alarm%
  (class object%
    (super-new)
    (define/public (alarm)
      (displayln "beeeeeeeep"))))
```

```
(define car-alarm%
  (class alarm%
    (super-new)
    (init-field proximity)
    (inherit alarm)
    (when (< proximity 10)
      (alarm))))

> (new car-alarm% [proximity 5])
beeeeeeeep
(object:car-alarm% ...)
```

(inherit/super *maybe-renamed* ...)

See `class*` and §5.2.3.2 “Inherited and Superclass Methods”; use outside the body of a `class*` form is a syntax error.

(inherit/inner *maybe-renamed* ...)

See `class*` and §5.2.3.2 “Inherited and Superclass Methods”; use outside the body of a `class*` form is a syntax error.

(rename-super *renamed* ...)

See `class*` and §5.2.3.2 “Inherited and Superclass Methods”; use outside the body of a `class*` form is a syntax error.

(rename-inner *renamed* ...)

See `class*` and §5.2.3.2 “Inherited and Superclass Methods”; use outside the body of a `class*` form is a syntax error.

(public* (*id expr*) ...)

Shorthand for `(begin (public id) ... (define id expr) ...)`.

(pubment* (*id expr*) ...)

Shorthand for `(begin (pubment id) ... (define id expr) ...)`.

```
(public-final* (id expr) ...)
```

Shorthand for `(begin (public-final id) ... (define id expr) ...)`.

```
(override* (id expr) ...)
```

Shorthand for `(begin (override id) ... (define id expr) ...)`.

```
(overment* (id expr) ...)
```

Shorthand for `(begin (overment id) ... (define id expr) ...)`.

```
(override-final* (id expr) ...)
```

Shorthand for `(begin (override-final id) ... (define id expr) ...)`.

```
(augment* (id expr) ...)
```

Shorthand for `(begin (augment id) ... (define id expr) ...)`.

```
(augride* (id expr) ...)
```

Shorthand for `(begin (augride id) ... (define id expr) ...)`.

```
(augment-final* (id expr) ...)
```

Shorthand for `(begin (augment-final id) ... (define id expr) ...)`.

```
(private* (id expr) ...)
```

Shorthand for `(begin (private id) ... (define id expr) ...)`.

```
(define/public id expr)
(define/public (id . formals) body ...+)
```

Shorthand for `(begin (public id) (define id expr))` or `(begin (public id) (define (id . formals) body ...+))`

```
(define/pubment id expr)
(define/pubment (id . formals) body ...+)
```

Shorthand for (begin (pubment *id*) (define *id expr*)) or (begin (pubment *id*) (define (*id* . *formals*) *body* ...+))

```
(define/public-final id expr)
(define/public-final (id . formals) body ...+)
```

Shorthand for (begin (public-final *id*) (define *id expr*)) or (begin (public-final *id*) (define (*id* . *formals*) *body* ...+))

```
(define/override id expr)
(define/override (id . formals) body ...+)
```

Shorthand for (begin (override *id*) (define *id expr*)) or (begin (override *id*) (define (*id* . *formals*) *body* ...+))

```
(define/overment id expr)
(define/overment (id . formals) body ...+)
```

Shorthand for (begin (overment *id*) (define *id expr*)) or (begin (overment *id*) (define (*id* . *formals*) *body* ...+))

```
(define/override-final id expr)
(define/override-final (id . formals) body ...+)
```

Shorthand for (begin (override-final *id*) (define *id expr*)) or (begin (override-final *id*) (define (*id* . *formals*) *body* ...+))

```
(define/augment id expr)
(define/augment (id . formals) body ...+)
```

Shorthand for (begin (augment *id*) (define *id expr*)) or (begin (augment *id*) (define (*id* . *formals*) *body* ...+))

```
(define/augride id expr)
(define/augride (id . formals) body ...+)
```

Shorthand for (begin (augride id) (define id expr)) or (begin (augride id) (define (id . formals) body ...+))

```
(define/augment-final id expr)
(define/augment-final (id . formals) body ...+)
```

Shorthand for (begin (augment-final id) (define id expr)) or (begin (augment-final id) (define (id . formals) body ...+))

```
(define/private id expr)
(define/private (id . formals) body ...+)
```

Shorthand for (begin (private id) (define id expr)) or (begin (private id) (define (id . formals) body ...+))

```
(class/derived original-datum
  (name-id super-expr (interface-expr ...) deserialize-id-expr)
  class-clause
  ...)
```

Like class*, but includes a sub-expression to be used as the source for all syntax errors within the class definition. For example, define-serializable-class expands to class/derived so that errors in the body of the class are reported in terms of define-serializable-class instead of class.

The *original-datum* is the original expression to use for reporting errors.

The *name-id* is used to name the resulting class; if it is #f, the class name is inferred.

The *super-expr*, *interface-expr*s, and *class-clause*s are as for class*.

If the *deserialize-id-expr* is not literally #f, then a serializable class is generated, and the result is two values instead of one: the class and a deserialize-info structure produced by make-deserialize-info. The *deserialize-id-expr* should produce a value suitable as the second argument to make-serialize-info, and it should refer to an export whose value is the deserialize-info structure.

Future optional forms may be added to the sequence that currently ends with *deserialize-id-expr*.

5.2.1 Initialization Variables

A class's initialization variables, declared with init, init-field, and init-rest, are instantiated for each object of a class. Initialization variables can be used in the initial

value expressions of fields, default value expressions for initialization arguments, and in initialization expressions. Only initialization variables declared with `init-field` can be accessed from methods; accessing any other initialization variable from a method is a syntax error.

The values bound to initialization variables are

- the arguments provided with `instantiate` or passed to `make-object`, if the object is created as a direct instance of the class; or,
- the arguments passed to the superclass initialization form or procedure, if the object is created as an instance of a derived class.

If an initialization argument is not provided for an initialization variable that has an associated *default-value-expr*, then the *default-value-expr* expression is evaluated to obtain a value for the variable. A *default-value-expr* is only evaluated when an argument is not provided for its variable. The environment of *default-value-expr* includes all of the initialization variables, all of the fields, and all of the methods of the class. If multiple *default-value-expr*s are evaluated, they are evaluated from left to right. Object creation and field initialization are described in detail in §5.3 "Creating Objects".

If an initialization variable has no *default-value-expr*, then the object creation or superclass initialization call must supply an argument for the variable, otherwise the `exn:fail:object` exception is raised.

Initialization arguments can be provided by name or by position. The external name of an initialization variable can be used with `instantiate` or with the superclass initialization form. Those forms also accept by-position arguments. The `make-object` procedure and the superclass initialization procedure accept only by-position arguments.

Arguments provided by position are converted into by-name arguments using the order of `init` and `init-field` clauses and the order of variables within each clause. When an `instantiate` form provides both by-position and by-name arguments, the converted arguments are placed before by-name arguments. (The order can be significant; see also §5.3 "Creating Objects".)

Unless a class contains an `init-rest` clause, when the number of by-position arguments exceeds the number of declared initialization variables, the order of variables in the superclass (and so on, up the superclass chain) determines the by-name conversion.

If a class expression contains an `init-rest` clause, there must be only one, and it must be last. If it declares a variable, then the variable receives extra by-position initialization arguments as a list (similar to a dotted "rest argument" in a procedure). An `init-rest` variable can receive by-position initialization arguments that are left over from a by-name conversion for a derived class. When a derived class's superclass initialization provides even more by-position arguments, they are prefixed onto the by-position arguments accumulated so far.

If too few or too many by-position initialization arguments are provided to an object creation or superclass initialization, then the exn:fail:object exception is raised. Similarly, if extra by-position arguments are provided to a class with an init-rest clause, the exn:fail:object exception is raised.

Unused (by-name) arguments are to be propagated to the superclass, as described in §5.3 "Creating Objects". Multiple initialization arguments can use the same name if the class derivation contains multiple declarations (in different classes) of initialization variables with the name. See §5.3 "Creating Objects" for further details.

See also §5.2.3.3 "Internal and External Names" for information about internal and external names.

5.2.2 Fields

Each field, init-field, and non-method define-values clause in a class declares one or more new fields for the class. Fields declared with field or init-field are public. Public fields can be accessed and mutated by subclasses using inherit-field. Public fields are also accessible outside the class via class-field-accessor and mutable via class-field-mutator (see §5.4 "Field and Method Access"). Fields declared with define-values are accessible only within the class.

A field declared with init-field is both a public field and an initialization variable. See §5.2.1 "Initialization Variables" for information about initialization variables.

An inherit-field declaration makes a public field defined by a superclass directly accessible in the class expression. If the indicated field is not defined in the superclass, the exn:fail:object exception is raised when the class expression is evaluated. Every field in a superclass is present in a derived class, even if it is not declared with inherit-field in the derived class. The inherit-field clause does not control inheritance, but merely controls lexical scope within a class expression.

When an object is first created, all of its fields have the #<undefined> value (see §3.18 "Void and Undefined"). The fields of a class are initialized at the same time that the class's initialization expressions are evaluated; see §5.3 "Creating Objects" for more information.

See also §5.2.3.3 "Internal and External Names" for information about internal and external names.

5.2.3 Methods

Method Definitions

Each `public`, `override`, `augment`, `pubment`, `overment`, `augride`, `public-final`, `override-final`, `augment-final`, and `private` clause in a class declares one or more method names. Each method name must have a corresponding *`method-definition`*. The order of `public`, etc., clauses and their corresponding definitions (among themselves, and with respect to other clauses in the class) does not matter.

As shown in the grammar for `class*`, a method definition is syntactically restricted to certain procedure forms, as defined by the grammar for *`method-procedure`*; in the last two forms of *`method-procedure`*, the body `id` must be one of the `id`s bound by `let-values` or `letrec-values`. A *`method-procedure`* expression is not evaluated directly. Instead, for each method, a class-specific method procedure is created; it takes an initial object argument, in addition to the arguments the procedure would accept if the *`method-procedure`* expression were evaluated directly. The body of the procedure is transformed to access methods and fields through the object argument.

A method declared with `public`, `pubment`, or `public-final` introduces a new method into a class. The method must not be present already in the superclass, otherwise the `exn:fail:object` exception is raised when the class expression is evaluated. A method declared with `public` can be overridden in a subclass that uses `override`, `overment`, or `override-final`. A method declared with `pubment` can be augmented in a subclass that uses `augment`, `augride`, or `augment-final`. A method declared with `public-final` cannot be overridden or augmented in a subclass.

A method declared with `override`, `overment`, or `override-final` overrides a definition already present in the superclass. If the method is not already present, the `exn:fail:object` exception is raised when the class expression is evaluated. A method declared with `override` can be overridden again in a subclass that uses `override`, `overment`, or `override-final`. A method declared with `overment` can be augmented in a subclass that uses `augment`, `augride`, or `augment-final`. A method declared with `override-final` cannot be overridden further or augmented in a subclass.

A method declared with `augment`, `augride`, or `augment-final` augments a definition already present in the superclass. If the method is not already present, the `exn:fail:object` exception is raised when the class expression is evaluated. A method declared with `augment` can be augmented further in a subclass that uses `augment`, `augride`, or `augment-final`. A method declared with `augride` can be overridden in a subclass that uses `override`, `overment`, or `override-final`. (Such an override merely replaces the augmentation, not the method that is augmented.) A method declared with `augment-final` cannot be overridden or augmented further in a subclass.

A method declared with `private` is not accessible outside the class expression, cannot be overridden, and never overrides a method in the superclass.

When a method is declared with `override`, `overment`, or `override-final`, then the superclass implementation of the method can be called using `super` form.

When a method is declared with `pubment`, `augment`, or `overment`, then a subclass augmenting method can be called using the `inner` form. The only difference between `public-final` and `pubment` without a corresponding `inner` is that `public-final` prevents the declaration of augmenting methods that would be ignored.

A method declared with `abstract` must be declared without an implementation. Subclasses may implement abstract methods via the `override`, `overment`, or `override-final` forms. Any class that contains or inherits any abstract methods is considered abstract and cannot be instantiated.

```
(super id arg ...)
(super id arg ... . arg-list-expr)
```

Always accesses the superclass method, independent of whether the method is overridden again in subclasses. Using the `super` form outside of `class*` is a syntax error. Each *arg* is as for `#%app`: either *arg-expr* or *keyword arg-expr*.

The second form is analogous to using `apply` with a procedure; the *arg-list-expr* must not be a parenthesized expression.

```
(inner default-expr id arg ...)
(inner default-expr id arg ... . arg-list-expr)
```

If the object's class does not supply an augmenting method, then *default-expr* is evaluated, and the *arg* expressions are not evaluated. Otherwise, the augmenting method is called with the *arg* results as arguments, and *default-expr* is not evaluated. If no `inner` call is evaluated for a particular method, then augmenting methods supplied by subclasses are never used. Using the `inner` form outside of `class*` is an syntax error.

The second form is analogous to using `apply` with a procedure; the *arg-list-expr* must not be a parenthesized expression.

Inherited and Superclass Methods

Each `inherit`, `inherit/super`, `inherit/inner`, `rename-super`, and `rename-inner` clause declares one or more methods that are defined in the class, but must be present in the superclass. The `rename-super` and `rename-inner` declarations are rarely used, since `inherit/super` and `inherit/inner` provide the same access. Also, superclass and augmenting methods are typically accessed through `super` and `inner` in a class that also declares the methods, instead of through `inherit/super`, `inherit/inner`, `rename-super`, or `rename-inner`.

Method names declared with `inherit`, `inherit/super`, or `inherit/inner` access overriding declarations, if any, at run time. Method names declared with `inherit/super` can also be used with the `super` form to access the superclass implementation, and method names declared with `inherit/inner` can also be used with the `inner` form to access an augmenting method, if any.

Method names declared with `rename-super` always access the superclass's implementation at run-time. Methods declared with `rename-inner` access a subclass's augmenting method, if any, and must be called with the form

```
(id (lambda () default-expr) arg ...)
```

so that a `default-expr` is available to evaluate when no augmenting method is available. In such a form, `lambda` is a literal identifier to separate the `default-expr` from the `arg`. When an augmenting method is available, it receives the results of the `arg` expressions as arguments.

Methods that are present in the superclass but not declared with `inherit`, `inherit/super`, or `inherit/inner` or `rename-super` are not directly accessible in the class (though they can be called with `send`). Every public method in a superclass is present in a derived class, even if it is not declared with `inherit` in the derived class; the `inherit` clause does not control inheritance, but merely controls lexical scope within a class expression.

If a method declared with `inherit`, `inherit/super`, `inherit/inner`, `rename-super`, or `rename-inner` is not present in the superclass, the `exn:fail:object` exception is raised when the class expression is evaluated.

Internal and External Names

Each method declared with `public`, `override`, `augment`, `pubment`, `overment`, `augride`, `public-final`, `override-final`, `augment-final`, `inherit`, `inherit/super`, `inherit/inner`, `rename-super`, and `rename-inner` can have separate internal and external names when `(internal-id external-id)` is used for declaring the method. The internal name is used to access the method directly within the class expression (including within `super` or `inner` forms), while the external name is used with `send` and `generic` (see §5.4 "Field and Method Access"). If a single `id` is provided for a method declaration, the identifier is used for both the internal and external names.

Method inheritance, overriding, and augmentation are based on external names only. Separate internal and external names are required for `rename-super` and `rename-inner` (for historical reasons, mainly).

Each `init`, `init-field`, `field`, or `inherit-field` variable similarly has an internal and an external name. The internal name is used within the class to access the variable, while the external name is used outside the class when providing initialization arguments (e.g., to

instantiate), inheriting a field, or accessing a field externally (e.g., with class-field-accessor). As for methods, when inheriting a field with inherit-field, the external name is matched to an external field name in the superclass, while the internal name is bound in the class expression.

A single identifier can be used as an internal identifier and an external identifier, and it is possible to use the same identifier as internal and external identifiers for different bindings. Furthermore, within a single class, a single name can be used as an external method name, an external field name, and an external initialization argument name. Overall, each internal identifier must be distinct from all other internal identifiers, each external method name must be distinct from all other method names, each external field name must be distinct from all other field names, and each initialization argument name must be distinct from all other initialization argument names.

By default, external names have no lexical scope, which means, for example, that an external method name matches the same syntactic symbol in all uses of send. The define-local-member-name and define-member-name forms introduce scoped external names.

When a class expression is compiled, identifiers used in place of external names must be symbolically distinct (when the corresponding external names are required to be distinct), otherwise a syntax error is reported. When no external name is bound by define-member-name, then the actual external names are guaranteed to be distinct when class expression is evaluated. When any external name is bound by define-member-name, the exn:fail:object exception is raised by class if the actual external names are not distinct.

(define-local-member-name *id* ...)

Unless it appears as the top-level definition, binds each *id* so that, within the scope of the definition, each use of each *id* as an external name is resolved to a hidden name generated by the define-local-member-name declaration. Thus, methods, fields, and initialization arguments declared with such external-name *id*s are accessible only in the scope of the define-local-member-name declaration. As a top-level definition, define-local-member-name binds *id* to its symbolic form.

The binding introduced by define-local-member-name is a syntax binding that can be exported and imported with modules. Each evaluation of a define-local-member-name declaration generates a distinct hidden name (except as a top-level definition). The interface->method-names procedure does not expose hidden names.

Examples:

```
(define-values (r o)
  (let ()
    (define-local-member-name m)
    (define c% (class object%
                 (define/public (m) 10)
                 (super-new)))
```

```
    (define o (new c%))

    (values (send o m)
            o)))

> r
10
> (send o m)
send: no such method
  method name: m
  class name: c%
```

(define-member-name *id key-expr*)

Maps a single external name to an external name that is determined by an expression. The value of *key-expr* must be the result of either a member-name-key expression or a generate-member-key call.

(member-name-key *identifier*)

Produces a representation of the external name for id in the environment of the member-name-key expression.

(generate-member-key) → member-name-key?

Produces a hidden name, just like the binding for define-local-member-name.

(member-name-key? *v*) → boolean?
  v : any/c

Returns #t for values produced by member-name-key and generate-member-key, #f otherwise.

(member-name-key=? *a-key b-key*) → boolean?
  a-key : member-name-key?
  b-key : member-name-key?

Produces #t if member-name keys *a-key* and *b-key* represent the same external name, #f otherwise.

(member-name-key-hash-code *a-key*) → integer?
  a-key : member-name-key?

Produces an integer hash code consistent with member-name-key=? comparisons, analogous to equal-hash-code.

Examples:

```
(define (make-c% key)
  (define-member-name m key)
  (class object%
    (define/public (m) 10)
    (super-new)))

> (send (new (make-c% (member-name-key m))) m)
10
> (send (new (make-c% (member-name-key p))) m)
send: no such method
  method name: m
  class name: eval:57:0
> (send (new (make-c% (member-name-key p))) p)
10

(define (fresh-c%)
  (let ([key (generate-member-key)])
    (values (make-c% key) key)))

(define-values (fc% key) (fresh-c%))

> (send (new fc%) m)
send: no such method
  method name: m
  class name: eval:57:0
> (let ()
    (define-member-name p key)
    (send (new fc%) p))
10
```

5.3 Creating Objects

The `make-object` procedure creates a new object with by-position initialization arguments, the `new` form creates a new object with by-name initialization arguments, and the `instantiate` form creates a new object with both by-position and by-name initialization arguments.

All fields in the newly created object are initially bound to the special `#<undefined>` value (see §3.18 “Void and Undefined”). Initialization variables with default value expressions (and no provided value) are also initialized to `#<undefined>`. After argument values are assigned to initialization variables, expressions in `field` clauses, `init-field` clauses with no provided argument, `init` clauses with no provided argument, private field definitions, and other expressions are evaluated. Those expressions are evaluated as they appear in the class expression, from left to right.

Sometime during the evaluation of the expressions, superclass-declared initializations must be evaluated once by using the `super-make-object` procedure, `super-new` form, or `super-instantiate` form.

By-name initialization arguments to a class that have no matching initialization variable are implicitly added as by-name arguments to a `super-make-object`, `super-new`, or `super-instantiate` invocation, after the explicit arguments. If multiple initialization arguments are provided for the same name, the first (if any) is used, and the unused arguments are propagated to the superclass. (Note that converted by-position arguments are always placed before explicit by-name arguments.) The initialization procedure for the `object%` class accepts zero initialization arguments; if it receives any by-name initialization arguments, then `exn:fail:object` exception is raised.

If the end of initialization is reached for any class in the hierarchy without invoking the superclass's initialization, the `exn:fail:object` exception is raised. Also, if superclass initialization is invoked more than once, the `exn:fail:object` exception is raised.

Fields inherited from a superclass are not initialized until the superclass's initialization procedure is invoked. In contrast, all methods are available for an object as soon as the object is created; the overriding of methods is not affected by initialization (unlike objects in C++).

```
(make-object class init-v ...) → object?
  class : class?
  init-v : any/c
```

Creates an instance of *class*. The *init-v*s are passed as initialization arguments, bound to the initialization variables of *class* for the newly created object as described in §5.2.1 "Initialization Variables". If *class* is not a class, the `exn:fail:contract` exception is raised.

```
(new class-expr (id by-name-expr) ...)
```

Creates an instance of the value of *class-expr* (which must be a class), and the value of each *by-name-expr* is provided as a by-name argument for the corresponding *id*.

```
(instantiate class-expr (by-pos-expr ...) (id by-name-expr) ...)
```

Creates an instance of the value of *class-expr* (which must be a class), and the values of the *by-pos-expr*s are provided as by-position initialization arguments. In addition, the value of each *by-name-expr* is provided as a by-name argument for the corresponding *id*.

```
super-make-object
```

Produces a procedure that takes by-position arguments an invokes superclass initialization. See §5.3 "Creating Objects" for more information.

```
(super-instantiate (by-pos-expr ...) (id by-expr ...) ...)
```

Invokes superclass initialization with the specified by-position and by-name arguments. See §5.3 “Creating Objects” for more information.

```
(super-new (id by-name-expr ...) ...)
```

Invokes superclass initialization with the specified by-name arguments. See §5.3 “Creating Objects” for more information.

5.4 Field and Method Access

In expressions within a class definition, the initialization variables, fields, and methods of the class are all part of the environment. Within a method body, only the fields and other methods of the class can be referenced; a reference to any other class-introduced identifier is a syntax error. Elsewhere within the class, all class-introduced identifiers are available, and fields and initialization variables can be mutated with `set!`.

5.4.1 Methods

Method names used within a class can only be used in the procedure position of an application expression; any other use is a syntax error.

To allow methods to be applied to lists of arguments, a method application can have the following form:

```
(method-id arg ... . arg-list-expr)
```

This form calls the method in a way analogous to `(apply method-id arg ... arg-list-expr)`. The `arg-list-expr` must not be a parenthesized expression.

Methods are called from outside a class with the `send`, `send/apply`, and `send/keyword-apply` forms.

```
(send obj-expr method-id arg ...)
(send obj-expr method-id arg ... . arg-list-expr)
```

Evaluates *obj-expr* to obtain an object, and calls the method with (external) name *method-id* on the object, providing the *arg* results as arguments. Each *arg* is as for `#%app`: either *arg-expr* or *keyword arg-expr*. In the second form, *arg-list-expr* cannot be a parenthesized expression.

If *obj-expr* does not produce an object, the exn:fail:contract exception is raised. If the object has no public method named *method-id*, the exn:fail:object exception is raised.

```
(send/apply obj-expr method-id arg ... arg-list-expr)
```

Like the dotted form of send, but *arg-list-expr* can be any expression.

```
(send/keyword-apply obj-expr method-id
                    keyword-list-expr value-list-expr
                    arg ... arg-list-expr)
```

Like send/apply, but with expressions for keyword and argument lists like keyword-apply.

```
(dynamic-send obj
              method-name
              v ...
              #:<kw> kw-arg ...) → any
  obj : object?
  method-name : symbol?
  v : any/c
  kw-arg : any/c
```

Calls the method on *obj* whose name matches *method-name*, passing along all given *v*s and *kw-arg*s.

```
(send* obj-expr msg ...+)

msg = (method-id arg ...)
    | (method-id arg ... . arg-list-expr)
```

Calls multiple methods (in order) of the same object. Each *msg* corresponds to a use of send.

For example,

```
(send* edit (begin-edit-sequence)
            (insert "Hello")
            (insert #\newline)
            (end-edit-sequence))
```

is the same as

```
(let ([o edit])
```

```
  (send o begin-edit-sequence)
  (send o insert "Hello")
  (send o insert #\newline)
  (send o end-edit-sequence))
```

```
(send+ obj-expr msg ...)

msg = (method-id arg ...)
    | (method-id arg ... . arg-list-expr)
```

Calls methods (in order) starting with the object produced by *obj-expr*. Each method call will be invoked on the result of the last method call, which is expected to be an object. Each *msg* corresponds to a use of `send`.

This is the functional analogue of `send*`.

Examples:

```
(define point%
  (class object%
    (super-new)
    (init-field [x 0] [y 0])
    (define/public (move-x dx)
      (new this% [x (+ x dx)]))
    (define/public (move-y dy)
      (new this% [y (+ y dy)]))))

> (send+ (new point%)
         (move-x 5)
         (move-y 7)
         (move-x 12))
(object:point% ...)
```

```
(with-method ((id (obj-expr method-id)) ...)
  body ...+)
```

Extracts methods from an object and binds a local name that can be applied directly (in the same way as declared methods within a class) for each method. Each *obj-expr* must produce an object, which must have a public method named by the corresponding *method-id*. The corresponding *id* is bound so that it can be applied directly (see §5.4.1 “Methods”).

Example:

```
(let ([s (new stack%)])
  (with-method ([push (s push!)]
                [pop (s pop!)])
```

```
      (push 10)
      (push 9)
      (pop)))
```

is the same as

```
(let ([s (new stack%)])
  (send s push! 10)
  (send s push! 9)
  (send s pop!))
```

5.4.2 Fields

```
(get-field id obj-expr)
```

Extracts the field with (external) name *id* from the value of *obj-expr*.

If *obj-expr* does not produce an object, the exn:fail:contract exception is raised. If the object has no *id* field, the exn:fail:object exception is raised.

```
(set-field! id obj-expr expr)
```

Sets the field with (external) name *id* from the value of *obj-expr* to the value of *expr*.

If *obj-expr* does not produce an object, the exn:fail:contract exception is raised. If the object has no *id* method, the exn:fail:object exception is raised.

```
(field-bound? id obj-expr)
```

Produces #t if the object result of *obj-expr* has a field with (external) name *id*, #f otherwise.

If *obj-expr* does not produce an object, the exn:fail:contract exception is raised.

```
(class-field-accessor class-expr field-id)
```

Returns an accessor procedure that takes an instance of the class produced by *class-expr* and returns the value of the object's field with (external) name *field-id*.

If *class-expr* does not produce a class, the exn:fail:contract exception is raised. If the class has no *field-id* field, the exn:fail:object exception is raised.

```
(class-field-mutator class-expr field-id)
```

Returns a mutator procedure that takes an instance of the class produced by *class-expr* and a value, and sets the value of the object's field with (external) name *field-id* to the given value. The result is `#<void>`.

If *class-expr* does not produce a class, the `exn:fail:contract` exception is raised. If the class has no *field-id* field, the `exn:fail:object` exception is raised.

5.4.3 Generics

A *generic* can be used instead of a method name to avoid the cost of relocating a method by name within a class.

```
(generic class-or-interface-expr id)
```

Produces a generic that works on instances of the class or interface produced by *class-or-interface-expr* (or an instance of a class/interface derived from `class-or-interface`) to call the method with (external) name *id*.

If *class-or-interface-expr* does not produce a class or interface, the `exn:fail:contract` exception is raised. If the resulting class or interface does not contain a method named *id*, the `exn:fail:object` exception is raised.

```
(send-generic obj-expr generic-expr arg ...)
(send-generic obj-expr generic-expr arg ... . arg-list-expr)
```

Calls a method of the object produced by *obj-expr* as indicated by the generic produced by *generic-expr*. Each *arg* is as for `#%app`: either *arg-expr* or *keyword arg-expr*. The second form is analogous to calling a procedure with `apply`, where *arg-list-expr* is not a parenthesized expression.

If *obj-expr* does not produce a object, or if *generic-expr* does not produce a generic, the `exn:fail:contract` exception is raised. If the result of *obj-expr* is not an instance of the class or interface encapsulated by the result of *generic-expr*, the `exn:fail:object` exception is raised.

```
(make-generic type method-name) → generic?
  type : (or/c class? interface?)
  method-name : symbol?
```

Like the `generic` form, but as a procedure that accepts a symbolic method name.

5.5 Mixins

```
(mixin (interface-expr ...) (interface-expr ...)
  class-clause ...)
```

Produces a *mixin*, which is a procedure that encapsulates a class extension, leaving the superclass unspecified. Each time that a mixin is applied to a specific superclass, it produces a new derived class using the encapsulated extension.

The given class must implement interfaces produced by the first set of *interface-expr*s. The result of the procedure is a subclass of the given class that implements the interfaces produced by the second set of *interface-expr*s. The *class-clause*s are as for `class*`, to define the class extension encapsulated by the mixin.

Evaluation of a `mixin` form checks that the *class-clause*s are consistent with both sets of *interface-expr*s.

5.6 Traits

```
(require racket/trait)
```

The bindings documented in this section are provided by the `racket/trait` library, not `racket/base` or `racket`.

A *trait* is a collection of methods that can be converted to a mixin and then applied to a class. Before a trait is converted to a mixin, the methods of a trait can be individually renamed, and multiple traits can be merged to form a new trait.

```
(trait trait-clause ...)

trait-clause = (public maybe-renamed ...)
             | (pubment maybe-renamed ...)
             | (public-final maybe-renamed ...)
             | (override maybe-renamed ...)
             | (overment maybe-renamed ...)
             | (override-final maybe-renamed ...)
             | (augment maybe-renamed ...)
             | (augride maybe-renamed ...)
             | (augment-final maybe-renamed ...)
             | (inherit maybe-renamed ...)
             | (inherit/super maybe-renamed ...)
             | (inherit/inner maybe-renamed ...)
             | method-definition
             | (field field-declaration ...)
             | (inherit-field maybe-renamed ...)
```

Creates a trait. The body of a `trait` form is similar to the body of a `class*` form, but restricted to non-private method definitions. In particular, the grammar of *maybe-renamed*, *method-definition*, and *field-declaration* are the same as for `class*`, and every *method-definition* must have a corresponding declaration (one of `public`, `over-`

ride, etc.). As in class, uses of method names in direct calls, super calls, and inner calls depend on bringing method names into scope via inherit, inherit/super, inherit/inner, and other method declarations in the same trait; an exception, compared to class is that overment binds a method name only in the corresponding method, and not in other methods of the same trait. Finally, macros such as public* and define/public work in trait as in class.

External identifiers in trait, trait-exclude, trait-exclude-field, trait-alias, trait-rename, and trait-rename-field forms are subject to binding via define-member-name and define-local-member-name. Although private methods or fields are not allowed in a trait form, they can be simulated by using a public or field declaration and a name whose scope is limited to the trait form.

```
(trait? v) → boolean?
  v : any/c
```

Returns #t if v is a trait, #f otherwise.

```
(trait->mixin tr) → (class? . -> . class?)
  tr : trait?
```

Converts a trait to a mixin, which can be applied to a class to produce a new class. An expression of the form

```
(trait->mixin
 (trait
   trait-clause ...))
```

is equivalent to

```
(lambda (%)
  (class %
    trait-clause ...
    (super-new)))
```

Normally, however, a trait's methods are changed and combined with other traits before converting to a mixin.

```
(trait-sum tr ...+) → trait?
  tr : trait?
```

Produces a trait that combines all of the methods of the given trs. For example,

```
(define t1
  (trait
```

```
    (define/public (m1) 1)))
(define t2
  (trait
    (define/public (m2) 2)))
(define t3 (trait-sum t1 t2))
```

creates a trait `t3` that is equivalent to

```
(trait
  (define/public (m1) 1)
  (define/public (m2) 2))
```

but `t1` and `t2` can still be used individually or combined with other traits.

When traits are combined with `trait-sum`, the combination drops `inherit`, `inherit/super`, `inherit/inner`, and `inherit-field` declarations when a definition is supplied for the same method or field name by another trait. The `trait-sum` operation fails (the `exn:fail:contract` exception is raised) if any of the traits to combine define a method or field with the same name, or if an `inherit/super` or `inherit/inner` declaration to be dropped is inconsistent with the supplied definition. In other words, declaring a method with `inherit`, `inherit/super`, or `inherit/inner`, does not count as defining the method; at the same time, for example, a trait that contains an `inherit/super` declaration for a method `m` cannot be combined with a trait that defines `m` as `augment`, since no class could satisfy the requirements of both `augment` and `inherit/super` when the trait is later converted to a mixin and applied to a class.

```
(trait-exclude trait-expr id)
```

Produces a new trait that is like the trait result of *`trait-expr`*, but with the definition of a method named by *`id`* removed; as the method definition is removed, either an `inherit`, `inherit/super`, or `inherit/inner` declaration is added:

- A method declared with `public`, `pubment`, or `public-final` is replaced with an `inherit` declaration.
- A method declared with `override` or `override-final` is replaced with an `inherit/super` declaration.
- A method declared with `augment`, `augride`, or `augment-final` is replaced with an `inherit/inner` declaration.
- A method declared with `overment` is not replaced with any `inherit` declaration.

If the trait produced by *`trait-expr`* has no method definition for *`id`*, the `exn:fail:contract` exception is raised.

```
(trait-exclude-field trait-expr id)
```

Produces a new trait that is like the trait result of *trait-expr*, but with the definition of a field named by *id* removed; as the field definition is removed, an `inherit-field` declaration is added.

```
(trait-alias trait-expr id new-id)
```

Produces a new trait that is like the trait result of *trait-expr*, but the definition and declaration of the method named by *id* is duplicated with the name *new-id*. The consistency requirements for the resulting trait are the same as for `trait-sum`, otherwise the `exn:fail:contract` exception is raised. This operation does not rename any other use of *id*, such as in method calls (even method calls to `identifier` in the cloned definition for *new-id*).

```
(trait-rename trait-expr id new-id)
```

Produces a new trait that is like the trait result of *trait-expr*, but all definitions and references to methods named *id* are replaced by definitions and references to methods named by *new-id*. The consistency requirements for the resulting trait are the same as for `trait-sum`, otherwise the `exn:fail:contract` exception is raised.

```
(trait-rename-field trait-expr id new-id)
```

Produces a new trait that is like the trait result of *trait-expr*, but all definitions and references to fields named *id* are replaced by definitions and references to fields named by *new-id*. The consistency requirements for the resulting trait are the same as for `trait-sum`, otherwise the `exn:fail:contract` exception is raised.

5.7 Object and Class Contracts

```
(class/c maybe-opaque member-spec ...)
```

```
maybe-opaque =
             | #:opaque

member-spec = method-spec
             | (field field-spec ...)
             | (init field-spec ...)
             | (init-field field-spec ...)
             | (inherit method-spec ...)
             | (inherit-field field-spec ...)
             | (super method-spec ...)
             | (inner method-spec ...)
             | (override method-spec ...)
             | (augment method-spec ...)
             | (augride method-spec ...)
             | (absent absent-spec ...)

method-spec = method-id
             | (method-id method-contract)

 field-spec = field-id
             | (field-id contract-expr)

absent-spec = method-id
             | (field field-id ...)
```

Produces a contract for a class.

There are two major categories of contracts listed in a `class/c` form: external and internal contracts. External contracts govern behavior when an object is instantiated from a class or when methods or fields are accessed via an object of that class. Internal contracts govern behavior when method or fields are accessed within the class hierarchy. This separation allows for stronger contracts for class clients and weaker contracts for subclasses.

Method contracts must contain an additional initial argument which corresponds to the implicit `this` parameter of the method. This allows for contracts which discuss the state of the object when the method is called (or, for dependent contracts, in other parts of the contract). Alternative contract forms, such as `->m`, are provided as a shorthand for writing method contracts.

Methods and fields listed in an `absent` clause must *not* be present in the class.

A class contract can be specified to be *opaque* with the `#:opaque` keyword. An opaque class contract will only accept a class that defines exactly the methods and fields specified by the contract. A contract error is raised if the contracted class contains any methods or fields that are not specified.

The external contracts are as follows:

- A method contract without a tag describes the behavior of the implementation of *method-id* on method sends to an object of the contracted class. This contract will continue to be checked in subclasses until the contracted class's implementation is no longer the entry point for dynamic dispatch.
- A field contract, tagged with `field`, describes the behavior of the value contained in that field when accessed via an object of that class. Since fields may be mutated, these contracts are checked on any external access and/or mutation of the field.
- An initialization argument contract, tagged with `init`, describes the expected behavior of the value paired with that name during class instantiation. The same name can be provided more than once, in which case the first such contract in the `class/c` form is applied to the first value tagged with that name in the list of initialization arguments, and so on.
- The contracts listed in an `init-field` section are treated as if each contract appeared in an `init` section and a `field` section.

The internal contracts are as follows:

- A method contract, tagged with `inherit`, describes the behavior of the method when invoked directly (i.e., via `inherit`) in any subclass of the contracted class. This contract, like external method contracts, applies until the contracted class's method implementation is no longer the entry point for dynamic dispatch.
- A field contract, tagged with `inherit-field`, describes the behavior of the value contained in that field when accessed directly (i.e., via `inherit-field`) in any subclass of the contracted class. Since fields may be mutated, these contracts are checked on any access and/or mutation of the field that occurs in such subclasses.
- A method contract, tagged with `super`, describes the behavior of *method-id* when called by the `super` form in a subclass. This contract only affects `super` calls in subclasses which call the contract class's implementation of *method-id*.
- A method contract, tagged with `inner`, describes the behavior the class expects of an augmenting method in a subclass. This contract affects any implementations of *method-id* in subclasses which can be called via `inner` from the contracted class. This means a subclass which implements *method-id* via `augment` or `overment` stop future subclasses from being affected by the contract, since further extension cannot be reached via the contracted class.
- A method contract, tagged with `override`, describes the behavior expected by the contracted class for *method-id* when called directly (i.e. by the application `(`*method-id* `...)`). This form can only be used if overriding the method in subclasses will change the entry point to the dynamic dispatch chain (i.e., the method has never been augmentable).

- A method contract, tagged with either `augment` or `augride`, describes the behavior provided by the contracted class for *`method-id`* when called directly from subclasses. These forms can only be used if the method has previously been augmentable, which means that no augmenting or overriding implementation will change the entry point to the dynamic dispatch chain. `augment` is used when subclasses can augment the method, and `augride` is used when subclasses can override the current augmentation.

```
(absent method-id ...)
```

See `class/c`; use outside of a `class/c` form is a syntax error.

```
(->m dom ... range)
```

Similar to `->`, except that the domain of the resulting contract contains one more element than the stated domain, where the first (implicit) argument is contracted with `any/c`. This contract is useful for writing simpler method contracts when no properties of `this` need to be checked.

```
(->*m (mandatory-dom ...) (optional-dom ...) rest range)
```

Similar to `->*`, except that the mandatory domain of the resulting contract contains one more element than the stated domain, where the first (implicit) argument is contracted with `any/c`. This contract is useful for writing simpler method contracts when no properties of `this` need to be checked.

```
(case->m (-> dom ... rest range) ...)
```

Similar to `case->`, except that the mandatory domain of each case of the resulting contract contains one more element than the stated domain, where the first (implicit) argument is contracted with `any/c`. This contract is useful for writing simpler method contracts when no properties of `this` need to be checked.

```
(->dm (mandatory-dependent-dom ...)
      (optional-dependent-dom ...)
      dependent-rest
      pre-cond
      dep-range)
```

Similar to `->d`, except that the mandatory domain of the resulting contract contains one more element than the stated domain, where the first (implicit) argument is contracted with `any/c`. In addition, `this` is appropriately bound in the body of the contract. This contract is useful for writing simpler method contracts when no properties of `this` need to be checked.

```
(object/c member-spec ...)
```

```
member-spec = method-spec
            | (field field-spec ...)

method-spec = method-id
            | (method-id method-contract)

 field-spec = field-id
            | (field-id contract-expr)
```

Produces a contract for an object.

Unlike the older form `object-contract`, but like `class/c`, arbitrary contract expressions are allowed. Also, method contracts for `object/c` follow those for `class/c`. An object wrapped with `object/c` behaves as if its class had been wrapped with the equivalent `class/c` contract.

```
(instanceof/c class-contract) → contract?
  class-contract : contract?
```

Produces a contract for an object, where the object is an instance of a class that conforms to *class-contract*.

```
(object-contract member-spec ...)
```

```
           member-spec = (method-id method-contract)
                       | (field field-id contract-expr)

        method-contract = (-> dom ... range)
                       | (->* (mandatory-dom ...)
                              (optional-dom ...)
                              rest
                              range)
                       | (->d (mandatory-dependent-dom ...)
                              (optional-dependent-dom ...)
                              dependent-rest
                              pre-cond
                              dep-range)

                   dom = dom-expr
                       | keyword dom-expr

                 range = range-expr
                       | (values range-expr ...)
                       | any

          mandatory-dom = dom-expr
                       | keyword dom-expr

           optional-dom = dom-expr
                       | keyword dom-expr

                  rest =
                       | #:rest rest-expr

mandatory-dependent-dom = [id dom-expr]
                       | keyword [id dom-expr]

 optional-dependent-dom = [id dom-expr]
                       | keyword [id dom-expr]

         dependent-rest =
                       | #:rest id rest-expr

              pre-cond =
                       | #:pre-cond boolean-expr

             dep-range = any
                       | [id range-expr] post-cond
                       | (values [id range-expr] ...) post-cond

             post-cond =
                       | #:post-cond boolean-expr
```

Produces a contract for an object.

Each of the contracts for a method has the same semantics as the corresponding function contract, but the syntax of the method contract must be written directly in the body of the object-contract—much like the way that methods in class definitions use the same syntax as regular function definitions, but cannot be arbitrary procedures. Unlike the method contracts for `class/c`, the implicit `this` argument is not part of the contract. To allow for the use of `this` in dependent contracts, `->d` contracts implicitly bind `this` to the object itself.

```
mixin-contract : contract?
```

A function contract that recognizes mixins. It guarantees that the input to the function is a class and the result of the function is a subclass of the input.

```
(make-mixin-contract type ...) → contract?
  type : (or/c class? interface?)
```

Produces a function contract that guarantees the input to the function is a class that implements/subclasses each *type*, and that the result of the function is a subclass of the input.

```
(is-a?/c type) → flat-contract?
  type : (or/c class? interface?)
```

Accepts a class or interface and returns a flat contract that recognizes objects that instantiate the class/interface.

```
(implementation?/c interface) → flat-contract?
  interface : interface?
```

Returns a flat contract that recognizes classes that implement *interface*.

```
(subclass?/c class) → flat-contract?
  class : class?
```

Returns a flat contract that recognizes classes that are subclasses of *class*.

5.8 Object Equality and Hashing

By default, objects that are instances of different classes or that are instances of a non-transparent class are `equal?` only if they are `eq?`. Like transparent structures, two objects that are instances of the same transparent class (i.e., every superclass of the class has `#f` as its inspector) are `equal?` when their field values are `equal?`.

To customize the way that a class instance is compared to other instances by `equal?`, implement the `equal<%>` interface.

```
equal<%> : interface?
```

The `equal<%>` interface includes three methods, which are analogous to the functions provided for a structure type with `prop:equal+hash`:

- `equal-to?` — Takes two arguments. The first argument is an object that is an instance of the same class (or a subclass that does not re-declare its implementation of `equal<%>`) and that is being compared to the target object. The second argument is an `equal?`-like procedure of two arguments that should be used for recursive equality testing. The result should be a true value if the object and the first argument of the method are equal, `#f` otherwise.
- `equal-hash-code-of` — Takes one argument, which is a procedure of one argument that should be used for recursive hash-code computation. The result should be an exact integer representing the target object's hash code.
- `equal-secondary-hash-code-of` — Takes one argument, which is a procedure of one argument that should be used for recursive hash-code computation. The result should be an exact integer representing the target object's secondary hash code.

The `equal<%>` interface is unusual in that declaring the implementation of the interface is different from inheriting the interface. Two objects can be equal only if they are instances of classes whose most specific ancestor to explicitly implement `equal<%>` is the same ancestor.

See `prop:equal+hash` for more information on equality comparisons and hash codes. The `equal<%>` interface is implemented with `interface*` and `prop:equal+hash`.

5.9 Object Serialization

```
(define-serializable-class* class-id superclass-expr
                            (interface-expr ...)
  class-clause ...)
```

Binds *class-id* to a class, where *superclass-expr*, the *interface-expr*s, and the *class-clause*s are as in `class*`.

This form can only be used at the top level, either within a module or outside. The *class-id* identifier is bound to the new class, and `deserialize-info:`*class-id* is also defined; if the definition is within a module, then the latter is provided from the module.

Serialization for the class works in one of two ways:

- If the class implements the built-in interface `externalizable<%>`, then an object is serialized by calling its `externalize` method; the result can be anything that is serializable (but, obviously, should not be the object itself). Deserialization creates an instance of the class with no initialization arguments, and then calls the object's `internalize` method with the result of `externalize` (or, more precisely, a deserialized version of the serialized result of a previous call).

  To support this form of serialization, the class must be instantiable with no initialization arguments. Furthermore, cycles involving only instances of the class (and other such classes) cannot be serialized.

- If the class does not implement `externalizable<%>`, then every superclass of the class must be either serializable or transparent (i.e,. have `#f` as its inspector). Serialization and deserialization are fully automatic, and may involve cycles of instances.

  To support cycles of instances, deserialization may create an instance of the call with all fields as the undefined value, and then mutate the object to set the field values. Serialization support does not otherwise make an object's fields mutable.

In the second case, a serializable subclass can implement `externalizable<%>`, in which case the `externalize` method is responsible for all serialization (i.e., automatic serialization is lost for instances of the subclass). In the first case, all serializable subclasses implement `externalizable<%>`, since a subclass implements all of the interfaces of its parent class.

In either case, if an object is an immediate instance of a subclass (that is not itself serializable), the object is serialized as if it was an immediate instance of the serializable class. In particular, overriding declarations of the `externalize` method are ignored for instances of non-serializable subclasses.

```
(define-serializable-class class-id superclass-expr
  class-clause ...)
```

Like `define-serializable-class*`, but without interface expressions (analogous to `class`).

```
externalizable<%> : interface?
```

The `externalizable<%>` interface includes only the `externalize` and `internalize` methods. See `define-serializable-class*` for more information.

5.10 Object Printing

To customize the way that a class instance is printed by `print`, `write` and `display`, implement the `printable<%>` interface.

```
printable<%> : interface?
```

The `printable<%>` interface includes only the `custom-print`, `custom-write`, and `custom-display` methods. The `custom-print` method accepts two arguments: the destination port and the current `quasiquote` depth as an exact nonnegative integer. The `custom-write` and `custom-display` methods each accepts a single argument, which is the destination port to `write` or `display` the object.

Calls to the `custom-print`, `custom-write`, or `custom-display` methods are like calls to a procedure attached to a structure type through the `prop:custom-write` property. In particular, recursive printing can trigger an escape from the call.

See `prop:custom-write` for more information. The `printable<%>` interface is implemented with `interface*` and `prop:custom-write`.

```
writable<%> : interface?
```

Like `printable<%>`, but includes only the `custom-write` and `custom-display` methods. A `print` request is directed to `custom-write`.

5.11 Object, Class, and Interface Utilities

```
(object? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is an object, `#f` otherwise.

Examples:

```
> (object? (new object%))
#t
> (object? object%)
#f
> (object? "clam chowder")
#f
```

```
(class? v) → boolean?
  v : any/c
```

Returns #t if v is a class, #f otherwise.

Examples:

```
> (class? object%)
#t
> (class? (class object% (super-new)))
#t
> (class? (new object%))
#f
> (class? "corn chowder")
#f
```

```
(interface? v) → boolean?
  v : any/c
```

Returns #t if v is an interface, #f otherwise.

Examples:

```
> (interface? (interface () empty cons first rest))
#t
> (interface? object%)
#f
> (interface? "gazpacho")
#f
```

```
(generic? v) → boolean?
  v : any/c
```

Returns #t if v is a generic, #f otherwise.

Examples:

```
(define c%
  (class object%
    (super-new)
    (define/public (m x)
      (+ 3.14 x))))

> (generic? (generic c% m))
#t
> (generic? c%)
#f
> (generic? "borscht")
#f
```

```
(object=? a b) → boolean?
  a : object?
  b : object?
```

Determines if two objects are the same object, or not; this procedure uses `eq?`, but also works properly with contracts.

Examples:

```
(define obj-1 (new object%))

> (define/contract obj-2 (object/c) obj-1)

> (object=? obj-1 obj-1)
#t
> (object=? (new object%) obj-1)
#f
> (object=? obj-1 obj-2)
#t
> (object=? obj-1 (new (class object% (super-new))))
#f
```

```
(object->vector object [opaque-v]) → vector?
  object : object?
  opaque-v : any/c = #f
```

Returns a vector representing *object* that shows its inspectable fields, analogous to `struct->vector`.

Examples:

```
> (object->vector (new object%))
'#(object:object% ...)
> (object->vector (new (class object%
                         (super-new)
                         (field [x 5] [y 10]))))
'#(object:eval:82:0 ...)
```

```
(class->interface class) → interface?
  class : class?
```

Returns the interface implicitly defined by *class*.

Example:

```
> (class->interface object%)
#<interface:object%>
```

```
(object-interface object) → interface?
  object : object?
```

Returns the interface implicitly defined by the class of *object*.

Example:

```
> (object-interface (new object%))
#<interface:object%>
```

```
(is-a? v type) → boolean?
  v : any/c
  type : (or/c interface? class?)
```

Returns #t if *v* is an instance of a class *type* or a class that implements an interface *type*, #f otherwise.

Examples:

```
(define point<%> (interface () get-x get-y))

(define 2d-point%
  (class* object% (point<%>)
    (super-new)
    (field [x 0] [y 0])
    (define/public (get-x) x)
    (define/public (get-y) y)))

> (is-a? (new 2d-point%) 2d-point%)
#t
> (is-a? (new 2d-point%) point<%>)
#t
> (is-a? (new object%) 2d-point%)
#f
> (is-a? (new object%) point<%>)
#f
```

```
(subclass? v class) → boolean?
  v : any/c
  class : class?
```

Returns #t if *v* is a class derived from (or equal to) *class*, #f otherwise.

Examples:

```
> (subclass? (class object% (super-new)) object%)
#t
> (subclass? object% (class object% (super-new)))
#f
> (subclass? object% object%)
#t
```

```
(implementation? v interface) → boolean?
  v : any/c
  interface : interface?
```

Returns `#t` if *v* is a class that implements *interface*, `#f` otherwise.

Examples:

```
(define i<%> (interface () go))

(define c%
  (class* object% (i<%>)
    (super-new)
    (define/public (go) 'go)))

> (implementation? c% i<%>)
#t
> (implementation? object% i<%>)
#f
```

```
(interface-extension? v interface) → boolean?
  v : any/c
  interface : interface?
```

Returns `#t` if *v* is an interface that extends *interface*, `#f` otherwise.

Examples:

```
(define point<%> (interface () get-x get-y))

(define colored-point<%> (interface (point<%>) color))

> (interface-extension? colored-point<%> point<%>)
#t
> (interface-extension? point<%> colored-point<%>)
#f
> (interface-extension? (interface () get-x get-y get-z) point<%>)
#f
```

```
(method-in-interface? sym interface) → boolean?
  sym : symbol?
  interface : interface?
```

Returns `#t` if *interface* (or any of its ancestor interfaces) includes a member with the name *sym*, `#f` otherwise.

Examples:

```
(define i<%> (interface () get-x get-y))

> (method-in-interface? 'get-x i<%>)
#t
> (method-in-interface? 'get-z i<%>)
#f
```

```
(interface->method-names interface) → (listof symbol?)
  interface : interface?
```

Returns a list of symbols for the method names in *interface*, including methods inherited from superinterfaces, but not including methods whose names are local (i.e., declared with `define-local-member-names`).

Examples:

```
(define i<%> (interface () get-x get-y))

> (interface->method-names i<%>)
'(get-y get-x)
```

```
(object-method-arity-includes? object
                               sym
                               cnt)   → boolean?
  object : object?
  sym : symbol?
  cnt : exact-nonnegative-integer?
```

Returns `#t` if *object* has a method named *sym* that accepts *cnt* arguments, `#f` otherwise.

Examples:

```
(define c%
  (class object%
    (super-new)
    (define/public (m x [y 0])
      (+ x y))))

> (object-method-arity-includes? (new c%) 'm 1)
#t
> (object-method-arity-includes? (new c%) 'm 2)
#t
```

```
> (object-method-arity-includes? (new c%) 'm 3)
#f
> (object-method-arity-includes? (new c%) 'n 1)
#f
```

```
(field-names object) → (listof symbol?)
  object : object?
```

Returns a list of all of the names of the fields bound in *object*, including fields inherited from superinterfaces, but not including fields whose names are local (i.e., declared with `define-local-member-names`).

Examples:

```
> (field-names (new object%))
'()
> (field-names (new (class object% (super-
new) (field [x 0] [y 0]))))
'(y x)
```

```
(object-info object) → (or/c class? #f) boolean?
  object : any/c
```

Returns two values, analogous to the return values of `struct-info`:

- *class*: a class or `#f`; the result is `#f` if the current inspector does not control any class for which the *object* is an instance.
- *skipped?*: `#f` if the first result corresponds to the most specific class of *object*, `#t` otherwise.

```
(class-info class)
    symbol?
    exact-nonnegative-integer?
    (listof symbol?)
 → (any/c exact-nonnegative-integer? . -> . any/c)
    (any/c exact-nonnegative-integer? any/c . -> . any/c)
    (or/c class? #f)
    boolean?
  class : class?
```

Returns seven values, analogous to the return values of `struct-type-info`:

- *name*: the class's name as a symbol;

- *field-cnt*: the number of fields (public and private) defined by the class;
- *field-name-list*: a list of symbols corresponding to the class's public fields; this list can be larger than *field-cnt* because it includes inherited fields;
- *field-accessor*: an accessor procedure for obtaining field values in instances of the class; the accessor takes an instance and a field index between 0 (inclusive) and *field-cnt* (exclusive);
- *field-mutator*: a mutator procedure for modifying field values in instances of the class; the mutator takes an instance, a field index between 0 (inclusive) and *field-cnt* (exclusive), and a new field value;
- *super-class*: a class for the most specific ancestor of the given class that is controlled by the current inspector, or #f if no ancestor is controlled by the current inspector;
- *skipped?*: #f if the sixth result is the most specific ancestor class, #t otherwise.

```
(struct exn:fail:object exn:fail ()
        #:extra-constructor-name make-exn:fail:object)
```

Raised for `class`-related failures, such as attempting to call a method that is not supplied by an object.

5.12 Surrogates

```
(require racket/surrogate)
```

The bindings documented in this section are provided by the `racket/surrogate` library, not `racket/base` or `racket`.

The `racket/surrogate` library provides an abstraction for building an instance of the *proxy design pattern*. The pattern consists of two objects, a *host* and a *surrogate* object. The host object delegates method calls to its surrogate object. Each host has a dynamically assigned surrogate, so an object can completely change its behavior merely by changing the surrogate.

```
(surrogate method-spec ...)

method-spec = (method-id arg-spec ...)
            | (override method-id arg-spec ...)
            | (override-final method-id (lambda () default-expr)
                              arg-spec ...)

  arg-spec = (id ...)
           | id
```

If neither `override` nor `override-final` is specified for a *`method-id`*, then `override` is assumed.

The `surrogate` form produces four values: a host mixin (a procedure that accepts and returns a class), a host interface, a surrogate class, and a surrogate interface.

The host mixin adds one additional field, `surrogate`, to its argument. It also adds a getter method, `get-surrogate`, and a setter method, `set-surrogate`, for changing the field. The `set-surrogate` method accepts instances of the class returned by the `surrogate` form or `#f`, and it updates the field with its argument; then, `set-surrogate` calls the `on-disable-surrogate` on the previous value of the field and `on-enable-surrogate` for the new value of the field. The `get-surrogate` method returns the current value of the field.

The host mixin has a single overriding method for each *`method-id`* in the `surrogate` form. Each of these methods is defined with a `case-lambda` with one arm for each *`arg-spec`*. Each arm has the variables as arguments in the *`arg-spec`*. The body of each method tests the `surrogate` field. If it is `#f`, the method just returns the result of invoking the super or inner method. If the `surrogate` field is not `#f`, the corresponding method of the object in the field is invoked. This method receives the same arguments as the original method, plus two extras. The extra arguments come at the beginning of the argument list. The first is the original object. The second is a procedure that calls the super or inner method (i.e., the method of the class that is passed to the mixin or an extension, or the method in an overriding class), with the arguments that the procedure receives.

The host interface has the names `set-surrogate`, `get-surrogate`, and each of the *`method-id`*s in the original form.

The surrogate class has a single public method for each *`method-id`* in the `surrogate` form. These methods are invoked by classes constructed by the mixin. Each has a corresponding method signature, as described in the above paragraph. Each method just passes its argument along to the super procedure it receives.

Note: if you derive a class from the surrogate class, do not both call the `super` argument and the super method of the surrogate class itself. Only call one or the other, since the default methods call the `super` argument.

Finally, the interface contains all of the names specified in surrogate's argument, plus `on-enable-surrogate` and `on-disable-surrogate`. The class returned by `surrogate` implements this interface.

6 Units

§14 “Units (Components)” in ***The Racket Guide*** introduces units.

Units organize a program into separately compilable and reusable components. The imports and exports of a unit are grouped into a *signature*, which can include “static” information (such as macros) in addition to placeholders for run-time values. Units with suitably matching signatures can be *linked* together to form a larger unit, and a unit with no imports can be *invoked* to execute its body.

```
(require racket/unit)
```

The bindings documented in this section are provided by the `racket/unit` and `racket` libraries, but not `racket/base`. The `racket/unit` module name can be used as a language name with `#lang`; see §6.10 “Single-Unit Modules”.

6.1 Creating Units

```
(unit
  (import tagged-sig-spec ...)
  (export tagged-sig-spec ...)
  init-depends-decl
  unit-body-expr-or-defn
  ...)

  tagged-sig-spec = sig-spec
                  | (tag id sig-spec)

        sig-spec = sig-id
                 | (prefix id sig-spec)
                 | (rename sig-spec (id id) ...)
                 | (only sig-spec id ...)
                 | (except sig-spec id ...)

init-depends-decl =
                 | (init-depend tagged-sig-id ...)

    tagged-sig-id = sig-id
                  | (tag id sig-id)
```

Produces a unit that encapsulates its *unit-body-expr-or-defn*s. Expressions in the `unit` body can refer to identifiers bound by the *sig-spec*s of the `import` clause, and the body must include one definition for each identifier of a *sig-spec* in the `export` clause. An identifier that is exported cannot be `set!`ed in either the defining unit or in importing units, although the implicit assignment to initialize the variable may be visible as a mutation.

Each import or export *sig-spec* ultimately refers to a *sig-id*, which is an identifier that is bound to a signature by `define-signature`.

In a specific import or export position, the set of identifiers bound or required by a particular *sig-id* can be adjusted in a few ways:

- `(prefix` *id sig-spec*`)` as an import binds the same as *sig-spec*, except that each binding is prefixed with *id*. As an export, this form causes definitions using the *id* prefix to satisfy the exports required by *sig-spec*.
- `(rename` *sig-spec* `(`*id id*`) ...)` as an import binds the same as *sig-spec*, except that the first *id* is used for the binding instead of the second *id* (where *sig-spec* by itself must imply a binding that is `bound-identifier=?` to second *id*). As an export, this form causes a definition for the first *id* to satisfy the export named by the second *id* in *sig-spec*.
- `(only` *sig-spec id ...*`)` as an import binds the same as *sig-spec*, but restricted to just the listed *id*s (where *sig-spec* by itself must imply a binding that is `bound-identifier=?` to each *id*). This form is not allowed for an export.
- `(except` *sig-spec id ...*`)` as an import binds the same as *sig-spec*, but excluding all listed *id*s (where *sig-spec* by itself must imply a binding that is `bound-identifier=?` to each *id*). This form is not allowed for an export.

As suggested by the grammar, these adjustments to a signature can be nested arbitrarily.

A unit's declared imports are matched with actual supplied imports by signature. That is, the order in which imports are supplied to a unit when linking is irrelevant; all that matters is the signature implemented by each supplied import. One actual import must be provided for each declared import. Similarly, when a unit implements multiple signatures, the order of the export signatures does not matter.

To support multiple imports or exports for the same signature, an import or export can be tagged using the form `(tag` *id sig-spec*`)`. When an import declaration of a unit is tagged, then one actual import must be given the same tag (with the same signature) when the unit is linked. Similarly, when an export declaration is tagged for a unit, then references to that particular export must explicitly use the tag.

A unit is prohibited syntactically from importing two signatures that are not distinct, unless they have different tags; two signatures are *distinct* only if they share no ancestor through `extends`. The same syntactic constraint applies to exported signatures. In addition, a unit is prohibited syntactically from importing the same identifier twice (after renaming and other transformations on a *sig-spec*), exporting the same identifier twice (again, after renaming), or exporting an identifier that is imported.

When units are linked, the bodies of the linked units are executed in an order that is specified at the linking site. An optional `(init-depend` *tagged-sig-id ...*`)` declaration constrains the allowed orders of linking by specifying that the current unit must be initialized

after the unit that supplies the corresponding import. Each *tagged-sig-id* in an `init-depend` declaration must have a corresponding import in the `import` clause.

```
(define-signature id extension-decl
  (sig-elem ...))

extension-decl =
               | extends sig-id

     sig-elem = id
              | (define-syntaxes (id ...) expr)
              | (define-values (id ...) expr)
              | (define-values-for-export (id ...) expr)
              | (contracted [id contract] ...)
              | (open sig-spec)
              | (struct id (field ...) struct-option ...)
              | (sig-form-id . datum)

        field = id
              | [id #:mutable]

 struct-option = #:mutable
               | #:constructor-name constructor-id
               | #:extra-constructor-name constructor-id
               | #:omit-constructor
               | #:omit-define-syntaxes
               | #:omit-define-values
```

Binds an identifier to a signature that specifies a group of bindings for import or export:

- Each *id* in a signature declaration means that a unit implementing the signature must supply a variable definition for the *id*. That is, *id* is available for use in units importing the signature, and *id* must be defined by units exporting the signature.
- Each `define-syntaxes` form in a signature declaration introduces a macro that is available for use in any unit that imports the signature. Free variables in the definition's *expr* refer to other identifiers in the signature first, or the context of the `define-signature` form if the signature does not include the identifier.
- Each `define-values` form in a signature declaration introduces code that effectively prefixes every unit that imports the signature. Free variables in the definition's *expr* are treated the same as for `define-syntaxes`.
- Each `define-values-for-export` form in a signature declaration introduces code that effectively suffixes every unit that exports the signature. Free variables in the definition's *expr* are treated the same as for `define-syntaxes`.

- Each `contracted` form in a signature declaration means that a unit exporting the signature must supply a variable definition for each *id* in that form. If the signature is imported, then uses of *id* inside the unit are protected by the appropriate contracts using the unit as the negative blame. If the signature is exported, then the exported values are protected by the appropriate contracts which use the unit as the positive blame, but internal uses of the exported identifiers are not protected. Variables in the *contract* expressions are treated the same as for `define-syntaxes`.
- Each `(open` *sig-spec*`)` adds to the signature everything specified by *sig-spec*.
- Each `(struct` *id* `(`*field* `...)` *struct-option* `...)` adds all of the identifiers that would be bound by `(struct` *id* `(`*field* `...)` `field-option ...)`, where the extra option `#:omit-constructor` omits the constructor identifier.
- Each `(`*sig-form-id* `.` *datum*`)` extends the signature in a way that is defined by *sig-form-id*, which must be bound by `define-signature-form`. One such binding is for `struct/ctc`.

When a `define-signature` form includes an `extends` clause, then the define signature automatically includes everything in the extended signature. Furthermore, any implementation of the new signature can be used as an implementation of the extended signature.

`(open` *sig-spec*`)`

Allowed only in a *sig-elem*; see `define-signature`.

`(define-values-for-export (`*id* `...)` *expr*`)`

Allowed only in a *sig-elem*; see `define-signature`.

`(contracted [`*id contract*`] ...)`

Allowed only in a *sig-elem*; see `define-signature`.

`(only` *sig-spec id* `...)`

Allowed only in a *sig-spec*; see `unit`.

`(except` *sig-spec id* `...)`

Allowed only in a *sig-spec*; see `unit`.

```
(rename sig-spec (id id) ...)
```

Allowed only in a *sig-spec*; see `unit`.

```
(prefix id sig-spec)
```

Allowed only in a *sig-spec*; see `unit`.

```
(import tagged-sig-spec ...)
```

Allowed only in certain forms; see, for example, `unit`.

```
(export tagged-sig-spec ...)
```

Allowed only in certain forms; see, for example, `unit`.

```
(link linkage-decl ...)
```

Allowed only in certain forms; see, for example, `compound-unit`.

```
(tag id sig-spec)
(tag id sig-id)
```

Allowed only in certain forms; see, for example, `unit`.

```
(init-depend tagged-sig-id ...)
```

Allowed only in a `init-depend-decl`; see `unit`.

```
extends
```

Allowed only within `define-signature`.

6.2 Invoking Units

```
(invoke-unit unit-expr)
(invoke-unit unit-expr (import tagged-sig-spec ...))
```

Invokes the unit produced by *unit-expr*. For each of the unit's imports, the `invoke-unit` expression must contain a *tagged-sig-spec* in the `import` clause; see `unit` for the grammar of *tagged-sig-spec*. If the unit has no imports, the `import` clause can be omitted.

When no *tagged-sig-spec*s are provided, *unit-expr* must produce a unit that expects no imports. To invoke the unit, all bindings are first initialized to the `#<undefined>` value. Next, the unit's body definitions and expressions are evaluated in order; in the case of a definition, evaluation sets the value of the corresponding variable(s). Finally, the result of the last expression in the unit is the result of the `invoke-unit` expression.

Each supplied *tagged-sig-spec* takes bindings from the surrounding context and turns them into imports for the invoked unit. The unit need not declare an import for every provided *tagged-sig-spec*, but one *tagged-sig-spec* must be provided for each declared import of the unit. For each variable identifier in each provided *tagged-sig-spec*, the value of the identifier's binding in the surrounding context is used for the corresponding import in the invoked unit.

```
(define-values/invoke-unit unit-expr
  (import tagged-sig-spec ...)
  (export tagged-sig-spec ...))
```

Like `invoke-unit`, but the values of the unit's exports are copied to new bindings.

The unit produced by *unit-expr* is linked and invoked as for `invoke-unit`. In addition, the `export` clause is treated as a kind of import into the local definition context. That is, for every binding that would be available in a unit that used the `export` clause's *tagged-sig-spec* as an import, a definition is generated for the context of the `define-values/invoke-unit` form.

6.3 Linking Units and Creating Compound Units

```
(compound-unit
  (import link-binding ...)
  (export tagged-link-id ...)
  (link linkage-decl ...))

  link-binding = (link-id : tagged-sig-id)

tagged-link-id = (tag id link-id)
               | link-id

  linkage-decl = ((link-binding ...) unit-expr tagged-link-id ...)
```

Links several units into one new compound unit without immediately invoking any of the linked units. The *unit-expr*s in the `link` clause determine the units to be linked in creating the compound unit. The *unit-expr*s are evaluated when the `compound-unit` form is evaluated.

The `import` clause determines the imports of the compound unit. Outside the compound unit, these imports behave as for a plain unit; inside the compound unit, they are propagated to some of the linked units. The `export` clause determines the exports of the compound unit. Again, outside the compound unit, these exports are treated the same as for a plain unit; inside the compound unit, they are drawn from the exports of the linked units. Finally, the left-hand and right-hand parts of each declaration in the `link` clause specify how the compound unit's imports and exports are propagated to the linked units.

Individual elements of an imported or exported signature are not available within the compound unit. Instead, imports and exports are connected at the level of whole signatures. Each specific import or export (i.e., an instance of some signature, possibly tagged) is given a *link-id* name. Specifically, a *link-id* is bound by the `import` clause or the left-hand part of a declaration in the `link` clause. A bound *link-id* is referenced in the right-hand part of a declaration in the `link` clause or by the `export` clause.

The left-hand side of a `link` declaration gives names to each expected export of the unit produced by the corresponding *unit-expr*. The actual unit may export additional signatures, and it may export an extension of a specific signature instead of just the specified one. If the unit does not export one of the specified signatures (with the specified tag, if any), the `exn:fail:contract` exception is raised when the `compound-unit` form is evaluated.

The right-hand side of a `link` declaration specifies the imports to be supplied to the unit produced by the corresponding *unit-expr*. The actual unit may import fewer signatures, and it may import a signature that is extended by the specified one. If the unit imports a signature (with a particular tag) that is not included in the supplied imports, the `exn:fail:contract` exception is raised when the `compound-unit` form is evaluated. Each *link-id* supplied as an import must be bound either in the `import` clause or in some declaration within the `link` clause.

The order of declarations in the `link` clause determines the order of invocation of the linked units. When the compound unit is invoked, the unit produced by the first *unit-expr* is invoked first, then the second, and so on. If the order specified in the `link` clause is inconsistent with `init-depend` declarations of the actual units, then the `exn:fail:contract` exception is raised when the `compound-unit` form is evaluated.

6.4 Inferred Linking

```
(define-unit unit-id
  (import tagged-sig-spec ...)
  (export tagged-sig-spec ...)
  init-depends-decl
  unit-body-expr-or-defn
  ...)
```

Binds *unit-id* to both a unit and static information about the unit.

Evaluating a reference to a *unit-id* bound by `define-unit` produces a unit, just like evaluating an `id` bound by `(define id (unit ...))`. In addition, however, *unit-id* can be used in `compound-unit/infer`. See `unit` for information on *tagged-sig-spec*, *init-depends-decl*, and *unit-body-expr-or-defn*.

```
(compound-unit/infer
  (import tagged-infer-link-import ...)
  (export tagged-infer-link-export ...)
  (link infer-linkage-decl ...))

tagged-infer-link-import = tagged-sig-id
                         | (link-id : tagged-sig-id)

tagged-infer-link-export = (tag id infer-link-export)
                         | infer-link-export

       infer-link-export = link-id
                         | sig-id

      infer-linkage-decl = ((link-binding ...) unit-id
                                               tagged-link-id ...)
                         | unit-id
```

Like `compound-unit`. Syntactically, the difference between `compound-unit` and `compound-unit/infer` is that the `unit-expr` for a linked unit is replaced with a *unit-id*, where a *unit-id* is bound by `define-unit` (or one of the other unit-binding forms that we introduce later in this section). Furthermore, an import can name just a *sig-id* without locally binding a *link-id*, and an export can be based on a *sig-id* instead of a *link-id*, and a declaration in the `link` clause can be simply a *unit-id* with no specified exports or imports.

The `compound-unit/infer` form expands to `compound-unit` by adding `sig-ids` as needed to the `import` clause, by replacing *sig-id*s in the `export` clause by *link-id*s, and by completing the declarations of the `link` clause. This completion is based on static information associated with each *unit-id*. Links and exports can be inferred when all signatures exported by the linked units are distinct from each other and from all imported

signatures, and when all imported signatures are distinct. Two signatures are *distinct* only if they share no ancestor through `extends`.

The long form of a `link` declaration can be used to resolve ambiguity by giving names to some of a unit's exports and supplying specific bindings for some of a unit's imports. The long form need not name all of a unit's exports or supply all of a unit's imports if the remaining parts can be inferred.

Like `compound-unit`, the `compound-unit/infer` form produces a (compound) unit without statically binding information about the result unit's imports and exports. That is, `compound-unit/infer` consumes static information, but it does not generate it. Two additional forms, `define-compound-unit` and `define-compound-unit/infer`, generate static information (where the former does not consume static information).

```
(define-compound-unit id
  (import link-binding ...)
  (export tagged-link-id ...)
  (link linkage-decl ...))
```

Like `compound-unit`, but binds static information about the compound unit like `define-unit`.

```
(define-compound-unit/infer id
  (import link-binding ...)
  (export tagged-infer-link-export ...)
  (link infer-linkage-decl ...))
```

Like `compound-unit/infer`, but binds static information about the compound unit like `define-unit`.

```
(define-unit-binding unit-id
  unit-expr
  (import tagged-sig-spec ...+)
  (export tagged-sig-spec ...+)
  init-depends-decl)
```

Like `define-unit`, but the unit implementation is determined from an existing unit produced by *unit-expr*. The imports and exports of the unit produced by *unit-expr* must be consistent with the declared imports and exports, otherwise the `exn:fail:contract` exception is raised when the `define-unit-binding` form is evaluated.

```
(invoke-unit/infer unit-spec)

unit-spec = unit-id
          | (link link-unit-id ...)
```

Like `invoke-unit`, but uses static information associated with *`unit-id`* to infer which imports must be assembled from the current context. If given a link form containing multiple *`link-unit-id`*s, then the units are first linked via `define-compound-unit/infer`.

```
(define-values/invoke-unit/infer maybe-exports unit-spec)

maybe-exports =
              | (export tagged-sig-spec ...)

    unit-spec = unit-id
              | (link link-unit-id ...)
```

Like `define-values/invoke-unit`, but uses static information associated with *`unit-id`* to infer which imports must be assembled from the current context and which exports should be bound by the definition. If given a link form containing multiple *`link-unit-id`*s, then the units are first linked via `define-compound-unit/infer`.

6.5 Generating A Unit from Context

```
(unit-from-context tagged-sig-spec)
```

Creates a unit that implements an interface using bindings in the enclosing environment. The generated unit is essentially the same as

```
(unit
  (import)
  (export tagged-sig-spec)
  (define id expr) ...)
```

for each `id` that must be defined to satisfy the exports, and each corresponding `expr` produces the value of `id` in the environment of the `unit-from-context` expression. (The unit cannot be written as above, however, since each `id` definition within the unit shadows the binding outside the `unit` form.)

See `unit` for the grammar of *`tagged-sig-spec`*.

```
(define-unit-from-context id tagged-sig-spec)
```

Like `unit-from-context`, in that a unit is constructed from the enclosing environment, and like `define-unit`, in that *`id`* is bound to static information to be used later with inference.

6.6 Structural Matching

```
(unit/new-import-export
  (import tagged-sig-spec ...)
  (export tagged-sig-spec ...)
  init-depends-decl
  ((tagged-sig-spec ...) unit-expr tagged-sig-spec))
```

Similar to `unit`, except the body of the unit is determined by an existing unit produced by *unit-expr*. The result is a unit whose implementation is *unit-expr*, but whose imports, exports, and initialization dependencies are as in the `unit/new-import-export` form (instead of as in the unit produced by *unit-expr*).

The final clause of the `unit/new-import-export` form determines the connection between the old and new imports and exports. The connection is similar to the way that `compound-unit` propagates imports and exports; the difference is that the connection between `import` and the right-hand side of the link clause is based on the names of elements in signatures, rather than the names of the signatures. That is, a *tagged-sig-spec* on the right-hand side of the link clause need not appear as a *tagged-sig-spec* in the `import` clause, but each of the bindings implied by the linking *tagged-sig-spec* must be implied by some *tagged-sig-spec* in the `import` clause. Similarly, each of the bindings implied by an `export` *tagged-sig-spec* must be implied by some left-hand-side *tagged-sig-spec* in the linking clause.

```
(define-unit/new-import-export unit-id
  (import tagged-sig-spec ...)
  (export tagged-sig-spec ...)
  init-depends-decl
  ((tagged-sig-spec ...) unit-expr tagged-sig-spec))
```

Like `unit/new-import-export`, but binds static information to *unit-id* like `define-unit`.

```
(unit/s
  (import tagged-sig-spec ...)
  (export tagged-sig-spec ...)
  init-depends-decl
  unit-id)
```

Like `unit/new-import-export`, but the linking clause is inferred, so *unit-id* must have the appropriate static information.

```
(define-unit/s name-id
  (import tagged-sig-spec ...)
  (export tagged-sig-spec ...)
  init-depends-decl
  unit-id)
```

Like `unit/s`, but binds static information to *name-id* like `define-unit`.

6.7 Extending the Syntax of Signatures

```
(define-signature-form sig-form-id expr)
(define-signature-form (sig-form-id id) body ...+)
```

Binds *sig-form-id* for use within a `define-signature` form.

In the first form, the result of *expr* must be a transformer procedure. In the second form, *sig-form-id* is bound to a transformer procedure whose argument is *id* and whose body is the *body*s. The result of the transformer must be a list of syntax objects, which are substituted for a use of *sig-form-id* in a `define-signature` expansion. (The result is a list so that the transformer can produce multiple declarations; `define-signature` has no splicing `begin` form.)

```
(struct/ctc id ([field contract-expr] ...) struct-option ...)

        field = id
              | [id #:mutable]

struct-option = #:mutable
              | #:omit-constructor
              | #:omit-define-syntaxes
              | #:omit-define-values
```

For use with `define-signature`. The `struct/ctc` form works similarly to `struct`, but the constructor, predicate, field accessors, and field mutators are contracted appropriately.

6.8 Unit Utilities

```
(unit? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a unit, `#f` otherwise.

```
(provide-signature-elements sig-spec ...)
```

Expands to a `provide` of all identifiers implied by the *sig-spec*s. See `unit` for the grammar of *sig-spec*.

6.9 Unit Contracts

```
(unit/c (import sig-block ...) (export sig-block ...))

sig-block = (tagged-sig-id [id contract] ...)
          | tagged-sig-id
```

A *unit contract* wraps a unit and checks both its imported and exported identifiers to ensure that they match the appropriate contracts. This allows the programmer to add contract checks to a single unit value without adding contracts to the imported and exported signatures.

The unit value must import a subset of the import signatures and export a superset of the export signatures listed in the unit contract. Any identifier which is not listed for a given signature is left alone. Variables used in a given *contract* expression first refer to other variables in the same signature, and then to the context of the `unit/c` expression.

```
(define-unit/contract unit-id
  (import sig-spec-block ...)
  (export sig-spec-block ...)
  init-depends-decl
  unit-body-expr-or-defn
  ...)

sig-spec-block = (tagged-sig-spec [id contract] ...)
               | tagged-sig-spec
```

The `define-unit/contract` form defines a unit compatible with link inference whose imports and exports are contracted with a unit contract. The unit name is used for the positive blame of the contract.

6.10 Single-Unit Modules

When `racket/unit` is used as a language name with `#lang`, the module body is treated as a unit body. The body must match the following *module-body* grammar:

```
module-body = require-decl ...
              (import tagged-sig-expr ...)
              (export tagged-sig-expr ...)
              init-depends-decl
              unit-body-expr-or-defn
              ...

require-decl = (require require-spec ...)
             | (begin require-decl ...)
```

```
        | derived-require-form
```

After any number of *require-decl*s, the content of the module is the same as a `unit` body.

The resulting unit is exported as *base*`@`, where *base* is derived from the enclosing module's name (i.e., its symbolic name, or its path without the directory and file suffix). If the module name ends in `-unit`, then *base* corresponds to the module name before `-unit`. Otherwise, the module name serves as *base*.

6.11 Single-Signature Modules

```
#lang racket/signature
```

The `racket/signature` language treats a module body as a unit signature.

The body must match the following *module-body* grammar:

```
module-body = (require require-spec ...) ... sig-spec ...
```

See §6.1 "Creating Units" for the grammar of *sig-spec*. Unlike the body of a `racket/unit` module, a `require` in a `racket/signature` module must be a literal use of `require`.

The resulting signature is exported as *base*`^`, where *base* is derived from the enclosing module's name (i.e., its symbolic name, or its path without the directory and file suffix). If the module name ends in `-sig`, then *base* corresponds to the module name before `-sig`. Otherwise, the module name serves as *base*.

6.12 Transformer Helpers

```
(require racket/unit-exptime)
```

The `racket/unit-exptime` library provides procedures that are intended for use by macro transformers. In particular, the library is typically imported using `for-syntax` into a module that defines macro with `define-syntax`.

```
(unit-static-signatures unit-identifier
                        err-syntax)
   (list/c (cons/c (or/c symbol? #f)
                   identifier?))
→  (list/c (cons/c (or/c symbol? #f)
                   identifier?))
 unit-identifier : identifier?
 err-syntax : syntax?
```

If *unit-identifier* is bound to static unit information via `define-unit` (or other such forms), the result is two values. The first value is for the unit's imports, and the second is for the unit's exports. Each result value is a list, where each list element pairs a symbol or `#f` with an identifier. The symbol or `#f` indicates the import's or export's tag (where `#f` indicates no tag), and the identifier indicates the binding of the corresponding signature.

If *unit-identifier* is not bound to static unit information, then the `exn:fail:syntax` exception is raised. In that case, the given *err-syntax* argument is used as the source of the error, where *unit-identifier* is used as the detail source location.

```
(signature-members sig-identifier
                                    (or/c identifier? #f)
                   err-syntax)   →  (listof identifier?)
                                    (listof identifier?)
                                    (listof identifier?)
  sig-identifier : identifier?
  err-syntax : syntax?
```

If *sig-identifier* is bound to static unit information via `define-signature` (or other such forms), the result is four values:

- an identifier or `#f` indicating the signature (of any) that is extended by the *sig-identifier* binding;
- a list of identifiers representing the variables supplied/required by the signature;
- a list of identifiers for variable definitions in the signature (i.e., variable bindings that are provided on import, but not defined by units that implement the signature); and
- a list of identifiers with syntax definitions in the signature.

If *sig-identifier* is not bound to a signature, then the `exn:fail:syntax` exception is raised. In that case, the given *err-syntax* argument is used as the source of the error, where *sig-identifier* is used as the detail source location.

7 Contracts

§7 “Contracts” in ***The Racket Guide*** introduces contracts.

The contract system guards one part of a program from another. Programmers specify the behavior of a module’s exports via `(provide (contract-out ....))`, and the contract system enforces those constraints.

```
(require racket/contract)
```

The bindings documented in this section are provided by the `racket/contract` and `racket` libraries, but not `racket/base`.

Contracts come in two forms: those constructed by the various operations listed in this section of the manual, and various ordinary Racket values that double as contracts, including

- symbols, booleans, characters, keywords, and `null`, which are treated as contracts that recognize themselves, using `eq?`,
- strings and byte strings, which are treated as contracts that recognize themselves using `equal?`,
- numbers, which are treated as contracts that recognize themselves using `=`,
- regular expressions, which are treated as contracts that recognize byte strings and strings that match the regular expression, and
- predicates: any procedure of arity 1 is treated as a predicate. During contract checking, it is applied to the values that appear and should return `#f` to indicate that the contract failed, and anything else to indicate it passed.

7.1 Data-structure Contracts

A *flat contract* can be fully checked immediately for a given value.

```
(flat-contract predicate) → flat-contract?
  predicate : (any/c . -> . any/c)
```

Constructs a flat contract from *predicate*. A value satisfies the contract if the predicate returns a true value.

```
(flat-named-contract type-name
                     predicate
                     [generator]) → flat-contract?
  type-name : any/c
  predicate : (or/c flat-contract? (any/c . -> . any))
  generator : (or/c #f (-> contract (-> int? any))) = #f
```

On predicates, behaves like flat-contract, but the first argument must be the (quoted) name of a contract used for error reporting. For example,

```
(flat-named-contract
 'odd-integer
 (lambda (x) (and (integer? x) (odd? x))))
```

turns the predicate into a contract with the name odd-integer.

On flat contracts, the new flat contract is the same as the old except for the name.

The generator argument adds a generator for the flat-named-contract. See contract-generate for more information.

any/c : flat-contract?

A flat contract that accepts any value.

When using this contract as the result portion of a function contract, consider using any instead; using any leads to better memory performance, but it also allows multiple results.

none/c : flat-contract?

A flat contract that accepts no values.

(or/c *contract* ...) → contract?
  contract : contract?

Takes any number of contracts and returns a contract that accepts any value that any one of the contracts accepts individually.

The or/c result tests any value by applying the contracts in order, from left to right, with the exception that it always moves the non-flat contracts (if any) to the end, checking them last. Thus, a contract such as (or/c (not/c real?) positive?) is guaranteed to only invoke the positive? predicate on real numbers.

If all of the arguments are procedures or flat contracts, the result is a flat contract. If only one of the arguments is a higher-order contract, the result is a contract that just checks the flat contracts and, if they don't pass, applies the higher-order contract.

If there are multiple higher-order contracts, or/c uses contract-first-order-passes? to distinguish between them. More precisely, when an or/c is checked, it first checks all of the flat contracts. If none of them pass, it calls contract-first-order-passes? with each of the higher-order contracts. If only one returns true, or/c uses that contract. If none of them return true, it signals a contract violation. If more than one returns true, it also signals a contract violation. For example, this contract

```
(or/c (-> number? number?)
      (-> string? string? string?))
```

does not accept a function like this one: `(lambda args ...)` since it cannot tell which of the two arrow contracts should be used with the function.

```
(and/c contract ...) → contract?
  contract : (or/c contract? (any/c . -> . any/c))
```

Takes any number of contracts and returns a contract that accepts any value that satisfies all of the contracts simultaneously.

If all of the arguments are procedures or flat contracts, the result is a flat contract.

The contract produced by `and/c` tests any value by applying the contracts in order, from left to right.

```
(not/c flat-contract) → flat-contract?
  flat-contract : (or/c flat-contract? (any/c . -> . any/c))
```

Accepts a flat contracts or a predicate and returns a flat contract that checks the inverse of the argument.

```
(=/c z) → flat-contract?
  z : real?
```

Returns a flat contract that requires the input to be a number and `=` to *z*.

```
(</c n) → flat-contract?
  n : real?
```

Returns a flat contract that requires the input to be a number and `<` than *n*.

```
(>/c n) → flat-contract?
  n : real?
```

Like `</c`, but for `>`.

```
(<=/c n) → flat-contract?
  n : real?
```

Like `</c`, but for `<=`.

```
(>=/c n) → flat-contract?
  n : real?
```

Like </c, but for >=.

```
(between/c n m) → flat-contract?
  n : real?
  m : real?
```

Returns a flat contract that requires the input to be a real number between *n* and *m* or equal to one of them.

```
(real-in n m) → flat-contract?
  n : real?
  m : real?
```

An alias for between/c.

```
(integer-in j k) → flat-contract?
  j : exact-integer?
  k : exact-integer?
```

Returns a flat contract that requires the input to be an exact integer between *j* and *k*, inclusive.

```
natural-number/c : flat-contract?
```

A flat contract that requires the input to be an exact non-negative integer.

```
(string-len/c len) → flat-contract?
  len : exact-nonnegative-integer?
```

Returns a flat contract that recognizes strings that have fewer than *len* characters.

```
false/c : flat-contract?
```

An alias #f for backwards compatibility.

```
printable/c : flat-contract?
```

A flat contract that recognizes values that can be written out and read back in with write and read.

```
(one-of/c v ...+) → flat-contract?
  v : any/c
```

Accepts any number of atomic values and returns a flat contract that recognizes those values, using eqv? as the comparison predicate. For the purposes of one-of/c, atomic values

are defined to be: characters, symbols, booleans, `null`, keywords, numbers, `#<void>`, and `#<undefined>`.

This is a backwards compatibility contract constructor. If neither `#<void>` nor `#<undefined>` are arguments, it simply passes its arguments to `or/c`.

```
(symbols sym ...+) → flat-contract?
  sym : symbol?
```

Accepts any number of symbols and returns a flat contract that recognizes those symbols.

This is a backwards compatibility constructor; it merely passes its arguments to `or/c`.

```
(vectorof  c
           [#:immutable immutable
            #:flat? flat?])        → contract?
  c : contract?
  immutable : (or/c #t #f 'dont-care) = 'dont-care
  flat? : boolean? = #f
```

Returns a contract that recognizes vectors. The elements of the vector must match *c*.

If the *flat?* argument is `#t`, then the resulting contract is a flat contract, and the *c* argument must also be a flat contract. Such flat contracts will be unsound if applied to mutable vectors, as they will not check future operations on the vector.

If the *immutable* argument is `#t` and the *c* argument is a flat contract, the result will be a flat contract. If the *c* argument is a chaperone contract, then the result will be a chaperone contract.

When a higher-order `vectorof` contract is applied to a vector, the result is not `eq?` to the input. The result will be a copy for immutable vectors and a chaperone or impersonator of the input for mutable vectors.

```
(vector-immutableof c) → contract?
  c : contract?
```

Returns the same contract as `(vectorof c #:immutable #t)`. This form exists for backwards compatibility.

```
(vector/c  c
           ...
           [#:immutable immutable
            #:flat? flat?])        → contract?
  c : contract?
  immutable : (or/c #t #f 'dont-care) = 'dont-care
  flat? : boolean? = #f
```

Returns a contract that recognizes vectors whose lengths match the number of contracts given. Each element of the vector must match its corresponding contract.

If the *flat?* argument is #t, then the resulting contract is a flat contract, and the *c* arguments must also be flat contracts. Such flat contracts will be unsound if applied to mutable vectors, as they will not check future operations on the vector.

If the *immutable* argument is #t and the *c* arguments are flat contracts, the result will be a flat contract. If the *c* arguments are chaperone contracts, then the result will be a chaperone contract.

When a higher-order vector/c contract is applied to a vector, the result is not eq? to the input. The result will be a copy for immutable vectors and a chaperone or impersonator of the input for mutable vectors.

```
(vector-immutable/c c ...) → contract?
  c : contract?
```

Returns the same contract as (vector/c c ... #:immutable #t). This form exists for reasons of backwards compatibility.

```
(box/c  c
        [#:immutable immutable
         #:flat? flat?])        → contract?
  c : contract?
  immutable : (or/c #t #f 'dont-care) = 'dont-care
  flat? : boolean? = #f
```

Returns a contract that recognizes boxes. The content of the box must match *c*.

If the *flat?* argument is #t, then the resulting contract is a flat contract, and the *c* argument must also be a flat contract. Such flat contracts will be unsound if applied to mutable boxes, as they will not check future operations on the box.

If the *immutable* argument is #t and the *c* argument is a flat contract, the result will be a flat contract. If the *c* argument is a chaperone contract, then the result will be a chaperone contract.

When a higher-order box/c contract is applied to a box, the result is not eq? to the input. The result will be a copy for immutable boxes and either a chaperone or impersonator of the input for mutable boxes.

```
(box-immutable/c c) → contract?
  c : contract?
```

Returns the same contract as (box/c c #:immutable #t). This form exists for reasons of backwards compatibility.

```
(listof c) → contract?
  c : (or/c contract? (any/c . -> . any/c))
```

Returns a contract that recognizes a list whose every element matches the contract `c`. Beware that when this contract is applied to a value, the result is not necessarily `eq?` to the input.

```
(non-empty-listof c) → contract?
  c : (or/c contract? (any/c . -> . any/c))
```

Returns a contract that recognizes non-empty lists whose elements match the contract `c`. Beware that when this contract is applied to a value, the result is not necessarily `eq?` to the input.

```
(cons/c car-c cdr-c) → contract?
  car-c : contract?
  cdr-c : contract?
```

Produces a contract that recognizes pairs whose first and second elements match *`car-c`* and *`cdr-c`*, respectively. Beware that when this contract is applied to a value, the result is not necessarily `eq?` to the input.

```
(list/c c ...) → contract?
  c : (or/c contract? (any/c . -> . any/c))
```

Produces a contract for a list. The number of elements in the list must match the number of arguments supplied to `list/c`, and each element of the list must match the corresponding contract. Beware that when this contract is applied to a value, the result is not necessarily `eq?` to the input.

```
(syntax/c c) → flat-contract?
  c : flat-contract?
```

Produces a flat contract that recognizes syntax objects whose `syntax-e` content matches `c`.

```
(struct/c struct-id contract-expr ...)
```

Produces a contract that recognizes instances of the structure type named by *`struct-id`*, and whose field values match the contracts produced by the *`contract-expr`*s.

Contracts for immutable fields must be either flat or chaperone contracts. Contracts for mutable fields may be impersonator contracts. If all fields are immutable and the *`contract-expr`*s evaluate to flat contracts, a flat contract is produced. If all the *`contract-expr`*s are chaperone contracts, a chaperone contract is produced. Otherwise, an impersonator contract is produced.

```
(struct/dc struct-id field-spec ...)

               field-spec = [field-id maybe-lazy contract-expr]
                          | [field-id (dep-field-id ...)
                                     maybe-lazy
                                     maybe-flat-or-impersonator
                                     maybe-dep-state
                                     contract-expr]

               maybe-lazy =
                          | #:lazy

maybe-flat-or-impersonator =
                          | #:flat
                          | #:impersonator

          maybe-dep-state =
                          | #:depends-on-state
```

Produces a contract that recognizes instances of the structure type named by *struct-id*, and whose field values match the contracts produced by the *field-spec*s.

If the *field-spec* lists the names of other fields, then the contract depends on values in those fields, and the *contract-expr* expression is evaluated each time a selector is applied, building a new contract for the fields based on the values of the *dep-field-id* fields. If the field is a dependent field, then it is assumed that the contract is a chaperone, but not always a flat contract (and theus the entire `struct/dc` contract is not a flat contract). If this is not the case, and the contract is always flat then the field must be annotated with the `#:flat`, or the field must be annotated with `#:chaperone` (in which case, it must be a mutable field).

If the `#:lazy` keyword appears, then the contract on the field is check lazily (only when a selector is applied); `#:lazy` contracts cannot be put on mutable fields.

If a dependent contract depends on some mutable state, then use the `#:depends-on-state` keyword argument (if a field's dependent contract depends on a mutable field, this keyword is automatically inferred). The presence of this keyword means that the contract expression is evaluated each time the corresponding field is accessed (or mutated, if it is a mutable field). Otherwise, the contract expression for a dependent field contract is evaluated when the contract is applied to a value.

Contracts for immutable fields must be either flat or chaperone contracts. Contracts for mutable fields may be impersonator contracts. If all fields are immutable and the *contract-expr*s evaluate to flat contracts, a flat contract is produced. If all the *contract-expr*s are chaperone contracts, a chaperone contract is produced. Otherwise, an impersonator contract is produced.

As an example, the function bst/c below returns a contract for binary search trees whose values are all between lo and hi. The lazy annotations ensure that this contract does not change the running time of operations that do not inspect the entire tree.

```
(struct bt (val left right))
(define (bst/c lo hi)
  (or/c #f
        (struct/dc bt
                   [val (between/c lo hi)]
                   [left (val) #:lazy (bst lo val)]
                   [right (val) #:lazy (bst val hi)])))
```

```
(parameter/c c) → contract?
  c : contract?
```

Produces a contract on parameters whose values must match contract.

```
(procedure-arity-includes/c n) → flat-contract?
  n : exact-nonnegative-integer?
```

Produces a contract for procedures that accept *n* argument (i.e,. the procedure? contract is implied).

```
(hash/c  key
         val
        [#:immutable immutable
         #:flat? flat?])        → contract?
  key : chaperone-contract?
  val : contract?
  immutable : (or/c #t #f 'dont-care) = 'dont-care
  flat? : boolean? = #f
```

Produces a contract that recognizes hash tables with keys and values as specified by the *key* and *val* arguments.

If the *flat?* argument is #t, then the resulting contract is a flat contract, and the *key* and *val* arguments must also be flat contracts. Such flat contracts will be unsound if applied to mutable hash tables, as they will not check future operations on the hash table.

If the *immutable* argument is #t and the *key* and *val* arguments are flat contracts, the result will be a flat contract. If either the domain or the range is a chaperone contract, then the result will be a chaperone contract.

If the *key* argument is a chaperone contract, then the resulting contract can only be applied to equal?-based hash tables. When a higher-order hash/c contract is applied to a hash table, the result is not eq? to the input. The result will be a copy for immutable hash tables, and either a chaperone or impersonator of the input for mutable hash tables.

```
(prompt-tag/c contract ...) → contract?
  contract : contract?
```

Takes any number of contracts and returns a contract that recognizes continuation prompt tags and will check any aborts or prompt handlers that use the contracted prompt tag.

Each *contract* will check the corresponding value passed to an abort-current-continuation and handled by the handler of a call to call-with-continuation-prompt.

If all of the *contract*s are chaperone contracts, the resulting contract will also be a chaperone contract. Otherwise, the contract is an impersonator contract.

Examples:

```
> (define/contract tag
    (prompt-tag/c (-> number? string?))
    (make-continuation-prompt-tag))

> (call-with-continuation-prompt
    (lambda ()
      (number->string
        (call-with-composable-continuation
          (lambda (k)
            (abort-current-continuation tag k)))))
    tag
    (lambda (k) (k "not a number")))
tag: contract violation
  expected: number?, given: "not a number"
  in: the 1st argument of
      ...
      (prompt-tag/c (-> number? string?))
  contract from: (definition tag)
  blaming: top-level
  at: eval:2.0
```

```
(continuation-mark-key/c contract) → contract?
  contract : contract?
```

Takes a single contract and returns a contract that recognizes continuation marks and will check any mappings of marks to values or any accesses of the mark value.

If the argument *contract* is a chaperone contract, the resulting contract will also be a chaperone contract. Otherwise, the contract is an impersonator contract.

Examples:

```
> (define/contract mark-key
    (continuation-mark-key/c (-> symbol? (listof symbol?)))
    (make-continuation-mark-key))

> (with-continuation-mark
    mark-key
    (lambda (s) (append s '(truffle fudge ganache)))
    (let ([mark-value (continuation-mark-set-first
                       (current-continuation-marks) mark-key)])
      (mark-value "chocolate-bar")))
mark-key: contract violation
  expected: symbol?, given: "chocolate-bar"
  in: the 1st argument of
      ...
      (continuation-mark-key/c
       (-> symbol? (listof symbol?)))
  contract from: (definition mark-key)
  blaming: top-level
  at: eval:2.0
```

(flat-rec-contract *id flat-contract-expr ...*)

Constructs a recursive flat contract. A *flat-contract-expr* can refer to *id* to refer recursively to the generated contract.

For example, the contract

```
(flat-rec-contract sexp
  (cons/c sexp sexp)
  number?
  symbol?)
```

is a flat contract that checks for (a limited form of) S-expressions. It says that a `sexp` is either two `sexp`s combined with `cons`, or a number, or a symbol.

Note that if the contract is applied to a circular value, contract checking will not terminate.

(flat-murec-contract ([*id flat-contract-expr ...*] ...) *body ...+*)

A generalization of `flat-rec-contract` for defining several mutually recursive flat contracts simultaneously. Each *id* is visible in the entire `flat-murec-contract` form, and the result of the final *body* is the result of the entire form.

any

Represents a contract that is always satisfied. In particular, it can accept multiple values. It can only be used in a result position of contracts like `->`. Using `any` elsewhere is a syntax error.

```
(promise/c expr)
```

Constructs a contract on a promise. The contract does not force the promise, but when the promise is forced, the contract checks that the result value meets the contract produced by *expr*.

7.2 Function Contracts

A *function contract* wraps a procedure to delay checks for its arguments and results. There are three primary function contract combinators that have increasing amounts of expressiveness and increasing additional overheads. The first `->` is the cheapest. It generates wrapper functions that can call the original function directly. Contracts built with `->*` require packaging up arguments as lists in the wrapper function and then using either `keyword-apply` or `apply`. Finally, `->i` is the most expensive (along with `->d`), because it requires delaying the evaluation of the contract expressions for the domain and range until the function itself is called or returns.

The `case->` contract is a specialized contract, designed to match `case-lambda` and `unconstrained-domain->` allows range checking without requiring that the domain have any particular shape (see below for an example use).

```
(-> dom ... range)

  dom = dom-expr
      | keyword dom-expr

range = range-expr
      | (values range-expr ...)
      | any
```

Produces a contract for a function that accepts a fixed number of arguments and returns either a fixed number of results or completely unspecified results (the latter when `any` is specified).

Each *dom-expr* is a contract on an argument to a function, and each *range-expr* is a contract on a result of the function.

For example,

```
(integer? boolean? . -> . integer?)
```

produces a contract on functions of two arguments. The first argument must be an integer, and the second argument must be a boolean. The function must produce an integer.

Using a `->` between two whitespace-delimited `.`s is the same as putting the `->` right after the enclosing opening parenthesis. See §2.4.3 “Lists and Racket Syntax” or §1.3.5 “Reading Pairs and Lists” for more information.

A domain specification may include a keyword. If so, the function must accept corresponding (mandatory) keyword arguments, and the values for the keyword arguments must match the corresponding contracts. For example:

```
(integer? #:x boolean? . -> . integer?)
```

is a contract on a function that accepts a by-position argument that is an integer and a `#:x` argument that is a boolean.

If `any` is used as the last sub-form for `->`, no contract checking is performed on the result of the function, and thus any number of values is legal (even different numbers on different invocations of the function).

If `(values range-expr ...)` is used as the last sub-form of `->`, the function must produce a result for each contract, and each value must match its respective contract.

```
(->* (mandatory-dom ...) optional-doms rest pre range post)

mandatory-dom = dom-expr
              | keyword dom-expr

optional-doms =
              | (optional-dom ...)

 optional-dom = dom-expr
              | keyword dom-expr

         rest =
              | #:rest rest-expr

          pre =
              | #:pre pre-cond-expr

        range = range-expr
              | (values range-expr ...)
              | any

         post =
              | #:post post-cond-expr
```

The `->*` contract combinator produces contracts for functions that accept optional arguments (either keyword or positional) and/or arbitrarily many arguments. The first clause of a `->*` contract describes the mandatory arguments, and is similar to the argument description of a `->` contract. The second clause describes the optional arguments. The range of description can either be `any` or a sequence of contracts, indicating that the function must return multiple values.

If present, the *rest-expr* contract governs the arguments in the rest parameter. Note that the *rest-expr* contract governs only the arguments in the rest parameter, not those in mandatory arguments. For example, this contract:

```
(->* () #:rest (cons/c integer? (listof integer?)) any)
```

does not match the function

```
(λ (x . rest) x)
```

because the contract insists that the function accept zero arguments (because there are no mandatory arguments listed in the contract). The `->i` contract does not know that the contract on the rest argument is going to end up disallowing empty argument lists.

The *pre-cond-expr* and *post-cond-expr* expressions are checked as the function is called and returns, respectively, and allow checking of the environment without an explicit connection to an argument (or a result).

As an example, the contract

```
(->* () (boolean? #:x integer?) #:rest (listof symbol?) symbol?)
```

matches functions that optionally accept a boolean, an integer keyword argument `#:x` and arbitrarily more symbols, and that return a symbol.

```
(->i (mandatory-dependent-dom ...)
     dependent-rest
     pre-condition
     dep-range
     post-condition)
(->i (mandatory-dependent-dom ...)
     (optional-dependent-dom ...)
     dependent-rest
     pre-condition
     dep-range
     post-condition)
```

```
mandatory-dependent-dom = id+ctc
                        | keyword id+ctc

 optional-dependent-dom = id+ctc
                        | keyword id+ctc

         dependent-rest =
                        | #:rest id+ctc

          pre-condition =
                        | #:pre (id ...) boolean-expr pre-condition
                        | #:pre/name (id ...) string boolean-expr pre-condition

        dependent-range = any
                        | id+ctc
                        | un+ctc
                        | (values id+ctc ...)
                        | (values un+ctc ...)

         post-condition =
                        | #:post (id ...) boolean-expr post-condition
                        | #:post/name (id ...) string boolean-expr post-condition

                id+ctc = [id contract-expr]
                        | [id (id ...) contract-expr]

                un+ctc = [_ contract-expr]
                        | [_ (id ...) contract-expr]
```

The `->i` contract combinator differs from the `->*` combinator in that the support pre- and post-condition clauses and in that each argument and result is named. These names can then be used in the subcontracts and in the pre-/post-condition clauses. In short, contracts now express dependencies among arguments and results.

The first sub-form of a `->i` contract covers the mandatory and the second sub-form covers the optional arguments. Following that is an optional rest-args contract, and an optional pre-condition. The pre-condition is introduced with the `#:pre` keyword followed by the list of names on which it depends. If the `#:pre/name` keyword is used, the string supplied is used as part of the error message; similarly with `#:post/name`.

The *dep-range* non-terminal specifies the possible result contracts. If it is `any`, then any value is allowed. Otherwise, the result contract pairs a name and a contract or a multiple values return with names and contracts. In the last two cases, the range contract may be optionally followed by a post-condition; the post-condition expression is not allowed if the range contract is `any`. Like the pre-condition, the post-condition must specify the variables on which it depends.

Consider this sample contract:

```
(->i ([x number?]
      [y (x) (>=/c x)])
     [result (x y) (and/c number? (>=/c (+ x y)))])
```

It specifies a function of two arguments, both numbers. The contract on the second argument (`y`) demands that it is greater than the first argument. The result contract promises a number that is greater than the sum of the two arguments. While the dependency specification for `y` signals that the argument contract depends on the value of the first argument, the dependency sequence for `result` indicates that the contract depends on both argument values. Since the contract for `x` does not depend on anything else, it does not come with any dependency sequence, not even `()`.

In general, an empty sequence is (nearly) equivalent to not adding a sequence at all except that the former is more expensive than the latter.

The contract expressions are not always evaluated in order. First, if there is no dependency for a given contract expression, the contract expression is evaluated at the time that the `->i` expression is evaluated rather than the time when the function is called or returns. These dependency-free contract expressions are evaluated in the order in which they are listed. Second, the dependent contract sub-expressions are evaluated when the contracted function is called or returns in some order that satisfies the dependencies. That is, if a contract for an argument depends on the value of some other contract, the former is evaluated first (so that the argument, with its contract checked, is available for the other). When there is no dependency between two arguments (or the result and an argument), then the contract that appears earlier in the source text is evaluated first.

Finally, if all of the identifier positions of the range contract are `_` s (underscores), then the range contract expressions are evaluated when the function is called and the underscore is not bound in the range, after the argument contracts are evaluated and checked. Otherwise, the range expressions are evaluated when the function returns.

If there are optional arguments that are not supplied, then the corresponding variables will be bound to a special value called `the-unsupplied-arg` value.

```
(->d (mandatory-dependent-dom ...)
     dependent-rest
     pre-condition
     dependent-range
     post-condition)
(->d (mandatory-dependent-dom ...)
     (optional-dependent-dom ...)
     dependent-rest
     pre-condition
     dependent-range
     post-condition)
```

```
mandatory-dependent-dom = [id dom-expr]
                        | keyword [id dom-expr]

 optional-dependent-dom = [id dom-expr]
                        | keyword [id dom-expr]

         dependent-rest =
                        | #:rest id rest-expr

          pre-condition =
                        | #:pre boolean-expr
                        | #:pre-cond boolean-expr

        dependent-range = any
                        | [_ range-expr]
                        | (values [_ range-expr] ...)
                        | [id range-expr]
                        | (values [id range-expr] ...)

         post-condition =
                        | #:post-cond boolean-expr
```

This contract is here for backwards compatibility; any new code should use `->i` instead.

This contract is similar to `->i`, but is "lax", meaning that it does not enforce contracts internally. For example, using this contract

```
(->d ([f (-> integer? integer?)])
     #:pre
     (zero? (f #f))
     any)
```

will allow `f` to be called with `#f`, trigger whatever bad behavior the author of `f` was trying to prohibit by insisting that `f`'s contract accept only integers.

The `#:pre-cond` and `#:post-cond` keywords are aliases for `#:pre` and `#:post` and are provided for backwards compatibility.

```
(case-> (-> dom-expr ... rest range) ...)

 rest =
      | #:rest rest-expr

range = range-expr
      | (values range-expr ...)
      | any
```

This contract form is designed to match `case-lambda`. Each argument to `case->` is a contract that governs a clause in the `case-lambda`. If the `#:rest` keyword is present, the corresponding clause must accept an arbitrary number of arguments. The *range* specification is just like that for `->` and `->*`.

(unconstrained-domain-> *range-expr* ...)

Constructs a contract that accepts a function, but makes no constraint on the function's domain. The *range-expr*s determine the number of results and the contract for each result.

Generally, this contract must be combined with another contract to ensure that the domain is actually known to be able to safely call the function itself.

For example, the contract

```
(provide
 (contract-out
  [f (->d ([size natural-number/c]
           [proc (and/c (unconstrained-domain-> number?)
                        (lambda (p)
                          (procedure-arity-includes? p size)))])
          ()
          [_ number?])]))
```

says that the function `f` accepts a natural number and a function. The domain of the function that `f` accepts must include a case for `size` arguments, meaning that `f` can safely supply `size` arguments to its input.

For example, the following is a definition of `f` that cannot be blamed using the above contract:

```
(define (f i g)
  (apply g (build-list i add1)))
```

predicate/c : contract?

Use this contract to indicate that some function is a predicate. It is semantically equivalent to `(-> any/c boolean?)`.

This contract also includes an optimization so that functions returning `#t` from `struct-predicate-procedure?` are just returned directly, without being wrapped. This contract is used by `provide/contract`'s `struct` subform so that struct predicates end up not being wrapped.

the-unsupplied-arg : unsupplied-arg?

Used by `->i` (and `->d`) to bind optional arguments that are not supplied by a call site.

```
(unsupplied-arg? v) → boolean?
  v : any/c
```

A predicate to determine whether *v* is `the-unsupplied-arg`.

7.3 Parametric Contracts

```
(require racket/contract/parametric)
```

The most convenient way to use parametric contract is to use `contract-out`'s `#:exists` keyword. The `racket/contract/parametric` provides a few more, general-purpose parametric contracts.

```
(parametric->/c (x ...) c)
```

Creates a contract for parametric polymorphic functions. Each function is protected by *c*, where each *x* is bound in *c* and refers to a polymorphic type that is instantiated each time the function is applied.

At each application of a function, the `parametric->/c` contract constructs a new opaque wrapper for each *x*; values flowing into the polymorphic function (i.e. values protected by some *x* in negative position with respect to `parametric->/c`) are wrapped in the corresponding opaque wrapper. Values flowing out of the polymorphic function (i.e. values protected by some *x* in positive position with respect to `parametric->/c`) are checked for the appropriate wrapper. If they have it, they are unwrapped; if they do not, a contract violation is signaled.

Examples:

```
> (define/contract (check x y)
    (parametric->/c [X] (boolean? X . -> . X))
    (if (or (not x) (equal? y 'surprise))
        'invalid
        y))

> (check #t 'ok)
'ok
> (check #f 'ignored)
check: broke its contract
  promised a(n) X; produced: 'invalid
  in: the range of
      ...
```

```
      (parametric->/c (X) ...)
  contract from: (function check)
  blaming: (function check)
  at: eval:2.0
> (check #t 'surprise)
'surprise
```

```
(new-∀/c name) → contract?
  name : symbol?
```

Constructs a new universal contract.

Universal contracts accept all values when in negative positions (e.g., function inputs) and wrap them in an opaque struct, hiding the precise value. In positive positions (e.g. function returns), a universal contract accepts only values that were previously accepted in negative positions (by checking for the wrappers).

The name is used to identify the contract in error messages.

For example, this contract:

```
(let ([a (new-∀/c 'a)])
  (-> a a))
```

describes the identity function (or a non-terminating function) That is, the first use of the `a` appears in a negative position and thus inputs to that function are wrapped with an opaque struct. Then, when the function returns, it is checked to determine whether the result is wrapped, since the second `a` appears in a positive position.

The `new-∀/c` construct constructor is dual to `new-∃/c`.

```
(new-∃/c name) → contract?
  name : symbol?
```

Constructs a new existential contract.

Existential contracts accept all values when in positive positions (e.g., function returns) and wrap them in an opaque struct, hiding the precise value. In negative positions (e.g. function inputs), they accepts only values that were previously accepted in positive positions (by checking for the wrappers).

The name is used to identify the contract in error messages.

For example, this contract:

```
(let ([a (new-∃/c 'a)])
  (-> (-> a a)
      any/c))
```

describes a function that accepts the identity function (or a non-terminating function) and returns an arbitrary value. That is, the first use of the `a` appears in a positive position and thus inputs to that function are wrapped with an opaque struct. Then, when the function returns, it is checked to see if the result is wrapped, since the second `a` appears in a negative position.

The `new-∃/c` construct constructor is dual to `new-∀/c`.

7.4 Lazy Data-structure Contracts

```
(contract-struct id (field-id ...))
```

> **NOTE:** This library is deprecated. Use `struct` instead. Lazy struct contracts no longer require a separate struct declaration; instead `struct/dc` and `struct/c` work directly with `struct` and `define-struct`.

Like `struct`, but with two differences: they do not define field mutators, and they define two contract constructors: *`id`*`/c` and *`id`*`/dc`. The first is a procedure that accepts as many arguments as there are fields and returns a contract for struct values whose fields match the arguments. The second is a syntactic form that also produces contracts on the structs, but the contracts on later fields may depend on the values of earlier fields.

The generated contract combinators are *lazy*: they only verify the contract holds for the portion of some data structure that is actually inspected. More precisely, a lazy data structure contract is not checked until a selector extracts a field of a struct.

```
(id/dc field-spec ...)

field-spec = [field-id contract-expr]
           | [field-id (field-id ...) contract-expr]
```

In each *`field-spec`* case, the first *`field-id`* specifies which field the contract applies to; the fields must be specified in the same order as the original `contract-struct`. The first case is for when the contract on the field does not depend on the value of any other field. The second case is for when the contract on the field does depend on some other fields, and the parenthesized *`field-id`*s indicate which fields it depends on; these dependencies can only be to earlier fields.

```
(define-contract-struct id (field-id ...))
```

> **NOTE:** This library is deprecated. Use `struct` instead. Lazy struct contracts no longer require a separate struct declaration; instead `struct/dc` and `struct/c` work directly with `struct` and `define-struct`.

Like `contract-struct`, but where the constructor's name is `make-id`, much like `define-struct`.

7.5 Structure Type Property Contracts

```
(struct-type-property/c value-contract) → contract?
  value-contract : contract?
```

Produces a contract for a structure type property. When the contract is applied to a struct type property, it produces a wrapped struct type property that applies *value-contract* to the value associated with the property when it used to create a new struct type (via `struct`, `make-struct-type`, etc).

The struct type property's accessor function is not affected; if it is exported, it must be protected separately.

As an example, consider the following module. It creates a structure type property, `prop`, whose value should be a function mapping a structure instance to a numeric predicate. The module also exports `app-prop`, which extracts the predicate from a structure instance and applies it to a given value.

```
> (module propmod racket
    (require racket/contract)
    (define-values (prop prop? prop-ref)
      (make-struct-type-property 'prop))
    (define (app-prop x v)
      (((prop-ref x) x) v))
    (provide/contract
     [prop? (-> any/c boolean?)]
     [prop (struct-type-property/c
            (-> prop? (-> integer? boolean?)))]
     [app-prop (-> prop? integer? boolean?)])
    (provide prop-ref))
```

The `structmod` module creates a structure type named `s` with a single field; the value of `prop` is a function that extracts the field value from an instance. Thus the field ought to

be an integer predicate, but notice that `structmod` places no contract on `s` enforcing that constraint.

```
> (module structmod racket
    (require 'propmod)
    (struct s (f) #:property prop (lambda (s) (s-f s)))
    (provide (struct-out s)))

> (require 'propmod 'structmod)
```

First we create an `s` instance with a integer predicate, so the constraint on `prop` is in fact satisfied. The first call to `app-prop` is correct; the second simply violates the contract of `app-prop`.

```
> (define s1 (s even?))

> (app-prop s1 5)
#f
> (app-prop s1 'apple)
app-prop: contract violation
  expected: integer?, given: 'apple
  in: the 2nd argument of
      (-> prop? integer? boolean?)
  contract from: propmod
  blaming: top-level
  at: eval:2.0
```

We are able to create `s` instances with values other than integer predicates, but applying `app-prop` on them blames `structmod`, because the function associated with `prop`—that is, `(lambda (s) (s-f s))`—does not always produce a value satisfying `(-> integer? boolean?)`.

```
> (define s2 (s "not a fun"))

> (app-prop s2 5)
prop: contract violation
  expected a procedure that accepts 1 mandatory argument
without any keywords, given: "not a fun"
  in: the range of
      ...
      (struct-type-property/c
        (-> prop? (-> integer? boolean?)))
  contract from: propmod
  blaming: structmod
```

```
  at: eval:2.0
> (define s3 (s list))

> (app-prop s3 5)
prop: contract violation
  expected: boolean?, given: '(5)
  in: the range of
      the range of
      ...
      (struct-type-property/c
       (-> prop? (-> integer? boolean?)))
  contract from: propmod
  blaming: structmod
  at: eval:2.0
```

The fix would be to propagate the obligation inherited from `prop` to `s`:

```
(provide (contract-out
          [struct s ([f (-> integer? boolean?)])]))
```

Finally, if we directly apply the property accessor, `prop-ref`, and then misuse the resulting function, the `propmod` module is blamed:

```
> ((prop-ref s3) 'apple)
prop: broke its contract
  promised: prop?, produced: 'apple
  in: the 1st argument of
      ...
      (struct-type-property/c
       (-> prop? (-> integer? boolean?)))
  contract from: propmod
  blaming: propmod
  at: eval:2.0
```

The `propmod` module has an obligation to ensure a function associated with `prop` is applied only to values satisfying `prop?`. By directly providing `prop-ref`, it enables that constraint to be violated (and thus it is blamed), even though the bad application actually occurs elsewhere.

Generally there is no need to provide a structure type property accessor at all; it is typically only used by other functions within the module. But if it must be provided, it should be protected thus:

```
(provide (contract-out
          [prop-ref (-> prop? (-> prop? (-> integer? boolean?)))]))
```

7.6 Attaching Contracts to Values

```
(contract-out p/c-item ...)

        p/c-item = (struct id ((id contract-expr) ...))
                 | (struct (id identifier) ((id contract-expr) ...))
                 | (rename orig-id id contract-expr)
                 | (id contract-expr)
                 | #:∃ exists-variables
                 | #:exists exists-variables

exists-variables = identifier
                 | (identifier ...)
```

A *provide-spec* for use in `provide` (currently only for the same phase level as the `provide` form; for example, `contract-out` cannot be nested within `for-syntax`). Each *id* is provided from the module. In addition, clients of the module must live up to the contract specified by *contract-expr* for each export.

The `contract-out` form treats modules as units of blame. The module that defines the provided variable is expected to meet the positive (co-variant) positions of the contract. Each module that imports the provided variable must obey the negative (contra-variant) positions of the contract. Each *contract-expr* in a `contract-out` form is effectively moved to the end of the enclosing module, so a *contract-expr* can refer to variables that are defined later in the same module.

Only uses of the contracted variable outside the module are checked. Inside the module, no contract checking occurs.

The `rename` form of `contract-out` exports the first variable (the internal name) with the name specified by the second variable (the external name).

The `struct` form of `contract-out` provides a structure-type definition, and each field has a contract that dictates the contents of the fields. The structure-type definition must appear before the `provide` clause within the enclosing module. If the structure type has a parent, the second `struct` form (above) must be used, with the first name referring to the structure type to export and the second name referring to the parent structure type. Unlike a `struct` definition, however, all of the fields (and their contracts) must be listed. The contract on the fields that the sub-struct shares with its parent are only used in the contract for the sub-struct's constructor, and the selector or mutators for the super-struct are not provided. The exported structure-type name always doubles as a constructor, even if the original structure-type name does not act as a constructor.

The `#:∃` and `#:exists` clauses define new abstract contracts. The variables are bound in the remainder of the `contract-out` form to new contracts that hide the values they accept and ensure that the exported functions are treated parametrically.

The implementation of `contract-out` attaches uses `syntax-property` to attach properties to the code it generates that records the syntax of the contracts in the fully expanded program. Specifically, the symbol `'provide/contract-original-contract` is bound to vectors of two elements, the exported identifier and a syntax object for the expression that produces the contract controlling the export.

```
(provide/contract p/c-item ...)
```

A legacy shorthand for `(provide (contract-out p/c-item ...))`, except that a *contract-expr* within `provide/contract` is evaluated at the position of the `provide/contract` form instead of at the end of the enclosing module.

7.6.1 Nested Contract Boundaries

```
(require racket/contract/region)
```

```
(with-contract blame-id (wc-export ...) free-var-list ... body ...+)
(with-contract blame-id results-spec free-var-list ... body ...+)

    wc-export = (id contract-expr)

  result-spec = #:result contract-expr
              | #:results (contract-expr ...)

free-var-list =
              | #:freevar id contract-expr
              | #:freevars ([id contract-expr] ...)
```

Generates a local contract boundary.

The first `with-contract` form cannot appear in expression position. All names defined within the first `with-contract` form are visible externally, but those names listed in the *wc-export* list are protected with the corresponding contract. The *body* of the form allows definition/expression interleaving if its context does.

The second `with-contract` form must appear in expression position. The final *body* expression should return the same number of values as the number of contracts listed in the *result-spec*, and each returned value is contracted with its respective contract. The sequence of *body* forms is treated as for `let`.

The *blame-id* is used for the positive positions of contracts paired with exported *id*s. Contracts broken within the `with-contract` *body* will use the *blame-id* for their negative position.

If a free-var-list is given, then any uses of the free variables inside the *body* will be protected

with contracts that blame the context of the `with-contract` form for the positive positions and the `with-contract` form for the negative ones.

```
(define/contract id contract-expr free-var-list init-value-expr)
(define/contract (head args) contract-expr free-var-list body ...+)
```

Works like `define`, except that the contract *contract-expr* is attached to the bound value. For the definition of *head* and *args*, see `define`. For the definition of *free-var-list*, see `with-contract`.

The `define/contract` form treats the individual definition as a contract region. The definition itself is responsible for positive (co-variant) positions of the contract, and references to *id* outside of the definition must meet the negative positions of the contract. Since the contract boundary is between the definition and the surrounding context, references to *id* inside the `define/contract` form are not checked.

If a free-var-list is given, then any uses of the free variables inside the *body* will be protected with contracts that blame the context of the `define/contract` form for the positive positions and the `define/contract` form for the negative ones.

```
(define-struct/contract struct-id ([field contract-expr] ...)
                        struct-option ...)
(define-struct/contract (struct-id super-struct-id)
                        ([field contract-expr] ...)
                        struct-option ...)
```

Works like `define-struct`, except that the arguments to the constructor, accessors, and mutators are protected by contracts. For the definitions of *field* and *struct-option*, see `define-struct`.

The `define-struct/contract` form only allows a subset of the *struct-option* keywords: `#:mutable`, `#:transparent`, `#:auto-value`, `#:omit-define-syntaxes`, `#:property` and `#:omit-define-values`.

Examples:

```
> (define-struct/contract fish ([color number?]))

> (make-fish 5)
#<fish>
> (make-fish #f)
make-fish: contract violation
  expected: number?, given: #f
  in: the 1st argument of
      (-> number? symbol? any)
  contract from: (struct fish)
```

```
  blaming: top-level
> (define-struct/contract (salmon fish) ([ocean symbol?]))

> (make-salmon 5 'atlantic)
#<salmon>
> (make-salmon 5 #f)
make-salmon: contract violation
  expected: symbol?, given: #f
  in: the 2nd argument of
      (-> any/c symbol? symbol? any)
  contract from: (struct salmon)
  blaming: top-level
> (make-salmon #f 'pacific)
make-fish: contract violation
  expected: number?, given: #f
  in: the 1st argument of
      (-> number? symbol? any)
  contract from: (struct fish)
  blaming: top-level
```

```
current-contract-region
```

Bound by `define-syntax-parameter`, this contains information about the current contract region, used by the above forms to determine the candidates for blame assignment.

7.6.2 Low-level Contract Boundaries

```
(contract contract-expr to-protect-expr
          positive-blame-expr negative-blame-expr)
(contract contract-expr to-protect-expr
          positive-blame-expr negative-blame-expr
          value-name-expr source-location-expr)
```

The primitive mechanism for attaching a contract to a value. The purpose of `contract` is as a target for the expansion of some higher-level contract specifying form.

The `contract` expression adds the contract specified by *`contract-expr`* to the value produced by *`to-protect-expr`*. The result of a `contract` expression is the result of the *`to-protect-expr`* expression, but with the contract specified by *`contract-expr`* enforced on *`to-protect-expr`*.

The values of *`positive-blame-expr`* and *`negative-blame-expr`* indicate how to assign blame for positive and negative positions of the contract specified by *`contract-expr`*. They may be any value, and are formatted as by `display` for purposes of contract violation error messages.

If specified, *value-name-expr* indicates a name for the protected value to be used in error messages. If not supplied, or if *value-name-expr* produces #f, no name is printed. Otherwise, it is also formatted as by display.

If specified, *source-location-expr* indicates the source location reported by contract violations. The expression must produce a srcloc structure, syntax object, #f, or a list or vector in the format accepted by the third argument to datum->syntax.

7.7 Building New Contract Combinators

```
(require racket/contract/combinator)
```

Contracts are represented internally as functions that accept information about the contract (who is to blame, source locations, etc.) and produce projections (in the spirit of Dana Scott) that enforce the contract. A projection is a function that accepts an arbitrary value, and returns a value that satisfies the corresponding contract. For example, a projection that accepts only integers corresponds to the contract (flat-contract integer?), and can be written like this:

```
(define int-proj
  (lambda (x)
    (if (integer? x)
        x
        (signal-contract-violation))))
```

As a second example, a projection that accepts unary functions on integers looks like this:

```
(define int->int-proj
  (lambda (f)
    (if (and (procedure? f)
             (procedure-arity-includes? f 1))
        (lambda (x)
          (int-proj (f (int-proj x))))
        (signal-contract-violation))))
```

Although these projections have the right error behavior, they are not quite ready for use as contracts, because they do not accommodate blame, and do not provide good error messages. In order to accommodate these, contracts do not just use simple projections, but use functions that accept a *blame object* encapsulating the names of two parties that are the candidates for blame, as well as a record of the source location where the contract was established and the name of the contract. They can then, in turn, pass that information to raise-blame-error to signal a good error message.

Here is the first of those two projections, rewritten for use in the contract system:

```
(define (int-proj blame)
  (lambda (x)
    (if (integer? x)
        x
        (raise-blame-error
         blame
         val
         '(expected "<integer>," given: "~e")
         val))))
```

The new argument specifies who is to be blamed for positive and negative contract violations.

Contracts, in this system, are always established between two parties. One party provides some value according to the contract, and the other consumes the value, also according to the contract. The first is called the "positive" person and the second the "negative". So, in the case of just the integer contract, the only thing that can go wrong is that the value provided is not an integer. Thus, only the positive party can ever accrue blame. The `raise-blame-error` function always blames the positive party.

Compare that to the projection for our function contract:

```
(define (int->int-proj blame)
  (let ([dom (int-proj (blame-swap blame))]
        [rng (int-proj blame)])
    (lambda (f)
      (if (and (procedure? f)
               (procedure-arity-includes? f 1))
          (lambda (x)
            (rng (f (dom x))))
          (raise-blame-error
           blame
           val
           '(expected "a procedure of one argument," given: "~e")
           val)))))
```

In this case, the only explicit blame covers the situation where either a non-procedure is supplied to the contract or the procedure does not accept one argument. As with the integer projection, the blame here also lies with the producer of the value, which is why `raise-blame-error` is passed `blame` unchanged.

The checking for the domain and range are delegated to the `int-proj` function, which is supplied its arguments in the first two lines of the `int->int-proj` function. The trick here is that, even though the `int->int-proj` function always blames what it sees as positive, we can swap the blame parties by calling `blame-swap` on the given blame object, replacing the positive party with the negative party and vice versa.

This technique is not merely a cheap trick to get the example to work, however. The reversal of the positive and the negative is a natural consequence of the way functions behave. That is, imagine the flow of values in a program between two modules. First, one module defines a function, and then that module is required by another. So, far the function itself has to go from the original, providing module to the requiring module. Now, imagine that the providing module invokes the function, supplying it an argument. At this point, the flow of values reverses. The argument is traveling back from the requiring module to the providing module! And finally, when the function produces a result, that result flows back in the original direction. Accordingly, the contract on the domain reverses the positive and the negative blame parties, just like the flow of values reverses.

We can use this insight to generalize the function contracts and build a function that accepts any two contracts and returns a contract for functions between them.

This projection also goes further and uses `blame-add-context` to improve the error messages when a contract violation is detected.

```
(define (make-simple-function-contract dom-proj range-proj)
  (lambda (blame)
    (let ([dom (dom-proj (blame-add-context blame "the argument
of" #:swap? #t))]
          [rng (range-proj (blame-add-context blame "the range
of"))])
      (lambda (f)
        (if (and (procedure? f)
                 (procedure-arity-includes? f 1))
            (lambda (x)
              (rng (f (dom x))))
            (raise-blame-error
             blame
             val
             '(expected "a procedure of one argu-
ment," given: "~e")
             val))))))
```

Projections like the ones described above, but suited to other, new kinds of value you might make, can be used with the contract library primitives below.

```
(make-contract [#:name name
                #:first-order test
                #:projection proj]
                #:stronger stronger) → contract?
  name : any/c = 'anonymous-contract
  test : (-> any/c any/c) = (λ (x) #t)
  proj : (-> blame? (-> any/c any/c))
         (λ (b)
           (λ (x)
             (if (test x)
               x
       =       (raise-blame-error
                b x
                '(expected "<∼a>," given: "∼e")
                name x))))
  stronger : (-> contract? contract? boolean?)
(make-chaperone-contract [#:name name
                          #:first-order test
                          #:projection proj]
                          #:stronger stronger)
 → chaperone-contract?
  name : any/c = 'anonymous-chaperone-contract
  test : (-> any/c any/c) = (λ (x) #t)
  proj : (-> blame? (-> any/c any/c))
         (λ (b)
           (λ (x)
             (if (test x)
               x
       =       (raise-blame-error
                b x
                '(expected "<∼a>," given: "∼e")
                name x))))
  stronger : (-> contract? contract? boolean?)
```

```
(make-flat-contract [#:name name
                     #:first-order test
                     #:projection proj]
                     #:stronger stronger) → flat-contract?
  name : any/c = 'anonymous-flat-contract
  test : (-> any/c any/c) = (λ (x) #t)
  proj : (-> blame? (-> any/c any/c))
         (λ (b)
           (λ (x)
             (if (test x)
               x
       =       (raise-blame-error
                b x
                '(expected "<∼a>," given: "∼e")
                name x))))
  stronger : (-> contract? contract? boolean?)
```

These functions build simple higher-order contracts, chaperone contracts, and flat contracts, respectively. They both take the same set of three optional arguments: a name, a first-order predicate, and a blame-tracking projection.

The *name* argument is any value to be rendered using `display` to describe the contract when a violation occurs. The default name for simple higher-order contracts is `anonymous-contract`, for chaperone contracts is `anonymous-chaperone-contract`, and for flat contracts is `anonymous-flat-contract`.

The first-order predicate *test* can be used to determine which values the contract applies to; usually, this is the set of values for which the contract fails immediately without any higher-order wrapping. This test is used by `contract-first-order-passes?`, and indirectly by `or/c` to determine which of multiple higher-order contracts to wrap a value with. The default test accepts any value.

The projection *proj* defines the behavior of applying the contract. It is a curried function of two arguments: the first application accepts a blame object, and the second accepts a value to protect with the contract. The projection must either produce the value, suitably wrapped to enforce any higher-order aspects of the contract, or signal a contract violation using `raise-blame-error`. The default projection produces an error when the first-order test fails, and produces the value unchanged otherwise.

Projections for chaperone contracts must produce a value that passes `chaperone-of?` when compared with the original, uncontracted value. Projections for flat contracts must fail precisely when the first-order test does, and must produce the input value unchanged otherwise. Applying a flat contract may result in either an application of the predicate, or the projection, or both; therefore, the two must be consistent. The existence of a separate projection only serves to provide more specific error messages. Most flat contracts do not need to supply an explicit projection.

The *stronger* argument is used to implement contract-stronger?. The first argument is always the contract itself and the second argument is whatever was passed as the second argument to contract-stronger?.

Examples:

```
(define int/c
  (make-flat-contract #:name 'int/c #:first-order integer?))

> (contract int/c 1 'positive 'negative)
1
> (contract int/c "not one" 'positive 'negative)
eval:4:0: broke its contract
  promised: int/c, produced: "not one"
  in: int/c
  contract from: positive
  blaming: positive
> (int/c 1)
#t
> (int/c "not one")
#f
(define int->int/c
  (make-contract
   #:name 'int->int/c
   #:first-order
   (λ (x) (and (procedure? x) (procedure-arity-includes? x 1)))
   #:projection
   (λ (b)
     (let ([domain ((contract-projection int/c) (blame-swap b))]
           [range ((contract-projection int/c) b)])
       (λ (f)
         (if (and (procedure? f) (procedure-arity-includes? f 1))
           (λ (x) (range (f (domain x))))
           (raise-blame-error
            b f
            '(expected "a function of one argu-
ment," 'given: "~e")
            f)))))))

> (contract int->int/c "not fun" 'positive 'negative)
regexp-match: contract violation
  expected: (or/c string? bytes? path? input-port?)
  given: ''given:
  argument position: 2nd
  other arguments...:
   #rx"^ "
```

```
(define halve
  (contract int->int/c (λ (x) (/ x 2)) 'positive 'negative))

> (halve 2)
1
> (halve 1/2)
halve: contract violation
  expected: int/c, given: 1/2
  in: ...
      int->int/c
  contract from: positive
  blaming: negative
> (halve 1)
halve: broke its contract
  promised: int/c, produced: 1/2
  in: ...
      int->int/c
  contract from: positive
  blaming: positive
```

```
(build-compound-type-name c/s ...) → any
  c/s : any/c
```

Produces an S-expression to be used as a name for a contract. The arguments should be either contracts or symbols. It wraps parentheses around its arguments and extracts the names from any contracts it is supplied with.

```
(coerce-contract id x) → contract?
  id : symbol?
  x : any/c
```

Converts a regular Racket value into an instance of a contract struct, converting it according to the description of contracts.

If `x` is not one of the coercible values, `coerce-contract` signals an error, using the first argument in the error message.

```
(coerce-contracts id xs) → (listof contract?)
  id : symbol?
  xs : (listof any/c)
```

Coerces all of the arguments in 'xs' into contracts (via `coerce-contract/f`) and signals an error if any of them are not contracts. The error messages assume that the function named by `id` got `xs` as its entire argument list.

```
(coerce-chaperone-contract id x) → chaperone-contract?
  id : symbol?
  x : any/c
```

Like coerce-contract, but requires the result to be a chaperone contract, not an arbitrary contract.

```
(coerce-chaperone-contracts id x)
 → (listof/c chaperone-contract?)
  id : symbol?
  x : (listof any/c)
```

Like coerce-contracts, but requires the results to be chaperone contracts, not arbitrary contracts.

```
(coerce-flat-contract id x) → flat-contract?
  id : symbol?
  x : any/c
```

Like coerce-contract, but requires the result to be a flat contract, not an arbitrary contract.

```
(coerce-flat-contracts id x) → (listof/c flat-contract?)
  id : symbol?
  x : (listof any/c)
```

Like coerce-contracts, but requires the results to be flat contracts, not arbitrary contracts.

```
(coerce-contract/f x) → (or/c contract? #f)
  x : any/c
```

Like coerce-contract, but returns #f if the value cannot be coerced to a contract.

7.7.1 Blame Objects

```
(blame? x) → boolean?
  x : any/c
```

This predicate recognizes blame objects.

```
(blame-add-context  blame
                    context
                   [#:important important
                    #:swap? swap?])        → blame?
```

```
blame : blame?
context : (or/c string? #f)
important : (or/c string? #f) = #f
swap? : boolean? = #f
```

Adds some context information to blame error messages that explicates which portion of the contract failed (and that gets rendered by `raise-blame-error`).

The *context* argument describes one layer of the portion of the contract, typically of the form `"the 1st argument of"` (in the case of a function contract) or `"a conjunct of"` (in the case of an `and/c` contract).

For example, consider this contract violation:

```
> (define/contract f
    (list/c (-> integer? integer?))
    (list (λ (x) x)))

> ((car f) #f)
f: contract violation
  expected: integer?, given: #f
  in: the 1st argument of
      the 1st element of
      (list/c (-> integer? integer?))
  contract from: (definition f)
  blaming: top-level
  at: eval:2.0
```

It shows that the portion of the contract being violated is the first occurrence of `integer?`, because the `->` and the `list/c` combinators each internally called `blame-add-context` to add the two lines following "in" in the error message.

The *important* argument is used to build the beginning part of the contract violation. The last *important* argument that gets added to a blame object is used. The `class/c` contract adds an important argument, as does the `->` contract (when `->` knows the name of the function getting the contract).

The *swap?* argument has the effect of calling `blame-swap` while adding the layer of context, but without creating an extra blame object.

The context information recorded in blame structs keeps track of combinators that do not add information, and add the string `"..."` for them, so programmers at least see that there was some context they are missing in the error messages. Accordingly, since there are combinators that should not add any context (e.g., `recursive-contract`), passing `#f` as the context string argument avoids adding the `"..."` string.

```
(blame-positive b) → any/c
  b : blame?
(blame-negative b) → any/c
  b : blame?
```

These functions produce printable descriptions of the current positive and negative parties of a blame object.

```
(blame-contract b) → any/c
  b : blame?
```

This function produces a description of the contract associated with a blame object (the result of `contract-name`).

```
(blame-value b) → any/c
  b : blame?
```

This function produces the name of the value to which the contract was applied, or `#f` if no name was provided.

```
(blame-source b) → srcloc?
  b : blame?
```

This function produces the source location associated with a contract. If no source location was provided, all fields of the structure will contain `#f`.

```
(blame-swap b) → blame?
  b : blame?
```

This function swaps the positive and negative parties of a blame object. (See also `blame-add-context`.)

```
(blame-original? b) → boolean?
  b : blame?
(blame-swapped? b) → boolean?
  b : blame?
```

These functions report whether the current blame of a given blame object is the same as in the original contract invocation (possibly of a compound contract containing the current one), or swapped, respectively. Each is the negation of the other; both are provided for convenience and clarity.

```
(blame-replace-negative b neg) → blame?
  b : blame?
  neg : any/c
```

Produces a blame? object just like b except that it uses neg instead of the negative position b has.

```
(raise-blame-error b x fmt v ...) → none/c
  b : blame?
  x : any/c
  fmt : (or/c string?
              (listof (or/c string?
                            'given 'given:
                            'expected 'expected:)))
  v : any/c
```

Signals a contract violation. The first argument, b, records the current blame information, including positive and negative parties, the name of the contract, the name of the value, and the source location of the contract application. The second argument, x, is the value that failed to satisfy the contract.

The remaining arguments are a format string, fmt, and its arguments, v ..., specifying an error message specific to the precise violation.

If fmt is a list, then the elements are concatenated together (with spaces added, unless there are already spaces at the ends of the strings), after first replacing symbols with either their string counterparts, or replacing 'given with "produced" and 'expected with "promised", depending on whether or not the b argument has been swapped or not (see blame-swap).

```
(struct exn:fail:contract:blame exn:fail:contract (object)
        #:extra-constructor-name make-exn:fail:contract:blame)
  object : blame?
```

This exception is raised to signal a contract error. The object field contains a blame object associated with a contract violation.

```
(current-blame-format) → (-> blame? any/c string? string?)
(current-blame-format proc) → void?
  proc : (-> blame? any/c string? string?)
```

A parameter that is used when constructing a contract violation error. Its value is procedure that accepts three arguments:

- the blame object for the violation,
- the value that the contract applies to, and
- a message indicating the kind of violation.

The procedure then returns a string that is put into the contract error message. Note that the value is often already included in the message that indicates the violation.

Examples:

```
(define (show-blame-error blame value message)
  (string-append
   "Contract Violation!\n"
   (format "Guilty Party: ~a\n" (blame-positive blame))
   (format "Innocent Party: ~a\n" (blame-negative blame))
   (format "Contracted Value Name: ~a\n" (blame-value blame))
   (format "Contract Location: ~s\n" (blame-source blame))
   (format "Contract Name: ~a\n" (blame-contract blame))
   (format "Offending Value: ~s\n" value)
   (format "Offense: ~a\n" message)))

> (current-blame-format show-blame-error)

> (define/contract (f x)
    (-> integer? integer?)
    (/ x 2))

> (f 2)
1
> (f 1)
Contract Violation!
Guilty Party: (function f)
Innocent Party: top-level
Contracted Value Name: f
Contract Location: #(struct:srcloc eval 4 0 4 1)
Contract Name: (-> integer? integer?)
Offending Value: 1/2
Offense: promised: integer?, produced: 1/2

> (f 1/2)
Contract Violation!
Guilty Party: top-level
Innocent Party: (function f)
Contracted Value Name: f
Contract Location: #(struct:srcloc eval 4 0 4 1)
Contract Name: (-> integer? integer?)
Offending Value: 1/2
Offense: expected: integer?, given: 1/2
```

7.7.2 Contracts as structs

The property `prop:contract` allows arbitrary structures to act as contracts. The property `prop:chaperone-contract` allows arbitrary structures to act as chaperone contracts; `prop:chaperone-contract` inherits `prop:contract`, so chaperone contract structures may also act as general contracts. The property `prop:flat-contract` allows arbitrary structures to act as flat contracts; `prop:flat-contract` inherits both `prop:chaperone-contract` and `prop:procedure`, so flat contract structures may also act as chaperone contracts, as general contracts, and as predicate procedures.

```
prop:contract : struct-type-property?
prop:chaperone-contract : struct-type-property?
prop:flat-contract : struct-type-property?
```

These properties declare structures to be contracts or flat contracts, respectively. The value for `prop:contract` must be a contract property constructed by `build-contract-property`; likewise, the value for `prop:chaperone-contract` must be a chaperone contract property constructed by `build-chaperone-contract-property` and the value for `prop:flat-contract` must be a flat contract property constructed by `build-flat-contract-property`.

```
prop:contracted : struct-type-property?
impersonator-prop:contracted : impersonator-property?
```

These properties attach a contract value to the protected structure, chaperone, or impersonator value. The function `has-contract?` returns `#t` for values that have one of these properties, and `value-contract` extracts the value from the property (which is expected to be the contract on the value).

```
(build-flat-contract-property [#:name get-name
                               #:first-order get-first-order
                               #:projection get-projection
                               #:stronger stronger
                               #:generator generator])
 → flat-contract-property?
 get-name : (-> contract? any/c)
          = (λ (c) 'anonymous-flat-contract)
 get-first-order : (-> contract? (-> any/c boolean?))
                 = (λ (c) (λ (x) #t))
 get-projection : (-> contract? (-> blame? (-> any/c any/c)))
                  (λ (c)
                    (λ (b)
                      (λ (x)
                =       (if ((get-first-order c) x)
                            x
                            (raise-blame-error
                             b x "expected <~a>, given: ~e" (get-name c) x)))))
 stronger : (or/c (-> contract? contract? boolean?) #f) = #f
 generator : (or/c (-> number? (listof (list any/c contract?)) any/c) #f)
           = #f
(build-chaperone-contract-property
 [#:name get-name
  #:first-order get-first-order
  #:projection get-projection
  #:stronger stronger
  #:generator generator])
 → chaperone-contract-property?
 get-name : (-> contract? any/c)
          = (λ (c) 'anonymous-chaperone-contract)
 get-first-order : (-> contract? (-> any/c boolean?))
                 = (λ (c) (λ (x) #t))
 get-projection : (-> contract? (-> blame? (-> any/c any/c)))
                  (λ (c)
                    (λ (b)
                      (λ (x)
                =       (if ((get-first-order c) x)
                            x
                            (raise-blame-error
                             b x "expected <~a>, given: ~e" (get-name c) x)))))
 stronger : (or/c (-> contract? contract? boolean?) #f) = #f
 generator : (or/c (-> number? (listof (list any/c contract?)) any/c) #f)
           = #f
```

```
(build-contract-property [#:name get-name
                          #:first-order get-first-order
                          #:projection get-projection
                          #:stronger stronger
                          #:generator generator])
 → contract-property?
  get-name : (-> contract? any/c) = (λ (c) 'anonymous-contract)
  get-first-order : (-> contract? (-> any/c boolean?))
                  = (λ (c) (λ (x) #t))
  get-projection : (-> contract? (-> blame? (-> any/c any/c)))
                   (λ (c)
                     (λ (b)
                       (λ (x)
                 =       (if ((get-first-order c) x)
                             x
                             (raise-blame-error
                              b x "expected <~a>, given: ~e" (get-name c) x)))))
  stronger : (or/c (-> contract? contract? boolean?) #f) = #f
  generator : (or/c (-> number? (listof (list any/c contract?)) any/c) #f)
            = #f
```

These functions build the arguments for `prop:contract`, `prop:chaperone-contract`, and `prop:flat-contract`, respectively.

A *contract property* specifies the behavior of a structure when used as a contract. It is specified in terms of five accessors: *`get-name`*, which produces a description to `write` as part of a contract violation; *`get-first-order`*, which produces a first-order predicate to be used by `contract-first-order-passes?`; *`get-projection`*, which produces a blame-tracking projection defining the behavior of the contract; *`stronger`*, which is a predicate that determines whether this contract (passed in the first argument) is stronger than some other contract (passed in the second argument); and *`generator`*, which makes a random value that matches the contract, given a size bound and an environment from which to draw interesting values.

These accessors are passed as (optional) keyword arguments to `build-contract-property`, and are applied to instances of the appropriate structure type by the contract system. Their results are used analogously to the arguments of `make-contract`.

A *chaperone contract property* specifies the behavior of a structure when used as a chaperone contract. It is specified using `build-chaperone-contract-property`, and accepts exactly the same set of arguments as `build-contract-property`. The only difference is that the projection accessor must return a value that passes `chaperone-of?` when compared with the original, uncontracted value.

A *flat contract property* specifies the behavior of a structure when used as a flat contract. It is specified using `build-flat-contract-property`, and accepts exactly the same set

of arguments as build-contract-property. The only difference is that the projection accessor is expected not to wrap its argument in a higher-order fashion, analogous to the constraint on projections in make-flat-contract.

```
(contract-property? x) → boolean?
  x : any/c
(chaperone-contract-property? x) → boolean?
  x : any/c
(flat-contract-property? x) → boolean?
  x : any/c
```

These predicates detect whether a value is a contract property, chaperone contract property, or a flat contract property, respectively.

7.7.3 Obligation Information in Check Syntax

Check Syntax in DrRacket shows obligation information for contracts according to syntax-properties that the contract combinators leave in the expanded form of the program. These properties indicate where contracts appear in the source and where the positive and negative positions of the contracts appear.

To make Check Syntax show obligation information for your new contract combinators, use the following properties (some helper macros and functions are below):

- 'racket/contract:contract :
  (vector/c symbol? (listof syntax?) (listof syntax?))

  This property should be attached to the result of a transformer that implements a contract combinator. It signals to Check Syntax that this is where a contract begins.

  The first element in the vector should be a unique (in the sense of eq?) value that Check Syntax can use a tag to match up this contract with its subpieces (specified by the two following syntax properties).

  The second and third elements of the vector are syntax objects from pieces of the contract, and Check Syntax will color them. The first list should contain subparts that are the responsibility of parties (typically modules) that provide implementations of the contract. The second list should contain subparts that are the responsibility of clients.

  For example, in (->* () #:pre #t any/c #:post #t), the ->* and the #:post should be in the first list and #:pre in the second list.

- 'racket/contract:negative-position : symbol?

  This property should be attached to sub-expressions of a contract combinator that are expected to be other contracts. The value of the property should be the key (the first

element from the vector for the 'racket/contract:contract property) indicating which contract this is.

This property should be used when the expression's value is a contract that clients are responsible for.

- 'racket/contract:positive-position : symbol?

  This form is just like 'racket/contract:negative-position, except that it should be used when the expression's value is a contract that the original party should be responsible for.

- 'racket/contract:contract-on-boundary : symbol?

  The presence of this property tells Check Syntax that it should start coloring from this point. It expects the expression to be a contract (and, thus, to have the 'racket/contract:contract property); this property indicates that this contract is on a (module) boundary.

  (The value of the property is not used.)

- 'racket/contract:internal-contract : symbol?

  Like 'racket/contract:contract-on-boundary, the presence of this property triggers coloring, but this is meant for use when the party (module) containing the contract (regardless of whether or not this module exports anything matching the contract) can be blamed for violating the contract. This comes into play for ->i contracts, since the contract itself has access to values under contract via the dependency.

```
(define/final-prop header body ...)

header = main-id
       | (main-id id ...)
       | (main-id id ... . id)
```

The same as (define header body ...), except that uses of main-id in the header are annotated with the 'racket/contract:contract property (as above).

```
(define/subexpression-pos-prop header body ...)

header = main-id
       | (main-id id ...)
       | (main-id id ... . id)
```

The same as (define header body ...), except that uses of main-id in the header are annotated with the 'racket/contract:contract property (as above) and arguments are annotated with the 'racket/contract:positive-position property.

7.7.4 Utilities for Building New Combinators

```
(contract-stronger? x y) → boolean?
  x : contract?
  y : contract?
```

Returns #t if the contract x accepts either fewer or the same number of values as y does.

This function is conservative, so it may return #f when x does, in fact, accept fewer values.

Examples:

```
> (contract-stronger? integer? integer?)
#t
> (contract-stronger? (between/c 25 75) (between/c 0 100))
#t
> (contract-stronger? (between/c 0 100) (between/c 25 75))
#f
> (contract-stronger? (between/c -10 0) (between/c 0 10))
#f
> (contract-stronger? (λ (x) (and (real? x) (<= x (random 10))))
                      (λ (x) (and (real? x) (<= x (+ 100 (random 10))))))
#f
```

```
(contract-first-order-passes? contract v) → boolean?
  contract : contract?
  v : any/c
```

Returns a boolean indicating whether the first-order tests of *contract* pass for v.

If it returns #f, the contract is guaranteed not to hold for that value; if it returns #t, the contract may or may not hold. If the contract is a first-order contract, a result of #t guarantees that the contract holds.

```
(contract-first-order c) → (-> any/c boolean?)
  c : contract?
```

Produces the first-order test used by or/c to match values to higher-order contracts.

7.8 Contract Utilities

```
(contract? v) → boolean?
  v : any/c
```

Returns `#t` if its argument is a contract (i.e., constructed with one of the combinators described in this section or a value that can be used as a contract) and `#f` otherwise.

```
(chaperone-contract? v) → boolean?
  v : any/c
```

Returns `#t` if its argument is a contract that guarantees that it returns a value which passes `chaperone-of?` when compared to the original, uncontracted value.

```
(impersonator-contract? v) → boolean?
  v : any/c
```

Returns `#t` if its argument is a contract that is not a chaperone contract nor a flat contract.

```
(flat-contract? v) → boolean?
  v : any/c
```

Returns `#t` when its argument is a contract that can be checked immediately (unlike, say, a function contract).

For example, `flat-contract` constructs flat contracts from predicates, and symbols, booleans, numbers, and other ordinary Racket values (that are defined as contracts) are also flat contracts.

```
(flat-contract-predicate v) → (any/c . -> . any/c)
  v : flat-contract?
```

Extracts the predicate from a flat contract.

```
(contract-name c) → any/c
  c : contract?
```

Produces the name used to describe the contract in error messages.

```
(value-contract v) → contract?
  v : has-contract?
```

Returns the contract attached to `v`, if recorded. Otherwise it returns `#f`.

To support `value-contract` and `has-contract?` in your own contract combinators, use `prop:contracted` or `impersonator-prop:contracted`.

```
(has-contract? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a value that has a recorded contract attached to it.

See also `value-contract`.

```
(contract-projection c) → (-> blame? (-> any/c any/c))
  c : contract?
```

Produces the projection defining a contract's behavior on protected values.

```
(make-none/c sexp-name) → contract?
  sexp-name : any/c
```

Makes a contract that accepts no values, and reports the name *sexp-name* when signaling a contract violation.

```
(recursive-contract contract-expr)
(recursive-contract contract-expr type)
```

Delays the evaluation of its argument until the contract is checked, making recursive contracts possible. If *type* is given, it describes the expected type of contract and must be one of the keywords `#:impersonator`, `#:chaperone`, or `#:flat`. If *type* is not given, an impersonator contract is created.

```
(opt/c contract-expr)
```

This optimizes its argument contract expression by traversing its syntax and, for known contract combinators, fuses them into a single contract combinator that avoids as much allocation overhead as possible. The result is a contract that should behave identically to its argument, except faster (due to less allocation).

```
(define-opt/c (id id ...) expr)
```

This defines a recursive contract and simultaneously optimizes it. Semantically, it behaves just as if the `-opt/c` were not present, defining a function on contracts (except that the body expression must return a contract). But, it also optimizes that contract definition, avoiding extra allocation, much like `opt/c` does.

For example,

```
(define-contract-struct bt (val left right))

(define-opt/c (bst-between/c lo hi)
  (or/c null?
        (bt/c [val (real-in lo hi)]
              [left (val) (bst-between/c lo val)]
              [right (val) (bst-between/c val hi)])))

(define bst/c (bst-between/c -inf.0 +inf.0))
```

defines the bst/c contract that checks the binary search tree invariant. Removing the -opt/c also makes a binary search tree contract, but one that is (approximately) 20 times slower.

7.9 racket/contract/base

```
(require racket/contract/base)
```

The racket/contract/base module provides a subset of the exports of racket/contract module. In particular, it contains everything in the

- §7.1 “Data-structure Contracts”
- §7.2 “Function Contracts”
- §7.6 “Attaching Contracts to Values” and
- §7.8 “Contract Utilities” sections.

Unfortunately, using racket/contract/base does not yield a significantly smaller memory footprint than racket/contract, but it can still be useful to add contracts to libraries that racket/contract uses to implement some of the more sophisticated parts of the contract system.

7.10 Legacy Contracts

```
(make-proj-contract name proj first-order) → contract?
  name : any/c
         (or/c (-> any/c
                   any/c
                   (list/c any/c any/c)
                   contact?
                   (-> any/c any/c))
  proj :       (-> any/c
                   any/c
                   (list/c any/c any/c)
                   contact?
                   boolean?
                   (-> any/c any/c)))
  first-order : (-> any/c boolean?)
```

Builds a contract using an old interface.

Modulo errors, it is equivalent to:

```
(make-contract
 #:name name
 #:first-order first-order
 #:projection
 (cond
   [(procedure-arity-includes? proj 5)
    (lambda (blame)
      (proj (blame-positive blame)
            (blame-negative blame)
            (list (blame-source blame) (blame-value blame))
            (blame-contract blame)
            (not (blame-swapped? blame))))]
   [(procedure-arity-includes? proj 4)
    (lambda (blame)
      (proj (blame-positive blame)
            (blame-negative blame)
            (list (blame-source blame) (blame-value blame))
            (blame-contract blame)))]))
```

```
(raise-contract-error val
                      src
                      pos
                      name
                      fmt
                      arg ...) → any/c
  val : any/c
  src : any/c
  pos : any/c
  name : any/c
  fmt : string?
  arg : any/c
```

Calls `raise-blame-error` after building a `blame` struct from the *val*, *src*, *pos*, and *name* arguments. The *fmt* string and following arguments are passed to `format` and used as the string in the error message.

```
(contract-proc c)
    (->* (symbol? symbol? (or/c syntax? (list/c any/c any/c)))
 →       (boolean?)
         (-> any/c any))
  c : contract?
```

Constructs an old-style projection from a contract.

The resulting function accepts the information that is in a `blame` struct and returns a projection function that checks the contract.

7.11 Random generation

```
(contract-random-generate ctc fuel [fail]) → any/c
  ctc : contract?
  fuel : int?
  fail : (-> any/c) = (λ () (error ...))
```

Attempts to randomly generate a value which will match the contract. The fuel argument limits how hard the generator tries to generate a value matching the contract and is a rough limit of the size of the resulting value.

The generator may fail to generate a contract, either because some contracts do not have corresponding generators (for example, not all predicates have generators) or because there is not enough fuel. In either case, the thunk *fail* is invoked.

8 Pattern Matching

§12 “Pattern Matching” in ***The Racket Guide*** introduces pattern matching.

The `match` form and related forms support general pattern matching on Racket values. See also §3.7 “Regular Expressions” for information on regular-expression matching on strings, bytes, and streams.

```
(require racket/match)
```

The bindings documented in this section are provided by the `racket/match` and `racket` libraries, but not `racket/base`.

```
(match val-expr clause ...)

clause = [pat body ...+]
       | [pat (=> id) body ...+]
```

Finds the first *pat* that matches the result of *val-expr*, and evaluates the corresponding *body*s with bindings introduced by *pat* (if any). The last *body* in the matching clause is evaluated in tail position with respect to the `match` expression.

To find a match, the *clause*s are tried in order. If no *clause* matches, then the `exn:misc:match?` exception is raised.

An optional `(=> id)` between a *pat* and the *body*s is bound to a *failure procedure* of zero arguments. If this procedure is invoked, it escapes back to the pattern matching expression, and resumes the matching process as if the pattern had failed to match. The *body*s must not mutate the object being matched before calling the failure procedure, otherwise the behavior of matching is unpredictable.

The grammar of *pat* is as follows, where non-italicized identifiers are recognized symbolically (i.e., not by binding).

pat	`::= id`	match anything, bind identifier
	`| (var id)`	match anything, bind identifier
	`| _`	match anything
	`| literal`	match literal
	`| (quote datum)`	match `equal?` value
	`| (list lvp ...)`	match sequence of *lvp*s
	`| (list-rest lvp ... pat)`	match *lvp*s consed onto a *pat*
	`| (list-no-order pat ...)`	match *pat*s in any order
	`| (list-no-order pat ... lvp)`	match *pat*s in any order
	`| (vector lvp ...)`	match vector of *pat*s
	`| (hash-table (pat pat) ...)`	match hash table
	`| (hash-table (pat pat) ...+ ooo)`	match hash table
	`| (cons pat pat)`	match pair of *pat*s
	`| (mcons pat pat)`	match mutable pair of *pat*s

	| `(box pat)`	match boxed *pat*
	| `(struct-id pat ...)`	match *struct-id* instance
	| `(struct struct-id (pat ...))`	match *struct-id* instance
	| `(regexp rx-expr)`	match string
	| `(regexp rx-expr pat)`	match string, result with *pat*
	| `(pregexp px-expr)`	match string
	| `(pregexp px-expr pat)`	match string, result with *pat*
	| `(and pat ...)`	match when all *pat*s match
	| `(or pat ...)`	match when any *pat* match
	| `(not pat ...)`	match when no *pat* matches
	| `(app expr pat)`	match (*expr* value) to *pat*
	| `(? expr pat ...)`	match if (*expr* value) and *pat*s
	| `(quasiquote qp)`	match a quasipattern
	| *derived-pattern*	match using extension
literal	::= `#t`	match true
	| `#f`	match false
	| *string*	match `equal?` string
	| *bytes*	match `equal?` byte string
	| *number*	match `equal?` number
	| `char`	match `equal?` character
	| `keyword`	match `equal?` keyword
	| `regexp`	match `equal?` regexp literal
	| `pregexp`	match `equal?` pregexp literal
lvp	::= *pat ooo*	greedily match *pat* instances
	| *pat*	match *pat*
qp	::= *literal*	match literal
	| *id*	match symbol
	| `(qp ...)`	match sequences of *qp*s
	| `(qp ... . qp)`	match *qp*s ending *qp*
	| `(qp ooo . qp)`	match *qp*s beginning with repeated *qp*
	| `#(qp ...)`	match vector of *qp*s
	| `#&qp`	match boxed *qp*
	| `,pat`	match *pat*
	| `,@(list lvp ...)`	match *lvp*s, spliced
	| `,@(list-rest lvp ... pat)`	match *lvp*s plus *pat*, spliced
	| `,@'qp`	match list-matching *qp*, spliced
ooo	::= `...`	zero or more; `...` is literal
	| `___`	zero or more
	| `..k`	*k* or more
	| `__k`	*k* or more

In more detail, patterns match as follows:

- *id* (excluding the reserved names `_`, `...`, `..._`, `..k`, and `..k` for non-negative integers *k*) or (`var` *id*) — matches anything, and binds *id* to the matching values. If an *id* is used multiple times within a pattern, the corresponding matches must be the same

according to (match-equality-test), except that instances of an *id* in different or and not sub-patterns are independent.

Examples:

```
> (match '(1 2 3)
    [(list a b a) (list a b)]
    [(list a b c) (list c b a)])
'(3 2 1)
> (match '(1 (x y z) 1)
    [(list a b a) (list a b)]
    [(list a b c) (list c b a)])
'(1 (x y z))
```

- _ — matches anything, without binding any identifiers.

  Example:

```
> (match '(1 2 3)
    [(list _ _ a) a])
3
```

- #t, #f, *string*, *bytes*, *number*, *char*, or (quote *datum*) — matches an equal? constant.

  Example:

```
> (match "yes"
    ["no" #f]
    ["yes" #t])
#t
```

- (list *lvp* ...) — matches a list of elements. In the case of (list *pat* ...), the pattern matches a list with as many element as *pat*s, and each element must match the corresponding *pat*. In the more general case, each *lvp* corresponds to a "spliced" list of greedy matches.

  For spliced lists, ... and ___ are aliases for zero or more matches. The ..*k* and __*k* forms are also aliases, specifying *k* or more matches. Pattern variables that precede these splicing operators are bound to lists of matching forms.

  Examples:

```
> (match '(1 2 3)
    [(list a b c) (list c b a)])
'(3 2 1)
> (match '(1 2 3)
    [(list 1 a ...) a])
'(2 3)
> (match '(1 2 3)
    [(list 1 a ..3) a]
    [_ 'else])
```

```
'else
> (match '(1 2 3 4)
    [(list 1 a ..3) a]
    [_ 'else])
'(2 3 4)
> (match '(1 2 3 4 5)
    [(list 1 a ..3 5) a]
    [_ 'else])
'(2 3 4)
> (match '(1 (2) (2) (2) 5)
    [(list 1 (list a) ..3 5) a]
    [_ 'else])
'(2 2 2)
```

- `(list-rest` *`lvp`* `...` *`pat`*`)` — similar to a `list` pattern, but the final *`pat`* matches the "rest" of the list after the last *`lvp`*. In fact, the matched value can be a non-list chain of pairs (i.e., an "improper list") if *`pat`* matches non-list values.

  Examples:

```
> (match '(1 2 3 . 4)
    [(list-rest a b c d) d])
4
> (match '(1 2 3 . 4)
    [(list-rest a ... d) (list a d)])
'((1 2 3) 4)
```

- `(list-no-order` *`pat`* `...)` — similar to a `list` pattern, but the elements to match each *`pat`* can appear in the list in any order.

  Example:

```
> (match '(1 2 3)
    [(list-no-order 3 2 x) x])
1
```

- `(list-no-order` *`pat`* `...` *`lvp`*`)` — generalizes `list-no-order` to allow a pattern that matches multiple list elements that are interspersed in any order with matches for the other patterns.

  Example:

```
> (match '(1 2 3 4 5 6)
    [(list-no-order 6 2 y ...) y])
'(1 3 4 5)
```

- `(vector` *`lvp`* `...)` — like a `list` pattern, but matching a vector.

  Example:

```
> (match #(1 (2) (2) (2) 5)
    [(vector 1 (list a) ..3 5) a])
'(2 2 2)
```

- `(hash-table (pat pat) ...)` — similar to `list-no-order`, but matching against hash table's key–value pairs.

  Example:

  ```
  > (match #hash(("a" . 1) ("b" . 2))
      [(hash-table ("b" b) ("a" a)) (list b a)])
  '(2 1)
  ```

- `(hash-table (pat pat) ...+ ooo)` — Generalizes `hash-table` to support a final repeating pattern.

  Example:

  ```
  > (match #hash(("a" . 1) ("b" . 2))
      [(hash-table (key val) ...) key])
  '("b" "a")
  ```

- `(cons pat1 pat2)` — matches a pair value.

  Example:

  ```
  > (match (cons 1 2)
      [(cons a b) (+ a b)])
  3
  ```

- `(mcons pat1 pat2)` — matches a mutable pair value.

  Example:

  ```
  > (match (mcons 1 2)
      [(cons a b) 'immutable]
      [(mcons a b) 'mutable])
  'mutable
  ```

- `(box pat)` — matches a boxed value.

  Example:

  ```
  > (match #&1
      [(box a) a])
  1
  ```

- `(struct-id pat ...)` or `(struct struct-id (pat ...))` — matches an instance of a structure type named *struct-id*, where each field in the instance matches the corresponding *pat*. See also `struct*`.

  Usually, *struct-id* is defined with `struct`. More generally, *struct-id* must be bound to expansion-time information for a structure type (see §4.7 "Structure Type

Transformer Binding"), where the information includes at least a predicate binding and field accessor bindings corresponding to the number of field *pat*s. In particular, a module import or a `unit` import with a signature containing a `struct` declaration can provide the structure type information.

Examples:

```
(define-struct tree (val left right))

> (match (make-tree 0 (make-tree 1 #f #f) #f)
    [(tree a (tree b  _ _) _) (list a b)])
'(0 1)
```

- `(struct` *struct-id* `_)` — matches any instance of *struct-id*, without regard to contents of the fields of the instance.

- `(regexp` *rx-expr*`)` — matches a string that matches the regexp pattern produced by *rx-expr*; see §3.7 "Regular Expressions" for more information about regexps.

  Examples:

```
> (match "apple"
    [(regexp #rx"p+") 'yes]
    [_ 'no])
'yes
> (match "banana"
    [(regexp #rx"p+") 'yes]
    [_ 'no])
'no
```

- `(regexp` *rx-expr pat*`)` — extends the `regexp` form to further constrain the match where the result of `regexp-match` is matched against *pat*.

  Examples:

```
> (match "apple"
    [(regexp #rx"p+(.)" (list _ "l")) 'yes]
    [_ 'no])
'yes
> (match "append"
    [(regexp #rx"p+(.)" (list _ "l")) 'yes]
    [_ 'no])
'no
```

- `(pregexp` *rx-expr*`)` or `(regexp` *rx-expr pat*`)` — like the `regexp` patterns, but if *rx-expr* produces a string, it is converted to a pattern using `pregexp` instead of `regexp`.

- `(and` *pat* `...)` — matches if all of the *pat*s match. This pattern is often used as `(and` *id pat*`)` to bind *id* to the entire value that matches *pat*.

Example:

```
> (match '(1 (2 3) 4)
    [(list _ (and a (list _ ...)) _) a])
'(2 3)
```

- `(or pat ...)` — matches if any of the `pat`s match. **Beware**: the result expression can be duplicated once for each `pat`! Identifiers in `pat` are bound only in the corresponding copy of the result expression; in a module context, if the result expression refers to a binding, then all `pat`s must include the binding.

  Example:

```
> (match '(1 2)
    [(or (list a 1) (list a 2)) a])
1
```

- `(not pat ...)` — matches when none of the `pat`s match, and binds no identifiers.

  Examples:

```
> (match '(1 2 3)
    [(list (not 4) ...) 'yes]
    [_ 'no])
'yes
> (match '(1 4 3)
    [(list (not 4) ...) 'yes]
    [_ 'no])
'no
```

- `(app expr pat)` — applies `expr` to the value to be matched; the result of the application is matched against `pat`.

  Example:

```
> (match '(1 2)
    [(app length 2) 'yes])
'yes
```

- `(? expr pat ...)` — applies `expr` to the value to be matched, and checks whether the result is a true value; the additional `pat`s must also match (i.e., `?` combines a predicate application and an `and` pattern).

  Example:

```
> (match '(1 3 5)
    [(list (? odd?) ...) 'yes])
'yes
```

- `(quasiquote qp)` — introduces a quasipattern, in which identifiers match symbols. Like the `quasiquote` expression form, `unquote` and `unquote-splicing` escape back to normal patterns.

Example:

```
> (match '(1 2 3)
    [`(1 ,a ,(? odd? b)) (list a b)])
'(2 3)
```

- *derived-pattern* — matches a pattern defined by a macro extension via `define-match-expander`.

Note that the matching process may destructure the input multiple times, and may evaluate expressions embedded in patterns such as `(app expr pat)` in arbitrary order, or multiple times. Therefore, such expressions must be safe to call multiple times, or in an order other than they appear in the original program.

8.1 Additional Matching Forms

```
(match* (val-expr ...+) clause* ...)

clause* = [(pat ...+) body ...+]
        | [(pat ...+) (=> id) body ...+]
```

Matches a sequence of values against each clause in order, matching only when all patterns in a clause match. Each clause must have the same number of patterns as the number of *val-expr*s.

Example:

```
> (match* (1 2 3)
   [(_ (? number?) x) (add1 x)])
4
```

```
(match/values expr clause clause ...)
```

If *expr* evaluates to `n` values, then match all `n` values against the patterns in *clause* .... Each clause must contain exactly `n` patterns. At least one clause is required to determine how many values to expect from *expr*.

```
(match-lambda clause ...)
```

Equivalent to `(lambda (id) (match id clause ...))`.

```
(match-lambda* clause ...)
```

Equivalent to `(lambda lst (match lst clause ...))`.

```
(match-lambda** clause* ...)
```

Equivalent to `(lambda (args ...) (match* (args ...) clause* ...))`, where the number of `args ...` is computed from the number of patterns appearing in each of the *clause**.

```
(match-let ([pat expr] ...) body ...+)
```

Generalizes `let` to support pattern bindings. Each *expr* is matched against its corresponding *pat* (the match must succeed), and the bindings that *pat* introduces are visible in the *body*s.

Example:

```
> (match-let ([(list a b) '(1 2)]
              [(vector x ...) #(1 2 3 4)])
    (list b a x))
'(2 1 (1 2 3 4))
```

```
(match-let* ([pat expr] ...) body ...+)
```

Like `match-let`, but generalizes `let*`, so that the bindings of each *pat* are available in each subsequent *expr*.

Example:

```
> (match-let* ([(list a b) '(#(1 2 3 4) 2)]
               [(vector x ...) a])
    x)
'(1 2 3 4)
```

```
(match-let-values ([(pat ...) expr] ...) body ...+)
```

Like `match-let`, but generalizes `let-values`.

```
(match-let*-values ([(pat ...) expr] ...) body ...+)
```

Like `match-let*`, but generalizes `let*-values`.

```
(match-letrec ([pat expr] ...) body ...+)
```

Like `match-let`, but generalizes `letrec`.

```
(match-define pat expr)
```

Defines the names bound by *pat* to the values produced by matching against the result of *expr*.

Examples:

```
> (match-define (list a b) '(1 2))

> b
2
```

```
(match-define-values (pat pats ...) expr)
```

Like `match-define` but for when expr produces multiple values. Like match/values, it requires at least one pattern to determine the number of values to expect.

Examples:

```
> (match-define-values (a b) (values 1 2))

> b
2
```

```
(exn:misc:match? v) → boolean?
  v : any/c
```

A predicate for the exception raised in the case of a match failure.

8.2 Extending `match`

```
(define-match-expander id proc-expr)
(define-match-expander id proc-expr proc-expr)
```

Binds *id* to a *match expander*.

The first *proc-expr* sub-expression must evaluate to a transformer that produces a *pat* for `match`. Whenever *id* appears as the beginning of a pattern, this transformer is given, at expansion time, a syntax object corresponding to the entire pattern (including *id*). The pattern is replaced with the result of the transformer.

A transformer produced by a second *proc-expr* sub-expression is used when *id* is used in an expression context. Using the second *proc-expr*, *id* can be given meaning both inside and outside patterns.

For example, to extend the pattern matcher and destructure syntax lists,

```
(define-match-expander syntax-list
  (lambda (stx)
    (syntax-case stx ()
      [(_ elts ...)
       #'(? (lambda (x) (and (syntax? x)
                             (list? (syntax->list x))))
            (app syntax->list (list elts ...)))])))
(define (make-keyword-predicate keyword)
  (lambda (stx)
    (and (identifier? stx)
         (free-identifier=? stx keyword))))
(define or-keyword? (make-keyword-predicate #'or))
(define and-keyword? (make-keyword-predicate #'and))

> (match #'(or 3 4)
    [(syntax-list (? or-keyword?) b c)
     (list "OOORRR!" b c)]
    [(syntax-list (? and-keyword?) b c)
     (list "AAANND!" b c)])
'("OOORRR!" #<syntax:47:0 3> #<syntax:47:0 4>)
> (match #'(and 5 6)
    [(syntax-list (? or-keyword?) b c)
     (list "OOORRR!" b c)]
    [(syntax-list (? and-keyword?) b c)
     (list "AAANND!" b c)])
'("AAANND!" #<syntax:48:0 5> #<syntax:48:0 6>)
```

prop:match-expander : struct-type-property?

A structure type property to identify structure types that act as match expanders like the ones created by `define-match-expander`.

The property value must be an exact non-negative integer or a procedure of one or two arguments. In the former case, the integer designates a field within the structure that should contain a procedure; the integer must be between `0` (inclusive) and the number of non-automatic fields in the structure type (exclusive, not counting supertype fields), and the designated field must also be specified as immutable.

If the property value is a procedure of one argument, then the procedure serves as the transformer for match expansion. If the property value is a procedure of two arguments, then the first argument is the structure whose type has `prop:match-expander` property, and the second argument is a syntax object as for a match expander..

If the property value is a assignment transformer, then the wrapped procedure is extracted with `set!-transformer-procedure` before it is called.

```
prop:legacy-match-expander : struct-type-property?
```

Like prop:match-expander, but for the legacy match syntax.

```
(match-equality-test) → (any/c any/c . -> . any)
(match-equality-test comp-proc) → void?
  comp-proc : (any/c any/c . -> . any)
```

A parameter that determines the comparison procedure used to check whether multiple uses of an identifier match the "same" value. The default is equal?.

```
(match/derived val-expr original-datum clause ...)
(match*/derived (val-expr ...) original-datum clause* ...)
```

Like match and match* respectively, but includes a sub-expression to be used as the source for all syntax errors within the form. For example, match-lambda expands to match/derived so that errors in the body of the form are reported in terms of match-lambda instead of match.

8.3 Library Extensions

```
(== val comparator)
(== val)
```

A match expander which checks if the matched value is the same as *val* when compared by *comparator*. If *comparator* is not provided, it defaults to equal?.

Examples:

```
> (match (list 1 2 3)
    [(== (list 1 2 3)) 'yes]
    [_ 'no])
'yes
> (match (list 1 2 3)
    [(== (list 1 2 3) eq?) 'yes]
    [_ 'no])
'no
> (match (list 1 2 3)
    [(list 1 2 (== 3 =)) 'yes]
    [_ 'no])
'yes
```

```
(struct* struct-id ([field pat] ...))
```

A `match` pattern form that matches an instance of a structure type named *`struct-id`*, where the field *`field`* in the instance matches the corresponding *`pat`*.

Any field of *`struct-id`* may be omitted, and such fields can occur in any order.

Examples:

```
(define-struct tree (val left right))

> (match (make-tree 0 (make-tree 1 #f #f) #f)
    [(struct* tree ([val a]
                    [left (struct* tree ([right #f] [val b]))]))
     (list a b)])
'(0 1)
```

9 Control Flow

9.1 Multiple Values

See §1.1.3 “Multiple Return Values” for general information about multiple result values. In addition to `call-with-values` (described in this section), the `let-values`, `let*-values`, `letrec-values`, and `define-values` forms (among others) create continuations that receive multiple values.

```
(values v ...) → any
  v : any/c
```

Returns the given `v`s. That is, `values` returns as provided arguments.

Examples:

```
> (values 1)
1
> (values 1 2 3)
1
2
3
> (values)
```

```
(call-with-values generator receiver) → any
  generator : (-> any)
  receiver : procedure?
```

Calls *generator*, and passes the values that *generator* produces as arguments to *receiver*. Thus, `call-with-values` creates a continuation that accepts any number of values that *receiver* can accept. The *receiver* procedure is called in tail position with respect to the `call-with-values` call.

Examples:

```
> (call-with-values (lambda () (values 1 2)) +)
3
> (call-with-values (lambda () 1) (lambda (x y) (+ x y)))
#<procedure>: arity mismatch;
 the expected number of arguments does not match the given
number
  expected: 2
  given: 1
  arguments...:
   1
```

9.2 Exceptions

See §1.1.15 “Exceptions” for information on the Racket exception model. It is based on a proposal by Friedman, Haynes, and Dybvig [Friedman95].

Whenever a primitive error occurs in Racket, an exception is raised. The value that is passed to the current exception handler for a primitive error is always an instance of the `exn` structure type. Every `exn` structure value has a `message` field that is a string, the primitive error message. The default exception handler recognizes exception values with the `exn?` predicate and passes the error message to the current error display handler (see `error-display-handler`).

Primitive procedures that accept a procedure argument with a particular required arity (e.g., `call-with-input-file`, `call/cc`) check the argument's arity immediately, raising `exn:fail:contract` if the arity is incorrect.

9.2.1 Error Message Conventions

Racket's *error message convention* is to produce error messages with the following shape:

⟨*srcloc*⟩`:` ⟨*name*⟩`:` ⟨*message*⟩;
⟨*continued-message*⟩ ...
⟨*field*⟩`:` ⟨*detail*⟩
...

The message starts with an optional source location, ⟨*srcloc*⟩, which is followed by a colon and space when present. The message continues with an optional ⟨*name*⟩ that usually identifies the complaining function, syntactic form, or other entity, but may also refer to an entity being complained about; the ⟨*name*⟩ is also followed by a colon and space when present.

The ⟨*message*⟩ should be relatively short, and it should be largely independent of specific values that triggered the error. More detailed explanation that requires multiple lines should continue with each line indented by a single space, in which case ⟨*message*⟩ should end in a semi-colon (but the semi-colon should be omitted if ⟨*continued-message*⟩ is not present). Message text should be lowercase—using semi-colons to separate sentences if needed, although long explanations may be better deferred to extra fields.

Specific values that triggered the error or other helpful information should appear in separate ⟨*field*⟩ lines, each of which is indented by two spaces. If a ⟨*detail*⟩ is especially long or takes multiple lines, it should start on its own line after the ⟨*field*⟩ label, and each of its lines should be indented by three spaces. Field names should be all lowercase.

A ⟨*field*⟩ name should end with `...` if the field provides relatively detailed information that might be distracting in common cases but useful in others. For example, when a contract failure is reported for a particular argument of a function, other arguments to the function

might be shown in an “other arguments...” field. The intent is that fields whose names end in ... might be hidden by default in an environment such as DrRacket.

Make ⟨*field*⟩ names as short as possible, relying on ⟨*message*⟩ or ⟨*continued message*⟩ text to clarify the meaning for a field. For example, prefer “given” to “given turtle” as a field name, where ⟨*message*⟩ is something like “given turtle is too sleepy” to clarify that “given” refers to a turtle.

9.2.2 Raising Exceptions

```
(raise v [barrier?]) → any
  v : any/c
  barrier? : any/c = #t
```

Raises an exception, where *v* represents the exception being raised. The *v* argument can be anything; it is passed to the current exception handler.

If *barrier?* is true, then the call to the exception handler is protected by a continuation barrier, so that multiple returns/escapes are impossible. All exceptions raised by `racket` functions effectively use `raise` with a `#t` value for *barrier?*.

Breaks are disabled from the time the exception is raised until the exception handler obtains control, and the handler itself is `parameterize-breaked` to disable breaks initially; see §9.6 “Breaks” for more information on breaks.

Examples:

```
> (with-handlers ([number? (lambda (n)
                             (+ n 5))])
    (raise 18 #t))
23
> (define-struct (my-exception exn:fail:user) ())

> (with-handlers ([my-exception? (lambda (e)
                                   #f)])
    (+ 5 (raise (make-my-exception
                 "failed"
                 (current-continuation-marks)))))
#f
> (raise 'failed #t)
uncaught exception: failed
```

```
(error sym) → any
  sym : symbol?
(error msg v ...) → any
```

```
  msg : string?
  v : any/c
(error src format v ...) → any
  src : symbol?
  format : string?
  v : any/c
```

Raises the exception `exn:fail`, which contains an error string. The different forms produce the error string in different ways:

- `(error sym)` creates a message string by concatenating `"error: "` with the string form of *sym*. Use this form sparingly.
- `(error msg v ...)` creates a message string by concatenating *msg* with string versions of the *v*s (as produced by the current error value conversion handler; see `error-value->string-handler`). A space is inserted before each *v*. Use this form sparingly, because it does not conform well to Racket's error message conventions; consider `raise-arguments-error`, instead.
- `(error src frmat v ...)` creates a message string equivalent to the string created by

  ```
  (format (string-append "~s: " frmat) src v ...)
  ```

  When possible, use functions such as `raise-argument-error`, instead, which construct messages that follow Racket's error message conventions.

In all cases, the constructed message string is passed to `make-exn:fail`, and the resulting exception is raised.

Examples:

```
> (error 'failed)
error: failed
> (error "failed" 23 'pizza (list 1 2 3))
failed 23 pizza (1 2 3)
> (error 'method-a "failed because ~a" "no argument supplied")
method-a: failed because no argument supplied
```

```
(raise-user-error sym) → any
  sym : symbol?
(raise-user-error msg v ...) → any
  msg : string?
  v : any/c
(raise-user-error src format v ...) → any
  src : symbol?
  format : string?
  v : any/c
```

Like error, but constructs an exception with make-exn:fail:user instead of make-exn:fail. The default error display handler does not show a "stack trace" for exn:fail:user exceptions (see §9.5 "Continuation Marks"), so raise-user-error should be used for errors that are intended for end users.

```
(raise-argument-error name expected v) → any
  name : symbol?
  expected : string?
  v : any/c
(raise-argument-error name
                      expected
                      bad-pos
                      v ...)   → any
  name : symbol?
  expected : string?
  bad-pos : exact-nonnegative-integer?
  v : any/c
```

Creates an exn:fail:contract value and raises it as an exception. The *name* argument is used as the source procedure's name in the error message. The *expected* argument is used as a description of the expected contract (i.e., as a string, but the string is intended to contain a contract expression).

In the first form, *v* is the value received by the procedure that does not have the expected type.

In the second form, the bad argument is indicated by an index *bad-pos* (counting from 0), and all of the original arguments *v* are provided (in order). The resulting error message names the bad argument and also lists the other arguments. If *bad-pos* is not less than the number of *v*s, the exn:fail:contract exception is raised.

Examples:

```
> (define (feed-machine bits)
    (if (not (integer? bits))
      (raise-argument-error 'feed-machine "integer?" bits)
      "fed the machine"))

> (feed-machine 'turkey)
feed-machine: contract violation
  expected: integer?
  given: 'turkey
> (define (feed-cow animal)
    (if (not (eq? animal 'cow))
      (raise-argument-error 'feed-cow "'cow" animal)
      "fed the cow"))
```

```
> (feed-cow 'turkey)
feed-cow: contract violation
  expected: 'cow
  given: 'turkey
> (define (feed-animals cow sheep goose cat)
    (if (not (eq? goose 'goose))
      (raise-argument-error 'feed-animals "'goose" 2 cow sheep goose cat)
      "fed the animals"))

> (feed-animals 'cow 'sheep 'dog 'cat)
feed-animals: contract violation
  expected: 'goose
  given: 'dog
  argument position: 3rd
  other arguments...:
   'cow
   'sheep
   'cat
```

```
(raise-result-error name expected v) → any
  name : symbol?
  expected : string?
  v : any/c
(raise-result-error name
                    expected
                    bad-pos
                    v ...)   → any
  name : symbol?
  expected : string?
  bad-pos : exact-nonnegative-integer?
  v : any/c
```

Like raise-argument-error, but the error message describe *v* as a "result" instead of an "argument."

```
(raise-arguments-error name
                       message
                       field
                       v ...
                       ...)     → any
  name : symbol?
  message : string?
  field : string?
  v : any/c
```

Creates an exn:fail:contract value and raises it as an exception. The *name* is used as the source procedure's name in the error message. The *message* is the error message;

if *message* contains newline characters, each extra line should be suitably indented (with one extra space at the start of each line), but it should not end with a newline character. Each *field* must have a corresponding *v*, and the two are rendered on their own line in the error message, with each *v* formatted using the error value conversion handler (see error-value->string-handler).

Example:

```
> (raise-arguments-error 'eat
                         "fish is smaller than its given meal"
                         "fish" 12
                         "meal" 13)
eat: fish is smaller than its given meal
  fish: 12
  meal: 13
```

```
(raise-range-error name
                   type-description
                   index-prefix
                   index
                   in-value
                   lower-bound
                   upper-bound
                   alt-lower-bound) → any
  name : symbol?
  type-description : string?
  index-prefix : string?
  index : exact-integer?
  in-value : any/c
  lower-bound : exact-integer?
  upper-bound : exact-integer?
  alt-lower-bound : (or/c #f exact-integer?)
```

Creates an exn:fail:contract value and raises it as an exception to report an out-of-range error. The *type-description* string describes the value for which the index is meant to select an element, and *index-prefix* is a prefix for the word "index." The *index* argument is the rejected index. The *in-value* argument is the value for which the index was meant. The *lower-bound* and *upper-bound* arguments specify the valid range of indices, inclusive. If *alt-lower-bound* is not #f, and if *index* is between *alt-lower-bound* and *upper-bound*, then the error is report as *index* being less than the "starting" index *lower-bound*.

Examples:

```
> (raise-range-error 'vector-ref "vector" "starting
" 5 #(1 2 3) 0 3)
vector-ref: starting index is out of range
```

```
  starting index: 5
  valid range: [0, 3]
  vector: '#(1 2 3)
> (raise-range-error 'vector-ref "vector" "ending
" 5 #(1 2 3) 0 3)
vector-ref: ending index is out of range
  ending index: 5
  valid range: [0, 3]
  vector: '#(1 2 3)
> (raise-range-error 'vector-ref "vector" "ending
" 1 #(1 2 3) 2 3 0)
vector-ref: ending index is smaller than starting index
  ending index: 1
  starting index: 2
  valid range: [0, 3]
  vector: '#(1 2 3)
```

```
(raise-type-error name expected v) → any
  name : symbol?
  expected : string?
  v : any/c
(raise-type-error name expected bad-pos v ...) → any
  name : symbol?
  expected : string?
  bad-pos : exact-nonnegative-integer?
  v : any/c
```

Like `raise-argument-error`, but with Racket's old formatting conventions, and where *expected* is used as a "type" description instead of a contract expression. Use `raise-argument-error` or `raise-result-error`, instead.

```
(raise-mismatch-error name
                      message
                      v ...+
                      ...+)    → any
  name : symbol?
  message : string?
  v : any/c
```

Similar to `raise-arguments-error`, but using Racket's old formatting conventions, with a required *v* immediately after the first *message* string, and with further *message* strings that are spliced into the message without line breaks or space. Use `raise-arguments-error`, instead.

```
(raise-arity-error name arity-v [arg-v ...]) → any
  name : (or/c symbol? procedure?)
```

```
           (or/c exact-nonnegative-integer?
                 arity-at-least?
 arity-v :       (listof
                  (or/c exact-nonnegative-integer?
                        arity-at-least?)))
 arg-v : any/c = #f
```

Creates an `exn:fail:contract:arity` value and `raise`s it as an exception. The *name* is used for the source procedure's name in the error message.

The *arity-v* value must be a possible result from `procedure-arity`, except that it does not have to be normalized (see `procedure-arity?` for the details of normalized arities); `raise-arity-error` will normalize the arity and use the normalized form in the error message. If *name* is a procedure, its actual arity is ignored.

The *arg-v* arguments are the actual supplied arguments, which are shown in the error message (using the error value conversion handler; see `error-value->string-handler`); also, the number of supplied *arg-v*s is explicitly mentioned in the message.

```
(raise-syntax-error  name
                     message
                    [expr
                     sub-expr
                     extra-sources]) → any
  name : (or/c symbol? #f)
  message : string?
  expr : any/c = #f
  sub-expr : any/c = #f
  extra-sources : (listof syntax?) = null
```

Creates an `exn:fail:syntax` value and `raise`s it as an exception. Macros use this procedure to report syntax errors.

The *name* argument is usually `#f` when *expr* is provided; it is described in more detail below. The *message* is used as the main body of the error message; if *message* contains newline characters, each new line should be suitably indented (with one space at the start), and it should not end with a newline character.

The optional *expr* argument is the erroneous source syntax object or S-expression (but the expression `#f` cannot be represented by itself; it must be wrapped as a syntax object). The optional *sub-expr* argument is a syntax object or S-expression (again, `#f` cannot represent itself) within *expr* that more precisely locates the error. Both may appear in the generated error-message text if `error-print-source-location` is `#t`. Source location information in the error-message text is similarly extracted from *sub-expr* or *expr* when at least one is a syntax object and `error-print-source-location` is `#t`.

If *sub-expr* is provided and not #f, it is used (in syntax form) for the exprs field of the generated exception record, else the *expr* is used if provided and not #f. In either case, the syntax object is consed onto *extra-sources* to produce the exprs field, or *extra-sources* is used directly for exprs if neither *expr* nor *sub-expr* is provided and not #f.

The form name used in the generated error message is determined through a combination of the *name*, *expr*, and *sub-expr* arguments:

- When *name* is #f, and when *expr* is either an identifier or a syntax pair containing an identifier as its first element, then the form name from the error message is the identifier's symbol.
- When *name* is #f and when *expr* is not an identifier or a syntax pair containing an identifier as its first element, then the form name in the error message is "?".
- symbol: When *name* is a symbol, then the symbol is used as the form name in the generated error message.

9.2.3 Handling Exceptions

```
(call-with-exception-handler f thunk) → any
  f : (any/c . -> . any)
  thunk : (-> any)
```

Installs *f* as the exception handler for the dynamic extent of the call to *thunk*. If an exception is raised during the evaluation of *thunk* (in an extension of the current continuation that does not have its own exception handler), then *f* is applied to the raised value in the continuation of the raise call (but the continuation is normally extended with a continuation barrier; see §1.1.12 "Prompts, Delimited Continuations, and Barriers" and raise).

Any procedure that takes one argument can be an exception handler. Normally, an exception handler escapes from the context of the raise call via abort-current-continuation or some other escape mechanism. To propagate an exception to the "previous" exception handler—that is, the exception handler associated with the rest of the continuation after the point where the called exception handler was associated with the continuation—an exception handler can simply return a result instead of escaping, in which case the raise call propagates the value to the previous exception handler (still in the dynamic extent of the call to raise, and under the same barrier, if any). If an exception handler returns a result and no previous handler is available, the uncaught-exception handler is used.

A call to an exception handler is parameterize-breaked to disable breaks, and it is wrapped with call-with-exception-handler to install an exception handler that reports both the original and newly raised exceptions via the error display handler and then escapes via the error escape handler.

```
(uncaught-exception-handler) → (any/c . -> . any)
(uncaught-exception-handler f) → void?
  f : (any/c . -> . any)
```

A parameter that determines an *uncaught-exception handler* used by `raise` when the relevant continuation has no exception handler installed with `call-with-exception-handler` or `with-handlers`. Unlike exception handlers installed with `call-with-exception-handler`, the uncaught-exception handler must not return a value when called by `raise`; if it returns, an exception is raised (to be handled by an exception handler that reports both the original and newly raised exception).

The default uncaught-exception handler prints an error message using the current error display handler (see `error-display-handler`) and then escapes by calling the current error escape handler (see `error-escape-handler`). The call to each handler is `parameterized` to set `error-display-handler` to the default error display handler, and it is `parameterize-breaked` to disable breaks. The call to the error escape handler is further parameterized to set `error-escape-handler` to the default error escape handler; if the error escape handler returns, then the default error escape handler is called.

When the current error display handler is the default handler, then the error-display call is parameterized to install an emergency error display handler that logs an error (see `log-error`) and never fails.

```
(with-handlers ([pred-expr handler-expr] ...)
  body ...+)
```

Evaluates each *pred-expr* and *handler-expr* in the order that they are specified, and then evaluates the *body*s with a new exception handler during its dynamic extent.

The new exception handler processes an exception only if one of the *pred-expr* procedures returns a true value when applied to the exception, otherwise the exception handler is invoked from the continuation of the `with-handlers` expression (by raising the exception again). If an exception is handled by one of the *handler-expr* procedures, the result of the entire `with-handlers` expression is the return value of the handler.

When an exception is raised during the evaluation of *body*s, each predicate procedure *pred-expr* is applied to the exception value; if a predicate returns a true value, the corresponding *handler-expr* procedure is invoked with the exception as an argument. The predicates are tried in the order that they are specified.

Before any predicate or handler procedure is invoked, the continuation of the entire `with-handlers` expression is restored, but also `parameterize-breaked` to disable breaks. Thus, breaks are disabled by default during the predicate and handler procedures (see §9.6 “Breaks”), and the exception handler is the one from the continuation of the `with-handlers` expression.

The exn:fail? procedure is useful as a handler predicate to catch all error exceptions. Avoid using (lambda (x) #t) as a predicate, because the exn:break exception typically should not be caught (unless it will be re-raised to cooperatively break). Beware, also, of catching and discarding exceptions, because discarding an error message can make debugging unnecessarily difficult.

```
(with-handlers* ([pred-expr handler-expr] ...)
  body ...+)
```

Like with-handlers, but if a *handler-expr* procedure is called, breaks are not explicitly disabled, and the handler call is in tail position with respect to the with-handlers* form.

9.2.4 Configuring Default Handling

```
(error-escape-handler) → (-> any)
(error-escape-handler proc) → void?
  proc : (-> any)
```

A parameter for the *error escape handler*, which takes no arguments and escapes from the dynamic context of an exception. The default error escape handler escapes using (abort-current-continuation (default-continuation-prompt-tag) void).

The error escape handler is normally called directly by an exception handler, in a parameterization that sets the error display handler and error escape handler to the default handlers, and it is normally parameterize-breaked to disable breaks. To escape from a run-time error in a different context, use raise or error.

Due to a continuation barrier around exception-handling calls, an error escape handler cannot invoke a full continuation that was created prior to the exception, but it can abort to a prompt (see call-with-continuation-prompt) or invoke an escape continuation (see call-with-escape-continuation).

```
(error-display-handler) → (string? any/c . -> . any)
(error-display-handler proc) → void?
  proc : (string? any/c . -> . any)
```

A parameter for the *error display handler*, which is called by the default exception handler with an error message and the exception value. More generally, the handler's first argument is a string to print as an error message, and the second is a value representing a raised exception.

The default error display handler displays its first argument to the current error port (determined by the current-error-port parameter) and extracts a stack trace (see continuation-mark-set->context) to display from the second argument if it is an exn value but not an exn:fail:user value.

The default error display handler in DrRacket also uses the second argument to highlight source locations.

To report a run-time error, use raise or procedures like error, instead of calling the error display handler directly.

```
(error-print-width) → (and/c exact-integer? (>=/c 3))
(error-print-width width) → void?
  width : (and/c exact-integer? (>=/c 3))
```

A parameter whose value is used as the maximum number of characters used to print a Racket value that is embedded in a primitive error message.

```
(error-print-context-length) → exact-nonnegative-integer?
(error-print-context-length cnt) → void?
  cnt : exact-nonnegative-integer?
```

A parameter whose value is used by the default error display handler as the maximum number of lines of context (or "stack trace") to print; a single "..." line is printed if more lines are available after the first *cnt* lines. A 0 value for *cnt* disables context printing entirely.

```
(error-value->string-handler)
    (any/c exact-nonnegative-integer?
 →         . -> .
           string?)
(error-value->string-handler proc) → void?
         (any/c exact-nonnegative-integer?
  proc :        . -> .
                string?)
```

A parameter that determines the *error value conversion handler*, which is used to print a Racket value that is embedded in a primitive error message.

The integer argument to the handler specifies the maximum number of characters that should be used to represent the value in the resulting string. The default error value conversion handler prints the value into a string (using the current global port print handler; see global-port-print-handler). If the printed form is too long, the printed form is truncated and the last three characters of the return string are set to "...".

If the string returned by an error value conversion handler is longer than requested, the string is destructively "truncated" by setting the first extra position in the string to the null character. If a non-string is returned, then the string "..." is used. If a primitive error string needs to be generated before the handler has returned, the default error value conversion handler is used.

Calls to an error value conversion handler are parameterized to re-install the default error value conversion handler, and to enable printing of unreadable values (see print-unreadable).

```
(error-print-source-location) → boolean?
(error-print-source-location include?) → void?
  include? : any/c
```

A parameter that controls whether read and syntax error messages include source information, such as the source line and column or the expression. This parameter also controls the error message when a module-defined variable is accessed before its definition is executed; the parameter determines whether the message includes a module name. Only the message field of an `exn:fail:read`, `exn:fail:syntax`, or `exn:fail:contract:variable` structure is affected by the parameter. The default is `#t`.

9.2.5 Built-in Exception Types

```
(struct exn (message continuation-marks)
        #:extra-constructor-name make-exn
        #:transparent)
  message : string?
  continuation-marks : continuation-mark-set?
```

The base structure type for exceptions. The `message` field contains an error message, and the `continuation-marks` field contains the value produced by `(current-continuation-marks)` immediately before the exception was raised.

```
(struct exn:fail exn ()
        #:extra-constructor-name make-exn:fail
        #:transparent)
```

Raised for exceptions that represent errors, as opposed to `exn:break`.

```
(struct exn:fail:contract exn:fail ()
        #:extra-constructor-name make-exn:fail:contract
        #:transparent)
```

Raised for errors from the inappropriate run-time use of a function or syntactic form.

```
(struct exn:fail:contract:arity exn:fail:contract ()
        #:extra-constructor-name make-exn:fail:contract:arity
        #:transparent)
```

Raised when a procedure is applied to the wrong number of arguments.

```
(struct exn:fail:contract:divide-by-zero exn:fail:contract ()
        #:extra-constructor-name
        make-exn:fail:contract:divide-by-zero
        #:transparent)
```

Raised for division by exact zero.

```
(struct exn:fail:contract:non-fixnum-result exn:fail:contract ()
        #:extra-constructor-name
        make-exn:fail:contract:non-fixnum-result
        #:transparent)
```

Raised by functions like `fx+` when the result would not be a fixnum.

```
(struct exn:fail:contract:continuation exn:fail:contract ()
        #:extra-constructor-name
        make-exn:fail:contract:continuation
        #:transparent)
```

Raised when a continuation is applied where the jump would cross a continuation barrier.

```
(struct exn:fail:contract:variable exn:fail:contract (id)
        #:extra-constructor-name make-exn:fail:contract:variable
        #:transparent)
  id : symbol?
```

Raised for a reference to a not-yet-defined top-level variable or module-level variable.

```
(struct exn:fail:syntax exn:fail (exprs)
        #:extra-constructor-name make-exn:fail:syntax
        #:transparent)
  exprs : (listof syntax?)
```

Raised for a syntax error that is not a `read` error. The `exprs` indicate the relevant source expressions, least-specific to most-specific.

```
(struct exn:fail:syntax:unbound exn:fail:syntax ()
        #:extra-constructor-name make-exn:fail:syntax:unbound
        #:transparent)
```

Raised by `#%top` or `set!` for an unbound identifier within a module.

```
(struct exn:fail:read exn:fail (srclocs)
        #:extra-constructor-name make-exn:fail:read
        #:transparent)
  srclocs : (listof srcloc?)
```

Raised for a `read` error. The `srclocs` indicate the relevant source expressions.

```
(struct exn:fail:read:eof exn:fail:read ()
        #:extra-constructor-name make-exn:fail:read:eof
        #:transparent)
```

Raised for a `read` error, specifically when the error is due to an unexpected end-of-file.

```
(struct exn:fail:read:non-char exn:fail:read ()
        #:extra-constructor-name make-exn:fail:read:non-char
        #:transparent)
```

Raised for a `read` error, specifically when the error is due to an unexpected non-character (i.e., “special”) element in the input stream.

```
(struct exn:fail:filesystem exn:fail ()
        #:extra-constructor-name make-exn:fail:filesystem
        #:transparent)
```

Raised for an error related to the filesystem (such as a file not found).

```
(struct exn:fail:filesystem:exists exn:fail:filesystem ()
        #:extra-constructor-name make-exn:fail:filesystem:exists
        #:transparent)
```

Raised for an error when attempting to create a file that exists already.

```
(struct exn:fail:filesystem:version exn:fail:filesystem ()
        #:extra-constructor-name make-exn:fail:filesystem:version
        #:transparent)
```

Raised for a version-mismatch error when loading an extension.

```
(struct exn:fail:network exn:fail ()
        #:extra-constructor-name make-exn:fail:network
        #:transparent)
```

Raised for TCP and UDP errors.

```
(struct exn:fail:out-of-memory exn:fail ()
        #:extra-constructor-name make-exn:fail:out-of-memory
        #:transparent)
```

Raised for an error due to insufficient memory, in cases where sufficient memory is at least available for raising the exception.

```
(struct exn:fail:unsupported exn:fail ()
        #:extra-constructor-name make-exn:fail:unsupported
        #:transparent)
```

Raised for an error due to an unsupported feature on the current platform or configuration.

```
(struct exn:fail:user exn:fail ()
        #:extra-constructor-name make-exn:fail:user
        #:transparent)
```

Raised for errors that are intended to be seen by end users. In particular, the default error printer does not show the program context when printing the error message.

```
(struct exn:break exn (continuation)
        #:extra-constructor-name make-exn:break
        #:transparent)
  continuation : continuation?
```

Raised asynchronously (when enabled) in response to a break request. The `continuation` field can be used by a handler to resume the interrupted computation.

```
prop:exn:srclocs : struct-type-property?
```

A property that identifies structure types that provide a list of `srcloc` values. The property is normally attached to structure types used to represent exception information.

The property value must be a procedure that accepts a single value—the structure type instance from which to extract source locations—and returns a list of `srcloc`s. Some error display handlers use only the first returned location.

As an example,

```
#lang racket

;; We create a structure that supports the
;; prop:exn:srcloc protocol.  It carries
;; with it the location of the syntax that
;; is guilty.
(define-struct (exn:fail:he-who-shall-not-be-named
                exn:fail)
  (a-srcloc)
  #:property prop:exn:srclocs
  (lambda (a-struct)
    (match a-struct
```

```
      [(struct exn:fail:he-who-shall-not-be-named
         (msg marks a-srcloc))
       (list a-srcloc)])))

;; We can play with this by creating a form that
;; looks at identifiers, and only flags specific ones.
(define-syntax (skeeterize stx)
  (syntax-case stx ()
    [(_ expr)
     (cond
       [(and (identifier? #'expr)
             (eq? (syntax-e #'expr) 'voldemort))
        (quasisyntax/loc stx
          (raise (make-exn:fail:he-who-shall-not-be-named
                  "oh dear don't say his name"
                  (current-continuation-marks)
                  (srcloc '#,(syntax-source #'expr)
                          '#,(syntax-line #'expr)
                          '#,(syntax-column #'expr)
                          '#,(syntax-position #'expr)
                          '#,(syntax-span #'expr)))))]
       [else
        ;; Otherwise, leave the expression alone.
        #'expr])]))

(define (f x)
  (* (skeeterize x) x))

(define (g voldemort)
  (* (skeeterize voldemort) voldemort))

;; Examples:
(f 7)
(g 7)
;; The error should highlight the use
;; of the one-who-shall-not-be-named
;; in g.
```

```
(exn:srclocs? v) → boolean?
  v : any/c
```

Returns `#t` if *v* has the `prop:exn:srclocs` property, `#f` otherwise.

```
(exn:srclocs-accessor v)
 → (exn:srclocs? . -> . (listof srcloc))
  v : exn:srclocs?
```

Returns the srcloc-getting procedure associated with v.

```
(struct srcloc (source line column position span)
        #:extra-constructor-name make-srcloc
        #:transparent)
  source : any/c
  line : (or/c exact-positive-integer? #f)
  column : (or/c exact-nonnegative-integer? #f)
  position : (or/c exact-positive-integer? #f)
  span : (or/c exact-nonnegative-integer? #f)
```

The fields of a srcloc instance are as follows:

- source — An arbitrary value identifying the source, often a path (see §14.1 “Paths”).
- line — The line number (counts from 1) or #f (unknown).
- column — The column number (counts from 0) or #f (unknown).
- position — The starting position (counts from 1) or #f (unknown).
- span — The number of covered positions (counts from 0) or #f (unknown).

9.3 Delayed Evaluation

```
(require racket/promise)
```

The bindings documented in this section are provided by the racket/promise and racket libraries, but not racket/base.

A *promise* encapsulates an expression to be evaluated on demand via force. After a promise has been forced, every later force of the promise produces the same result.

```
(promise? v) → boolean?
  v : any/c
```

Returns #t if v is a promise, #f otherwise.

```
(delay body ...+)
```

Creates a promise that, when forced, evaluates the *body*s to produce its value. The result is then cached, so further uses of force produce the cached value immediately. This includes multiple values and exceptions.

```
(lazy body ...+)
```

Like `delay`, if the last *body* produces a promise when forced, then this promise is `force`d, too, to obtain a value. In other words, this form creates a composable promise, where the computation of its body is "attached" to the computation of the following promise, and a single `force` iterates through the whole chain, tail-calling each step.

Note that the last *body* of this form must produce a single value, but the value can itself be a `delay` promise that returns multiple values.

The `lazy` form is useful for implementing lazy libraries and languages, where tail calls can be wrapped in a promise.

```
(force v) → any
  v : any/c
```

If *v* is a promise, then the promise is forced to obtain a value. If the promise has not been forced before, then the result is recorded in the promise so that future `force`s on the promise produce the same value (or values). If forcing the promise raises an exception, then the exception is similarly recorded so that forcing the promise will raise the same exception every time.

If *v* is `force`d again before the original call to `force` returns, then the `exn:fail` exception is raised.

If *v* is not a promise, then it is returned as the result.

```
(promise-forced? promise) → boolean?
  promise : promise?
```

Returns `#t` if *promise* has been forced.

```
(promise-running? promise) → boolean?
  promise : promise?
```

Returns `#t` if *promise* is currently being forced. (Note that a promise can be either running or forced but not both.)

9.3.1 Additional Promise Kinds

```
(delay/name body ...+)
```

Creates a "call-by-name" promise that is similar to `delay`-promises, except that the resulting value is not cached. This kind of promise is essentially a thunk that is wrapped in a way that `force` recognizes.

If a `delay/name` promise forces itself, no exception is raised, the promise is never considered "running" or "forced" in the sense of `promise-running?` and `promise-forced?`.

```
(delay/strict body ...+)
```

Creates a “strict” promise: it is evaluated immediately, and the result is wrapped in a promise value. Note that the body can evaluate to multiple values, and forcing the resulting promise will return these values.

```
(delay/sync body ...+)
```

Produces a promise where an attempt to `force` the promise by a thread other than one currently running the promise causes the `force` to block until a result is available. This kind of promise is also a synchronizable event for use with `sync`; `sync`ing on the promise does not `force` it, but merely waits until a value is forced by another thread.

If a promise created by `delay/sync` is forced on a thread that is already running the promise, an exception is raised in the same way as for promises created with `delay`.

```
(delay/thread body/option ...+)

body/option = body
            | #:group thread-group-expr
```

Like `delay/sync`, but begins the computation immediately on a newly created thread. The thread is created under the thread group specified by *thread-group-expr*, which defaults to `(make-thread-group)`. A `#:group` specification can appear at most once.

Exceptions raised by the *body*s are caught as usual and raised only when the promise is `force`d.

```
(delay/idle body/option ...+)

body/option = body
            | #:wait-for wait-evt-expr
            | #:work-while while-evt-expr
            | #:tick tick-secs-expr
            | #:use use-ratio-expr
```

Like `delay/thread`, but with the following differences:

- the computation does not start until the event produced by *wait-evt-expr* is ready, where the default is `(system-idle-evt)`;
- the computation thread gets to work only when the process is otherwise idle as determined by *while-evt-expr*, which also defaults to `(system-idle-evt)`;

- the thread is allowed to run only periodically: out of every *tick-secs-expr* (defaults to 0.2) seconds, the thread is allowed to run *use-ratio-expr* (defaults to 0.12) of the time proportionally; i.e., the thread runs for (* *tick-secs-expr* *use-ratio-expr*) seconds.

If the promise is forced before the computation is done, it runs the rest of the computation immediately without waiting on events or periodically restricting evaluation.

A #:wait-for, #:work-while, #:tick, or #:use specification can appear at most once.

9.4 Continuations

See §1.1.1 “Sub-expression Evaluation and Continuations” and §1.1.12 “Prompts, Delimited Continuations, and Barriers” for general information about continuations. Racket’s support for prompts and composable continuations most closely resembles Dorai Sitaram’s % and fcontrol operator [Sitaram93].

Racket installs a continuation barrier around evaluation in the following contexts, preventing full-continuation jumps into the evaluation context protected by the barrier:

- applying an exception handler, an error escape handler, or an error display handler (see §9.2 “Exceptions”);
- applying a macro transformer (see §11.4 “Syntax Transformers”), evaluating a compile-time expression, or applying a module name resolver (see §13.4.1 “Resolving Module Names”);
- applying a custom-port procedure (see §12.1.9 “Custom Ports”), an event guard procedure (see §10.2.1 “Events”), or a parameter guard procedure (see §10.3.2 “Parameters”);
- applying a security-guard procedure (see §13.6 “Security Guards”);
- applying a will procedure (see §15.3 “Wills and Executors”); or
- evaluating or loading code from the stand-alone Racket command line (see §17.1 “Running Racket or GRacket”).

In addition, extensions of Racket may install barriers in additional contexts. Finally, call-with-continuation-barrier applies a thunk barrier between the application and the current continuation.

```
(call-with-continuation-prompt proc
                               [prompt-tag
                                handler]
                               arg ...)   → any
```

```
proc : procedure?
prompt-tag : continuation-prompt-tag?
           = (default-continuation-prompt-tag)
handler : (or/c procedure? #f) = #f
arg : any/c
```

Applies *proc* to the given *arg*s with the current continuation extended by a prompt. The prompt is tagged by *prompt-tag*, which must be a result from either default-continuation-prompt-tag (the default) or make-continuation-prompt-tag. The result of *proc* is the result of the call-with-continuation-prompt call.

The *handler* argument specifies a handler procedure to be called in tail position with respect to the call-with-continuation-prompt call when the installed prompt is the target of an abort-current-continuation call with *prompt-tag*; the remaining arguments of abort-current-continuation are supplied to the handler procedure. If *handler* is #f, the default handler accepts a single *abort-thunk* argument and calls (call-with-continuation-prompt *abort-thunk prompt-tag* #f); that is, the default handler reinstalls the prompt and continues with a given thunk.

```
(abort-current-continuation prompt-tag
                            v ...+)    → any
  prompt-tag : any/c
  v : any/c
```

Resets the current continuation to that of the nearest prompt tagged by *prompt-tag* in the current continuation; if no such prompt exists, the exn:fail:contract:continuation exception is raised. The *v*s are delivered as arguments to the target prompt's handler procedure.

The protocol for *v*s supplied to an abort is specific to the *prompt-tag*. When abort-current-continuation is used with (default-continuation-prompt-tag), generally, a single thunk should be supplied that is suitable for use with the default prompt handler. Similarly, when call-with-continuation-prompt is used with (default-continuation-prompt-tag), the associated handler should generally accept a single thunk argument.

```
(make-continuation-prompt-tag) → continuation-prompt-tag?
(make-continuation-prompt-tag sym) → continuation-prompt-tag?
  sym : symbol?
```

Creates a prompt tag that is not equal? to the result of any other value (including prior or future results from make-continuation-prompt-tag). The optional *sym* argument, if supplied, is used when printing the prompt tag.

```
(default-continuation-prompt-tag) → continuation-prompt-tag?
```

Returns a constant prompt tag for which a prompt is installed at the start of every thread's continuation; the handler for each thread's initial prompt accepts any number of values and returns. The result of `default-continuation-prompt-tag` is the default tag for any procedure that accepts a prompt tag.

```
(call-with-current-continuation proc
                                [prompt-tag]) → any
  proc : (continuation? . -> . any)
  prompt-tag : continuation-prompt-tag?
             = (default-continuation-prompt-tag)
```

Captures the current continuation up to the nearest prompt tagged by *prompt-tag*; if no such prompt exists, the `exn:fail:contract:continuation` exception is raised. The truncated continuation includes only continuation marks and `dynamic-wind` frames installed since the prompt.

The capture continuation is delivered to *proc*, which is called in tail position with respect to the `call-with-current-continuation` call.

If the continuation argument to *proc* is ever applied, then it removes the portion of the current continuation up to the nearest prompt tagged by *prompt-tag* (not including the prompt; if no such prompt exists, the `exn:fail:contract:continuation` exception is raised), or up to the nearest continuation frame (if any) shared by the current and captured continuations—whichever is first. While removing continuation frames, `dynamic-wind` *post-thunk*s are executed. Finally, the (unshared portion of the) captured continuation is appended to the remaining continuation, applying `dynamic-wind` *pre-thunk*s.

The arguments supplied to an applied procedure become the result values for the restored continuation. In particular, if multiple arguments are supplied, then the continuation receives multiple results.

If, at application time, a continuation barrier would be introduced by replacing the current continuation with the applied one, then the `exn:fail:contract:continuation` exception is raised.

A continuation can be invoked from the thread (see §10.1 "Threads") other than the one where it was captured.

```
(call/cc proc [prompt-tag]) → any
  proc : (continuation? . -> . any)
  prompt-tag : continuation-prompt-tag?
             = (default-continuation-prompt-tag)
```

The `call/cc` binding is an alias for `call-with-current-continuation`.

```
(call-with-composable-continuation proc
                                   [prompt-tag]) → any
```

```
  proc : (continuation? . -> . any)
  prompt-tag : continuation-prompt-tag?
             = (default-continuation-prompt-tag)
```

Similar to call-with-current-continuation, but applying the resulting continuation procedure does not remove any portion of the current continuation. Instead, application always extends the current continuation with the captured continuation (without installing any prompts other than those captured in the continuation).

When call-with-composable-continuation is called, if a continuation barrier appears in the continuation before the closest prompt tagged by *prompt-tag*, the exn:fail:contract:continuation exception is raised (because attempting to apply the continuation would always fail).

```
(call-with-escape-continuation proc) → any
  proc : (continuation? . -> . any)
```

Like call-with-current-continuation, but *proc* is not called in tail position, and the continuation procedure supplied to *proc* can only be called during the dynamic extent of the call-with-escape-continuation call.

Due to the limited applicability of its continuation, call-with-escape-continuation can be implemented more efficiently than call-with-current-continuation.

A continuation obtained from call-with-escape-continuation is actually a kind of prompt. Escape continuations are provided mainly for backwards compatibility, since they pre-date general prompts in Racket, and because call/ec is often an easy replacement for call/cc to improve performance.

```
(call/ec proc) → any
  proc : (continuation? . -> . any)
```

The call/ec binding is an alias for call-with-escape-continuation.

```
(let/cc k body ...+)
```

Equivalent to (call/cc (lambda (k) body ...)).

```
(let/ec k body ...+)
```

Equivalent to (call/ec (lambda (k) body ...)).

```
(call-with-continuation-barrier thunk) → any
  thunk : (-> any)
```

Applies *thunk* with a continuation barrier between the application and the current continuation. The results of *thunk* are the results of the call-with-continuation-barrier call.

```
(continuation-prompt-available? prompt-tag
                                [cont])      → any
  prompt-tag : continuation-prompt-tag?
  cont : continuation? = (call/cc values)
```

Returns #t if *cont*, which must be a continuation, includes a prompt tagged by *prompt-tag*, #f otherwise.

```
(continuation? v) → boolean?
  v : any/c
```

Return #t if *v* is a continuation as produced by call-with-current-continuation, call-with-composable-continuation, or call-with-escape-continuation, #f otherwise.

```
(continuation-prompt-tag? v) → boolean?
  v : any/c
```

Returns #t if *v* is a continuation prompt tag as produced by default-continuation-prompt-tag or make-continuation-prompt-tag.

```
(dynamic-wind pre-thunk
              value-thunk
              post-thunk) → any
  pre-thunk : (-> any)
  value-thunk : (-> any)
  post-thunk : (-> any)
```

Applies its three thunk arguments in order. The value of a dynamic-wind expression is the value returned by *value-thunk*. The *pre-thunk* procedure is invoked before calling *value-thunk* and *post-thunk* is invoked after *value-thunk* returns. The special properties of dynamic-wind are manifest when control jumps into or out of the *value-thunk* application (either due to a prompt abort or a continuation invocation): every time control jumps into the *value-thunk* application, *pre-thunk* is invoked, and every time control jumps out of *value-thunk*, *post-thunk* is invoked. (No special handling is performed for jumps into or out of the *pre-thunk* and *post-thunk* applications.)

When dynamic-wind calls *pre-thunk* for normal evaluation of *value-thunk*, the continuation of the *pre-thunk* application calls *value-thunk* (with dynamic-wind's special jump handling) and then *post-thunk*. Similarly, the continuation of the *post-thunk* application returns the value of the preceding *value-thunk* application to the continuation of the entire dynamic-wind application.

When *pre-thunk* is called due to a continuation jump, the continuation of *pre-thunk*

- jumps to a more deeply nested *pre-thunk*, if any, or jumps to the destination continuation; then
- continues with the context of the *pre-thunk*'s dynamic-wind call.

Normally, the second part of this continuation is never reached, due to a jump in the first part. However, the second part is relevant because it enables jumps to escape continuations that are contained in the context of the dynamic-wind call. Furthermore, it means that the continuation marks (see §9.5 "Continuation Marks") and parameterization (see §10.3.2 "Parameters") for *pre-thunk* correspond to those of the dynamic-wind call that installed *pre-thunk*. The *pre-thunk* call, however, is parameterize-breaked to disable breaks (see also §9.6 "Breaks").

Similarly, when *post-thunk* is called due to a continuation jump, the continuation of *post-thunk* jumps to a less deeply nested *post-thunk*, if any, or jumps to a *pre-thunk* protecting the destination, if any, or jumps to the destination continuation, then continues from the *post-thunk*'s dynamic-wind application. As for *pre-thunk*, the parameterization of the original dynamic-wind call is restored for the call, and the call is parameterize-breaked to disable breaks.

In both cases, the target for a jump is recomputed after each *pre-thunk* or *post-thunk* completes. When a prompt-delimited continuation (see §1.1.12 "Prompts, Delimited Continuations, and Barriers") is captured in a *post-thunk*, it might be delimited and instantiated in such a way that the target of a jump turns out to be different when the continuation is applied than when the continuation was captured. There may even be no appropriate target, if a relevant prompt or escape continuation is not in the continuation after the restore; in that case, the first step in a *pre-thunk* or *post-thunk*'s continuation can raise an exception.

Examples:

```
> (let ([v (let/ec out
             (dynamic-wind
              (lambda () (display "in "))
              (lambda ()
                (display "pre ")
                (display (call/cc out))
                #f)
              (lambda () (display "out "))))])
    (when v (v "post ")))
in pre out in post out

> (let/ec k0
    (let/ec k1
      (dynamic-wind
```

```
       void
       (lambda () (k0 'cancel))
       (lambda () (k1 'cancel-canceled)))))
'cancel-canceled
> (let* ([x (make-parameter 0)]
         [l null]
         [add (lambda (a b)
                (set! l (append l (list (cons a b))))])
    (let ([k (parameterize ([x 5])
               (dynamic-wind
                   (lambda () (add 1 (x)))
                   (lambda () (parameterize ([x 6])
                                (let ([k+e (let/cc k (cons k void))])
                                  (add 2 (x))
                                  ((cdr k+e))
                                  (car k+e))))
                   (lambda () (add 3 (x)))))])
      (parameterize ([x 7])
        (let/cc esc
          (k (cons void esc)))))
    l)
'((1 . 5) (2 . 6) (3 . 5) (1 . 5) (2 . 6) (3 . 5))
```

9.4.1 Additional Control Operators

```
(require racket/control)
```

The bindings documented in this section are provided by the `racket/control` library, not `racket/base` or `racket`.

The `racket/control` library provides various control operators from the research literature on higher-order control operators, plus a few extra convenience forms. These control operators are implemented in terms of `call-with-continuation-prompt`, `call-with-composable-continuations`, etc., and they generally work sensibly together. Many are redundant; for example, `reset` and `prompt` are aliases.

```
(call/prompt proc [prompt-tag handler] arg ...) → any
  proc : procedure?
  prompt-tag : continuation-prompt-tag?
             = (default-continuation-prompt-tag)
  handler : (or/c procedure? #f) = #f
  arg : any/c
```

The `call/prompt` binding is an alias for `call-with-continuation-prompt`.

```
(abort/cc prompt-tag v ...+) → any
  prompt-tag : any/c
  v : any/c
```

The abort/cc binding is an alias for abort-current-continuation.

```
(call/comp proc [prompt-tag]) → any
  proc : (continuation? . -> . any)
  prompt-tag : continuation-prompt-tag?
             = (default-continuation-prompt-tag)
```

The call/comp binding is an alias for call-with-composable-continuation.

```
(abort v ...) → any
  v : any/c
```

Returns the vs to a prompt using the default continuation prompt tag and the default abort handler.

That is, (abort v ...) is equivalent to

```
(abort-current-continuation
 (default-continuation-prompt-tag)
 (lambda () (values v ...)))
```

Example:

```
> (prompt
    (printf "start here\n")
    (printf "answer is ~a\n" (+ 2 (abort 3))))
start here
3
```

```
(% expr)
(% expr handler-expr)
(% expr handler-expr #:tag tag-expr)
(fcontrol v #:tag prompt-tag) → any
  v : any/c
  prompt-tag : (default-continuation-prompt-tag)
```

Sitaram's operators [Sitaram93].

The essential reduction rules are:

```
(% val proc) => val
(% E[(fcontrol val)] proc) => (proc val (lambda (x) E[x]))
  ; where E has no %
```

When *handler-expr* is omitted, % is the same as `prompt`. If *prompt-tag* is provided, % uses specific prompt tags like `prompt-at`.

Examples:

```
> (% (+ 2 (fcontrol 5))
     (lambda (v k)
       (k v)))
7
> (% (+ 2 (fcontrol 5))
     (lambda (v k)
       v))
5
```

```
(prompt expr ...+)
(control id expr ...+)
```

Among the earliest operators for higher-order control [Felleisen88a, Felleisen88, Sitaram90].

The essential reduction rules are:

```
(prompt val) => val
(prompt E[(control k expr)]) => (prompt ((lambda (k) expr)
                                         (lambda (v) E[v])))
  ; where E has no prompt
```

Examples:

```
> (prompt
    (+ 2 (control k (k 5))))
7
> (prompt
    (+ 2 (control k 5)))
5
> (prompt
    (+ 2 (control k (+ 1 (control k1 (k1 6))))))
7
> (prompt
    (+ 2 (control k (+ 1 (control k1 (k 6))))))
8
> (prompt
    (+ 2 (control k (control k1 (control k2 (k2 6))))))
6
```

```
(prompt-at prompt-tag-expr expr ...+)
(control-at prompt-tag-expr id expr ...+)
```

Like `prompt` and `control`, but using specific prompt tags:

```
(prompt-at tag val) => val
(prompt-at tag E[(control-at tag k expr)]) => (prompt-at tag
                                                ((lambda (k) expr)
                                                 (lambda (v) E[v])))
  ; where E has no prompt-at for tag
```

```
(reset expr ...+)
(shift id expr ...+)
```

Danvy and Filinski's operators [Danvy90].

The essential reduction rules are:

```
(reset val) => val
(reset E[(shift k expr)]) => (reset ((lambda (k) expr)
                                     (lambda (v) (reset E[v]))))
  ; where E has no reset
```

The `reset` and `prompt` forms are interchangeable.

```
(reset-at prompt-tag-expr expr ...+)
(shift-at prompt-tag-expr identifier expr ...+)
```

Like `reset` and `shift`, but using the specified prompt tags.

```
(prompt0 expr ...+)
(reset0 expr ...+)
(control0 id expr ...+)
(shift0 id expr ...+)
```

Generalizations of `prompt`, etc. [Shan04].

The essential reduction rules are:

```
(prompt0 val) => val
(prompt0 E[(control0 k expr)]) => ((lambda (k) expr)
                                   (lambda (v) E[v]))
(reset0 val) => val
(reset0 E[(shift0 k expr)]) => ((lambda (k) expr)
                                (lambda (v) (reset0 E[v])))
```

The `reset0` and `prompt0` forms are interchangeable. Furthermore, the following reductions apply:

```
(prompt E[(control0 k expr)]) => (prompt ((lambda (k) expr)
                                          (lambda (v) E[v])))
(reset E[(shift0 k expr)]) => (reset ((lambda (k) expr)
                                      (lambda (v) (reset0 E[v]))))
(prompt0 E[(control k expr)]) => (prompt0 ((lambda (k) expr)
                                           (lambda (v) E[v])))
(reset0 E[(shift k expr)]) => (reset0 ((lambda (k) expr)
                                       (lambda (v) (reset E[v]))))
```

That is, both the `prompt`/`reset` and `control`/`shift` sites must agree for `0`-like behavior, otherwise the non-`0` behavior applies.

```
(prompt0-at prompt-tag-expr expr ...+)
(reset0-at prompt-tag-expr expr ...+)
(control0-at prompt-tag-expr id expr ...+)
(shift0-at prompt-tag-expr id expr ...+)
```

Variants of `prompt0`, etc., that accept a prompt tag.

```
(spawn proc) → any
  proc : ((any/c . -> . any) . -> . any)
```

The operators of Hieb and Dybvig [Hieb90].

The essential reduction rules are:

```
(prompt-at tag obj) => obj
(spawn proc) => (prompt tag (proc (lambda (x) (abort tag x))))
(prompt-at tag E[(abort tag proc)])
  => (proc (lambda (x) (prompt-at tag E[x])))
  ; where E has no prompt-at for tag
```

```
(splitter proc) → any
        (((-> any) . -> . any)
  proc :  ((continuation? . -> . any) . -> . any)
          . -> . any)
```

The operator of Queinnec and Serpette [Queinnec91].

The essential reduction rules are:

```
(splitter proc) => (prompt-at tag
                     (proc (lambda (thunk)
```

```
                              (abort tag thunk))
                            (lambda (proc)
                              (control0-at tag k (proc k)))))
(prompt-at tag E[(abort tag thunk)]) => (thunk)
  ; where E has no prompt-at for tag
(prompt-at tag E[(control0-at tag k expr)]) => ((lambda (k) expr)
                                                (lambda (x) E[x]))
  ; where E has no prompt-at for tag
```

```
(new-prompt) → any
(set prompt-expr expr ...+)
(cupto prompt-expr id expr ...+)
```

The operators of Gunter et al. [Gunter95].

In this library, `new-prompt` is an alias for `make-continuation-prompt-tag`, `set` is an alias for `prompt0-at`, and `cupto` is an alias for `control0-at`.

9.5 Continuation Marks

See §1.1.11 “Continuation Frames and Marks” and §1.1.12 “Prompts, Delimited Continuations, and Barriers” for general information about continuation marks.

The list of continuation marks for a key k and a continuation C that extends C_0 is defined as follows:

- If C is an empty continuation, then the mark list is `null`.
- If C’s first frame contains a mark m for k, then the mark list for C is `(cons m lst)`, where `lst` is the mark list for k in C_0.
- If C’s first frame does not contain a mark keyed by k, then the mark list for C is the mark list for C_0.

The `with-continuation-mark` form installs a mark on the first frame of the current continuation (see §2.19 “Continuation Marks: `with-continuation-mark`”). Procedures such as `current-continuation-marks` allow inspection of marks.

Whenever Racket creates an exception record for a primitive exception, it fills the `continuation-marks` field with the value of `(current-continuation-marks)`, thus providing a snapshot of the continuation marks at the time of the exception.

When a continuation procedure returned by `call-with-current-continuation` or `call-with-composable-continuation` is invoked, it restores the captured continuation,

and also restores the marks in the continuation's frames to the marks that were present when `call-with-current-continuation` or `call-with-composable-continuation` was invoked.

```
(continuation-marks cont [prompt-tag]) → continuation-mark-set?
  cont : (or/c continuation? thread? #f)
  prompt-tag : continuation-prompt-tag?
             = (default-continuation-prompt-tag)
```

Returns an opaque value containing the set of continuation marks for all keys in the continuation *cont* (or the current continuation of *cont* if it is a thread) up to the prompt tagged by *prompt-tag*. If *cont* is `#f`, the resulting set of continuation marks is empty. If *cont* is an escape continuation (see §1.1.12 "Prompts, Delimited Continuations, and Barriers"), then the current continuation must extend *cont*, or the `exn:fail:contract` exception is raised. If *cont* was not captured with respect to *prompt-tag* and does not include a prompt for *prompt-tag*, the `exn:fail:contract` exception is raised. If *cont* is a dead thread, the result is an empty set of continuation marks.

```
(current-continuation-marks [prompt-tag])
 → continuation-mark-set?
  prompt-tag : continuation-prompt-tag?
             = (default-continuation-prompt-tag)
```

Returns an opaque value containing the set of continuation marks for all keys in the current continuation up to *prompt-tag*. In other words, it produces the same value as

```
(call-with-current-continuation
  (lambda (k)
    (continuation-marks k prompt-tag))
  prompt-tag)
```

```
(continuation-mark-set->list  mark-set
                              key-v
                             [prompt-tag]) → list?
  mark-set : continuation-mark-set?
  key-v : any/c
  prompt-tag : continuation-prompt-tag?
             = (default-continuation-prompt-tag)
```

Returns a newly-created list containing the marks for *key-v* in *mark-set*, which is a set of marks returned by `current-continuation-marks`. The result list is truncated at the first point, if any, where continuation frames were originally separated by a prompt tagged with *prompt-tag*.

```
(make-continuation-mark-key) → continuation-mark-key?
(make-continuation-mark-key sym) → continuation-mark-key?
  sym : symbol?
```

Creates a continuation mark key that is not `equal?` to the result of any other value (including prior and future results from `make-continuation-mark-key`). The continuation mark key can be used as the key argument for `with-continuation-mark` or accessor procedures like `continuation-mark-set-first`. The mark key can be chaperoned or impersonated, unlike other values that are used as the mark key.

The optional *sym* argument, if provided, is used when printing the continuation mark.

```
(continuation-mark-set->list* mark-set
                              key-list
                             [none-v
                              prompt-tag]) → (listof vector?)
  mark-set : continuation-mark-set?
  key-list : (listof any/c)
  none-v : any/c = #f
  prompt-tag : continuation-prompt-tag?
             = (default-continuation-prompt-tag)
```

Returns a newly-created list containing vectors of marks in *mark-set* for the keys in *key-list*, up to *prompt-tag*. The length of each vector in the result list is the same as the length of *key-list*, and a value in a particular vector position is the value for the corresponding key in *key-list*. Values for multiple keys appear in a single vector only when the marks are for the same continuation frame in *mark-set*. The *none-v* argument is used for vector elements to indicate the lack of a value.

```
(continuation-mark-set-first mark-set
                             key-v
                            [none-v
                             prompt-tag]) → any
  mark-set : (or/c continuation-mark-set? #f)
  key-v : any/c
  none-v : any/c = #f
  prompt-tag : continuation-prompt-tag?
             = (default-continuation-prompt-tag)
```

Returns the first element of the list that would be returned by `(continuation-mark-set->list (or mark-set (current-continuation-marks prompt-tag)) key-v prompt-tag)`, or *none-v* if the result would be the empty list. Typically, this result can be computed more quickly using `continuation-mark-set-first` than using `continuation-mark-set->list`.

```
(call-with-immediate-continuation-mark key-v
                                       proc
                                      [default-v]) → any
  key-v : any/c
  proc : (any/c . -> . any)
  default-v : any/c = #f
```

Calls *proc* with the value associated with *key-v* in the first frame of the current continuation (i.e., a value that would be replaced if the call to `call-with-immediate-continuation-mark` were replaced with a `with-continuation-mark` form using *key-v* as the key expression). If no such value exists in the first frame, *default-v* is passed to *proc*. The *proc* is called in tail position with respect to the `call-with-immediate-continuation-mark` call.

This function could be implemented with a combination of `with-continuation-mark`, `current-continuation-marks`, and `continuation-mark-set->list`, but `call-with-immediate-continuation-mark` is implemented more efficiently; it inspects only the first frame of the current continuation.

```
(continuation-mark-key? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a mark key created by `make-continuation-mark-key`, `#f` otherwise.

```
(continuation-mark-set? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a mark set created by `continuation-marks` or `current-continuation-marks`, `#f` otherwise.

```
(continuation-mark-set->context mark-set) → list?
  mark-set : continuation-mark-set?
```

Returns a list representing an approximate "stack trace" for *mark-set*'s continuation. The list contains pairs, where the `car` of each pair contains either `#f` or a symbol for a procedure name, and the `cdr` of each pair contains either `#f` or a `srcloc` value for the procedure's source location (see §12.1.4 "Counting Positions, Lines, and Columns"); the `car` and `cdr` are never both `#f`.

Conceptually, the stack-trace list is the result of `continuation-mark-set->list` with *mark-set* and Racket's private key for procedure-call marks. The implementation may be different, however, and the results may merely approximate the correct answer. Thus, while the result may contain useful hints to humans about the context of an expression, it is not reliable enough for programmatic use.

A stack trace is extracted from an exception and displayed by the default error display handler (see `error-display-handler`) for exceptions other than `exn:fail:user` (see `raise-user-error` in §9.2.2 "Raising Exceptions").

Examples:

```
> (define (extract-current-continuation-marks key)
    (continuation-mark-set->list
     (current-continuation-marks)
     key))
```

```
> (with-continuation-mark 'key 'mark
    (extract-current-continuation-marks 'key))
'(mark)
> (with-continuation-mark 'key1 'mark1
    (with-continuation-mark 'key2 'mark2
      (list
       (extract-current-continuation-marks 'key1)
       (extract-current-continuation-marks 'key2))))
'((mark1) (mark2))
> (with-continuation-mark 'key 'mark1
    (with-continuation-mark 'key 'mark2 ; replaces previous mark
      (extract-current-continuation-marks 'key)))
'(mark2)
> (with-continuation-mark 'key 'mark1
    (list ; continuation extended to evaluate the argument
     (with-continuation-mark 'key 'mark2
        (extract-current-continuation-marks 'key))))
'((mark2 mark1))
> (let loop ([n 1000])
    (if (zero? n)
        (extract-current-continuation-marks 'key)
        (with-continuation-mark 'key n
          (loop (sub1 n)))))
'(1)
```

9.6 Breaks

A *break* is an asynchronous exception, usually triggered through an external source controlled by the user, or through the `break-thread` procedure. A break exception can only occur in a thread while breaks are enabled. When a break is detected and enabled, the `exn:break` exception is raised in the thread sometime afterward; if breaking is disabled when `break-thread` is called, the break is suspended until breaking is again enabled for the thread. While a thread has a suspended break, additional breaks are ignored.

Breaks are enabled through the `break-enabled` parameter-like procedure, and through the `parameterize-break` form, which is analogous to `parameterize`. The `break-enabled` procedure does not represent a parameter to be used with `parameterize`, because changing the break-enabled state of a thread requires an explicit check for breaks, and this check is incompatible with the tail evaluation of a `parameterize` expression's body.

Certain procedures, such as `semaphore-wait/enable-break`, enable breaks temporarily while performing a blocking action. If breaks are enabled for a thread, and if a break is triggered for the thread but not yet delivered as an `exn:break` exception, then the break

is guaranteed to be delivered before breaks can be disabled in the thread. The timing of exn:break exceptions is not guaranteed in any other way.

Before calling a with-handlers predicate or handler, an exception handler, an error display handler, an error escape handler, an error value conversion handler, or a pre-thunk or post-thunk for a dynamic-wind, the call is parameterize-breaked to disable breaks. Furthermore, breaks are disabled during the transitions among handlers related to exceptions, during the transitions between pre-thunks and post-thunks for dynamic-wind, and during other transitions for a continuation jump. For example, if breaks are disabled when a continuation is invoked, and if breaks are also disabled in the target continuation, then breaks will remain disabled from the time of the invocation until the target continuation executes unless a relevant dynamic-wind pre-thunk or post-thunk explicitly enables breaks.

If a break is triggered for a thread that is blocked on a nested thread (see call-in-nested-thread), and if breaks are enabled in the blocked thread, the break is implicitly handled by transferring it to the nested thread.

When breaks are enabled, they can occur at any point within execution, which makes certain implementation tasks subtle. For example, assuming breaks are enabled when the following code is executed,

```
(with-handlers ([exn:break? (lambda (x) (void))])
  (semaphore-wait s))
```

then it is *not* the case that a #<void> result means the semaphore was decremented or a break was received, exclusively. It is possible that *both* occur: the break may occur after the semaphore is successfully decremented but before a #<void> result is returned by semaphore-wait. A break exception will never damage a semaphore, or any other built-in construct, but many built-in procedures (including semaphore-wait) contain internal sub-expressions that can be interrupted by a break.

In general, it is impossible using only semaphore-wait to implement the guarantee that either the semaphore is decremented or an exception is raised, but not both. Racket therefore supplies semaphore-wait/enable-break (see §10.2.3 “Semaphores”), which does permit the implementation of such an exclusive guarantee:

```
(parameterize-break #f
  (with-handlers ([exn:break? (lambda (x) (void))])
    (semaphore-wait/enable-break s)))
```

In the above expression, a break can occur at any point until breaks are disabled, in which case a break exception is propagated to the enclosing exception handler. Otherwise, the break can only occur within semaphore-wait/enable-break, which guarantees that if a break exception is raised, the semaphore will not have been decremented.

To allow similar implementation patterns over blocking port operations, Racket provides `read-bytes-avail!/enable-break`, `write-bytes-avail/enable-break`, and other procedures.

```
(break-enabled) → boolean?
(break-enabled on?) → void?
  on? : any/c
```

Gets or sets the break enabled state of the current thread. If *on?* is not supplied, the result is `#t` if breaks are currently enabled, `#f` otherwise. If *on?* is supplied as `#f`, breaks are disabled, and if *on?* is a true value, breaks are enabled.

```
(parameterize-break boolean-expr body ...+)
```

Evaluates *boolean-expr* to determine whether breaks are initially enabled while evaluating the *body*s in sequence. The result of the `parameter-break` expression is the result of the last `expr`.

Like `parameterize` (see §10.3.2 “Parameters”), a fresh thread cell (see §10.3.1 “Thread Cells”) is allocated to hold the break-enabled state of the continuation, and calls to `break-enabled` within the continuation access or modify the new cell. Unlike parameters, the break setting is not inherited by new threads.

```
(current-break-parameterization) → break-parameterization?
```

Analogous to `(current-parameterization)` (see §10.3.2 “Parameters”); it returns a break-parameterization (effectively, a thread cell) that holds the current continuation’s break-enable state.

```
(call-with-break-parameterization break-param
                                  thunk)      → any
  break-param : break-parameterization?
  thunk : (-> any)
```

Analogous to `(call-with-parameterization parameterization thunk)` (see §10.3.2 “Parameters”), calls *thunk* in a continuation whose break-enabled state is in *break-param*. The *thunk* is *not* called in tail position with respect to the `call-with-break-parameterization` call.

9.7 Exiting

```
(exit [v]) → any
  v : any/c = #t
```

Passes `v` to the current exit handler. If the exit handler does not escape or terminate the thread, `#<void>` is returned.

```
(exit-handler) → (any/c . -> . any)
(exit-handler proc) → void?
  proc : (any/c . -> . any)
```

A parameter that determines the current *exit handler*. The exit handler is called by `exit`.

The default exit handler in the Racket executable takes any argument and shuts down the OS-level Racket process. The argument is used as the OS-level exit code if it is an exact integer between `1` and `255` (which normally means "failure"); otherwise, the exit code is `0`, (which normally means "success").

```
(executable-yield-handler) → ((integer-in 0 255) . -> . any)
(executable-yield-handler proc) → void?
  proc : ((integer-in 0 255) . -> . any)
```

A parameter that determines a procedure to be called as the Racket process is about to exit normally. The procedure associated with this parameter is not call when `exit` (or, more precisely, the default exit handler) is used to exit early. The argument to the handler is the status code that is returned to the system on exit. The default executable-yield handler simply returns `#<void>`.

The `scheme/gui/base` library sets this parameter to wait until all frames are closed, timers stopped, and queued events handled in the main eventspace. See `scheme/gui/base` for more information.

10 Concurrency and Parallelism

Racket supports multiple threads of control within a program, thread-local storage, some primitive synchronization mechanisms, and a framework for composing synchronization abstractions. In addition, the `racket/future` and `racket/place` libraries provide support for parallelism to improve performance.

10.1 Threads

See §1.1.13 “Threads” for basic information on the Racket thread model. See also §10.4 “Futures”.

When a thread is created, it is placed into the management of the current custodian and added to the current thread group. A thread can have any number of custodian managers added through `thread-resume`.

A thread that has not terminated can be garbage collected (see §1.1.7 “Garbage Collection”) if it is unreachable and suspended or if it is unreachable and blocked on only unreachable events through `semaphore-wait`, `semaphore-wait/enable-break`, `channel-put`, `channel-get`, `sync`, `sync/enable-break`, or `thread-wait`.

In GRacket, a handler thread for an eventspace is blocked on an internal semaphore when its event queue is empty. Thus, the handler thread is collectible when the eventspace is unreachable and contains no visible windows or running timers.

All constant-time procedures and operations provided by Racket are thread-safe because they are *atomic*. For example, `set!` assigns to a variable as an atomic action with respect to all threads, so that no thread can see a “half-assigned” variable. Similarly, `vector-set!` assigns to a vector atomically. The `hash-set!` procedure is not atomic, but the table is protected by a lock; see §3.13 “Hash Tables” for more information. Port operations are generally not atomic, but they are thread-safe in the sense that a byte consumed by one thread from an input port will not be returned also to another thread, and procedures like `port-commit-peeked` and `write-bytes-avail` offer specific concurrency guarantees.

10.1.1 Creating Threads

```
(thread thunk) → thread?
  thunk : (-> any)
```

Calls *thunk* with no arguments in a new thread of control. The `thread` procedure returns immediately with a *thread descriptor* value. When the invocation of *thunk* returns, the thread created to invoke *thunk* terminates.

```
(thread? v) → thread?
  v : any/c
```

Returns `#t` if *v* is a thread descriptor, `#f` otherwise.

```
(current-thread) → thread?
```

Returns the thread descriptor for the currently executing thread.

```
(thread/suspend-to-kill thunk) → thread
  thunk : (-> any)
```

Like `thread`, except that "killing" the thread through `kill-thread` or `custodian-shutdown-all` merely suspends the thread instead of terminating it.

```
(call-in-nested-thread thunk [cust]) → any
  thunk : (->any)
  cust : custodian? = (current-custodian)
```

Creates a nested thread managed by *cust* to execute *thunk*. (The nested thread's current custodian is inherited from the creating thread, independent of the *cust* argument.) The current thread blocks until *thunk* returns, and the result of the `call-in-nested-thread` call is the result returned by *thunk*.

The nested thread's exception handler is initialized to a procedure that jumps to the beginning of the thread and transfers the exception to the original thread. The handler thus terminates the nested thread and re-raises the exception in the original thread.

If the thread created by `call-in-nested-thread` dies before *thunk* returns, the `exn:fail` exception is raised in the original thread. If the original thread is killed before *thunk* returns, a break is queued for the nested thread.

If a break is queued for the original thread (with `break-thread`) while the nested thread is running, the break is redirected to the nested thread. If a break is already queued on the original thread when the nested thread is created, the break is moved to the nested thread. If a break remains queued on the nested thread when it completes, the break is moved to the original thread.

10.1.2 Suspending, Resuming, and Killing Threads

```
(thread-suspend thd) → void?
  thd : thread?
```

Immediately suspends the execution of *thd* if it is running. If the thread has terminated or is already suspended, `thread-suspend` has no effect. The thread remains suspended (i.e., it does not execute) until it is resumed with `thread-resume`. If the current custodian does not solely manage *thd* (i.e., some custodian of *thd* is not the current custodian or a subordinate), the `exn:fail:contract` exception is raised, and the thread is not suspended.

```
(thread-resume thd [benefactor]) → void?
  thd : thread?
  benefactor : (or/c thread? custodian? #f) = #f
```

Resumes the execution of *thd* if it is suspended and has at least one custodian (possibly added through *benefactor*, as described below). If the thread has terminated, or if the thread is already running and *benefactor* is not supplied, or if the thread has no custodian and *benefactor* is not supplied, then `thread-resume` has no effect. Otherwise, if *benefactor* is supplied, it triggers up to three additional actions:

- If *benefactor* is a thread, whenever it is resumed from a suspended state in the future, then *thd* is also resumed. (Resuming *thd* may trigger the resumption of other threads that were previously attached to *thd* through `thread-resume`.)
- New custodians may be added to *thd*'s set of managers. If *benefactor* is a thread, then all of the thread's custodians are added to *thd*. Otherwise, *benefactor* is a custodian, and it is added to *thd* (unless the custodian is already shut down). If *thd* becomes managed by both a custodian and one or more of its subordinates, the redundant subordinates are removed from *thd*. If *thd* is suspended and a custodian is added, then *thd* is resumed only after the addition.
- If *benefactor* is a thread, whenever it receives a new managing custodian in the future, then *thd* also receives the custodian. (Adding custodians to *thd* may trigger adding the custodians to other threads that were previously attached to *thd* through `thread-resume`.)

```
(kill-thread thd) → void?
  thd : thread?
```

Terminates the specified thread immediately, or suspends the thread if *thd* was created with `thread/suspend-to-kill`. Terminating the main thread exits the application. If *thd* has already terminated, `kill-thread` does nothing. If the current custodian does not manage *thd* (and none of its subordinates manages *thd*), the `exn:fail:contract` exception is raised, and the thread is not killed or suspended.

Unless otherwise noted, procedures provided by Racket (and GRacket) are kill-safe and suspend-safe; that is, killing or suspending a thread never interferes with the application of procedures in other threads. For example, if a thread is killed while extracting a character from an input port, the character is either completely consumed or not consumed, and other threads can safely use the port.

```
(break-thread thd) → void?
  thd : thread?
```

Registers a break with the specified thread. If breaking is disabled in *thd*, the break will be ignored until breaks are re-enabled (see §9.6 “Breaks”).

```
(sleep [secs]) → void?
  secs : (>=/c 0) = 0
```

Causes the current thread to sleep until at least *secs* seconds have passed after it starts sleeping. A zero value for *secs* simply acts as a hint to allow other threads to execute. The value of *secs* can be a non-integer to request a sleep duration to any precision; the precision of the actual sleep time is unspecified.

```
(thread-running? thd) → any
  thd : thread?
```

Returns `#t` if *thd* has not terminated and is not suspended, `#f` otherwise.

```
(thread-dead? thd) → any
  thd : thread?
```

Returns `#t` if *thd* has terminated, `#f` otherwise.

10.1.3 Synchronizing Thread State

```
(thread-wait thd) → void?
  thd : thread?
```

Blocks execution of the current thread until *thd* has terminated. Note that `(thread-wait (current-thread))` deadlocks the current thread, but a break can end the deadlock (if breaking is enabled; see §9.6 “Breaks”).

```
(thread-dead-evt thd) → evt?
  thd : thread?
```

Returns a synchronizable event (see §10.2.1 “Events”) that is ready if and only if *thd* has terminated. Unlike using *thd* directly, however, a reference to the event does not prevent *thd* from being garbage collected (see §1.1.7 “Garbage Collection”). For a given *thd*, `thread-dead-evt` always returns the same (i.e., `eq?`) result.

```
(thread-resume-evt thd) → evt?
  thd : thread?
```

Returns a synchronizable event (see §10.2.1 “Events”) that becomes ready when *thd* is running. (If *thd* has terminated, the event never becomes ready.) If *thd* runs and is then

suspended after a call to thread-resume-evt, the result event remains ready; after each suspend of thd a fresh event is generated to be returned by thread-resume-evt. The result of the event is thd, but if thd is never resumed, then reference to the event does not prevent thd from being garbage collected (see §1.1.7 "Garbage Collection").

```
(thread-suspend-evt thd) → evt?
  thd : thread?
```

Returns a synchronizable event (see §10.2.1 "Events") that becomes ready when thd is suspended. (If thd has terminated, the event will never unblock.) If thd is suspended and then resumes after a call to thread-suspend-evt, the result event remains ready; after each resume of thd created a fresh event to be returned by thread-suspend-evt.

10.1.4 Thread Mailboxes

Each thread has a *mailbox* through which it can receive arbitrary messages. In other words, each thread has a built-in asynchronous channel.

See also §10.2.4 "Buffered Asynchronous Channels".

```
(thread-send thd v [fail-thunk]) → any
  thd : thread?
  v : any/c
  fail-thunk : (or/c (-> any) #f)
             = (lambda () (raise-mismatch-error ....))
```

Queues v as a message to thd without blocking. If the message is queued, the result is #<void>. If thd stops running—as in thread-running?—before the message is queued, then fail-thunk is called (through a tail call) if it is a procedure to produce the result, or #f is returned if fail-thunk is #f.

```
(thread-receive) → any/c
```

Receives and dequeues a message queued for the current thread, if any. If no message is available, thread-receive blocks until one is available.

```
(thread-try-receive) → any/c
```

Receives and dequeues a message queued for the current thread, if any, or returns #f immediately if no message is available.

```
(thread-receive-evt) → evt?
```

Returns a constant synchronizable event (see §10.2.1 "Events") that becomes ready when the synchronizing thread has a message to receive. The event result is the event itself.

```
(thread-rewind-receive lst) → void?
  lst : list?
```

Pushes the elements of *lst* back onto the front of the current thread's queue. The elements are pushed one by one, so that the first available message is the last element of *lst*.

10.2 Synchronization

Racket's synchronization toolbox spans three layers:

- synchronizable events — a general framework for synchronization;
- channels — a primitive that can be used, in principle, to build most other kinds of synchronizable events (except the ones that compose events); and
- semaphores — a simple and especially cheap primitive for synchronization.

10.2.1 Events

A *synchronizable event* (or just *event* for short) works with the sync procedure to coordinate synchronization among threads. Certain kinds of objects double as events, including ports and threads. Other kinds of objects exist only for their use as events.

At any point in time, an event is either *ready* for synchronization, or it is not; depending on the kind of event and how it is used by other threads, an event can switch from not ready to ready (or back), at any time. If a thread synchronizes on an event when it is ready, then the event produces a particular *synchronization result*.

Synchronizing an event may affect the state of the event. For example, when synchronizing a semaphore, then the semaphore's internal count is decremented, just as with semaphore-wait. For most kinds of events, however (such as a port), synchronizing does not modify the event's state.

The following act as events in Racket. An extension or embedding application can extend the set of primitive events — in particular, an eventspace in GRacket is an event — and new structure types can generate events (see prop:evt).

- *semaphore* — a semaphore is ready when semaphore-wait would not block. The synchronization result of *semaphore* is *semaphore* itself.
- *semaphore-peek* — a semaphore-peek event returned by semaphore-peek-evt applied to *semaphore* is ready exactly when *semaphore* is ready. The synchronization result of *semaphore-peek* is *semaphore-peek* itself.

- *channel* — a channel returned by make-channel is ready when channel-get would not block. The channel's result as an event is the same as the channel-get result.

- *channel-put* — an event returned by channel-put-evt applied to *channel* is ready when channel-put would not block on *channel*. The synchronization result of *channel-put* is *channel-put* itself.

- *async-channel* — a channel returned by make-async-channel is ready when async-channel-get would not block. The channel's result as an event is the same as the async-channel-get result.

- *async-channel-put* — an event returned by async-channel-put-evt applied to *async-channel* is ready when async-channel-put would not block on *async-channel*. The synchronization result of *async-channel-put* is *async-channel-put* itself.

- *input-port* — an input port is ready as an event when read-byte would not block. The synchronization result of *input-port* is *input-port* itself.

- *output-port* — an output port is ready when write-bytes-avail would not block or when the port contains buffered characters and write-bytes-avail* can flush part of the buffer (although write-bytes-avail might block). The synchronization result of *output-port* is *output-port* itself.

- *closed* — an event produced by port-closed-evt applied to *port* is ready after *port* is closed. The synchronization result of *closed* is *closed* itself.

- *progress* — an event produced by port-progress-evt applied to *input-port* is ready after any subsequent read from *input-port*. The synchronization result of *progress* is *progress* itself.

- *tcp-listener* — a TCP listener is ready when tcp-accept would not block. The synchronization result of *listener* is *listener* itself.

- *thd* — a thread is ready when thread-wait would not block. The synchronization result of *thread* is *thread* itself.

- *thread-dead* — an event returned by thread-dead-evt applied to thd is ready when thd has terminated. The synchronization result of *thread-dead* is *thread-dead* itself.

- *thread-resume* — an event returned by thread-resume-evt applied to thd is ready when thd subsequently resumes execution (if it was not already running). The event's result is thd.

- *thread-suspend* — an event returned by thread-suspend-evt applied to thd is ready when thd subsequently suspends execution (if it was not already suspended). The event's result is thd.

- *thread-receive* — an event returned by thread-receive-evt is ready when the synchronizing thread has a message to receive. The synchronization result of *thread-receive* is *thread-receive* itself.

- *alarm* — an event returned by alarm-evt is ready after a particular date and time. The synchronization result of *alarm* is *alarm* itself.

- *subprocess* — a subprocess is ready when subprocess-wait would not block. The synchronization result of *subprocess* is *subprocess* itself.

- *will-executor* — a will executor is ready when will-execute would not block. The synchronization result of *will-executor* is *will-executor* itself.

- *custodian-box* — a custodian box is ready when its custodian is shut down. The synchronization result of *custodian-box* is *custodian-box* itself.

- *udp* — an event returned by udp-send-evt or udp-receive!-evt is ready when a send or receive on the original socket would block, respectively. The synchronization result of *udp* is *udp* itself.

- *log-receiver* — a log receiver as produced by make-log-receiver is ready when a logged message is available. The event's result is a vector, as described with make-log-receiver.

- *choice* — an event returned by choice-evt is ready when one or more of the *evt*s supplied to choice-evt are ready. If the choice event is chosen, one of its ready *evt*s is chosen pseudo-randomly, and the result is the chosen *evt*'s result.

- *wrap* — an event returned by wrap-evt applied to *evt* and *proc* is ready when *evt* is ready. The event's result is obtained by a call to *proc* (with breaks disabled) on the result of evt.

- *handle* — an event returned by handle-evt applied to *evt* and *proc* is ready when *evt* is ready. The event's result is obtained by a tail call to *proc* on the result of *evt*.

- *guard* — an event returned by guard-evt applied to *thunk* generates a new event every time that *guard* is used with sync (or whenever it is part of a choice event used with sync, etc.); the generated event is the result of calling *thunk* when the synchronization begins; if *thunk* returns a non-event, then *thunk*'s result is replaced with an event that is ready and whose result is *guard*.

- *nack-guard* — an event returned by nack-guard-evt applied to *proc* generates a new event every time that *nack-guard* is used with sync (or whenever it is part of a choice event used with sync, etc.); the generated event is the result of calling *proc* with a NACK ("negative acknowledgment") event when the synchronization begins; if *proc* returns a non-event, then *proc*'s result is replaced with an event that is ready and whose result is *nack-guard*.

  If the event from *proc* is not ultimately chosen as the unblocked event, then the NACK event supplied to *proc* becomes ready with a #<void> value. This NACK event

becomes ready when the event is abandoned because some other event is chosen, because the synchronizing thread is dead, or because control escaped from the call to `sync` (even if *`nack-guard`*'s *`proc`* has not yet returned a value). If the event returned by *`proc`* is chosen, then the NACK event never becomes ready.

- *`poll-guard`* — an event returned by `poll-guard-evt` applied to *`proc`* generates a new event every time that `poll-guard` is used with `sync` (or whenever it is part of a choice event used with `sync`, etc.); the generated event is the result of calling *`proc`* with a boolean: `#t` if the event will be used for a poll, `#f` for a blocking synchronization.

  If `#t` is supplied to *`proc`*, if breaks are disabled, if the polling thread is not terminated, and if polling the resulting event produces a result, the event will certainly be chosen for its result.

- *`struct`* — a structure whose type has the `prop:evt` property identifies/generates an event through the property.

- `always-evt` — a constant event that is always ready. The synchronization result of `always-evt` is `always-evt` itself.

- `never-evt` — a constant event that is never ready.

- *`idle`* — an event produced by `system-idle-evt` is ready when, if this event were replaced by `never-evt`, no thread in the system would be available to run. In other words, all threads must be suspended or blocked on events with timeouts that have not yet expired. The event's result is `#<void>`.

- *`place-channel`* — a place channel is ready when `place-channel-get` would not block. The channel's result as an event is the same as the `place-channel-get` result.

- *`place-dead`* — an event returned by `(place-dead-evt p)` is ready when *`p`* has terminated. The synchronization result of *`place-dead`* is *`place-dead`* itself.

```
(evt? v) → boolean?
  v : any/c
```

Returns `#t` if *`v`* is a synchronizable event, `#f` otherwise.

```
(sync evt ...+) → any
  evt : evt?
```

Blocks as long as none of the synchronizable events *`evt`*s are ready, as defined above.

When at least one *`evt`* is ready, its synchronization result (often *`evt`* itself) is returned. If multiple *`evt`*s are ready, one of the *`evt`*s is chosen pseudo-randomly for the result; the `current-evt-pseudo-random-generator` parameter sets the random-number generator that controls this choice.

```
(sync/timeout timeout evt ...+) → any
  timeout : (or/c #f (and/c real? (not/c negative?)) (-> any))
  evt : evt?
```

Like sync if *timeout* is #f. If *timeout* is a real number, then the result is #f if *timeout* seconds pass without a successful synchronization. If *timeout* is a procedure, then it is called in tail position if polling the *evt*s discovers no ready events.

A zero value for *timeout* is equivalent to (lambda () #f). In either case, each *evt* is checked at least once before returning #f or calling *timeout*.

See also alarm-evt for an alternative timeout mechanism.

```
(sync/enable-break evt ...+) → any
  evt : evt?
```

Like sync, but breaking is enabled (see §9.6 “Breaks”) while waiting on the *evt*s. If breaking is disabled when sync/enable-break is called, then either all *evt*s remain unchosen or the exn:break exception is raised, but not both.

```
(sync/timeout/enable-break timeout evt ...+) → any
  timeout : (or/c #f (and/c real? (not/c negative?)) (-> any))
  evt : evt?
```

Like sync/enable-break, but with a timeout as for sync/timeout.

```
(choice-evt evt ...) → evt?
  evt : evt?
```

Creates and returns a single event that combines the *evt*s. Supplying the result to sync is the same as supplying each *evt* to the same call.

```
(wrap-evt evt wrap) → evt?
  evt : (and/c evt? (not/c handle-evt?))
  wrap : (any/c . -> . any)
```

Creates an event that is in a ready when *evt* is ready, but whose result is determined by applying *wrap* to the result of *evt*. The call to *wrap* is parameterize-breaked to disable breaks initially. The *evt* cannot be an event created by handle-evt or any combination of choice-evt involving an event from handle-evt.

```
(handle-evt evt handle) → handle-evt?
  evt : (and/c evt? (not/c handle-evt?))
  handle : (any/c . -> . any)
```

Like wrap, except that *handle* is called in tail position with respect to the synchronization request, and without breaks explicitly disabled.

```
(guard-evt generator) → evt?
  generator : (-> evt?)
```

Creates a value that behaves as an event, but that is actually an event generator. For details, see the overview.

```
(nack-guard-evt generator) → evt?
  generator : (evt? . -> . evt?)
```

Creates a value that behaves as an event, but that is actually an event generator; the generator procedure receives an event that becomes ready with a #<void> value if the generated event was not ultimately chosen. For details, see the overview.

```
(poll-guard-evt generator) → evt?
  generator : (boolean? . -> . evt?)
```

Creates a value that behaves as an event, but that is actually an event generator; the generator procedure receives a boolean indicating whether the event is used for polling. For details, see the overview.

```
always-evt : evt?
```

A constant event that is always ready, with itself as its result.

```
never-evt : evt?
```

A constant event that is never ready.

```
(system-idle-evt) → evt?
```

Returns an event that is ready when the system is otherwise idle; see the overview for more information. The result of the system-idle-evt procedure is always the same event.

```
(alarm-evt msecs) → evt
  msecs : nonnegative-number?
```

Returns a synchronizable event that is not ready when (current-inexact-milliseconds) would return a value that is less than *msecs*, and it is ready when (current-inexact-milliseconds) would return a value that is more than *msecs*.

```
(handle-evt? evt) → boolean?
  evt : evt?
```

Returns `#t` if *evt* was created by `handle-evt` or by `choice-evt` applied to another event for which `handle-evt?` produces `#t`. Such events are illegal as an argument to `handle-evt` or `wrap-evt`, because they cannot be wrapped further. For any other event, `handle-evt?` produces `#f`, and the event is a legal argument to `handle-evt` or `wrap-evt` for further wrapping.

`prop:evt` : `struct-type-property?`

A structure type property that identifies structure types whose instances can serve as synchronizable events. The property value can be any of the following:

- An event *evt*: In this case, using the structure as an event is equivalent to using *evt*.
- A procedure *proc* of one argument: In this case, the structure is similar to an event generated by `guard-evt`, except that the would-be guard procedure *proc* receives the structure as an argument, instead of no arguments.
- An exact, non-negative integer between `0` (inclusive) and the number of non-automatic fields in the structure type (exclusive, not counting supertype fields): The integer identifies a field in the structure, and the field must be designated as immutable. If the field contains an object or an event-generating procedure of one argument, the event or procedure is used as above. Otherwise, the structure acts as an event that is never ready.

Instances of a structure type with the `prop:input-port` or `prop:output-port` property are also synchronizable by virtue of being a port. If the structure type has more than one of `prop:evt`, `prop:input-port`, and `prop:output-port`, then the `prop:evt` value (if any) takes precedence for determining the instance's behavior as an event, and the `prop:input-port` property takes precedence over `prop:output-port` for synchronization.

Examples:

```
> (define-struct wt (base val)
                  #:property prop:evt (struct-field-index base))

> (define sema (make-semaphore))

> (sync/timeout 0 (make-wt sema #f))
#f
> (semaphore-post sema)

> (sync/timeout 0 (make-wt sema #f))
#<semaphore>
> (semaphore-post sema)
```

```
> (sync/timeout 0 (make-wt (lambda (self) (wt-val self)) sema))
#<semaphore>
> (semaphore-post sema)

> (define my-wt (make-wt (lambda (self) (wrap-evt
                                         (wt-val self)
                                         (lambda (x) self)))
                         sema))

> (sync/timeout 0 my-wt)
#<wt>
> (sync/timeout 0 my-wt)
#f
```

```
(current-evt-pseudo-random-generator)
 → pseudo-random-generator?
(current-evt-pseudo-random-generator generator) → void?
  generator : pseudo-random-generator?
```

A parameter that determines the pseudo-random number generator used by `sync` for events created by `choice-evt`.

10.2.2 Channels

A *channel* both synchronizes a pair of threads and passes a value from one to the other. Channels are synchronous; both the sender and the receiver must block until the (atomic) transaction is complete. Multiple senders and receivers can access a channel at once, but a single sender and receiver is selected for each transaction.

Channel synchronization is *fair*: if a thread is blocked on a channel and transaction opportunities for the channel occur infinitely often, then the thread eventually participates in a transaction.

For buffered asynchronous channels, see §10.2.4 “Buffered Asynchronous Channels”.

```
(channel? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a channel, `#f` otherwise.

```
(make-channel) → channel?
```

Creates and returns a new channel. The channel can be used with `channel-get`, with `channel-try-get`, or as a synchronizable event (see §10.2.1 “Events”) to receive a value through the channel. The channel can be used with `channel-put` or through the result of `channel-put-evt` to send a value through the channel.

```
(channel-get ch) → any
  ch : channel?
```

Blocks until a sender is ready to provide a value through *ch*. The result is the sent value.

```
(channel-try-get ch) → any
  ch : channel?
```

Receives and returns a value from *ch* if a sender is immediately ready, otherwise returns `#f`.

```
(channel-put ch v) → void?
  ch : channel?
  v : any/c
```

Blocks until a receiver is ready to accept the value *v* through *ch*.

```
(channel-put-evt ch v) → channel-put-evt?
  ch : channel?
  v : any/c
```

Returns a fresh synchronizable event for use with `sync`. The event is ready when `(channel-put ch v)` would not block, and the event's synchronization result is the event itself.

```
(channel-put-evt? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a channel-put event produced by `channel-put-evt`, `#f` otherwise.

10.2.3 Semaphores

A *semaphore* has an internal counter; when this counter is zero, the semaphore can block a thread's execution (through `semaphore-wait`) until another thread increments the counter (using `semaphore-post`). The maximum value for a semaphore's internal counter is platform-specific, but always at least `10000`.

A semaphore's counter is updated in a single-threaded manner, so that semaphores can be used for reliable synchronization. Semaphore waiting is *fair*: if a thread is blocked on a semaphore and the semaphore's internal value is non-zero infinitely often, then the thread is eventually unblocked.

In addition to its use with semaphore-specific procedures, semaphores can be used as events; see §10.2.1 "Events".

```
(semaphore? v) → boolean?
  v : any/c
```

Returns #t if *v* is a semaphore, #f otherwise.

```
(make-semaphore [init]) → semaphore?
  init : exact-nonnegative-integer? = 0
```

Creates and returns a new semaphore with the counter initially set to *init*. If *init* is larger than a semaphore's maximum internal counter value, the exn:fail exception is raised.

```
(semaphore-post sema) → void?
  sema : semaphore?
```

Increments the semaphore's internal counter and returns #<void>. If the semaphore's internal counter has already reached its maximum value, the exn:fail exception is raised.

```
(semaphore-wait sema) → void?
  sema : semaphore?
```

Blocks until the internal counter for semaphore *sema* is non-zero. When the counter is non-zero, it is decremented and semaphore-wait returns #<void>.

```
(semaphore-try-wait? sema) → boolean?
  sema : semaphore?
```

Like semaphore-wait, but semaphore-try-wait? never blocks execution. If *sema*'s internal counter is zero, semaphore-try-wait? returns #f immediately without decrementing the counter. If *sema*'s counter is positive, it is decremented and #t is returned.

```
(semaphore-wait/enable-break sema) → void?
  sema : semaphore?
```

Like semaphore-wait, but breaking is enabled (see §9.6 "Breaks") while waiting on *sema*. If breaking is disabled when semaphore-wait/enable-break is called, then either the semaphore's counter is decremented or the exn:break exception is raised, but not both.

```
(semaphore-peek-evt sema) → semaphore-peek-evt?
  sema : semaphore?
```

Creates and returns a new synchronizable event (for use with sync, for example) that is ready when *sema* is ready, but synchronizing the event does not decrement *sema*'s internal count.

```
(semaphore-peek-evt? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a semaphore wrapper produced by `semaphore-peek-evt`, `#f` otherwise.

```
(call-with-semaphore  sema
                      proc
                     [try-fail-thunk]
                      arg ...)        → any
  sema : semaphore?
  proc : procedure?
  try-fail-thunk : (or/c (-> any) #f) = #f
  arg : any/c
```

Waits on *sema* using `semaphore-wait`, calls *proc* with all *arg*s, and then posts to *sema*. A continuation barrier blocks full continuation jumps into or out of *proc* (see §1.1.12 “Prompts, Delimited Continuations, and Barriers”), but escape jumps are allowed, and *sema* is posted on escape. If *try-fail-thunk* is provided and is not `#f`, then `semaphore-try-wait?` is called on *sema* instead of `semaphore-wait`, and *try-fail-thunk* is called if the wait fails.

```
(call-with-semaphore/enable-break  sema
                                   proc
                                  [try-fail-thunk]
                                   arg ...)        → any
  sema : semaphore?
  proc : procedure?
  try-fail-thunk : (or/c (-> any) #f) = #f
  arg : any/c
```

Like `call-with-semaphore`, except that `semaphore-wait/enable-break` is used with *sema* in non-try mode. When *try-fail-thunk* is provided and not `#f`, then breaks are enabled around the use of `semaphore-try-wait?` on *sema*.

10.2.4 Buffered Asynchronous Channels

```
(require racket/async-channel)
```

The bindings documented in this section are provided by the `racket/async-channel` library, not `racket/base` or `racket`.

See also §10.1.4 “Thread Mailboxes”.

```
(async-channel? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is an asynchronous channel, `#f` otherwise.

```
(make-async-channel [limit]) → async-channel?
  limit : (or/c exact-positive-integer? #f) = #f
```

Returns an asynchronous channel with a buffer limit of *limit* items. A get operation blocks when the channel is empty, and a put operation blocks when the channel has *limit* items already. If *limit* is `#f`, the channel buffer has no limit (so a put never blocks).

The asynchronous channel value can be used directly with `sync`. The channel blocks until `async-channel-get` would return a value, and the unblock result is the received value.

```
(async-channel-get ach) → any/c
  ach : async-channel?
```

Blocks until at least one value is available in *ach*, and then returns the first of the values that were put into `async-channel`.

```
(async-channel-try-get ach) → any/c
  ach : async-channel?
```

If at least one value is immediately available in *ach*, returns the first of the values that were put into *ach*. If `async-channel` is empty, the result is `#f`.

```
(async-channel-put ach v) → void?
  ach : async-channel?
  v : any/c
```

Puts *v* into *ach*, blocking if *ach*'s buffer is full until space is available.

```
(async-channel-put-evt ach v) → evt?
  ach : async-channel?
  v : any/c
```

Returns a synchronizable event that is blocked while `(async-channel-put ach v)` would block. The unblock result is the event itself. See also `sync`.

Examples:

```
(define (server input-channel output-channel)
  (thread (lambda ()
            (define (get)
              (async-channel-get input-channel))
            (define (put x)
              (async-channel-put output-channel x))
```

```
            (define (do-large-computation)
              (sqrt 9))
            (let loop ([data (get)])
              (case data
                [(quit) (void)]
                [(add) (begin
                         (put (+ 1 (get)))
                         (loop (get)))]
                [(long) (begin
                          (put (do-large-computation))
                          (loop (get)))])))))

(define to-server (make-async-channel))

(define from-server (make-async-channel))

> (server to-server from-server)
#<thread>
> (async-channel? to-server)
#t
> (printf "Adding 1 to 4\n")
Adding 1 to 4

> (async-channel-put to-server 'add)

> (async-channel-put to-server 4)

> (printf "Result is ~a\n" (async-channel-get from-server))
Result is 5

> (printf "Ask server to do a long computation\n")
Ask server to do a long computation

> (async-channel-put to-server 'long)

> (printf "I can do other stuff\n")
I can do other stuff

> (printf "Ok, computation from server is ~a\n"
          (async-channel-get from-server))
Ok, computation from server is 3

> (async-channel-put to-server 'quit)
```

10.3 Thread-Local Storage

Thread cells provides primitive support for thread-local storage. Parameters combine thread cells and continuation marks to support thread-specific, continuation-specific binding.

10.3.1 Thread Cells

A *thread cell* contains a thread-specific value; that is, it contains a specific value for each thread, but it may contain different values for different threads. A thread cell is created with a default value that is used for all existing threads. When the cell's content is changed with `thread-cell-set!`, the cell's value changes only for the current thread. Similarly, `thread-cell-ref` obtains the value of the cell that is specific to the current thread.

A thread cell's value can be *preserved*, which means that when a new thread is created, the cell's initial value for the new thread is the same as the creating thread's current value. If a thread cell is non-preserved, then the cell's initial value for a newly created thread is the default value (which was supplied when the cell was created).

Within the current thread, the current values of all preserved threads cells can be captured through `current-preserved-thread-cell-values`. The captured set of values can be imperatively installed into the current thread through another call to `current-preserved-thread-cell-values`. The capturing and restoring threads can be different.

```
(thread-cell? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a thread cell, `#f` otherwise.

```
(make-thread-cell v [preserved?]) → thread-cell?
  v : any/c
  preserved? : any/c = #f
```

Creates and returns a new thread cell. Initially, `v` is the cell's value for all threads. If `preserved?` is true, then the cell's initial value for a newly created threads is the creating thread's value for the cell, otherwise the cell's value is initially `v` in all future threads.

```
(thread-cell-ref cell) → any
  cell : thread-cell?
```

Returns the current value of `cell` for the current thread.

```
(thread-cell-set! cell v) → any
  cell : thread-cell?
  v : any/c
```

Sets the value in *cell* to *v* for the current thread.

Examples:

```
> (define cnp (make-thread-cell '(nerve) #f))

> (define cp (make-thread-cell '(cancer) #t))

> (thread-cell-ref cnp)
'(nerve)
> (thread-cell-ref cp)
'(cancer)
> (thread-cell-set! cnp '(nerve nerve))

> (thread-cell-set! cp '(cancer cancer))

> (thread-cell-ref cnp)
'(nerve nerve)
> (thread-cell-ref cp)
'(cancer cancer)
> (define ch (make-channel))

> (thread (lambda ()
            (channel-put ch (thread-cell-ref cnp))
            (channel-put ch (thread-cell-ref cp))
            (channel-get ch)
            (channel-put ch (thread-cell-ref cp))))
#<thread>
> (channel-get ch)
'(nerve)
> (channel-get ch)
'(cancer cancer)
> (thread-cell-set! cp '(cancer cancer cancer))

> (thread-cell-ref cp)
'(cancer cancer cancer)
> (channel-put ch 'ok)

> (channel-get ch)
'(cancer cancer)
```

```
(current-preserved-thread-cell-values) → thread-cell-values?
(current-preserved-thread-cell-values thread-cell-vals) → void?
  thread-cell-vals : thread-cell-values?
```

When called with no arguments, this procedure produces a *thread-cell-vals* that represents the current values (in the current thread) for all preserved thread cells.

When called with a *thread-cell-vals* generated by a previous call to current-preserved-thread-cell-values, the values of all preserved thread cells (in the current thread) are set to the values captured in *thread-cell-vals*; if a preserved thread cell was created after *thread-cell-vals* was generated, then the thread cell's value for the current thread reverts to its initial value.

```
(thread-cell-values? v) → boolean?
  v : any/c
```

Returns #t if *v* is a set of thread cell values produced by current-preserved-thread-cell-values, #f otherwise.

10.3.2 Parameters

§4.13 "Dynamic Binding: `parameterize`" in ***The Racket Guide*** introduces parameters.

See §1.1.14 "Parameters" for basic information on the parameter model. Parameters correspond to *preserved thread fluids* in Scsh [Gasbichler02].

To parameterize code in a thread- and continuation-friendly manner, use `parameterize`. The `parameterize` form introduces a fresh thread cell for the dynamic extent of its body expressions.

When a new thread is created, the parameterization for the new thread's initial continuation is the parameterization of the creator thread. Since each parameter's thread cell is preserved, the new thread "inherits" the parameter values of its creating thread. When a continuation is moved from one thread to another, settings introduced with `parameterize` effectively move with the continuation.

In contrast, direct assignment to a parameter (by calling the parameter procedure with a value) changes the value in a thread cell, and therefore changes the setting only for the current thread. Consequently, as far as the memory manager is concerned, the value originally associated with a parameter through `parameterize` remains reachable as long the continuation is reachable, even if the parameter is mutated.

```
(make-parameter v [guard]) → parameter?
  v : any/c
  guard : (or/c (any/c . -> . any) #f) = #f
```

Returns a new parameter procedure. The value of the parameter is initialized to *v* in all threads. If *guard* is supplied, it is used as the parameter's guard procedure. A guard procedure takes one argument. Whenever the parameter procedure is applied to an argument, the argument is passed on to the guard procedure. The result returned by the guard procedure is used as the new parameter value. A guard procedure can raise an exception to reject a change to the parameter's value. The *guard* is not applied to the initial *v*.

```
(parameterize ([parameter-expr value-expr] ...)
  body ...+)

  parameter-expr : parameter?
```

§4.13 “Dynamic Binding: `parameterize`” in ***The Racket Guide*** introduces `parameterize`.

The result of a `parameterize` expression is the result of the last *body*. The *parameter-expr*s determine the parameters to set, and the *value-expr*s determine the corresponding values to install while evaluating the `body-expr`s. All of the *parameter-expr*s are evaluated first (and checked with `parameter?`), then all *value-expr*s are evaluated, and then the parameters are bound in the continuation to preserved thread cells that contain the values of the *value-expr*s. The last `body-expr` is in tail position with respect to the entire `parameterize` form.

Outside the dynamic extent of a `parameterize` expression, parameters remain bound to other thread cells. Effectively, therefore, old parameters settings are restored as control exits the `parameterize` expression.

If a continuation is captured during the evaluation of `parameterize`, invoking the continuation effectively re-introduces the parameterization, since a parameterization is associated to a continuation via a continuation mark (see §9.5 “Continuation Marks”) using a private key.

Examples:

```
> (parameterize ([exit-handler (lambda (x) 'no-exit)])
    (exit))

> (define p1 (make-parameter 1))

> (define p2 (make-parameter 2))

> (parameterize ([p1 3]
                 [p2 (p1)])
    (cons (p1) (p2)))
'(3 . 1)
> (let ([k (let/cc out
             (parameterize ([p1 2])
               (p1 3)
               (cons (let/cc k
                       (out k))
                     (p1)))))])
    (if (procedure? k)
        (k (p1))
        k))
'(1 . 3)
> (define ch (make-channel))
```

```
> (parameterize ([p1 0])
    (thread (lambda ()
              (channel-put ch (cons (p1) (p2))))))
#<thread>
> (channel-get ch)
'(0 . 2)
> (define k-ch (make-channel))

> (define (send-k)
    (parameterize ([p1 0])
      (thread (lambda ()
                (let/ec esc
                  (channel-put ch
                               ((let/cc k
                                  (channel-put k-ch k)
                                  (esc)))))))))

> (send-k)
#<thread>
> (thread (lambda () ((channel-get k-ch)
                      (let ([v (p1)])
                        (lambda () v)))))
#<thread>
> (channel-get ch)
1
> (send-k)
#<thread>
> (thread (lambda () ((channel-get k-ch) p1)))
#<thread>
> (channel-get ch)
0
```

```
(parameterize* ((parameter-expr value-expr) ...)
  body ...+)
```

Analogous to `let*` compared to `let`, `parameterize*` is the same as a nested series of single-parameter `parameterize` forms.

```
(make-derived-parameter parameter
                        guard
                        wrap)      → parameter?
  parameter : parameter?
  guard : (any/c . -> . any)
  wrap : (any/c . -> . any)
```

Returns a parameter procedure that sets or retrieves the same value as *parameter*, but with:

- *guard* applied when setting the parameter (before any guard associated with *parameter*), and
- *wrap* applied when obtaining the parameter's value.

See also chaperone-procedure, which can also be used to guard parameter procedures.

```
(parameter? v) → boolean?
  v : any/c
```

Returns #t if *v* is a parameter procedure, #f otherwise.

```
(parameter-procedure=? a b) → boolean?
  a : parameter?
  b : parameter?
```

Returns #t if the parameter procedures *a* and *b* always modify the same parameter with the same guards (although possibly with different chaperones), #f otherwise.

```
(current-parameterization) → parameterization?
```

Returns the current continuation's parameterization.

```
(call-with-parameterization parameterization
                            thunk)           → any
  parameterization : parameterization?
  thunk : (-> any)
```

Calls *thunk* (via a tail call) with *parameterization* as the current parameterization.

```
(parameterization? v) → boolean?
  v : any/c
```

Returns #t if *v* is a parameterization returned by current-parameterization, #f otherwise.

10.4 Futures

§18.10 "Parallelism with Futures" in ***The Racket Guide*** introduces futures.

```
(require racket/future)
```

The bindings documented in this section are provided by the racket/future and racket libraries, but not racket/base.

Currently, parallel support for future is enabled by default for Windows, Linux x86/x86_64, and Mac OS X x86/x86_64. To enable support for other platforms, use `--enable-futures` with `configure` when building Racket.

The future and touch functions from racket/future provide access to parallelism as supported by the hardware and operating system. In contrast to thread, which provides concurrency for arbitrary computations without parallelism, future provides parallelism for limited computations. A *future* executes its work in parallel (assuming that support for parallelism is available) until it detects an attempt to perform an operation that is too complex for the system to run safely in parallel. Similarly, work in a future is suspended if it depends in some way on the current continuation, such as raising an exception. A suspended computation for a future is resumed when touch is applied to the future.

"Safe" parallel execution of a future means that all operations provided by the system must be able to enforce contracts and produce results as documented. "Safe" does not preclude concurrent access to mutable data that is visible in the program. For example, a computation in a future might use `set!` to modify a shared variable, in which case concurrent assignment to the variable can be visible in other futures and threads. Furthermore, guarantees about the visibility of effects and ordering are determined by the operating system and hardware—which rarely support, for example, the guarantee of sequential consistency that is provided for thread-based concurrency. At the same time, operations that seem obviously safe may have a complex enough implementation internally that they cannot run in parallel. See also §18.10 "Parallelism with Futures" in ***The Racket Guide***.

A future never runs in parallel if all of the custodians that allow its creating thread to run are shut down. Such futures can execute through a call to touch, however.

```
(future thunk) → future?
  thunk : (-> any)
(touch f) → any
  f : future?
```

The future procedure returns a future value that encapsulates *thunk*. The touch function forces the evaluation of the *thunk* inside the given future, returning the values produced by *thunk*. After touch forces the evaluation of a *thunk*, the resulting values are retained by the future in place of *thunk*, and additional touches of the future return those values.

Between a call to future and touch for a given future, the given *thunk* may run speculatively in parallel to other computations, as described above.

```
> (let ([f (future (lambda () (+ 1 2)))])
    (list (+ 3 4) (touch f)))
'(7 3)
```

```
(futures-enabled?) → boolean?
```

Returns whether parallel support for futures is enabled in the current Racket configuration.

```
(current-future) → (or/c #f future?)
```

Returns the descriptor of the future whose thunk execution is the current continuation. If a future thunk itself uses `touch`, future-thunk executions can be nested, in which case the descriptor of the most immediately executing future is returned. If the current continuation is not a future-thunk execution, the result is `#f`.

```
(future? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a future value, `#f` otherwise.

```
(would-be-future thunk) → future?
  thunk : (-> any)
```

Returns a future that never runs in parallel, but that consistently logs all potentially "unsafe" operations during the execution of the future's thunk (i.e., operations that interfere with parallel execution).

With a normal future, certain circumstances might prevent the logging of unsafe operations. For example, when executed with debug-level logging,

```
(touch (future (lambda ()
                 (printf "hello1")
                 (printf "hello2")
                 (printf "hello3"))))
```

might log three messages, one for each `printf` invocation. However, if the `touch` is performed before the future has a chance to start running in parallel, the future thunk evaluates in the same manner as any ordinary thunk, and no unsafe operations are logged. Replacing `future` with `would-be-future` ensures the logging of all three calls to `printf`.

```
(processor-count) → exact-positive-integer?
```

Returns the number of parallel computation units (e.g., processors or cores) that are available on the current machine.

```
(make-fsemaphore init) → fsemaphore?
  init : exact-nonnegative-integer?
```

Creates and returns a new *future semaphore* with the counter initially set to *init*.

A future semaphore is similar to a plain semaphore, but future-semaphore operations can be performed safely in parallel (to synchronize parallel computations). In contrast, operations on plain semaphores are not safe to perform in parallel, and they therefore prevent a computation from continuing in parallel.

```
(fsemaphore? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is an future semaphore value, `#f` otherwise.

```
(fsemaphore-post fsema) → void?
  fsema : fsemaphore?
```

Increments the future semaphore's internal counter and returns `#<void>`.

```
(fsemaphore-wait fsema) → void?
  fsema : fsemaphore?
```

Blocks until the internal counter for *fsema* is non-zero. When the counter is non-zero, it is decremented and `fsemaphore-wait` returns `#<void>`.

```
(fsemaphore-try-wait? fsema) → boolean?
  fsema : fsemaphore?
```

Like `fsemaphore-wait`, but `fsemaphore-try-wait?` never blocks execution. If *fsema*'s internal counter is zero, `fsemaphore-try-wait?` returns `#f` immediately without decrementing the counter. If *fsema*'s counter is positive, it is decremented and `#t` is returned.

```
(fsemaphore-count fsema) → exact-nonnegative-integer?
  fsema : fsemaphore?
```

Returns *fsema*'s current internal counter value.

10.5 Futures Visualizer

§18.10 "Parallelism with Futures" in ***The Racket Guide*** introduces the future visualizer.

```
(require future-visualizer)
```

The *futures visualizer* is a graphical profiling tool for parallel programs written using `future`. The tool shows a timeline of a program's execution including all future-related events, as well as the overall amount of processor utilization at any point during the program's lifetime.

```
(visualize-futures e ...)
(visualize-futures-thunk thunk) → any
  thunk : (-> any)
```

The `visualize-futures` macro enables the collection of data required by the visualizer and displays a profiler window showing the corresponding trace. The `visualize-futures-thunk` provides similar functionality where program code is contained within *thunk*.

A typical program using profiling might look like the following:

```
(require racket/future
         future-visualizer)

(visualize-futures
 (let ([f (future (lambda () ...))])
   ...
   (touch f)))
```

The preceding program is equivalent to:

```
(require racket/future
         future-visualizer/trace
         future-visualizer)

(start-future-tracing!)
(let ([f (future (lambda () ...))])
  ...
  (touch f))
(stop-future-tracing!)
(show-visualizer)
```

```
(show-visualizer #:timeline timeline) → void?
  timeline : (listof indexed-future-event?)
```

Displays the visualizer window. If the function is called with no arguments, it must be preceded by the following sequence: a call to `start-future-tracing!`, program code that is being traced, and a call to `stop-future-tracing!` – in which case the visualizer will show data for all events logged in between those calls (via `timeline-events`). Note that `visualize-futures` and `visualize-futures-thunk` are simpler alternatives to using these primitives directly. The *timeline* argument can be used to show the visualizer for a previously-generated trace.

10.5.1 Execution Timeline

The *execution timeline*, shown in the top left-hand corner of the profiler window, displays a history of the program and all events associated with its futures, with OS-level threads or *processes* organized along the y-axis and time increasing along the x-axis. A coloring convention is used to distinguish between different types of events (see §10.6.1 "Future Performance Logging" for a full description of these event types):

- Blue dot: `'create`

- Green bar: `'start-work`, `'start-0-work`
- Orange dot: `'sync`
- Red dot: `'block`, `'touch`
- White dot: `'result`, `'end-work`
- Green dot: `'touch-pause`, `'touch-resume`

Mousing over any event connects it via purple lines to the sequence of events for its future. Additionally, orange dotted lines with arrowheads may be shown to indicate operations performed from one future to another (e.g. `'create` or `'touch` actions). To view details about two events simultaneously, a selection can be tacked by clicking the mouse.

The timeline displays vertical lines at 100-microsecond intervals. Note that though the time interval is fixed, the pixel distance between lines varies based on the event density for any given time range to prevent overlapping event circles.

```
(timeline-pict  events
                [#:x x
                 #:y y
                 #:width width
                 #:height height
                 #:selected-event-index selected-event-index])
 → pict?
  events : (listof indexed-future-event?)
  x : (or #f exact-nonnegative-integer?) = #f
  y : (or #f exact-nonnegative-integer?) = #f
  width : (or #f exact-nonnegative-integer?) = #f
  height : (or #f exact-nonnegative-integer?) = #f
  selected-event-index : (or #f exact-nonnegative-integer?) = #f
```

Returns a `pict` showing the execution timeline for the trace in *`events`*. The optional arguments *`x`*, *`y`*, *`width`*, and *`height`* can be used to obtain a specific area (in pixels) of the timeline image. The *`selected-event-index`* argument, if specified, shows the timeline image as if the user placed the mouse pointer over the `indexed-future-event` with the corresponding index.

10.5.2 Future Creation Tree

The *creation tree* shows a tree with a single node per future created by the program. This display can be particularly useful for programs which spawn futures in nested fashion (futures within futures). For any given future node, the children of that node represent futures which were created by that future (within the scope of its thunk). For all programs, the root

of the tree is a special node representing the main computation thread (the runtime thread), and is denoted *RTT*.

```
(creation-tree-pict  events
                    [#:x x
                     #:y y
                     #:width width
                     #:node-width node-width
                     #:height height
                     #:padding padding
                     #:zoom zoom])            → pict?
  events : (listof indexed-future-event?)
  x : (or #f exact-nonnegative-integer?) = #f
  y : (or #f exact-nonnegative-integer?) = #f
  width : (or #f exact-nonnegative-integer?) = #f
  node-width : (or #f exact-nonnegative-integer?) = #f
  height : (or #f exact-nonnegative-integer?) = #f
  padding : (or #f exact-nonnegative-integer?) = #f
  zoom : exact-nonnegative-integer? = 1
```

Returns a `pict` showing the future creation tree for the trace in *events*. The optional arguments *x*, *y*, *width*, and *height* can be used to obtain a specific area (in pixels) of the creation tree image. The *node-width* argument specifies (in pixels) the diameter of each node. The *padding* argument specifies the minimum space vertically between each depth and horizontally between siblings. The *zoom* argument specifies the zoom factor for the tree image in the range 1-5, where 5 returns a 500% zoom.

10.6 Futures Tracing

§18.10 “Parallelism with Futures” in ***The Racket Guide*** introduces the future visualizer.

```
(require future-visualizer/trace)
```

The *futures trace* module exposes low-level information about the execution of parallel programs written using `future`.

```
(trace-futures e ...)
(trace-futures-thunk thunk) → (listof indexed-future-event?)
  thunk : (-> any)
```

The `trace-futures` macro and `trace-futures-thunk` function track the execution of a program using futures and return the program trace as a list of `indexed-future-event` structures.

This program:

```
(require racket/future
```

```
          future-visualizer/trace)

(trace-futures
 (let ([f (future (lambda () ...))])
   ...
   (touch f)))
```

Is equivalent to:

```
(require racket/future
         future-visualizer/trace)

(start-future-tracing!)
(let ([f (future (lambda () ...))])
  ...
  (touch f))
(stop-future-tracing!)
(timeline-events)
```

```
(start-future-tracing!) → void?
(stop-future-tracing!) → void?
(timeline-events) → (listof indexed-future-event?)
```

The start-future-tracing! procedure enables the collection of future-related execution data. This function should be called immediately prior to executing code the programmer wishes to profile.

The stop-future-tracing! procedure must be used to indicate the end of code the programmer wishes to trace. Tracing works by simply using a log receiver to record all future-related log events; this procedure logs a special message that is well-known to the log receiver to mean 'stop recording'.

The timeline-events procedure returns the program trace as a list of indexed-future-event structures.

```
(struct indexed-future-event (index event)
        #:extra-constructor-name make-indexed-future-event)
  index : exact-nonnegative-integer?
  event : future-event?
```

Represents an individual log message in a program trace. Because multiple future-event structures may contain identical timestamps, the index field ranks them in the order in which they were recorded in the log output.

10.6.1 Future Performance Logging

Racket traces use logging (see §14.5 “Logging”) extensively to report information about how futures are evaluated. Logging output is useful for debugging the performance of programs that use futures.

Though textual log output can be viewed directly (or retrieved in code via `trace-futures`), it is much easier to use the graphical profiler tool provided by `future-visualizer`.

In addition to its string message, each event logged for a future has a data value that is an instance of a `future-event` prefab structure:

```
(struct future-event (future-id
                      proc-id
                      action
                      time-id
                      prim-name
                      user-data)
        #:extra-constructor-name make-future-event
        #:prefab)
  future-id : (or exact-nonnegative-integer? #f)
  proc-id : exact-nonnegative-integer?
  action : symbol?
  time-id : real?
  prim-name : (or symbol? #f)
  user-data : (or #f symbol? exact-nonnegative-integer?)
```

The `future-id` field is an exact integer that identifies a future, or it is `#f` when `action` is `'missing`. The `future-id` field is particularly useful for correlating logged events.

The `proc-id` fields is an exact, non-negative integer that identifies a parallel process. Process 0 is the main Racket process, where all expressions other than future thunks evaluate.

The `time-id` field is an inexact number that represents time in the same way as `current-inexact-milliseconds`.

The `action` field is a symbol:

- `'create`: a future was created.
- `'complete`: a future’s thunk evaluated successfully, so that `touch` will produce a value for the future immediately.
- `'start-work` and `'end-work`: a particular process started and ended working on a particular future.

- `'start-0-work`: like `'start-work`, but for a future thunk that for some structural reason could not be started in a process other than 0 (e.g., the thunk requires too much local storage to start).
- `'start-overflow-work`: like `'start-work`, where the future thunk's work was previously stopped due to an internal stack overflow.
- `'sync`: blocking (processes other than 0) or initiation of handing (process 0) for an "unsafe" operation in a future thunk's evaluation; the operation must run in process 0.
- `'block`: like `'sync`, but for a part of evaluation that must be delayed until the future is `touch`ed, because the evaluation may depend on the current continuation.
- `'touch` (never in process 0): like `'sync` or `'block`, but for a `touch` operation within a future thunk.
- `'overflow` (never in process 0): like `'sync` or `'block`, but for the case that a process encountered an internal stack overflow while evaluating a future thunk.
- `'result` or `'abort`: waiting or handling for `'sync`, `'block`, or `'touch` ended with a value or an error, respectively.
- `'suspend` (never in process 0): a process blocked by `'sync`, `'block`, or `'touch` abandoned evaluation of a future; some other process may pick up the future later.
- `'touch-pause` and `'touch-resume` (in process 0, only): waiting in `touch` for a future whose thunk is being evaluated in another process.
- `'missing`: one or more events for the process were lost due to internal buffer limits before they could be reported, and the `time-id` field reports an upper limit on the time of the missing events; this kind of event is rare.

Assuming no `'missing` events, then `'start-work`, `'start-0-work`, `'start-overflow-work` is always paired with `'end-work`; `'sync`, `'block`, and `'touch` are always paired with `'result`, `'abort`, or `'suspend`; and `'touch-pause` is always paired with `'touch-resume`.

In process 0, some event pairs can be nested within other event pairs: `'sync`, `'block`, or `'touch` with `'result` or `'abort`; and `'touch-pause` with `'touch-resume`.

An `'block` in process 0 is generated when an unsafe operation is handled. This type of event will contain a symbol in the `unsafe-op-name` field that is the name of the operation. In all other cases, this field contains `#f`.

The `prim-name` field will always be `#f` unless the event occurred on process 0 and its `action` is either `'block` or `'sync`. If these conditions are met, `prim-name` will contain the name of the Racket primitive which required the future to synchronize with the runtime thread (represented as a symbol).

The `user-data` field may take on a number of different values depending on both the `action` and `prim-name` fields:

- `'touch` on process 0: contains the integer ID of the future being touched.
- `'sync` and `prim-name` = `|allocate memory|`: The size (in bytes) of the requested allocation.
- `'sync` and `prim-name` = `jit_on_demand`: The runtime thread is performing a JIT compilation on behalf of the future `future-id`. The field contains the name of the function being JIT compiled (as a symbol).
- `'create`: A new future was created. The field contains the integer ID of the newly created future.

10.7 Places

§18.11 "Parallelism with Places" in ***The Racket Guide*** introduces places.

```
(require racket/place)
```

The bindings documented in this section are provided by the `racket/place` and `racket` libraries, but not `racket/base`.

Places enable the development of parallel programs that take advantage of machines with multiple processors, cores, or hardware threads.

Currently, parallel support for places is enabled only for Racket 3m (which is the main variant of Racket), and only by default for Windows, Linux x86/x86_64, and Mac OS X x86/x86_64. To enable support for other platforms, use `--enable-places` with `configure` when building Racket. The `place-enabled?` function reports whether places run in parallel.

A *place* is a parallel task that is effectively a separate instance of the Racket virtual machine. Places communicate through *place channels*, which are endpoints for a two-way buffered communication.

To a first approximation, place channels support only immutable, transparent values as messages. In addition, place channels themselves can be sent across channels to establish new (possibly more direct) lines of communication in addition to any existing lines. Finally, mutable values produced by `shared-flvector`, `make-shared-flvector`, `shared-fxvector`, `make-shared-fxvector`, `shared-bytes`, and `make-shared-bytes` can be sent across place channels; mutation of such values is visible to all places that share the value, because they are allowed in a *shared memory space*. See `place-message-allowed?`.

A place channel can be used as a synchronizable event (see §10.2.1 "Events") to receive a value through the channel. A place can also receive messages with `place-channel-get`, and messages can be sent with `place-channel-put`.

Two place channels are `equal?` if they are endpoints for the same underlying channels while both or neither is a place descriptor. Place channels can be `equal?` without being `eq?` after being sent messages through a place channel.

Constraints on messages across a place channel—and therefore on the kinds of data that

places share—enable greater parallelism than `future`, even including separate garbage collection of separate places. At the same time, the setup and communication costs for places can be higher than for futures.

For example, the following expression launches two places, echoes a message to each, and then waits for the places to terminate:

```
(let ([pls (for/list ([i (in-range 2)])
             (dynamic-place "place-worker.rkt" 'place-main))])
  (for ([i (in-range 2)]
        [p pls])
    (place-channel-put p i)
    (printf "~a\n" (place-channel-get p)))
  (map place-wait pls))
```

The `"place-worker.rkt"` module must export the `place-main` function that each place executes, where `place-main` must accept a single place channel argument:

```
#lang racket
(provide place-main)

(define (place-main pch)
  (place-channel-put pch (format "Hello from place ~a"
                                 (place-channel-get pch))))
```

```
(place-enabled?) → boolean?
```

Returns `#t` if Racket is configured so that `dynamic-place` and `place` create places that can run in parallel, `#f` if `dynamic-place` and `place` are simulated using `thread`.

```
(place? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a *place descriptor* value, `#f` otherwise. Every place descriptor is also a place channel.

```
(place-channel? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is place channel, `#f` otherwise.

```
(dynamic-place module-path start-name) → place?
  module-path : (or/c module-path? path?)
  start-name : symbol?
```

Creates a place to run the procedure that is identified by *module-path* and *start-name*. The result is a place descriptor value that represents the new parallel task; the place descriptor is returned immediately. The place descriptor value is also a place channel that permits communication with the place.

The module indicated by *module-path* must export a function with the name `start-proc`. The function must accept a single argument, which is a place channel that corresponds to the other end of communication for the place descriptor returned by `place`.

When the place is created, the initial exit handler terminates the place, using the argument to the exit handler as the place's *completion value*. Use `(exit v)` to immediately terminate a place with the completion value `v`. Since a completion value is limited to an exact integer between `0` and `255`, any other value for `v` is converted to `0`.

If the function indicated by *module-path* and `start-proc` returns, then the place terminates with the completion value `0`.

In the created place, the `current-input-port` parameter is set to an empty input port, while the values of the `current-output-port` and `current-error-port` parameters are connected to the current ports in the creating place. If the output ports are file-stream ports, then the connected ports in the places share the underlying stream, otherwise a thread in the creating place pumps bytes to the current ports in the creating place.

The *module-path* argument must not be a module path of the form `(quote sym)` unless the module is predefined (see `module-predefined?`).

```
(dynamic-place* module-path
                start-name
               [#:in in
                #:out out
                #:err err]) → place?
                              (or/c output-port? #f)
                              (or/c input-port? #f)
                              (or/c input-port? #f)
  module-path : (or/c module-path? path?)
  start-name : symbol?
  in : (or/c input-port? #f) = #f
  out : (or/c output-port? #f) = (current-output-port)
  err : (or/c output-port? #f) = (current-error-port)
```

Like `dynamic-place`, but accepts specific ports to the new place's ports, and returns a created port when `#f` is supplied for a port. The *in*, *out*, and *err* ports are connected to the `current-input-port`, `current-output-port`, and `current-error-port` ports, respectively, for the place. Any of the ports can be `#f`, in which case a file-stream port (for an operating-system pipe) is created and returned by `dynamic-place*`. The *err* argument can be `'stdout`, in which case the same file-stream port or that is supplied as standard output is also used for standard error. For each port or `'stdout` that is provided, no pipe is created

and the corresponding returned value is `#f`.

The caller of `dynamic-place*` is responsible for closing all returned ports; none are closed automatically.

The `dynamic-place*` procedure returns four values:

- a place descriptor value representing the created place;
- an output port piped to the place's standard input, or `#f` if *in* was a port;
- an input port piped from the place's standard output, or `#f` if *out* was a port;
- an input port piped from the place's standard error, or `#f` if *err* was a port or `'stdout`.

```
(place id body ...+)
```

Creates a place that evaluates *body* expressions with *id* bound to a place channel. The *body*s close only over *id* plus the top-level bindings of the enclosing module, because the *body*s are lifted to a function that is exported by the module. The result of `place` is a place descriptor, like the result of `dynamic-place`.

```
(place* maybe-port ...
        id
        body ...+)

maybe-port =
           | #:in in-expr
           | #:out out-expr
           | #:err err-expr
```

Like `place`, but supports optional `#:in`, `#:out`, and `#:err` expressions (at most one of each) to specify ports in the same way and with the same defaults as `dynamic-place*`. The result of a `place*` form is also the same as for `dynamic-place*`.

```
(place-wait p) → exact-integer?
  p : place?
```

Returns the completion value of the place indicated by *p*, blocking until the place has terminated.

If any pumping threads were created to connect a non-file-stream port to the ports in the place for *p* (see `dynamic-place`), `place-wait` returns only when the pumping threads have completed.

```
(place-dead-evt p) → evt?
  p : place?
```

Returns a synchronizable event (see §10.2.1 “Events”) that is ready if and only if *p* has terminated.

If any pumping threads were created to connect a non-file-stream port to the ports in the place for *p* (see dynamic-place), the event returned by place-dead-evt may become ready even if a pumping thread is still running.

```
(place-kill p) → void?
  p : place?
```

Immediately terminates the place, setting the place’s completion value to 1 if the place does not have a completion value already.

```
(place-break p) → void?
  p : place?
```

Sends place *p* a break signal; see §9.6 “Breaks”.

```
(place-channel) → place-channel? place-channel?
```

Returns two place channels. Data sent through the first channel can be received through the second channel, and data sent through the second channel can be received from the first.

Typically, one place channel is used by the current place to send messages to a destination place; the other place channel is sent to the destination place (via an existing place channel).

```
(place-channel-put pch v) → void
  pch : place-channel?
  v : place-message-allowed?
```

Sends a message *v* on channel *pch*. Since place channels are asynchronous, place-channel-put calls are non-blocking.

See place-message-allowed? form information on automatic coercions in *v*, such as converting a mutable string to an immutable string.

```
(place-channel-get pch) → place-message-allowed?
  pch : place-channel?
```

Returns a message received on channel *pch*, blocking until a message is available.

```
(place-channel-put/get pch v) → any/c
  pch : place-channel?
  v : any/c
```

Sends an immutable message *v* on channel *pch* and then waits for a message (perhaps a reply) on the same channel.

```
(place-message-allowed? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is allowed as a message on a place channel, `#f` otherwise.

If `(place-enabled?)` returns `#f`, then the result is always `#t` and no conversions are performed on *v* as a message. Otherwise, the following kinds of data are allowed as messages:

- numbers, characters, booleans, and `#<void>`;
- symbols, where the `eq?`ness of uninterned symbols is preserved within a single message, but not across messages;
- strings and byte strings, where mutable strings and byte strings are automatically replaced by immutable variants;
- paths (for any platform);
- pairs, lists, vectors, and immutable prefab structures containing message-allowed values, where a mutable vector is automatically replaced by an immutable vector;
- hash tables where mutable hash tables are automatically replaced by immutable variants;
- place channels, where a place descriptor is automatically replaced by a plain place channel;
- file-stream ports and TCP ports, where the underlying representation (such as a file descriptor, socket, or handle) is duplicated and attached to a fresh port in the receiving place;
- C pointers as created or accessed via `ffi/unsafe`; and
- values produced by `shared-flvector`, `make-shared-flvector`, `shared-fxvector`, `make-shared-fxvector`, `shared-bytes`, and `make-shared-bytes`.

10.8 Distributed Places

See also §18.12 “Distributed Places” in ***The Racket Guide***.

```
(require racket/place/distributed)
```

Distributed places support programs whose computation may span physical machines. The design relies on machine *nodes* that perform computation. The programmer configures a new distributed system using a declarative syntax and callbacks. A node begins life with one initial place: the *message router*. After a node has been configured, its message router

is activated by calling the message-router function. The message router listens on a TCP port for incoming connections from other nodes in the distributed system. Places can be spawned within the node by sending place-spawn request messages to the node's message router.

The distributed places implementation relies on two assumptions:

- The user's ".ssh/config" and ".ssh/authorized_keys" files are configured correctly to allow passwordless connection to remote hosts via public key authentication.
- Distributed places does not support the specification of ssh usernames. If a non-default ssh username is required the ".ssh/config" file should be used to specifiy the username.
- All machines run the same version of Racket. Futures versions of distributed places may use the zo binary data format for serialization.

The following example illustrates a configuration and use of distributed places that starts a new node on the current machine and passes it a "Hello World" string:

Example:

```
> (module hello-world-example racket/base
    (require racket/place/distributed
             racket/place)

    (provide hello-world)

    (define (hello-world)
      (place ch
        (printf "hello-world received: ~a\n" (place-channel-
get ch))
        (place-channel-put ch "Hello World\n")
        (printf "hello-world sent: Hello World\n")))


    (module+ main
      ; 1) spawns a node running at "localhost" and listenting on
port
      ; 6344 for incomming connections.
      ; 2) connects to the node running at localhost:6344 and cre-
ates a
      ; place on that node by calling the hello-world procedure
from
      ; the current module.
      ; 3) returns a remote-node% instance (node) and a
```

```
      ; remote-connection% instance (pl) for communicating with
the
      ; new node and place
      (define-values (node pl)
        (spawn-node-supervise-place-thunk-at "localhost"
                                             #:listen-port 6344
                                             (quote-module-
path "..")
                                             'hello-world))

      ; starts a message router which adds three event-
container<%>s to
      ; its list of events to handle: the node and two after-
seconds
      ; event containers . Two seconds after the launch of the
message-router, a
      ; message will be sent to the pl place.  After six seconds,
the
      ; program and all spawned nodes and places will terminate.
      (message-router
        node
        (after-seconds 2
          (*channel-put pl "Hello")
          (printf "message-router received: ~a\n" (*channel-
get pl)))

        (after-seconds 6
          (exit 0)))))
```

```
(message-router ec ...+) → void?
  ec : (is-a?/c event-container<%>)
```

Waits in an endless loop for one of many events to become ready. The `message-router` procedure constructs a `node%` instance to serve as the message router for the node. The `message-router` procedure then adds all the declared `event-container<%>`s to the `node%` and finally calls the never ending loop `sync-events` method, which handles events for the node.

```
(spawn-node-with-dynamic-place-at
 hostname
 instance-module-path
 instance-place-function-name
 [#:listen-port port
  #:initial-message initial-message
  #:racket-path racket-path
  #:ssh-bin-path ssh-path
  #:distributed-launch-path launcher-path
  #:restart-on-exit restart-on-exit])
→ (is-a?/c remote-connection%)
 hostname : string?
 instance-module-path : module-path?
 instance-place-function-name : symbol?
 port : port-no? = DEFAULT-ROUTER-PORT
 initial-message : any = #f
 racket-path : string-path? = (racket-path)
 ssh-path : string-path? = (ssh-bin-path)
 launcher-path : string-path?
               = (path->string distributed-launch-path)
 restart-on-exit : any/c = #f
```

Spawns a new remote node at *hostname* with one instance place specified by the *instance-module-path* and *instance-place-function-name*; that is, the result of (dynamic-place *instance-module-path instance-place-function-name*). is called in the new place of the new node.

The result is a remote-node% instance, not a remote-connection%. Use get-first-place on the result to obtain a remote-connection%.

The *restart-on-exit* argument can be #t to instruct the remote-connection% instance to respawn the place on the remote node should it exit or terminate at any time. It can also be a procedure of zero arguments to implement the restart procedure, or it can be an object that support a restart method that takes a place argument.

```
(spawn-node-supervise-dynamic-place-at
 hostname
 instance-module-path
 instance-place-function-name
 [#:listen-port port
  #:initial-message initial-message
  #:racket-path racket-path
  #:ssh-bin-path ssh-path
  #:distributed-launch-path launcher-path
  #:restart-on-exit restart-on-exit])
→ (is-a?/c remote-node%)
  (is-a?/c remote-connection%)
```

```
 hostname : string?
 instance-module-path : module-path?
 instance-place-function-name : symbol?
 port : port-no? = DEFAULT-ROUTER-PORT
 initial-message : any = #f
 racket-path : string-path? = (racket-path)
 ssh-path : string-path? = (ssh-bin-path)
 launcher-path : string-path?
               = (path->string distributed-launch-path)
 restart-on-exit : any/c = #f
```

Like `spawn-node-with-dynamic-place-at`, but the result is two values: the new `remote-node%` and its `remote-connection%` instance.

```
(spawn-node-with-place-thunk-at
 hostname
 instance-module-path
 instance-thunk-function-name
 [#:listen-port port
  #:initial-message initial-message
  #:racket-path racket-path
  #:ssh-bin-path ssh-path
  #:distributed-launch-path launcher-path
  #:restart-on-exit restart-on-exit])
→ (is-a?/c remote-connection%)
 hostname : string?
 instance-module-path : module-path?
 instance-thunk-function-name : symbol?
 port : port-no? = DEFAULT-ROUTER-PORT
 initial-message : any = #f
 racket-path : string-path? = (racket-path)
 ssh-path : string-path? = (ssh-bin-path)
 launcher-path : string-path?
               = (path->string distributed-launch-path)
 restart-on-exit : any/c = #f
```

Like `spawn-node-with-dynamic-place-at`, but in the new node, `(dynamic-place instance-module-path instance-thunk-function-name)` is called to create a place and return the newly constructed the place descriptor. That is, the *instance-thunk-function-name* function should use `dynamic-place` or `place` to create an initial place in the new node.

```
(spawn-node-supervise-place-thunk-at
 hostname
 instance-module-path
 instance-thunk-function-name
 [#:listen-port port
  #:initial-message initial-message
  #:racket-path racket-path
  #:ssh-bin-path ssh-path
  #:distributed-launch-path launcher-path
  #:restart-on-exit restart-on-exit])
  (is-a?/c remote-node%)
→ (is-a?/c remote-connection%)
 hostname : string?
 instance-module-path : module-path?
 instance-thunk-function-name : symbol?
 port : port-no? = DEFAULT-ROUTER-PORT
 initial-message : any = #f
 racket-path : string-path? = (racket-path)
 ssh-path : string-path? = (ssh-bin-path)
 launcher-path : string-path?
               = (path->string distributed-launch-path)
 restart-on-exit : any/c = #f
```

Like spawn-node-with-place-thunk-at, but the result is two values (like spawn-node-supervise-dynamic-place-at): the new remote-node% and its remote-connection% instance.

```
(spawn-remote-racket-node
 hostname
 [#:listen-port port
  #:racket-path racket-path
  #:ssh-bin-path ssh-path
  #:distributed-launch-path launcher-path]
  #:use-current-ports use-current-ports)
→ (is-a?/c remote-node%)
 hostname : string?
 port : port-no? = DEFAULT-ROUTER-PORT
 racket-path : string-path? = (racket-path)
 ssh-path : string-path? = (ssh-bin-path)
 launcher-path : string-path?
               = (path->string distributed-launch-path)
 use-current-ports : #f
```

Spawns a new remote node at *hostname* and returns a remote-node% handle.

```
(create-place-node  hostname
                   [#:listen-port port
                    #:racket-path racket-path
                    #:ssh-bin-path ssh-path
                    #:distributed-launch-path launcher-path]
                    #:use-current-ports use-current-ports)
 → (is-a?/c remote-node%)
  hostname : string?
  port : port-no? = DEFAULT-ROUTER-PORT
  racket-path : string-path? = (racket-path)
  ssh-path : string-path? = (ssh-bin-path)
  launcher-path :  string-path?
                = (path->string distributed-launch-path)
  use-current-ports : #t
```

Like spawn-remote-racket-node, but the current-output-port and current-error-port are used as the standard ports for the spawned process instead of new pipe ports.

```
(supervise-dynamic-place-at  remote-node
                             instance-module-path
                             instance-place-function-name
                            [#:restart-on-exit restart-on-exit])
 → (is-a?/c remote-connection%)
  remote-node : (is-a?/c remote-node%)
  instance-module-path : module-path?
  instance-place-function-name : symbol?
  restart-on-exit : any/c = #f
```

Creates a new place on *remote-node* by using dynamic-place to invoke *instance-place-function-name* from the module *instance-module-path*.

```
(supervise-place-thunk-at  remote-node
                           instance-module-path
                           instance-thunk-function-name
                          [#:restart-on-exit restart-on-exit])
 → (is-a?/c remote-connection%)
  remote-node : (is-a?/c remote-node%)
  instance-module-path : module-path?
  instance-thunk-function-name : symbol?
  restart-on-exit : any/c = #f
```

Creates a new place at *remote-node* by executing the thunk exported as *instance-thunk-function-name* from the module *instance-module-path*. The function should use dynamic-place or place to create a place in the new node.

```
(supervise-process-at  hostname
                       commandline-argument ...+
                      [#:listen-port port])
 → (is-a?/c remote-process%)
  hostname : string?
  commandline-argument : string?
  port : port-no? = DEFAULT-ROUTER-PORT
```

Spawns an attached external process at host *hostname*.

```
 (supervise-named-dynamic-place-at
  remote-node
  place-name
  instance-module-path
  instance-place-function-name
 [#:restart-on-exit restart-on-exit])
 → (is-a?/c remote-connection%)
  remote-node : (is-a?/c remote-node%)
  place-name : symbol?
  instance-module-path : module-path?
  instance-place-function-name : symbol?
  restart-on-exit : any/c = #f
```

Creates a new place on the *remote-node* by using dynamic-place to invoke *instance-place-function-name* from the module *instance-module-path*. The *place-name* symbol is used to establish later connections to the named place.

```
 (supervise-named-place-thunk-at
  remote-node
  place-name
  instance-module-path
  instance-thunk-function-name
 [#:restart-on-exit restart-on-exit])
 → (is-a?/c remote-connection%)
  remote-node : (is-a?/c remote-node%)
  place-name : symbol?
  instance-module-path : module-path?
  instance-thunk-function-name : symbol?
  restart-on-exit : any/c = #f
```

Creates a new place on the *remote-node* by executing the thunk *instance-thunk-function-name* from the module *instance-module-path*; that is, the function should use dynamic-place or place to create a place. The *place-name* symbol is used to establish later connections to the named place.

```
(supervise-thread-at  remote-node
                      instance-module-path
                      instance-thunk-function-name
                     [#:restart-on-exit restart-on-exit])
 → (is-a?/c remote-connection%)
  remote-node : (is-a?/c remote-node%)
  instance-module-path : module-path?
  instance-thunk-function-name : symbol?
  restart-on-exit : any/c = #f
```

Creates a new threadon the *remote-node* by using `dynamic-require` to invoke `instance-place-function-name` from the module *instance-module-path*.

```
(restart-every  seconds
                [#:retry retry
                 #:on-final-fail on-final-fail])
 → (is-a/c respawn-and-fire%)
  seconds : (number?)
  retry : (or/c number? #f) = #f
  on-final-fail : (or/c #f (-> any/c)) = #f
```

Returns a `restarter%` instance that should be supplied to a `#:restart-on-exit` argument.

```
(every-seconds seconds-expr body ....)
```

Returns a `respawn-and-fire%` instance that should be supplied to a `message-router`. The `respawn-and-fire%` instance executes *body*s once every *N* seconds, where *N* is the result of *seconds-expr*.

```
(after-seconds seconds-expr body ....)
```

Returns a `after-seconds%` instance that should be supplied to a `message-router`. The `after-seconds%` instance executes the *body*s after a delay of *N* seconds from the start of the event loop, where *N* is the result of *seconds-expr*.

```
(connect-to-named-place node name)
 → (is-a?/c remote-connection%)
  node : (is-a?/c remote-node%)
  name : symbol?
```

Connects to a named place on the *node* named *name* and returns a `remote-connection%` object.

```
(log-message severity msg) → void?
  severity : (or/c 'fatal 'error 'warning 'info 'debug)
  msg : string?
```

Logs a message at the root node.

```
event-container<%> : interface?
```

All objects that are supplied to the `message-router` must implement the `event-container<%>` interface. The `message-router` calls the `register` method on each supplied `event-container<%>` to obtain a list of events on which the event loop should wait.

```
(send an-event-container register events) → (listof events?)
  events : (listof events?)
```

Returns the list of events inside the `event-container<%>` that should be waited on by the `message-router`.

The following classes all implement `event-container<%>` and can be supplied to a `message-router`: `spawned-process%`, `place-socket-bridge%`, `node%`, `remote-node%`, `remote-connection%`, `place% connection%`, `respawn-and-fire%`, and `after-seconds%`.

```
spawned-process% : class?
  superclass: object%
  extends: event-container<%>
```

```
(send a-spawned-process get-pid) → exact-positive-integer?
```

```
 (new spawned-process%
     [cmdline-list cmdline-list]
    [[parent parent]])
 → (is-a?/c spawned-process%)
  cmdline-list : (listof (or/c string? path?))
  parent : (is-a?/c remote-node%) = #f
```

The *cmdline-list* is a list of command line arguments of type `string` and/or `path`.

The *parent* argument is a `remote-node%` instance that will be notified when the process dies via a `(send parent process-died this)` call.

```
place-socket-bridge% : class?
  superclass: object%
  extends: event-container<%>
```

```
(send a-place-socket-bridge get-sc-id)
 → exact-positive-integer?
```

```
 (new place-socket-bridge%
    [pch pch]
    [sch sch]
    [id id])
 → (is-a?/c place-socket-bridge%)
  pch : place-channel?
  sch : (is-a?/c socket-connection%)
  id : exact-positive-integer?
```

The *pch* argument is a `place-channel`. Messages received on *pch* are forwarded to the socket-connection% *sch* via a `dcgm` message. e.g. `(sconn-write-flush sch (dcgm DCGM-TYPE-INTER-DCHANNEL id id msg))` The *id* is a `exact-positive-integer` that identifies the socket-connection subchannel for this inter-node place connection.

```
socket-connection% : class?
  superclass: object%
  extends: event-container<%>
```

```
 (new socket-connection%
    [[host host]
     [port port]
     [retry-times retry-times]
     [delay delay]
     [background-connect? background-connect?]
     [in in]
     [out out]
     [remote-node remote-node]])
 → (is-a?/c socket-connection%)
  host : (or/c string? #f) = #f
  port : (or/c port-no? #f) = #f
  retry-times : exact-nonnegative-integer? = 30
  delay : number? = 1
  background-connect? : any/c = #f
```

```
  in : (or/c input-port? #f) = #f
  out : (or/c output-port #f) = #f
  remote-node : (or/c (is-a?/c remote-node%) #f) = #f
```

When a *host* and *port* are supplied a new tcp connection is established. If a input-port? and output-port? are supplied as *in* and *out*, the ports are used as a connection to the remote host. The *retry-times* argument specifies how many times to retry the connection attempt should it fail to connect and defaults to 30 retry attempts. Often a remote node is still booting up when a connection is attempted and the connection needs to be retried several times. The *delay* argument specifies how many seconds to wait between retry attempts. The *background-connect?* argument defaults to #t and specifies that the constructor should retry immediately and that connecion establishment should occur in the background. Finally, the *remote-node* argument specifies the remote-node% instance that should be notified should the connection fail.

```
node% : class?
  superclass: object%
  extends: event-container<%>
```

The node% instance controls a distributed places node. It launches places and routes inter-node place messages in the distributed system. The message-router form constructs a node% instance under the hood. Newly spawned nodes also have a node% instance in their initial place that serves as the node's message router.

```
(new node% [[listen-port listen-port]]) → (is-a?/c node%)
  listen-port : tcp-listen-port? = #f
```

Constructs a node% that will listen on *listen-port* for inter-node connections.

```
(send a-node sync-events) → void?
```

Starts the never ending event loop for this distributed places node.

```
remote-node% : class?
  superclass: object%
  extends: event-container<%>
```

The node% instance controls a distributed places node. It launches compute places and routes inter-node place messages in the distributed system. This is the remote api to a distributed places node. Instances of remote-node% are returned by spawn-remote-racket-node, spawn-node-supervise-dynamic-place-at, and spawn-node-supervise-place-thunk-at.

```
(new remote-node%
   [[listen-port listen-port]
    [restart-on-exit restart-on-exit]])
 → (is-a?/c remote-node%)
  listen-port : tcp-listen-port? = #f
  restart-on-exit : any/c = #f
```

Constructs a node% that will listen on *listen-port* for inter-node connections.

When set to true the *restart-on-exit* parameter causes the specified node to be restarted when the ssh session spawning the node dies.

```
(send a-remote-node get-first-place)
 → (is-a?/c remote-connection%)
```

Returns the remote-connection% object instance for the first place spawned on this node.

```
(send a-remote-node get-first-place-channel) → place-
channel?
```

Returns the communication channel for the first place spawned on this node.

```
(send a-remote-node get-log-prefix) → string?
```

Returns (format "PLACE ~a:~a" host-name listen-port)

```
(send a-remote-node launch-place
  place-exec
 [#:restart-on-exit restart-on-exit
  #:one-sided-place? one-sided-place?])
 → (is-a?/c remote-connection%)
  place-exec : list?
  restart-on-exit : any/c = #f
  one-sided-place? : any/c = #f
```

Launches a place on the remote node represented by this remote-node% instance.

The *place-exec* argument describes how the remote place should be launched, and it should have one of the following shapes:

- (list 'place *place-module-path place-thunk*)
- (list 'dynamic-place *place-module-path place-func*)

The difference between these two launching methods is that the 'place version of *place-exec* expects a thunk to be exported by the module place-module-path. Executing the thunk is expected to create a new place and return a place

descriptor to the newly created place. The 'dynamic-place version of *place-exec* expects place-func to be a function taking a single argument, the initial channel argument, and calls dynamic-place on behalf of the user and creates the new place from the place-module-path and place-func.

The *restart-on-exit* argument is treated in the same way as for spawn-node-with-dynamic-place-at.

The *one-sided-place?* argument is an internal use argument for launching remote places from within a place using the old design pattern.

```
(send a-remote-node remote-connect name) → remote-
connection%
  name : string?
```

Connects to a named place on the remote node represented by this remote-node% instance.

```
(send a-remote-node send-exit) → void?
```

Sends a message instructing the remote node represented by this remote-node% instance to exit immediately

```
(node-send-exit remote-node%) → void?
  remote-node% : node
```

Sends node a message telling it to exit immediately.

```
(node-get-first-place remote-node%)
 → (is-a?/c remote-connection%)
  remote-node% : node
```

Returns the remote-connection% instance of the first place spawned at this node

```
(distributed-place-wait remote-connection%) → void?
  remote-connection% : place
```

Waits for place to terminate.

```
remote-connection% : class?
  superclass: object%
  extends: event-container<%>
```

The remote-connection% instance provides a remote api to a place running on a remote distributed places node. It launches a places or connects to a named place and routes inter-node place messages to the remote place.

```
(new remote-connection%
   [node node]
   [place-exec place-exec]
   [name name]
   [restart-on-exit restart-on-exit]
   [one-sided-place? one-sided-place?]
   [on-channel on-channel])
 → (is-a?/c remote-connection%)
 node : (is-a?/c remote-node%)
 place-exec : list?
 name : string?
 restart-on-exit : #f
 one-sided-place? : #f
 on-channel : #f
```

Constructs a remote-connection% instance.

The *place-exec* argument describes how the remote place should be launched in the same way as for launch-place in remote-node%.

The *restart-on-exit* argument is treated in the same way as for spawn-node-with-dynamic-place-at.

The *one-sided-place?* argument is an internal use argument for launching remote places from within a place using the old design pattern.

See set-on-channel! for description of *on-channel* argument.

```
(send a-remote-connection set-on-channel! callback) → void?
  callback : (-> channel msg void?)
```

Installs a handler function that handles messages from the remote place. The setup/distributed-docs module uses this callback to handle job completion messages.

```
place% : class?
  superclass: object%
  extends: event-container<%>
```

The place% instance represents a place launched on a distributed places node at that node. It launches a compute places and routes inter-node place messages to the place.

```
(new place%
   [node node]
   [place-exec place-exec]
   [ch-id ch-id]
   [sc sc]
  [[on-place-dead on-place-dead]]) → (is-a?/c place%)
```

```
  node : (is-a?/c remote-connection%)
  place-exec : list?
  ch-id : exact-positive-integer?
  sc : (is-a?/c socket-connection%)
  on-place-dead : (-> event void?) = default-on-place-dead
```

Constructs a remote-connection% instance. The *place-exec* argument describes how the remote place should be launched in the same way as for launch-place in remote-node%. The *ch-id* and *sc* arguments are internally used to establish routing between the remote node spawning this place and the place itself. The *on-place-dead* callback handles the event when the newly spawned place terminates.

```
(send a-place wait-for-die) → void?
```

Blocks and waits for the subprocess representing the remote-node% to exit.

```
connection% : class?
  superclass: object%
  extends: event-container<%>
```

The connection% instance represents a connection to a named-place instance running on the current node. It routes inter-node place messages to the named place.

```
 (new connection%
    [node node]
    [name name]
    [ch-id ch-id]
    [sc sc])        → (is-a?/c connection%)
  node : (is-a?/c remote-node%)
  name : string?
  ch-id : exact-positive-integer?
  sc : (is-a?/c socket-connection%)
```

Constructs a remote-connection% instance. The place-exec argument describes how the remote place should be launched in the same way as for launch-place in remote-node%. The *ch-id* and *sc* arguments are internally used to establish routing between the remote node and this named-place.

```
respawn-and-fire% : class?
  superclass: object%
  extends: event-container<%>
```

The respawn-and-fire% instance represents a thunk that should execute every n seconds.

```
(new respawn-and-fire%
   [seconds seconds]
   [thunk thunk])
→ (is-a?/c respawn-and-fire%)
 seconds : (and/c real? (not/c negative?))
 thunk : (-> void?)
```

Constructs a `respawn-and-fire%` instance that when placed inside a `message-router` construct causes the supplied thunk to execute every `n` seconds.

```
after-seconds% : class?
  superclass: object%
  extends: event-container<%>
```

The `after-seconds%` instance represents a thunk that should execute after `n` seconds.

```
(new after-seconds%
   [seconds seconds]
   [thunk thunk])
→ (is-a?/c after-seconds%)
 seconds : (and/c real? (not/c negative?))
 thunk : (-> void?)
```

Constructs an `after-seconds%` instance that when placed inside a `message-router` construct causes the supplied thunk to execute after `n` seconds.

```
restarter% : class?
  superclass: after-seconds%
  extends: event-container<%>
```

The `restarter%` instance represents a restart strategy.

```
(new restarter%
    [seconds seconds]
   [[retry retry]
    [on-final-fail on-final-fail]])
→ (is-a?/c restarter%)
 seconds : number?
 retry : (or/c number? #f) = #f
 on-final-fail : (or/c #f (-> any/c)) = #f
```

Constructs an `restarter%` instance that when supplied to a `#:restart-on-exit` argument, attempts to restart the process every *seconds*. The *retry*

argument specifies how many time to attempt to restart the process before giving up. If the process stays alive for `(* 2 seconds)` the attempted retries count is reset to `0`. The *on-final-fail* thunk is called when the number of retries is exceeded

```
distributed-launch-path : path?
```

Contains the local path to the distributed places launcher. The distributed places launcher is the bootsrap file that launches the message router on a new node.

```
(ssh-bin-path) → string?
```

Returns the path to the ssh binary on the local system in string form.

Example:

```
> (ssh-bin-path)
#<path:/usr/bin/ssh>
```

```
(racket-path) → path?
```

Returns the path to the currently executing racket binary on the local system.

```
(build-distributed-launch-path collects-path) → string?
  collects-path : path-string?
```

Returns the path to the distributed places launch file. The function can take an optional argument specifying the path to the collects directory.

```
(spawn-node-at  hostname
                [#:listen-port port
                 #:racket-path racket-path
                 #:ssh-bin-path ssh-path
                 #:distributed-launch-path launcher-path])
 → channel?
  hostname : string?
  port : port-no? = DEFAULT-ROUTER-PORT
  racket-path : string-path? = (racket-path)
  ssh-path : string-path? = (ssh-bin-path)
  launcher-path :  string-path?
                 = (path->string distributed-launch-path)
```

Spawns a node in the background using a Racket thread and returns a channel that becomes ready with a `remote-node%` once the node has spawned successfully

```
(spawn-nodes/join nodes-descs) → void?
  nodes-descs : list?
```

Spawns a list of nodes by calling `(lambda (x) (apply keyword-apply spawn-node-at x))` for each node description in *nodes-descs* and then waits for each node to spawn.

```
(*channel-put ch msg) → void?
  ch : (or/c place-channel? async-bi-channel?
             channel? (is-a?/c remote-connection%))
  msg : any
```

Sends *msg* over *ch* channel.

```
(*channel-get ch) → any
  ch : (or/c place-channel? async-bi-channel?
             channel? (is-a?/c remote-connection%))
```

Returns a message received on *ch* channel.

```
(*channel? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is one of `place-channel?`, `async-bi-channel?`, `channel?`, or `(is-a?/c remote-connection%)`.

```
(send-new-place-channel-to-named-dest ch
                                      src-id
                                      dest-list)
 → place-channel?
  ch : *channel?
  src-id : any
  dest-list : (listof string? port-no? string?)
```

Creates and returns a new place channel connection to a named place at *dest-list*. The *dest-list* argument is a list of a remote-hostname remote-port and named-place name. The channel *ch* should be a connection to a `message-router`.

```
(mr-spawn-remote-node  mrch
                       host
                      [#:listen-port listen-port
                       #:solo solo])              → void?
  mrch : *channel?
  host : string?
  listen-port : port-no? = DEFAULT-ROUTER-PORT
  solo : boolean? = #f
```

Sends a message to a message router over *mrch* channel asking the message router to spawn a new node at *host* listening on port *listen-port*. If the #:solo keyword argument is supplied the new node is not folded into the complete network with other nodes in the distributed system.

```
(mr-supervise-named-dynamic-place-at mrch
                                     dest
                                     name
                                     path
                                     func) → void?
  mrch : *channel?
  dest : (listof string? port-no?)
  name : string?
  path : string?
  func : symbol?
```

Sends a message to a message router over *mrch* channel asking the message router to spawn a named place at *dest* named *name*. The place is spawned at the remote node by calling dynamic place with module-path *path* and function *func*. The *dest* parameter should be a list of remote-hostname and remote-port.

```
(mr-connect-to mrch dest name) → void?
  mrch : *channel?
  dest : (listof string? port-no?)
  name : string?
```

Sends a message to a message router over *mrch* channel asking the message router to create a new connection to the named place named *name* at *dest*. The *dest* parameter should be a list of remote-hostname and remote-port.

```
(start-message-router/thread [#:listen-port listen-port
                              #:nodes nodes])
 → thread? channel?
  listen-port : port-no? = DEFAULT-ROUTER-PORT
  nodes : list? = null
```

Starts a message router in a racket thread connected to *nodes*, listening on port *listen-port*, and returns a channel? connection to the message router.

```
(port-no? no) → boolean?
  no : (and/c exact-nonnegative-integer?  (integer-in 0 65535))
```

Returns #t if *no* is a exact-nonnegative-integer? between 0 and 65535.

```
DEFAULT-ROUTER-PORT : port-no?
```

The default port for distributed places message router.

```
named-place-typed-channel% : class?
  superclass: object%
```

```
(new named-place-typed-channel% [ch ch])
 → (is-a?/c named-place-typed-channel%)
  ch : place-channel?
```

The *ch* argument is a `place-channel`.

```
(send a-named-place-typed-channel get type) → any
  type : symbol?
```

Returns the first message received on `ch` that has the type *type*. Messages are lists and their type is the first item of the list which should be a `symbol?`. Messages of other types that are received are queued for later `get` requests.

```
(tc-get type ch) → void?
  type : symbol?
  ch : place-channel?
```

Gets a message of type *type* from the `named-place-typed-channel%` *ch*.

```
(write-flush datum port) → void?
  datum : any
  port : port?
```

Writes *datum* to *port* and then flushes *port*.

```
(printf/f format args ...) → void?
  format : string?
  args : any
```

Calls `printf` followed by a call to `flush-output`.

```
(displayln/f item) → void?
  item : any
```

Calls `displayln` followed by a call to `flush-output`.

Example:

```
> (write-flush "Hello World" (current-output-port))
"Hello World"
```

10.8.1 Define Remote Server

```
(require racket/place/define-remote-server)
```

```
(define-remote-server [name identifier?] rpc-forms ...+)
(define-named-remote-server [name identifier?] rpc-forms ...+)
```

The `define-remote-server` and `define-named-remote-server` forms are nearly identical. The `define-remote-server` form should be used with `supervise-dynamic-place-at` to build a private rpc server, while the `define-named-remote-server` form should be used with `supervise-named-dynamic-place-at` to build a rpc server inside a named place.

The `define-named-remote-server` form takes an identifier and a list of custom expressions as its arguments. From the identifier a function is created by prepending the `make-` prefix. This procedure takes a single argument a `place-channel`. In the example below, the `make-tuple-server` identifier is the `place-function-name` given to the `supervise-named-dynamic-place-at` form to spawn an rpc server. The server created by the `make-tuple-server` procedure sits in a loop waiting for rpc requests from the `define-rpc` functions documented below.

```
(define-state id value)
```

Expands to a `define`, which is closed over by the `define-rpc` functions to form local state.

```
(define-rpc (id args ...) body ...)
```

Expands to a client rpc function `name-id` which sends *id* and *args* ... to the rpc server `rpc-place` and waits for a response. `(define (name-id rpc-place args ...) body)`

```
(define-cast (id args ...) body ...)
```

Expands to a client rpc function `name-id` which sends *id* and *args* ... to the rpc server `rpc-place` but does not receive any response. A cast is a one-way communication technique. `(define (name-id rpc-place args ...) body)`

The `define-state` custom form translates into a simple `define` form, which is closed over by the `define-rpc` forms.

The `define-rpc` form is expanded into two parts. The first part is the client stubs that call the rpc functions. The client function name is formed by concatenating the `define-named-remote-server` identifier, `tuple-server`, with the RPC function name `set` to form `tuple-server-set`. The RPC client functions take a destination argument which is a `remote-connection%` descriptor and then the RPC function arguments. The RPC

client function sends the RPC function name, `set`, and the RPC arguments to the destination by calling an internal function `named-place-channel-put`. The RPC client then calls `named-place-channel-get` to wait for the RPC response.

The second expansion part of `define-rpc` is the server implementation of the RPC call. The server is implemented by a match expression inside the `make-tuple-server` function. The match clause for `tuple-server-set` matches on messages beginning with the `'set` symbol. The server executes the RPC call with the communicated arguments and sends the result back to the RPC client.

The `define-cast` form is similar to the `define-rpc` form except there is no reply message from the server to client

Example:

```
> (module tuple-server-example racket/base
    (require racket/match
             racket/place/define-remote-server)

    (define-named-remote-server tuple-server
      (define-state h (make-hash))
      (define-rpc (set k v)
        (hash-set! h k v)
        v)
      (define-rpc (get k)
        (hash-ref h k #f))
      (define-cast (hello)
        (printf "Hello from define-cast\n")
        (flush-output))))
```

Example:

```
> (module bank-server-example racket/base
    (require racket/match
             racket/place/define-remote-server)

    (define-remote-server bank
      (define-state accounts (make-hash))
      (define-rpc (new-account who)
         (match (hash-has-key? accounts who)
           [#t '(already-exists)]
           [else
             (hash-set! accounts who 0)
             (list 'created who)]))
      (define-rpc (remove who amount)
         (cond
```

```
            [(hash-ref accounts who (lambda () #f)) =>
               (lambda (balance)
                 (cond [(<= amount balance)
                        (define new-balance (- balance amount))
                        (hash-set! accounts who new-balance)
                        (list 'ok new-balance)]
                       [else
                        (list 'insufficient-funds balance)]))]
            [else
              (list 'invalid-account who)]))
      (define-rpc (add who amount)
        (cond
           [(hash-ref accounts who (lambda () #f)) =>
              (lambda (balance)
                (define new-balance (+ balance amount))
                (hash-set! accounts who new-balance)
                (list 'ok new-balance))]
           [else
             (list 'invalid-account who)])))
```

```
(log-to-parent msg [#:severity severity]) → void?
  msg : string?
  severity : symbol? = 'info
```

The log-to-parent procedure can be used inside a `define-remote-server` or `define-named-remote-server` form to send a logging message to the remote owner of the rpc server.

10.8.2 Async Bidirectional Channels

```
(require racket/place/private/async-bi-channel)
```

```
(make-async-bi-channel) → async-bi-channel?
```

Creates and returns a opaque structure, which is the async bidirectional channel.

```
(async-bi-channel? ch) → boolean?
  ch : any
```

A predicate that returns #t for async bidirectional channels.

```
(async-bi-channel-get ch) → any
  ch : async-bi-channel?
```

Returns the next available message from the async bidirectional channel *ch*.

```
(async-bi-channel-put ch msg) → void?
  ch : async-bi-channel?
  msg : any
```

Sends message *msg* to the remote end of the async bidirectional channel *ch*.

10.8.3 Distributed Places MPI

```
(require racket/place/distributed/rmpi)
```

```
(struct rmpi-comm (id cnt channels))
  id : exact-nonnegative-integer?
  cnt : exact-positive-integer?
  channels : vector?
```

The communicator struct rmpi-comm contains the rmpi process rank id, the quantity of processes in the communicator group, cnt, and a vector of place-channels, one for each process in the group.

```
(rmpi-id comm) → exact-nonnegative-integer?
  comm : rmpi-comm?
```

Takes a rmpi communicator structure, *comm*, and returns the node id of the RMPI process.

```
(rmpi-cnt comm) → exact-positive-integer?
  comm : rmpi-comm?
```

Takes a rmpi communicator structure, *comm*, and returns the count of the RMPI processes in the communicator group.

```
(rmpi-send comm dest val) → void?
  comm : rmpi-comm?
  dest : exact-nonnegative-integer?
  val : any
```

Sends *val* to destination rmpi process number *dest* using the RMPI communicator structure *comm*.

```
(rmpi-recv comm src) → any
  comm : rmpi-comm?
  src : exact-nonnegative-integer?
```

Receives a message from source rmpi process number *src* using the RMPI communicator structure *comm*.

```
                    rmpi-comm?
(rmpi-init ch) → (listof any)
                    (is-a?/c named-place-typed-channel%)
  ch : place-channel?
```

Creates the rmpi-comm structure instance using the named place's original place-channel *ch*. In addition to the communicator structure, rmpi-init returns a list of initial arguments and the original place-channel *ch* wrapped in a named-place-typed-channel%. The named-place-typed-channel% wrapper allows for the reception of list messages typed by an initial symbol.

```
(rmpi-broadcast comm src) → any
  comm : rmpi-comm?
  src : exact-nonnegative-integer?
(rmpi-broadcast comm src val) → any
  comm : rmpi-comm?
  src : exact-nonnegative-integer?
  val : any
```

Broadcasts *val* from *src* to all rmpi processes in the communication group using a hypercube algorithm. Receiving processes call (rmpi-broadcast *comm src*).

```
(rmpi-reduce comm dest op val) → any
  comm : rmpi-comm?
  dest : exact-nonnegative-integer?
  op : procedure?
  val : any
```

Reduces *val* using the *op* operator to *dest* rmpi node using a hypercube algorithm.

```
(rmpi-barrier comm) → void?
  comm : rmpi-comm?
```

Introduces a synchronization barrier for all rmpi processes in the communcication group *comm*.

```
(rmpi-allreduce comm op val) → any
  comm : rmpi-comm?
  op : procedure?
  val : any
```

Reduces *val* using the *op* operator to rmpi node 0 and then broadcasts the reduced value to all nodes in the communication group.

```
(rmpi-partition comm num) → exact-nonnegative-integer?
                          exact-nonnegative-integer?
  comm : rmpi-comm?
  num : exact-nonnegative-integer?
```

Partitions *num* into `rmpi-cnt` equal pieces and returns the offset and length for the `rmpi-id`th piece.

```
(rmpi-build-default-config
 #:racket-path racket-path
 #:distributed-launch-path distributed-launch-path
 #:mpi-module mpi-module
 #:mpi-func mpi-func
 #:mpi-args mpi-args)
→ hash?
 racket-path : string?
 distributed-launch-path : string?
 mpi-module : string?
 mpi-func : symbol?
 mpi-args : (listof any)
```

Builds a hash from keywords to keyword arguments for use with the `rmpi-launch function`.

```
(rmpi-launch default-node-config config) → void?
  default-node-config : hash?
  config : (listof (list/c string? port-no? symbol?
                           exact-nonnegative-integer?))
```

Launches distributed places nodes running `#:mpi-func` in `#:mpi-module` with `#:mpi-args`. The config is a list of node configs, where each node config consists of a hostname, port, named place symbol and rmpi id number, followed by and optional hash of keyword `#:racket-path`, `#:distributed-launch-path`, `#:mpi-module`, `#:mpi-func`, and `#:mpi-args` to keyword arguments. Missing optional keyword arguments will be taken from the *default-node-config* hash of keyword arguments.

```
(rmpi-finish comm tc) → void?
  comm : rmpi-comm?
  tc : (is-a?/c named-place-typed-channel%)
```

Rendezvous with the `rmpi-launch`, using the *tc* returned by `rmpi-launch`, to indicate that the RMPI module is done executing and that `rmpi-launch` can return control to its caller.

Example:

```
> (rmpi-launch
      (rmpi-build-default-config
        #:racket-path "/tmp/mplt/bin/racket"
        #:distributed-launch-path (build-distributed-launch-path
                                   "/tmp/mplt/collects")
        #:mpi-module "/tmp/mplt/kmeans.rkt"
        #:mpi-func   'kmeans-place
        #:mpi-args    (list "/tmp/mplt/color100.bin" #t 100
                            9 10 1e-07))

      (list (list "nodea.example.com" 6340 'kmeans_0 0)
            (list "nodeb.example.com" 6340 'kmeans_1 1)
            (list "nodec.example.com" 6340 'kmeans_2 2)
            (list "noded.example.com" 6340 'kmeans_3 3)))
```

11 Macros

§16 “Macros” in ***The Racket Guide*** introduces Macros.

See §1.2 “Syntax Model” for general information on how programs are parsed. In particular, the subsection §1.2.3.2 “Expansion Steps” describes how parsing triggers macros, and §1.2.3.5 “Transformer Bindings” describes how macro transformers are called.

11.1 Pattern-Based Syntax Matching

```
(syntax-case stx-expr (literal-id ...)
  clause ...)

     clause = [pattern result-expr]
            | [pattern fender-expr result-expr]

    pattern = _
            | id
            | (pattern ...)
            | (pattern ...+ . pattern)
            | (pattern ... pattern ellipsis pattern ...)
            | (pattern ... pattern ellipsis pattern ... . pattern)
            | #(pattern ...)
            | #(pattern ... pattern ellipsis pattern ...)
            | #s(key-datum pattern ...)
            | #s(key-datum pattern ... pattern ellipsis pattern ...)
            | (ellipsis stat-pattern)
            | const

stat-pattern = id
            | (stat-pattern ...)
            | (stat-pattern ...+ . stat-pattern)
            | #(stat-pattern ...)
            | const

   ellipsis = ...
```

Finds the first *pattern* that matches the syntax object produced by *stx-expr*, and for which the corresponding *fender-expr* (if any) produces a true value; the result is from the corresponding *result-expr*, which is in tail position for the `syntax-case` form. If no *clause* matches, then the `exn:fail:syntax` exception is raised; the exception is generated by calling `raise-syntax-error` with `#f` as the “name” argument, a string with a generic error message, and the result of *stx-expr*.

A syntax object matches a *pattern* as follows:

_

A _ pattern (i.e., an identifier with the same binding as _) matches any syntax object.

id

An *id* matches any syntax object when it is not bound to ... or _ and does not have the same binding as any *literal-id*. The *id* is further bound as *pattern variable* for the corresponding *fender-expr* (if any) and *result-expr*. A pattern-variable binding is a transformer binding; the pattern variable can be reference only through forms like `syntax`. The binding's value is the syntax object that matched the pattern with a *depth marker* of 0.

An *id* that has the same binding as a *literal-id* matches a syntax object that is an identifier with the same binding in the sense of `free-identifier=?`. The match does not introduce any pattern variables.

(*pattern* ...)

A (*pattern* ...) pattern matches a syntax object whose datum form (i.e., without lexical information) is a list with as many elements as sub-*pattern*s in the pattern, and where each syntax object that corresponds to an element of the list matches the corresponding sub-*pattern*.

Any pattern variables bound by the sub-*pattern*s are bound by the complete pattern; the bindings must all be distinct.

(*pattern* ...+ . *pattern*)

The last *pattern* must not be a (*pattern* ...), (*pattern* ...+ . *pattern*), (*pattern* ... *pattern ellipsis pattern* ...), or (*pattern* ... *pattern ellipsis pattern* ... . *pattern*) form.

Like the previous kind of pattern, but matches syntax objects that are not necessarily lists; for *n* sub-*pattern*s before the last sub-*pattern*, the syntax object's datum must be a pair such that *n*-1 `cdr`s produce pairs. The last sub-*pattern* is matched against the syntax object corresponding to the *n*th `cdr` (or the `datum->syntax` coercion of the datum using the nearest enclosing syntax object's lexical context and source location).

(*pattern* ... *pattern ellipsis pattern* ...)

Like the (*pattern* ...) kind of pattern, but matching a syntax object with any number (zero or more) elements that match the sub-*pattern* followed by *ellipsis* in the corresponding position relative to other sub-*pattern*s.

For each pattern variable bound by the sub-*pattern* followed by *ellipsis*, the larger pattern binds the same pattern variable to a list of values, one for each element of the syntax object matched to the sub-*pattern*, with an incremented depth marker. (The sub-*pattern* itself may contain *ellipsis*, leading to a pattern variables bound to lists of lists of syntax objects with a depth marker of 2, and so on.)

(*pattern* ... *pattern ellipsis pattern* ... . *pattern*)

Like the previous kind of pattern, but with a final sub-*pattern* as for (*pattern* ...+ . *pattern*). The final *pattern* never matches a syntax object whose datum is a pair.

#(*pattern* ...)

Like a (*pattern* ...) pattern, but matching a vector syntax object whose elements match the corresponding sub-*pattern*s.

#(*pattern* ... *pattern ellipsis pattern* ...)

Like a (*pattern* ... *pattern ellipsis pattern* ...) pattern, but matching a vector syntax object whose elements match the corresponding sub-*pattern*s.

#s(*key-datum pattern* ...)

Like a (*pattern* ...) pattern, but matching a prefab structure syntax object whose fields match the corresponding sub-*pattern*s. The *key-datum* must correspond to a valid first argument to make-prefab-struct.

#s(*key-datum pattern* ... *pattern ellipsis pattern* ...)

Like a (*pattern* ... *pattern ellipsis pattern* ...) pattern, but matching a prefab structure syntax object whose elements match the corresponding sub-*pattern*s.

(*ellipsis stat-pattern*)

Matches the same as *stat-pattern*, which is like a *pattern*, but identifiers with the binding `...` are treated the same as other *id*s.

const

A *const* is any datum that does not match one of the preceding forms; a syntax object matches a *const* pattern when its datum is `equal?` to the `quoted` *const*.

If *stx-expr* produces a syntax object that is tainted or armed, then any syntax object bound by a *pattern* are tainted—unless the binding corresponds to the whole syntax object produced by *stx-expr*, in which case it remains tainted or armed.

Examples:

```
> (require (for-syntax racket/base))

> (define-syntax (swap stx)
    (syntax-case stx ()
      [(_ a b) #'(let ([t a])
                   (set! a b)
                   (set! b t))]))

> (let ([x 5] [y 10])
    (swap x y)
    (list x y))
'(10 5)
> (syntax-case #'(ops 1 2 3 => +) (=>)
    [(_ x ... => op) #'(op x ...)])
#<syntax:554:0 (+ 1 2 3)>
> (syntax-case #'(let ([x 5] [y 9] [z 12])
                   (+ x y z))
                (let)
    [(let ([var expr] ...) body ...)
     (list #'(var ...)
           #'(expr ...))])
'(#<syntax:555:0 (x y z)> #<syntax:555:0 (5 9 12)>)
```

(`syntax-case*` *stx-expr* (*literal-id* ...) *id-compare-expr*
  clause ...)

Like `syntax-case`, but *id-compare-expr* must produce a procedure that accepts two arguments. A *literal-id* in a *pattern* matches an identifier for which the procedure

returns true when given the identifier to match (as the first argument) and the identifier in the *pattern* (as the second argument).

In other words, `syntax-case` is like `syntax-case*` with an *id-compare-expr* that produces `free-identifier=?`.

```
(with-syntax ([pattern stx-expr] ...)
  body ...+)
```

Similar to `syntax-case`, in that it matches a *pattern* to a syntax object. Unlike `syntax-case`, all *pattern*s are matched, each to the result of a corresponding *stx-expr*, and the pattern variables from all matches (which must be distinct) are bound with a single *body* sequence. The result of the `with-syntax` form is the result of the last *body*, which is in tail position with respect to the `with-syntax` form.

If any *pattern* fails to match the corresponding *stx-expr*, the `exn:fail:syntax` exception is raised.

A `with-syntax` form is roughly equivalent to the following `syntax-case` form:

```
(syntax-case (list stx-expr ...) ()
  [(pattern ...) (let () body ...+)])
```

However, if any individual *stx-expr* produces a non-syntax object, then it is converted to one using `datum->syntax` and the lexical context and source location of the individual *stx-expr*.

Examples:

```
> (define-syntax (hello stx)
    (syntax-case stx ()
      [(_ name place)
       (with-syntax ([print-name #'(printf "~a\n" 'name)]
                     [print-place #'(printf "~a\n" 'place)])
         #'(begin
             (define (name times)
               (printf "Hello\n")
               (for ([i (in-range 0 times)])
                    print-name))
             (define (place times)
               (printf "From\n")
               (for ([i (in-range 0 times)])
                    print-place))))]))

> (hello jon utah)
```

```
> (jon 2)
Hello
jon
jon

> (utah 2)
From
utah
utah

> (define-syntax (math stx)
    (define (make+1 expression)
      (with-syntax ([e expression])
        #'(+ e 1)))
    (syntax-case stx ()
      [(_ numbers ...)
       (with-syntax ([(added ...)
                      (map make+1
                           (syntax->list #'(numbers ...)))])
         #'(begin
             (printf "got ~a\n" added)
             ...))]))

> (math 3 1 4 1 5 9)
got 4
got 2
got 5
got 2
got 6
got 10
```

(syntax *template*)

```
     template = id
              | (template-elem ...)
              | (template-elem ...+ . template)
              | #(template-elem ...)
              | #s(key-datum template-elem ...)
              | (ellipsis stat-template)
              | const

template-elem = template ellipsis ...

stat-template = id
              | (stat-template ...)
              | (stat-template ... . stat-template)
              | #(stat-template ...)
              | #s(key-datum stat-template ...)
              | const

     ellipsis = ...
```

Constructs a syntax object based on a *template*, which can include pattern variables bound by `syntax-case` or `with-syntax`.

Template forms produce a syntax object as follows:

id

If *id* is bound as a pattern variable, then *id* as a template produces the pattern variable's match result. Unless the *id* is a sub-*template* that is replicated by *ellipsis* in a larger *template*, the pattern variable's value must be a syntax object with a depth marker of 0 (as opposed to a list of matches).

More generally, if the pattern variable's value has a depth marker *n*, then it can only appear within a template where it is replicated by at least *n* *ellipsis*es. In that case, the template will be replicated enough times to use each match result at least once.

If *id* is not bound as a pattern variable, then *id* as a template produces `(quote-syntax` *id*`)`.

`(`*template-elem* `...)`

Produces a syntax object whose datum is a list, and where the elements of the list correspond to syntax objects produced by the *template-elem*s.

A *template-elem* is a sub-*template* replicated by any number of *ellipsis*es:

- If the sub-*template* is replicated by no *ellipsis*es, then it generates a single syntax object to incorporate into the result syntax object.
- If the sub-*template* is replicated by one *ellipsis*, then it generates a sequence of syntax objects that is "inlined" into the resulting syntax object.

  The number of generated elements depends on the values of pattern variables referenced within the sub-*template*. There must be at least one pattern variable whose value has a depth marker less than the number of *ellipsis*es after the pattern variable within the sub-*template*.

  If a pattern variable is replicated by more *ellipsis*es in a *template* than the depth marker of its binding, then the pattern variable's result is determined normally for inner *ellipsis*es (up to the binding's depth marker), and then the result is replicated as necessary to satisfy outer *ellipsis*es.
- For each *ellipsis* after the first one, the preceding element (with earlier replicating *ellipsis*es) is conceptually wrapped with parentheses for generating output, and then the wrapping parentheses are removed in the resulting syntax object.

(*template-elem* ... . *template*)

Like the previous form, but the result is not necessarily a list; instead, the place of the empty list in the resulting syntax object's datum is taken by the syntax object produced by *template*.

#(*template-elem* ...)

Like the (*template-elem* ...) form, but producing a syntax object whose datum is a vector instead of a list.

#s(*key-datum template-elem* ...)

Like the (*template-elem* ...) form, but producing a syntax object whose datum is a prefab structure instead of a list. The *key-datum* must correspond to a valid first argument of `make-prefab-struct`.

(*ellipsis stat-template*)

Produces the same result as *stat-template*, which is like a *template*, but ... is treated like an *id* (with no pattern binding).

const

> A *const* template is any form that does not match the preceding cases, and it produces the result (quote-syntax *const*).

A (syntax *template*) form is normally abbreviated as #'*template*; see also §1.3.7 “Reading Quotes”. If *template* contains no pattern variables, then #'*template* is equivalent to (quote-syntax *template*).

(quasisyntax *template*)

Like syntax, but (unsyntax *expr*) and (unsyntax-splicing *expr*) escape to an expression within the *template*.

The *expr* must produce a syntax object (or syntax list) to be substituted in place of the unsyntax or unsyntax-splicing form within the quasiquoting template, just like unquote and unquote-splicing within quasiquote. (If the escaped expression does not generate a syntax object, it is converted to one in the same way as for the right-hand side of with-syntax.) Nested quasisyntaxes introduce quasiquoting layers in the same way as nested quasiquotes.

Also analogous to quasiquote, the reader converts #` to quasisyntax, #, to unsyntax, and #,@ to unsyntax-splicing. See also §1.3.7 “Reading Quotes”.

(unsyntax *expr*)

Illegal as an expression form. The unsyntax form is for use only with a quasisyntax template.

(unsyntax-splicing *expr*)

Illegal as an expression form. The unsyntax-splicing form is for use only with a quasisyntax template.

(syntax/loc *stx-expr template*)

Like syntax, except that the immediate resulting syntax object takes its source-location information from the result of *stx-expr* (which must produce a syntax object), unless the *template* is just a pattern variable, or both the source and position of *stx-expr* are #f.

(quasisyntax/loc *stx-expr template*)

Like quasisyntax, but with source-location assignment like syntax/loc.

```
(quote-syntax/prune id)
```

Like quote-syntax, but the lexical context of *id* is pruned via identifier-prune-lexical-context to including binding only for the symbolic name of *id* and for '#%top. Use this form to quote an identifier when its lexical information will not be transferred to other syntax objects (except maybe to '#%top for a top-level binding).

```
(syntax-rules (literal-id ...)
  [(id . pattern) template] ...)
```

Equivalent to

```
(lambda (stx)
  (syntax-case stx (literal-id ...)
    [(generated-id . pattern) (syntax-protect #'template)]  ...))
```

where each *generated-id* binds no identifier in the corresponding *template*.

```
(syntax-id-rules (literal-id ...)
  [pattern template] ...)
```

Equivalent to

```
(make-set!-transformer
 (lambda (stx)
   (syntax-case stx (literal-id ...)
     [pattern (syntax-protect #'template)]  ...)))
```

```
(define-syntax-rule (id . pattern) template)
```

Equivalent to

```
(define-syntax id
  (syntax-rules ()
   [(id . pattern) template]))
```

but with syntax errors potentially phrased in terms of *pattern*.

```
...
```

The ... transformer binding prohibits ... from being used as an expression. This binding is useful only in syntax patterns and templates, where it indicates repetitions of a pattern or template. See syntax-case and syntax.

```
_
```

The _ transformer binding prohibits _ from being used as an expression. This binding is useful only in syntax patterns, where it indicates a pattern that matches any syntax object. See `syntax-case`.

```
(syntax-pattern-variable? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a value that, as a transformer-binding value, makes the bound variable as pattern variable in `syntax` and other forms. To check whether an identifier is a pattern variable, use `syntax-local-value` to get the identifier's transformer value, and then test the value with `syntax-pattern-variable?`.

The `syntax-pattern-variable?` procedure is provided `for-syntax` by `racket/base`.

11.2 Syntax Object Content

```
(syntax? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a syntax object, `#f` otherwise. See also §1.2.2 "Syntax Objects".

```
(identifier? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a syntax object and `(syntax-e stx)` produces a symbol.

```
(syntax-source stx) → any
  stx : syntax?
```

Returns the source for the syntax object *stx*, or `#f` if none is known. The source is represented by an arbitrary value (e.g., one passed to `read-syntax`), but it is typically a file path string. Source-location information is dropped for a syntax object that is marshaled as part of compiled code; see also `current-compile`.

```
(syntax-line stx) → (or/c exact-positive-integer? #f)
  stx : syntax?
```

Returns the line number (positive exact integer) for the start of the syntax object in its source, or `#f` if the line number or source is unknown. The result is `#f` if and only if `(syntax-column stx)` produces `#f`. See also §12.1.4 "Counting Positions, Lines, and Columns", and see `syntax-source` for information about marshaling compiled syntax objects.

```
(syntax-column stx) → (or/c exact-nonnegative-integer? #f)
  stx : syntax?
```

Returns the column number (non-negative exact integer) for the start of the syntax object in its source, or `#f` if the source column is unknown. The result is `#f` if and only if `(syntax-line stx)` produces `#f`. See also §12.1.4 “Counting Positions, Lines, and Columns”, and see `syntax-source` for information about marshaling compiled syntax objects.

```
(syntax-position stx) → (or/c exact-positive-integer? #f)
  stx : syntax?
```

Returns the character position (positive exact integer) for the start of the syntax object in its source, or `#f` if the source position is unknown. See also §12.1.4 “Counting Positions, Lines, and Columns”, and see `syntax-source` for information about marshaling compiled syntax objects.

```
(syntax-span stx) → (or/c exact-nonnegative-integer? #f)
  stx : syntax?
```

Returns the span (non-negative exact integer) in characters of the syntax object in its source, or `#f` if the span is unknown. See also `syntax-source` for information about marshaling compiled syntax objects.

```
(syntax-original? stx) → boolean?
  stx : syntax?
```

Returns `#t` if *stx* has the property that `read-syntax` attaches to the syntax objects that they generate (see §11.7 “Syntax Object Properties”), and if *stx*’s lexical information does not indicate that the object was introduced by a syntax transformer (see §1.2.2 “Syntax Objects”). The result is `#f` otherwise. This predicate can be used to distinguish syntax objects in an expanded expression that were directly present in the original expression, as opposed to syntax objects inserted by macros.

```
(syntax-source-module stx [source?])
 → (or/c module-path-index? symbol? path? resolved-module-path? #f)
  stx : syntax?
  source? : any/c = #f
```

Returns an indication of the module whose source contains *stx*, or `#f` if *stx* has no source module. If *source?* is `#f`, then result is a module path index or symbol (see §13.4.2 “Compiled Modules and References”) or a resolved module path; if *source?* is true, the result is a path or symbol corresponding to the loaded module’s source in the sense of `current-module-declare-source`.

```
(syntax-e stx) → any
  stx : syntax?
```

Unwraps the immediate datum structure from a syntax object, leaving nested syntax structure (if any) in place. The result of `(syntax-e stx)` is one of the following:

- a symbol
- a syntax pair (described below)
- the empty list
- an immutable vector containing syntax objects
- an immutable box containing syntax objects
- an immutable hash table containing syntax object values (but not necessarily syntax object keys)
- an immutable prefab structure containing syntax objects
- some other kind of datum—usually a number, boolean, or string—that is interned when `datum-intern-literal` would convert the value

A *syntax pair* is a pair containing a syntax object as its first element, and either the empty list, a syntax pair, or a syntax object as its second element.

A syntax object that is the result of `read-syntax` reflects the use of delimited `.` in the input by creating a syntax object for every pair of parentheses in the source, and by creating a pair-valued syntax object *only* for parentheses in the source. See §1.3.5 “Reading Pairs and Lists” for more information.

If `stx` is tainted or armed, then any syntax object in the result of `(syntax-e stx)` is tainted, and multiple calls to `syntax-e` may return values that are not `eq?`. For a `stx` that is not armed, the results from multiple calls to `syntax-e` of `stx` are `eq?`.

```
(syntax->list stx) → (or/c list? #f)
  stx : syntax?
```

Returns a list of syntax objects or `#f`. The result is a list of syntax objects when `(syntax->datum stx)` would produce a list. In other words, syntax pairs in `(syntax-e stx)` are flattened.

If `stx` is tainted or armed, then any syntax object in the result of `(syntax->list stx)` is tainted.

```
(syntax->datum stx) → any
  stx : syntax?
```

Returns a datum by stripping the lexical information, source-location information, properties, and tamper status from *stx*. Inside of pairs, (immutable) vectors, (immutable) boxes, immutable hash table values (not keys), and immutable prefab structures, syntax objects are recursively stripped.

The stripping operation does not mutate *stx*; it creates new pairs, vectors, boxes, hash tables, and prefab structures as needed to strip lexical and source-location information recursively.

```
(datum->syntax ctxt v [srcloc prop ignored]) → syntax?
  ctxt : (or/c syntax? #f)
  v : any/c
            (or/c syntax? #f
                  (list/c any/c
                          (or/c exact-positive-integer? #f)
                          (or/c exact-nonnegative-integer? #f)
                          (or/c exact-positive-integer? #f)
  srcloc :                (or/c exact-nonnegative-integer? #f))
                  (vector/c any/c
                           (or/c exact-positive-integer? #f)
                           (or/c exact-nonnegative-integer? #f)
                           (or/c exact-positive-integer? #f)
                           (or/c exact-nonnegative-integer? #f)))
         = #f
  prop : (or/c syntax? #f) = #f
  ignored : (or/c syntax? #f) = #f
```

Converts the datum *v* to a syntax object. The contents of pairs, vectors, and boxes, the fields of prefab structures, and the values of immutable hash tables are recursively converted. The keys of prefab structures and the keys of immutable hash tables are not converted. Mutable vectors and boxes are replaced by immutable vectors and boxes. For any kind of value other than a pair, vector, box, immutable hash table, immutable prefab structure, or syntax object, conversion means wrapping the value with lexical information, source-location information, and properties after the value is interned via `datum-intern-literal`.

Converted objects in *v* are given the lexical context information of *ctxt* and the source-location information of *srcloc*. If *v* is not already a syntax object, then the resulting immediate syntax object is given the properties (see §11.7 “Syntax Object Properties”) of *prop*; if *v* is a pair, vector, box, immutable hash table, or immutable prefab structure, recursively converted values are not given properties or certificates. If *ctxt* is tainted or armed, then the resulting syntax object from `datum->syntax` is tainted.

Any of *ctxt*, *srcloc*, or *prop* can be #f, in which case the resulting syntax has no lexical context, source information, and/or new properties.

If *srcloc* is not #f or a syntax object, it must be a list or vector of five elements:

```
(list source-name line column position span)
or (vector source-name line column position span)
```

where source-name-v is an arbitrary value for the source name; line is an integer for the source line, or #f; column is an integer for the source column, or #f; position is an integer for the source position, or #f; and span is an integer for the source span, or #f. The line and column values must both be numbers or both be #f, otherwise the exn:fail:contract exception is raised.

Graph structure is not preserved by the conversion of *v* to a syntax object. Instead, *v* is essentially unfolded into a tree. If *v* has a cycle through pairs, vectors, boxes, immutable hash tables, and immutable prefab structures, then the exn:fail:contract exception is raised.

The *ignored* argument is allowed for backward compatibility and has no effect on the returned syntax object.

```
(datum-intern-literal v) → any/c
  v : any/c
```

Converts some values to be consistent with an interned result produced by the default reader in read-syntax mode.

If *v* is a number, character, string, byte string, or regular expression, then the result is a value that is equal? to *v* and eq? to a potential result of the default reader. (Note that mutable strings and byte strings are interned as immutable strings and byte strings.)

If *v* is an uninterned or an unreadable symbol, the result is still *v*, since an interned symbol would not be equal? to *v*.

The conversion process does not traverse compound values. For example, if *v* is a pair containing strings, then the strings within *v* are not interned.

If *v1* and *v2* are equal? but not eq?, then it is possible that (datum-intern-literal *v1*) will return *v1* and—sometime after *v1* becomes unreachable as determined by the garbage collector (see §1.1.7 “Garbage Collection”)—(datum-intern-literal *v2*) can still return *v2*. In other words, datum-intern-literal may adopt a given value as an interned representative, but if a former representative becomes otherwise unreachable, then datum-intern-literal may adopt a new representative.

```
(syntax-shift-phase-level stx shift) → syntax?
  stx : syntax?
  shift : (or/c exact-integer? #f)
```

Returns a syntax object that is like *stx*, but with all of its top-level and module binding shifted by *shift* phase levels. If *shift* is #f, then only bindings at phase level 0 are shifted to the label phase level. If *shift* is 0, then the result is *stx*.

```
(generate-temporaries stx-pair) → (listof identifier?)
  stx-pair : (or syntax? list?)
```

Returns a list of identifiers that are distinct from all other identifiers. The list contains as many identifiers as *stx-pair* contains elements. The *stx-pair* argument must be a syntax pair that can be flattened into a list. The elements of *stx-pair* can be anything, but string, symbol, keyword (possibly wrapped as syntax), and identifier elements will be embedded in the corresponding generated name, which is useful for debugging purposes.

The generated identifiers are built with interned symbols (not gensyms); see also §1.4.15 “Printing Compiled Code”.

```
(identifier-prune-lexical-context  id-stx
                                   [syms]) → identifier?
  id-stx : identifier?
  syms : (listof symbol?) = (list (syntax-e id-stx))
```

Returns an identifier with the same binding as *id-stx*, but without lexical information from *id-stx* that does not apply to the symbols in *syms*, where even further extension of the lexical information drops information for other symbols. In particular, transferring the lexical context via datum->syntax from the result of this function to a symbol other than one in *syms* produces an identifier with no binding.

See also quote-syntax/prune.

```
(identifier-prune-to-source-module id-stx) → identifier?
  id-stx : identifier?
```

Returns an identifier with its lexical context minimized to that needed for syntax-source-module. The minimized lexical context does not include any bindings.

```
(syntax-recertify new-stx
                  old-stx
                  inspector
                  key)      → syntax?
  new-stx : syntax?
  old-stx : syntax?
  inspector : inspector?
  key : any/c
```

For backward compatibility only; returns *new-stx*.

11.3 Syntax Object Bindings

```
(bound-identifier=? a-id b-id [phase-level]) → boolean?
  a-id : syntax?
  b-id : syntax?
  phase-level : (or/c exact-integer? #f)
              = (syntax-local-phase-level)
```

Returns `#t` if the identifier `a-id` would bind `b-id` (or vice versa) if the identifiers were substituted in a suitable expression context at the phase level indicated by `phase-level`, `#f` otherwise. A `#f` value for `phase-level` corresponds to the label phase level.

Examples:

```
> (define-syntax (check stx)
    (syntax-case stx ()
      [(_ x y)
       (if (bound-identifier=? #'x #'y)
           #'(let ([y 'wrong]) (let ([x 'binds]) y))
           #'(let ([y 'no-binds]) (let ([x 'wrong]) y)))]))

> (check a a)
'binds
> (check a b)
'no-binds
> (define-syntax-rule (check-a x) (check a x))

> (check-a a)
'no-binds
```

```
(free-identifier=?  a-id
                    b-id
                   [a-phase-level
                    b-phase-level]) → boolean?
  a-id : syntax?
  b-id : syntax?
  a-phase-level : (or/c exact-integer? #f)
                = (syntax-local-phase-level)
  b-phase-level : (or/c exact-integer? #f) = a-phase-level
```

Returns `#t` if `a-id` and `b-id` access the same local binding, module binding, or top-level binding—perhaps via rename transformers—at the phase levels indicated by `a-phase-level` and `b-phase-level`, respectively. A `#f` value for `a-phase-level` or `b-phase-level` corresponds to the label phase level.

"Same module binding" means that the identifiers refer to the same original definition site, and not necessarily to the same `require` or `provide` site. Due to renaming in `require` and

`provide`, or due to a transformer binding to a rename transformer, the identifiers may return distinct results with `syntax-e`.

Examples:

```
> (define-syntax (check stx)
    (syntax-case stx ()
      [(_ x)
       (if (free-identifier=? #'car #'x)
           #'(list 'same: x)
           #'(list 'different: x))]))

> (check car)
'(same: #<procedure:car>)
> (check mcar)
'(different: #<procedure:mcar>)
> (let ([car list])
    (check car))
'(different: #<procedure:list>)
> (require (rename-in racket/base [car kar]))

> (check kar)
'(same: #<procedure:car>)
```

```
(free-transformer-identifier=? a-id b-id) → boolean?
  a-id : syntax?
  b-id : syntax?
```

Same as `(free-identifier=? a-id b-id (add1 (syntax-local-phase-level)))`.

```
(free-template-identifier=? a-id b-id) → boolean?
  a-id : syntax?
  b-id : syntax?
```

Same as `(free-identifier=? a-id b-id (sub1 (syntax-local-phase-level)))`.

```
(free-label-identifier=? a-id b-id) → boolean?
  a-id : syntax?
  b-id : syntax?
```

Same as `(free-identifier=? a-id b-id #f)`.

```
(check-duplicate-identifier ids) → (or/c identifier? #f)
  ids : (listof identifier?)
```

Compares each identifier in *ids* with every other identifier in the list with bound-identifier=?. If any comparison returns #t, one of the duplicate identifiers is returned (the first one in *ids* that is a duplicate), otherwise the result is #f.

```
(identifier-binding id-stx [phase-level])
    (or/c 'lexical
          #f
          (listof module-path-index?
                  symbol?
 →                module-path-index?
                  symbol?
                  (or/c 0 1)
                  (or/c exact-integer? #f)
                  (or/c exact-integer? #f)))
  id-stx : syntax?
  phase-level : (or/c exact-integer? #f)
              = (syntax-local-phase-level)
```

Returns one of three kinds of values, depending on the binding of *id-stx* at the phase level indicated by *phase-level* (where a #f value for *phase-level* corresponds to the label phase level):

- The result is 'lexical if *id-stx* has a local binding. If 'lexical is produced for any *phase-level* value, then it is produced for all *phase-level* values.
- The result is a list of seven items when *id-stx* has a module binding: (list *source-mod source-id nominal-source-mod nominal-source-id source-phase import-phase nominal-export-phase*).
  - *source-mod* is a module path index (see §13.4.2 "Compiled Modules and References") that indicates the defining module.
  - *source-id* is a symbol for the identifier's name at its definition site in the source module. This can be different from the local name returned by syntax->datum for several reasons: the identifier is renamed on import, it is renamed on export, or it is implicitly renamed because the identifier (or its import) was generated by a macro invocation.
  - *nominal-source-mod* is a module path index (see §13.4.2 "Compiled Modules and References") that indicates the module `required` into the context of *id-stx* to provide its binding. It can be different from *source-mod* due to a re-export in *nominal-source-mod* of some imported identifier. If the same binding is imported in multiple ways, an arbitrary representative is chosen.
  - *nominal-source-id* is a symbol for the identifier's name as exported by *nominal-source-mod*. It can be different from *source-id* due to a renaming `provide`, even if *source-mod* and *nominal-source-mod* are the same.
  - *source-phase* is 1 if the source definition is for-syntax, 0 otherwise.

  - *import-phase* is 0 if the binding import of *nominal-source-mode* is a plain require, 1 if it is from a for-syntax import, etc.
  - *nominal-export-phase* is the phase level of the export from *nominal-source-mod*.

- The result is #f if *id-stx* has a top-level binding (or, equivalently, if it is unbound).

If *id-stx* is bound to a rename-transformer, the result from identifier-binding is for the identifier in the transformer, so that identifier-binding is consistent with free-identifier=?.

```
(identifier-transformer-binding id-stx)
    (or/c 'lexical
          #f
          (listof module-path-index?
                  symbol?
 →                module-path-index?
                  symbol?
                  (or/c 0 1)
                  (or/c exact-integer? #f)
                  (or/c exact-integer? #f)))
  id-stx : syntax?
```

Same as (identifier-binding *id-stx* (add1 (syntax-local-phase-level))).

```
(identifier-template-binding id-stx)
    (or/c 'lexical
          #f
          (listof module-path-index?
                  symbol?
 →                module-path-index?
                  symbol?
                  (or/c 0 1)
                  (or/c exact-integer? #f)
                  (or/c exact-integer? #f)))
  id-stx : syntax?
```

Same as (identifier-binding *id-stx* (sub1 (syntax-local-phase-level))).

```
(identifier-label-binding id-stx)
```

```
   (or/c 'lexical
         #f
         (listof module-path-index?
                 symbol?
→                module-path-index?
                 symbol?
                 (or/c 0 1)
                 (or/c exact-integer? #f)
                 (or/c exact-integer? #f)))
 id-stx : syntax?
```

Same as `(identifier-binding id-stx #f)`.

11.4 Syntax Transformers

```
(set!-transformer? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a value created by `make-set!-transformer` or an instance of a structure type with the `prop:set!-transformer` property, `#f` otherwise.

```
(make-set!-transformer proc) → set!-transformer?
  proc : (syntax? . -> . syntax?)
```

Creates an assignment transformer that cooperates with `set!`. If the result of `make-set!-transformer` is bound to `id` as a transformer binding, then `proc` is applied as a transformer when `id` is used in an expression position, or when it is used as the target of a `set!` assignment as `(set! id expr)`. When the identifier appears as a `set!` target, the entire `set!` expression is provided to the transformer.

Example:

```
> (let ([x 1]
        [y 2])
    (let-syntax ([x (make-set!-transformer
                     (lambda (stx)
                       (syntax-case stx (set!)
                         ; Redirect mutation of x to y
                         [(set! id v) #'(set! y v)]
                         ; Normal use of x really gets x
                         [id (identifier? #'id)  #'x])))])
      (begin
        (set! x 3)
        (list x y))))
'(1 3)
```

```
(set!-transformer-procedure transformer)
 → (syntax? . -> . syntax?)
  transformer : set!-transformer?
```

Returns the procedure that was passed to `make-set!-transformer` to create *transformer* or that is identified by the `prop:set!-transformer` property of *transformer*.

```
prop:set!-transformer : struct-type-property?
```

A structure type property to identify structure types that act as assignment transformers like the ones created by `make-set!-transformer`.

The property value must be an exact integer or procedure of one or two arguments. In the former case, the integer designates a field within the structure that should contain a procedure; the integer must be between `0` (inclusive) and the number of non-automatic fields in the structure type (exclusive, not counting supertype fields), and the designated field must also be specified as immutable.

If the property value is a procedure of one argument, then the procedure serves as a syntax transformer and for `set!` transformations. If the property value is a procedure of two arguments, then the first argument is the structure whose type has `prop:set!-transformer` property, and the second argument is a syntax object as for a syntax transformer and for `set!` transformations; `set!-transformer-procedure` applied to the structure produces a new function that accepts just the syntax object and calls the procedure associated through the property. Finally, if the property value is an integer, the target identifier is extracted from the structure instance; if the field value is not a procedure of one argument, then a procedure that always calls `raise-syntax-error` is used, instead.

If a value has both the `prop:set!-transformer` and `prop:rename-transformer` properties, then the latter takes precedence. If a structure type has the `prop:set!-transformer` and `prop:procedure` properties, then the former takes precedence for the purposes of macro expansion.

```
(rename-transformer? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a value created by `make-rename-transformer` or an instance of a structure type with the `prop:rename-transformer` property, `#f` otherwise.

```
(make-rename-transformer id-stx
                         [delta-introduce]) → rename-transformer?
  id-stx : syntax?
  delta-introduce : (identifier? . -> . identifier?)
                  = (lambda (id) id)
```

Creates a rename transformer that, when used as a transformer binding, acts as a transformer that inserts the identifier *id-stx* in place of whatever identifier binds the transformer, including in non-application positions, in `set!` expressions.

Such a transformer could be written manually, but the one created by `make-rename-transformer` triggers special cooperation with the parser and other syntactic forms when *id* is bound to the rename transformer:

- The parser installs a `free-identifier=?` and `identifier-binding` equivalence between *id* and *id-stx*, as long as *id-stx* does not have a true value for the `'not-free-identifier=?` syntax property.
- A `provide` of *id* provides the binding indicated by *id-stx* instead of *id*, as long as *id-stx* does not have a true value for the `'not-free-identifier=?` syntax property and as long as *id-stx* has a binding.
- If `provide` exports *id*, it uses a symbol-valued `'nominal-id` property of *id-stx* to specify the "nominal source identifier" of the binding as reported by `identifier-binding`.
- If *id-stx* has a true value for the `'not-provide-all-defined` syntax property, then *id* (or its target) is not exported by `all-defined-out`.
- The `syntax-local-value` and `syntax-local-make-delta-introducer` functions recognize rename-transformer bindings and consult their targets.

```
(rename-transformer-target transformer) → identifier?
  transformer : rename-transformer?
```

Returns the identifier passed to `make-rename-transformer` to create *transformer* or as indicated by a `prop:rename-transformer` property on *transformer*.

```
prop:rename-transformer : struct-type-property?
```

A structure type property to identify structure types that act as rename transformers like the ones created by `make-rename-transformer`.

The property value must be an exact integer or an identifier syntax object. In the former case, the integer designates a field within the structure that should contain an identifier; the integer must be between `0` (inclusive) and the number of non-automatic fields in the structure type (exclusive, not counting supertype fields), and the designated field must also be specified as immutable.

If the property value is an identifier, the identifier serves as the target for renaming, just like the first argument to `make-rename-transformer`. If the property value is an integer, the target identifier is extracted from the structure instance; if the field value is not an identifier, then an identifier `?` with an empty context is used, instead.

```
(local-expand  stx
               context-v
               stop-ids
              [intdef-ctx]) → syntax?
  stx : syntax?
  context-v : (or/c 'expression 'top-level 'module 'module-begin list?)
  stop-ids : (or/c (listof identifier?) #f)
  intdef-ctx :  (or/c internal-definition-context?
                      (and/c pair?
                             (listof internal-definition-context?))
                      #f)
             = #f
```

Expands *stx* in the lexical context of the expression currently being expanded. The *context-v* argument is used as the result of `syntax-local-context` for immediate expansions; a list indicates an internal-definition context, and more information on the form of the list is below.

When an identifier in *stop-ids* is encountered by the expander in a sub-expression, expansions stops for the sub-expression. If *stop-ids* is a non-empty list and does not contain just `module*`, then `begin`, `quote`, `set!`, `lambda`, `case-lambda`, `let-values`, `letrec-values`, `if`, `begin0`, `with-continuation-mark`, `letrec-syntaxes+values`, `#%app`, `#%expression`, `#%top`, and `#%variable-reference` are added to *stop-ids*. If `#%app`, `#%top`, or `#%datum` appears in *stop-ids*, then application, top-level variable reference, and literal data expressions without the respective explicit form are not wrapped with the explicit form. If *stop-ids* is `#f` instead of a list, then *stx* is expanded only as long as the outermost form of *stx* is a macro (i.e., expansion does not proceed to sub-expressions). A fully expanded form can include the bindings listed in §1.2.3.1 "Fully Expanded Programs" plus the `letrec-syntaxes+values` form and `#%expression` in any expression position.

When `#%plain-module-begin` is not itself in *stop-ids* and `module*` is in *stop-ids*, then the `#%plain-module-begin` transformer refrains from expanding `module*` sub-forms. Otherwise, the `#%plain-module-begin` transformer detects and expands sub-forms (such as `define-values`) independent of the correspond identifier's presence in *stop-ids*.

The optional *intdef-ctx* argument must be either `#f`, the result of `syntax-local-make-definition-context`, or a list of such results. In the latter two cases, lexical information for internal definitions is added to *stx* before it is expanded (in reverse order relative to the list). The lexical information is also added to the expansion result (because the expansion might introduce bindings or references to internal-definition bindings).

For a particular internal-definition context, generate a unique value and put it into a list for *context-v*. To allow liberal expansion of `define` forms, the generated value should be an instance of a structure with a true value for `prop:liberal-define-context`. If the internal-definition context is meant to be self-contained, the list for *context-v* should

contain only the generated value; if the internal-definition context is meant to splice into an immediately enclosing context, then when `syntax-local-context` produces a list, `cons` the generated value onto that list.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see `syntax-transforming?`), otherwise the `exn:fail:contract` exception is raised.

Examples:

```
> (define-syntax-rule (do-print x ...)
    (printf x ...))

> (define-syntax-rule (hello x)
    (do-print "hello ~a" x))

> (define-syntax (show stx)
    (syntax-case stx ()
      [(_ x)
       (let ([partly (local-expand #'(hello x)
                                   'expression
                                   (list #'do-print))]
             [fully (local-expand #'(hello x)
                                  'expression
                                  #f)])
         (printf "partly expanded: ~s\n" (syntax->datum partly))
         (printf "fully expanded: ~s\n" (syntax->datum fully))
         fully)]))

> (show 1)
partly expanded: (do-print "hello ~a" 1)
fully expanded: (printf "hello ~a" 1)
hello 1
```

```
(syntax-local-expand-expression stx) → syntax? syntax?
  stx : syntax?
```

Like `local-expand` given `'expression` and an empty stop list, but with two results: a syntax object for the fully expanded expression, and a syntax object whose content is opaque. The latter can be used in place of the former (perhaps in a larger expression produced by a macro transformer), and when the macro expander encounters the opaque object, it substitutes the fully expanded expression without re-expanding it; the `exn:fail:syntax` exception is raised if the expansion context includes bindings or marks that were not present for the original expansion, in which case re-expansion might produce different results. Consistent use of `syntax-local-expand-expression` and the opaque object thus avoids quadratic expansion times when local expansions are nested.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(local-transformer-expand  stx
                           context-v
                           stop-ids
                          [intdef-ctx]) → syntax?
  stx : syntax?
  context-v : (or/c 'expression 'top-level 'module 'module-begin list?)
  stop-ids : (or/c (listof identifier?) #f)
  intdef-ctx : (or/c internal-definition-context? #f) = #f
```

Like local-expand, but *stx* is expanded as a transformer expression instead of a run-time expression, and any lifted expressions—from calls to syntax-local-lift-expression during the expansion of *stx*—are captured into a let-values form in the result.

```
(local-expand/capture-lifts  stx
                             context-v
                             stop-ids
                            [intdef-ctx
                             lift-ctx]) → syntax?
  stx : syntax?
  context-v : (or/c 'expression 'top-level 'module 'module-begin list?)
  stop-ids : (or/c (listof identifier?) #f)
  intdef-ctx : (or/c internal-definition-context? #f) = #f
  lift-ctx : any/c = (gensym 'lifts)
```

Like local-expand, but the result is a syntax object that represents a begin expression. Lifted expressions—from calls to syntax-local-lift-expression during the expansion of *stx*—appear with their identifiers in define-values forms, and the expansion of *stx* is the last expression in the begin. The *lift-ctx* value is reported by syntax-local-lift-context during local expansion. The lifted expressions are not expanded, but instead left as provided in the begin form.

```
(local-transformer-expand/capture-lifts  stx
                                         context-v
                                         stop-ids
                                        [intdef-ctx
                                         lift-ctx]) → syntax?
  stx : syntax?
  context-v : (or/c 'expression 'top-level 'module 'module-begin list?)
  stop-ids : (or/c (listof identifier?) #f)
  intdef-ctx : (or/c internal-definition-context? #f) = #f
  lift-ctx : any/c = (gensym 'lifts)
```

Like local-expand/capture-lifts, but *stx* is expanded as a transformer expression instead of a run-time expression. Lifted expressions are reported as define-values forms (in the transformer environment).

```
(internal-definition-context? v) → boolean?
  v : any/c
```

Returns #t if *v* is an internal-definition context, #f otherwise.

```
(syntax-local-make-definition-context [intdef-ctx])
 → internal-definition-context?
  intdef-ctx : (or/c internal-definition-context? #f) = #f
```

Creates an opaque internal-definition context value to be used with local-expand and other functions. A transformer should create one context for each set of internal definitions to be expanded, and use it when expanding any form whose lexical context should include the definitions. After discovering an internal define-values or define-syntaxes form, use syntax-local-bind-syntaxes to add bindings to the context. Finally, the transformer must call internal-definition-context-seal after all bindings have been added; if an unsealed internal-definition context is detected in a fully expanded expression, the exn:fail:contract exception is raised.

If *intdef-ctx* is not #f, then the new internal-definition context extends the given one. That is, expanding in the new internal-definition context can use bindings previously introduced into *intdef-ctx*.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(syntax-local-bind-syntaxes id-list
                            expr
                            intdef-ctx) → void?
  id-list : (listof identifier?)
  expr : (or/c syntax? #f)
  intdef-ctx : internal-definition-context?
```

Binds each identifier in *id-list* within the internal-definition context represented by *intdef-ctx*, where *intdef-ctx* is the result of syntax-local-make-definition-context. Supply #f for *expr* when the identifiers correspond to define-values bindings, and supply a compile-time expression when the identifiers correspond to define-syntaxes bindings; in the latter case, the number of values produced by the expression should match the number of identifiers, otherwise the exn:fail:contract:arity exception is raised.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(internal-definition-context-seal intdef-ctx) → void?
  intdef-ctx : internal-definition-context?
```

Indicates that no further bindings will be added to *intdef-ctx*, which must not be sealed already. See also syntax-local-make-definition-context.

```
(identifier-remove-from-definition-context id-stx
                                           intdef-ctx)
 → identifier?
  id-stx : identifier?
  intdef-ctx : (or/c internal-definition-context?
                     (listof internal-definition-context?))
```

Removes *intdef-ctx* (or each identifier in the list) from the lexical information of *id-stx*. This operation is useful for correlating an identifier that is bound in an internal-definition context with its binding before the internal-definition context was created.

If simply removing the contexts produces a different binding than completely ignoring the contexts (due to nested internal definition contexts, for example), then the resulting identifier is given a syntax mark to simulate a non-existent lexical context. The *intdef-ctx* argument can be a list because removing internal-definition contexts one at a time can produce a different intermediate binding than removing them all at once.

```
(syntax-local-value  id-stx
                    [failure-thunk
                     intdef-ctx])  → any
  id-stx : syntax?
  failure-thunk : (or/c (-> any) #f) = #f
  intdef-ctx : (or/c internal-definition-context?
                     #f)                          = #f
```

Returns the transformer binding value of *id-stx* in either the context associated with *intdef-ctx* (if not #f) or the context of the expression being expanded (if *intdef-ctx* is #f). If *intdef-ctx* is provided, it must be an extension of the context of the expression being expanded.

If *id-stx* is bound to a rename transformer created with make-rename-transformer, syntax-local-value effectively calls itself with the target of the rename and returns that result, instead of the rename transformer.

If *id-stx* has no transformer binding (via define-syntax, let-syntax, etc.) in that environment, the result is obtained by applying *failure-thunk* if not #f. If *failure-thunk* is false, the exn:fail:contract exception is raised.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the

exn:fail:contract exception is raised.

```
(syntax-local-value/immediate  id-stx
                               [failure-thunk
                                intdef-ctx])   → any
  id-stx : syntax?
  failure-thunk : (or/c (-> any) #f) = #f
  intdef-ctx : (or/c internal-definition-context?
                     #f)                       = #f
```

Like syntax-local-value, but the result is normally two values. If *id-stx* is bound to a rename transformer, the results are the rename transformer and the identifier in the transformer. If *id-stx* is not bound to a rename transformer, then the results are the value that syntax-local-value would produce and #f.

Beware that `provide` on an *id* bound to a rename transformer may export the target of the rename instead of *id*. See make-rename-transformer for more information.

If *id-stx* has no transformer binding, then *failure-thunk* is called (and it can return any number of values), or an exception is raised if *failure-thunk* is #f.

```
(syntax-local-lift-expression stx) → identifier?
  stx : syntax?
```

Returns a fresh identifier, and cooperates with the `module`, `letrec-syntaxes+values`, `define-syntaxes`, `begin-for-syntax`, and top-level expanders to bind the generated identifier to the expression *stx*.

A run-time expression within a module is lifted to the module's top level, just before the expression whose expansion requests the lift. Similarly, a run-time expression outside of a module is lifted to a top-level definition. A compile-time expression in a `letrec-syntaxes+values` or `define-syntaxes` binding is lifted to a `let` wrapper around the corresponding right-hand side of the binding. A compile-time expression within `begin-for-syntax` is lifted to a `define` declaration just before the requesting expression within the `begin-for-syntax`.

Other syntactic forms can capture lifts by using local-expand/capture-lifts or local-transformer-expand/capture-lifts.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(syntax-local-lift-values-expression n stx)
 → (listof identifier?)
  n : exact-nonnegative-integer?
  stx : syntax?
```

Like syntax-local-lift-expression, but binds the result to *n* identifiers, and returns a list of the *n* identifiers.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(syntax-local-lift-context) → any/c
```

Returns a value that represents the target for expressions lifted via syntax-local-lift-expression. That is, for different transformer calls for which this procedure returns the same value (as determined by eq?), lifted expressions for the two transformer are moved to the same place. Thus, the result is useful for caching lift information to avoid redundant lifts.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(syntax-local-lift-module-end-declaration stx) → void?
  stx : syntax?
```

Cooperates with the module form to insert *stx* as a top-level declaration at the end of the module currently being expanded. If the current expression being transformed is in phase level 0 and not in the module top-level, then *stx* is eventually expanded in an expression context. If the current expression being transformed is in a higher phase level (i.e., nested within some number of begin-for-syntaxes within a module top-level), then the lifted declaration is placed at the very end of the module (under a suitable number of begin-for-syntaxes), instead of merely the end of the enclosing begin-for-syntax.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

If the current expression being transformed is not within a module form (see syntax-transforming-module-expression?), then the exn:fail:contract exception is raised.

```
(syntax-local-lift-require raw-require-spec
                           stx)             → syntax?
  raw-require-spec : any/c
  stx : syntax?
```

Lifts a #%require form corresponding to *raw-require-spec* (either as a syntax object or datum) to the top-level or to the top of the module currently being expanded or to an enclosing begin-for-syntax..

The resulting syntax object is the same as *stx*, except that a fresh syntax mark is added. The same syntax mark is added to the lifted #%require form, so that the #%require form can bind uses of imported identifiers in the resulting syntax object (assuming that the lexical information of *stx* includes the binding environment into which the #%require is lifted).

If *raw-require-spec* and *stx* are part of the input to a transformer, then typically syntax-local-introduce should be applied to each before passing them to syntax-local-lift-require, and then syntax-local-introduce should be applied to the result of syntax-local-lift-require. Otherwise, marks added by the macro expander can prevent access to the new imports.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(syntax-local-lift-provide raw-provide-spec-stx) → void?
  raw-provide-spec-stx : syntax?
```

Lifts a #%provide form corresponding to *raw-provide-spec-stx* to the top of the module currently being expanded or to an enclosing begin-for-syntax.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

If the current expression being transformed is not within a module form (see syntax-transforming-module-expression?), then the exn:fail:contract exception is raised.

```
(syntax-local-name) → any/c
```

Returns an inferred name for the expression position being transformed, or #f if no such name is available. A name is normally a symbol or an identifier. See also §1.2.6 “Inferred Value Names”.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(syntax-local-context)
 → (or/c 'expression 'top-level 'module 'module-begin list?)
```

Returns an indication of the context for expansion that triggered a syntax transformer call. See §1.2.3.3 “Expansion Context” for more information on contexts.

The symbol results indicate that the expression is being expanded for an expression context, a top-level context, a module context, or a module-begin context.

A list result indicates expansion in an internal-definition context. The identity of the list’s first element (i.e., its eq?ness) reflects the identity of the internal-definition context; in particular two transformer expansions receive the same first value if and only if they are invoked

for the same internal-definition context. Later values in the list similarly identify internal-definition contexts that are still being expanded, and that required the expansion of nested internal-definition contexts.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(syntax-local-phase-level) → (or/c exact-integer? #f)
```

During the dynamic extent of a syntax transformer application by the expander, the result is the phase level of the form being expanded. Otherwise, the result is 0.

```
(syntax-local-module-exports mod-path)
 → (listof (cons/c (or/c exact-integer? #f) (listof symbol?)))
             (or/c module-path?
                   (and/c syntax?
  mod-path :               (lambda (stx)
                             (module-path? (syntax->datum stx)))))
```

Returns an association list from phase-level numbers (or #f for the label phase level) to lists of symbols, where the symbols are the names of provided bindings from *mod-path* at the corresponding phase level.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(syntax-local-submodules) → (listof symbol?)
```

Returns a list of submodule names that are declared via module (as opposed to module*) in the current expansion context.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(syntax-local-get-shadower id-stx) → identifier?
  id-stx : identifier?
```

Returns *id-stx* if no binding in the current expansion context shadows *id-stx* (ignoring unsealed internal-definition contexts and identifiers that had the 'unshadowable syntax property), if *id-stx* has no module bindings in its lexical information, and if the current expansion context is not a module context.

If a binding of inner-identifier shadows *id-stx*, the result is the same as (syntax-local-get-shadower inner-identifier), except that it has the location and properties

of *id-stx*. When searching for a shadowing binding, bindings from unsealed internal-definition contexts are ignored.

Otherwise, the result is the same as *id-stx* with its module bindings (if any) removed from its lexical information, and the lexical information of the current module context (if any) added.

Thus, the result is an identifier corresponding to the innermost shadowing of *id-stx* in the current context if it is shadowed, and a module-contextless version of *id-stx* otherwise.

If *id-stx* is tainted or armed, then the resulting identifier is tainted.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(syntax-local-certifier [active?])
  ((syntax?) (any/c (or/c procedure? #f))
→   . ->* . syntax?)
 active? : boolean? = #f
```

For backward compatibility only; returns a procedure that returns its first argument.

```
(syntax-transforming?) → boolean?
```

Returns #t during the dynamic extent of a syntax transformer application by the expander and while a module is being visited, #f otherwise.

```
(syntax-transforming-module-expression?) → boolean?
```

Returns #t during the dynamic extent of a syntax transformer application by the expander for an expression within a module form, #f otherwise.

```
(syntax-local-introduce stx) → syntax?
  stx : syntax?
```

Produces a syntax object that is like *stx*, except that a syntax mark for the current expansion is added (possibly canceling an existing mark in parts of *stx*). See §1.2.3.5 “Transformer Bindings” for information on syntax marks.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see syntax-transforming?), otherwise the exn:fail:contract exception is raised.

```
(make-syntax-introducer) → (syntax? . -> . syntax?)
```

Produces a procedure that behaves like `syntax-local-introduce`, but using a fresh syntax mark. Multiple applications of the same `make-syntax-introducer` result procedure use the same mark, and different result procedures use distinct marks.

```
(make-syntax-delta-introducer  ext-stx
                               base-stx
                              [phase-level])
 → (syntax? . -> . syntax?)
  ext-stx : syntax?
  base-stx : (or/c syntax? #f)
  phase-level : (or/c #f exact-integer?)
              = (syntax-local-phase-level)
```

Produces a procedure that behaves like `syntax-local-introduce`, but using the syntax marks of *ext-stx* that are not shared with *base-stx*. If *ext-stx* does not extend the set of marks in *base-stx* or if *base-stx* is `#f`, and if *ext-stx* has a module binding in the phase level indicated by *phase-level*, then any marks of *ext-stx* that would be needed to preserve its binding are not transferred in an introduction.

This procedure is potentially useful when *m-id* has a transformer binding that records some *orig-id*, and a use of *m-id* introduces a binding of *orig-id*. In that case, the syntax marks in the use of *m-id* since the binding of *m-id* should be transferred to the binding instance of *orig-id*, so that it captures uses with the same lexical context as the use of *m-id*.

More typically, however, `syntax-local-make-delta-introducer` should be used, since it cooperates with rename transformers.

If *ext-stx* is tainted or armed, then an identifier result from the created procedure is tainted.

```
(syntax-local-make-delta-introducer id)
 → (identifier? . -> . identifier?)
  id : identifier?
```

Determines the binding of *id*. If the binding is not a rename transformer, the result is an introducer as created by `make-syntax-delta-introducer` using *id* and the binding of *id* in the environment of expansion. If the binding is a rename transformer, then the introducer is one composed with the target of the rename transformer and its binding. Furthermore, the *delta-introduce* functions associated with the rename transformers (supplied as the second argument to `make-rename-transformer`) are composed (in first-to-last order) before the introducers created with `make-syntax-delta-introducer` (which are composed last-to-first).

The `exn:fail:contract` exception is raised if *id* or any identifier in its rename-transformer chain has no binding.

This procedure must be called during the dynamic extent of a syntax transformer application by the expander or while a module is visited (see `syntax-transforming?`), otherwise the `exn:fail:contract` exception is raised.

```
(syntax-local-transforming-module-provides?) → boolean?
```

Returns `#t` while a provide transformer is running (see `make-provide-transformer`) or while an `expand` sub-form of `#%provide` is expanded, `#f` otherwise.

```
(syntax-local-module-defined-identifiers)
 → (and/c hash? immutable?)
```

Can be called only while `syntax-local-transforming-module-provides?` returns `#t`.

It returns a hash table mapping a phase-level number (such as `0`) to a list of all definitions at that phase level within the module being expanded. This information is used for implementing `provide` sub-forms like `all-defined-out`.

Beware that the phase-level keys are absolute relative to the enclosing module, and not relative to the current transformer phase level as reported by `syntax-local-phase-level`.

```
(syntax-local-module-required-identifiers mod-path
                                          phase-level)
  (listof (cons/c (or/c exact-integer? #f)
→                 (listof identifier?)))
 mod-path : (or/c module-path? #f)
 phase-level : (or/c exact-integer? #f #t)
```

Can be called only while `syntax-local-transforming-module-provides?` returns `#t`.

It returns an association list mapping phase levels to lists of identifiers. Each list of identifiers includes all bindings imported (into the module being expanded) using the module path *`mod-path`*, or all modules if *`mod-path`* is `#f`. The association list includes all identifiers imported with a *`phase-level`* shift, of all shifts if *`phase-level`* is `#t`.

When an identifier is renamed on import, the result association list includes the identifier by its internal name. Use `identifier-binding` to obtain more information about the identifier.

Beware that the phase-level keys are absolute relative to the enclosing module, and not relative to the current transformer phase level as reported by `syntax-local-phase-level`.

```
prop:liberal-define-context : struct-type-property?
(liberal-define-context? v) → boolean?
  v : any/c
```

An instance of a structure type with a true value for the prop:liberal-define-context property can be used as an element of an internal-definition context representation in the result of syntax-local-context or the second argument of local-expand. Such a value indicates that the context supports *liberal expansion* of define forms into potentially multiple define-values and define-syntaxes forms. The 'module and 'module-body contexts implicitly allow liberal expansion.

The liberal-define-context? predicate returns #t if v is an instance of a structure with a true value for the prop:liberal-define-context property, #f otherwise.

11.4.1 require Transformers

```
(require racket/require-transform)
```

The bindings documented in this section are provided by the racket/require-transform library, not racket/base or racket.

A transformer binding whose value is a structure with the prop:require-transformer property implements a derived *require-spec* for require as a *require transformer*.

A require transformer is called with the syntax object representing its use as a *require-spec* within a require form, and the result must be two lists: a list of imports and a list of import-sources.

If the derived form contains a sub-form that is a *require-spec*, then it can call expand-import to transform the sub-*require-spec* to lists of imports and import sources.

See also define-require-syntax, which supports macro-style require transformers.

```
(expand-import stx) → (listof import?) (listof import-source?)
  stx : syntax?
```

Expands the given *require-spec* to lists of imports and import sources. The latter specifies modules to be instantiated or visited, so the modules that it represents should be a superset of the modules represented in the former list (so that a module will be instantiated or visited even if all of imports are eventually filtered from the former list).

```
(make-require-transformer proc) → require-transformer?
        (syntax? . -> . (values
  proc :                 (listof import?)
                         (listof import-source?)))
```

Creates a require transformer using the given procedure as the transformer.

```
prop:require-transformer : struct-type-property?
```

A property to identify require transformers. The property value must be a procedure that takes the structure and returns a transformer procedure; the returned transformer procedure takes a syntax object and returns import and import-source lists.

```
(require-transformer? v) → boolean?
  v : any/c
```

Returns `#t` if `v` has the `prop:require-transformer` property, `#f` otherwise.

```
(struct import (local-id
                src-sym
                src-mod-path
                mode
                req-mode
                orig-mode
                orig-stx)
        #:extra-constructor-name make-import)
  local-id : identifier?
  src-sym : symbol?
  src-mod-path : module-path?
  mode : (or/c exact-integer? #f)
  req-mode : (or/c exact-integer? #f)
  orig-mode : (or/c exact-integer? #f)
  orig-stx : syntax?
```

A structure representing a single imported identifier:

- `local-id` — the identifier to be bound within the importing module.
- `src-sym` — the external name of the binding as exported from its source module.
- `src-mod-path` — a module path (relative to the importing module) for the source of the imported binding.
- `orig-stx` — a syntax object for the source of the import, used for error reporting.
- `mode` — the phase level of the binding in the importing module.
- `req-mode` — the phase level shift of the import relative to the exporting module.
- `orig-mode` — the phase level of the binding as exported by the exporting module.

```
(struct import-source (mod-path-stx mode)
        #:extra-constructor-name make-import-source)
                (and/c syntax?
  mod-path-stx :        (lambda (x)
                          (module-path? (syntax->datum x))))
  mode : (or/c exact-integer? #f)
```

A structure representing an imported module, which must be instantiated or visited even if no binding is imported into a module.

- `mod-path-stx` — a module path (relative to the importing module) for the source of the imported binding.
- `mode` — the phase level shift of the import.

```
(current-require-module-path) → (or/c #f module-path-index?)
(current-require-module-path module-path) → void?
  module-path : (or/c #f module-path-index?)
```

A parameter that determines how relative `require`-level module paths are expanded to `#%require`-level module paths by `convert-relative-module-path` (which is used implicitly by all built-in `require` sub-forms).

When the value of `current-require-module-path` is `#f`, relative module paths are left as-is, which means that the `require` context determines the resolution of the module path.

The `require` form parameterizes `current-require-module-path` as `#f` while invoking sub-form transformers, while `relative-in` parameterizes to a given module path.

```
(convert-relative-module-path module-path)
    (or/c module-path?
          (and/c syntax?
                 (lambda (stx)
                   (module-path? (syntax-e stx)))))
→                (or/c module-path?
                       (and/c syntax?
                              (lambda (stx)
                                (module-path? (syntax-e stx)))))
  module-path :
```

Converts *module-path* according to `current-require-module-path`.

If *module-path* is not relative or if the value of `current-require-module-path` is `#f`, then *module-path* is returned. Otherwise, *module-path* is converted to an absolute module path that is equivalent to *module-path* relative to the value of `current-require-module-path`.

```
(syntax-local-require-certifier)
    ((syntax?) (or/c #f (syntax? . -> . syntax?))
→     . ->* . syntax?)
```

For backward compatibility only; returns a procedure that returns its first argument.

11.4.2 `provide` Transformers

```
(require racket/provide-transform)
```

The bindings documented in this section are provided by the `racket/provide-transform` library, not `racket/base` or `racket`.

A transformer binding whose value is a structure with the `prop:provide-transformer` property implements a derived *`provide-spec`* for `provide` as a *provide transformer*. A provide transformer is applied as part of the last phase of a module's expansion, after all other declarations and expressions within the module are expanded.

A transformer binding whose value is a structure with the `prop:provide-pre-transformer` property implements a derived *`provide-spec`* for `provide` as a *provide pre-transformer*. A provide pre-transformer is applied as part of the first phase of a module's expansion. Since it is used in the first phase, a provide pre-transformer can use functions such as `syntax-local-lift-expression` to introduce expressions and definitions in the enclosing module.

An identifier can have a transformer binding to a value that acts both as a provide transformer and provide pre-transformer. The result of a provide pre-transformer is *not* automatically re-expanded, so a provide pre-transformer can usefully expand to itself in that case.

A transformer is called with the syntax object representing its use as a *`provide-spec`* within a `provide` form and a list of symbols representing the export modes specified by enclosing *`provide-spec`*s. The result of a provide transformer must be a list of `export`s, while the result of a provide pre-transformer is a syntax object to be used as a *`provide-spec`* in the last phase of module expansion.

If a derived form contains a sub-form that is a *`provide-spec`*, then it can call `expand-export` or `pre-expand-export` to transform the sub-*`provide-spec`* sub-form.

See also `define-provide-syntax`, which supports macro-style provide transformers.

```
(expand-export stx modes) → (listof export?)
  stx : syntax?
  modes : (listof (or/c exact-integer? #f))
```

Expands the given *`provide-spec`* to a list of exports. The *`modes`* list controls the expansion of sub-*`provide-specs`*; for example, an identifier refers to a binding in the phase level of the enclosing `provide` form, unless the *`modes`* list specifies otherwise. Normally, *`modes`* is either empty or contains a single element.

```
(pre-expand-export stx modes) → syntax?
  stx : syntax?
  modes : (listof (or/c exact-integer? #f))
```

Expands the given *provide-spec* at the level of provide pre-transformers. The *modes* argument is the same as for `expand-export`.

```
(make-provide-transformer proc) → provide-transformer?
  proc : (syntax? (listof (or/c exact-integer? #f))
          . -> . (listof export?))
(make-provide-transformer proc pre-proc)
 → (and/c provide-transformer? provide-pre-transformer?)
  proc : (syntax? (listof (or/c exact-integer? #f))
          . -> . (listof export?))
  pre-proc : (syntax? (listof (or/c exact-integer? #f))
              . -> . syntax?)
```

Creates a provide transformer (i.e., a structure with the `prop:provide-transformer` property) using the given procedure as the transformer. If a *pre-proc* is provided, then the result is also a provide pre-transformer.

```
(make-provide-pre-transformer pre-proc)
 → provide-pre-transformer?
  pre-proc : (syntax? (listof (or/c exact-integer? #f))
              . -> . syntax?)
```

Like `make-provide-transformer`, but for a value that is a provide pre-transformer, only.

```
prop:provide-transformer : struct-type-property?
```

A property to identify provide transformers. The property value must be a procedure that takes the structure and returns a transformer procedure; the returned transformer procedure takes a syntax object and mode list and returns an export list.

```
prop:provide-pre-transformer : struct-type-property?
```

A property to identify provide pre-transformers. The property value must be a procedure that takes the structure and returns a transformer procedure; the returned transformer procedure takes a syntax object and mode list and returns a syntax object.

```
(provide-transformer? v) → boolean?
  v : any/c
```

Returns `#t` if *v* has the `prop:provide-transformer` property, `#f` otherwise.

```
(provide-pre-transformer? v) → boolean?
  v : any/c
```

Returns `#t` if *v* has the `prop:provide-pre-transformer` property, `#f` otherwise.

```
(struct export (local-id out-sym mode protect? orig-stx)
        #:extra-constructor-name make-export)
  local-id : identifier?
  out-sym : symbol?
  mode : (or/c exact-integer? #f)
  protect? : any/c
  orig-stx : syntax?
```

A structure representing a single imported identifier:

- `local-id` — the identifier that is bound within the exporting module.
- `out-sym` — the external name of the binding.
- `orig-stx` — a syntax object for the source of the export, used for error reporting.
- `protect?` — indicates whether the identifier should be protected (see §13.10 “Code Inspectors”).
- `mode` — the phase level of the binding in the exporting module.

```
(syntax-local-provide-certifier)
 → ((syntax?) (or/c #f (syntax? . -> . syntax?))
    . ->* . syntax?)
```

For backward compatibility only; returns a procedure that returns its first argument.

11.4.3 Keyword-Argument Conversion Introspection

```
(require racket/keyword-transform)
```

The bindings documented in this section are provided by the `racket/keyword-transform` library, not `racket/base` or `racket`.

```
(syntax-procedure-alias-property stx)
    (or/c #f
          (letrec ([val? (recursive-contract
 →                        (or/c (cons/c identifier? identifier?)
                                (cons/c val? val?)))])
            val?))
  stx : syntax?
```

```
(syntax-procedure-converted-arguments-property stx)
   (or/c #f
         (letrec ([val? (recursive-contract
→                        (or/c (cons/c identifier? identifier?)
                              (cons/c val? val?)))])
           val?))
  stx : syntax?
```

At expansion time, reports the value of a syntax property that can be attached to an identifier by the expansion of a keyword-application form during the same expansion time. See `lambda` for more information about the property.

The property value is normally a pair consisting of the original identifier and an identifier that appears in the expansion. Property-value merging via `syntax-track-origin` can make the value a pair of such values, and so on.

11.5 Syntax Parameters

```
(require racket/stxparam)
```

The bindings documented in this section are provided by the `racket/stxparam` library, not `racket/base` or `racket`.

```
(define-syntax-parameter id expr)
```

Binds *id* as syntax to a *syntax parameter*. The *expr* is an expression in the transformer environment that serves as the default value for the syntax parameter. The value is typically obtained by a transformer using `syntax-parameter-value`.

The *id* can be used with `syntax-parameterize` or `syntax-parameter-value` (in a transformer). If *expr* produces a procedure of one argument or a `make-set!-transformer` result, then *id* can be used as a macro. If *expr* produces a `rename-transformer` result, then *id* can be used as a macro that expands to a use of the target identifier, but `syntax-local-value` of *id* does not produce the target's value.

```
(syntax-parameterize ([id expr] ...) body-expr ...+)
```

See also `splicing-syntax-parameteriz`

Each *id* must be bound to a syntax parameter using `define-syntax-parameter`. Each *expr* is an expression in the transformer environment. During the expansion of the *body-expr*s, the value of each *expr* is bound to the corresponding *id*.

If an *expr* produces a procedure of one argument or a `make-set!-transformer` result, then its *id* can be used as a macro during the expansion of the *body-expr*s. If *expr* produces a `rename-transformer` result, then *id* can be used as a macro that expands to a use of the target identifier, but `syntax-local-value` of *id* does not produce the target's value.

11.5.1 Syntax Parameter Inspection

```
(require racket/stxparam-exptime)
```

```
(syntax-parameter-value id-stx) → any
  id-stx : syntax?
```

This procedure is intended for use in a transformer environment, where *id-stx* is an identifier bound in the normal environment to a syntax parameter. The result is the current value of the syntax parameter, as adjusted by `syntax-parameterize` form.

This binding is provided `for-syntax` by `racket/stxparam`, since it is normally used in a transformer. It is provided normally by `racket/stxparam-exptime`.

```
(make-parameter-rename-transformer id-stx) → any
  id-stx : syntax?
```

This procedure is intended for use in a transformer, where *id-stx* is an identifier bound to a syntax parameter. The result is a transformer that behaves as *id-stx*, but that cannot be used with `syntax-parameterize` or `syntax-parameter-value`.

Using `make-parameter-rename-transformer` is analogous to defining a procedure that calls a parameter. Such a procedure can be exported to others to allow access to the parameter value, but not to change the parameter value. Similarly, `make-parameter-rename-transformer` allows a syntax parameter to be used as a macro, but not changed.

The result of `make-parameter-rename-transformer` is not treated specially by `syntax-local-value`, unlike the result of `make-rename-transformer`.

This binding is provided `for-syntax` by `racket/stxparam`, since it is normally used in a transformer. It is provided normally by `racket/stxparam-exptime`.

11.6 Local Binding with Splicing Body

```
(require racket/splicing)
```

The bindings documented in this section are provided by the `racket/splicing` library, not `racket/base` or `racket`.

```
splicing-let
splicing-letrec
splicing-let-values
splicing-letrec-values
splicing-let-syntax
```

```
splicing-letrec-syntax
splicing-let-syntaxes
splicing-letrec-syntaxes
splicing-letrec-syntaxes+values
splicing-local
```

Like `let`, `letrec`, `let-values`, `letrec-values`, `let-syntax`, `letrec-syntax`, `let-syntaxes`, `letrec-syntaxes`, `letrec-syntaxes+values`, and `local`, except that in a definition context, the body forms are spliced into the enclosing definition context (in the same way as for `begin`).

Examples:

```
> (splicing-let-syntax ([one (lambda (stx) #'1)])
    (define o one))

> o
1
> one
one: undefined;
 cannot reference undefined identifier
```

When a splicing binding form occurs in a top-level context or module context, its local bindings are treated similarly to definitions. In particular, if a reference to one of the splicing form's bound variables is evaluated before the variable is initialized, an unbound variable error is raised, instead of the variable evaluating to the undefined value. Also, syntax bindings are evaluated every time the module is visited, instead of only once during compilation as in `let-syntax`, etc.

Example:

```
> (splicing-letrec ([x bad]
                    [bad 1])
    x)
bad.3: undefined;
 cannot reference undefined identifier
```

```
splicing-syntax-parameterize
```

Like `syntax-parameterize`, except that in a definition context, the body forms are spliced into the enclosing definition context (in the same way as for `begin`). In a definition context, the body of `splicing-syntax-parameterize` can be empty.

Note that require transformers and provide transformers are not affected by syntax parameterization. While all uses of `require` and `provide` will be spliced into the enclosing context, derived import or export specifications will expand as if they had not been inside of the `splicing-syntax-parameterize`.

Examples:

```
> (define-syntax-parameter place (lambda (stx) #'"Kansas"))

> (define-syntax-rule (where) `(at ,(place)))

> (where)
'(at "Kansas")
> (splicing-syntax-parameterize ([place (lambda (stx) #'"Oz")])
    (define here (where)))

> here
'(at "Oz")
```

11.7 Syntax Object Properties

Every syntax object has an associated *syntax property* list, which can be queried or extended with `syntax-property`. Properties are not preserved for a `syntax-quoted` syntax object in a compiled form that is marshaled to a byte string.

In `read-syntax`, the reader attaches a `'paren-shape` property to any pair or vector syntax object generated from parsing a pair `[` and `]` or `{` and `}`; the property value is `#\[` in the former case, and `#\{` in the latter case. The `syntax` form copies any `'paren-shape` property from the source of a template to corresponding generated syntax.

Both the syntax input to a transformer and the syntax result of a transformer may have associated properties. The two sets of properties are merged by the syntax expander: each property in the original and not present in the result is copied to the result, and the values of properties present in both are combined with `cons` (result value first, original value second).

Before performing the merge, however, the syntax expander automatically adds a property to the original syntax object using the key `'origin`. If the source syntax has no `'origin` property, it is set to the empty list. Then, still before the merge, the identifier that triggered the macro expansion (as syntax) is `cons`ed onto the `'origin` property so far. The `'origin` property thus records (in reverse order) the sequence of macro expansions that produced an expanded expression. Usually, the `'origin` value is an immutable list of identifiers. However, a transformer might return syntax that has already been expanded, in which case an `'origin` list can contain other lists after a merge. The `syntax-track-origin` procedure implements this tracking.

Besides `'origin` tracking for general macro expansion, Racket adds properties to expanded syntax (often using `syntax-track-origin`) to record additional expansion details:

- When a `begin` form is spliced into a sequence with internal definitions (see §1.2.3.7 “Internal Definitions”), `syntax-track-origin` is applied to every spliced element

from the `begin` body. The second argument to `syntax-track-origin` is the `begin` form, and the third argument is the `begin` keyword (extracted from the spliced form).

- When an internal `define-values` or `define-syntaxes` form is converted into a `letrec-syntaxes+values` form (see §1.2.3.7 "Internal Definitions"), `syntax-track-origin` is applied to each generated binding clause. The second argument to `syntax-track-origin` is the converted form, and the third argument is the `define-values` or `define-syntaxes` keyword form the converted form.
- When a `letrec-syntaxes+values` expression is fully expanded, syntax bindings disappear, and the result is either a `letrec-values` form (if the unexpanded form contained non-syntax bindings), or only the body of the `letrec-syntaxes+values` form (wrapped with `begin` if the body contained multiple expressions). To record the disappeared syntax bindings, a property is added to the expansion result: an immutable list of identifiers from the disappeared bindings, as a `'disappeared-binding` property.
- When a subtyping `struct` form is expanded, the identifier used to reference the base type does not appear in the expansion. Therefore, the `struct` transformer adds the identifier to the expansion result as a `'disappeared-use` property.
- When a reference to an unexported or protected identifier from a module is discovered, the `'protected` property is added to the identifier with a `#t` value.
- When `read-syntax` generates a syntax object, it attaches a property to the object (using a private key) to mark the object as originating from a read. The `syntax-original?` predicate looks for the property to recognize such syntax objects. (See §11.2 "Syntax Object Content" for more information.)

See §11.9.1 "Information on Expanded Modules" for information about properties generated by the expansion of a module declaration. See `lambda` and §1.2.6 "Inferred Value Names" for information about properties recognized when compiling a procedure. See `current-compile` for information on properties and byte codes.

```
(syntax-property stx key v) → syntax?
  stx : syntax?
  key : any/c
  v : any/c
(syntax-property stx key) → any
  stx : syntax?
  key : any/c
```

The three-argument form extends *stx* by associating an arbitrary property value *v* with the key *key*; the result is a new syntax object with the association (while *stx* itself is unchanged).

The two-argument form returns an arbitrary property value associated to *stx* with the key *key*, or `#f` if no value is associated to *stx* for *key*.

```
(syntax-property-symbol-keys stx) → list?
  stx : syntax?
```

Returns a list of all symbols that as keys have associated properties in *stx*. Uninterned symbols (see §3.6 “Symbols”) are not included in the result list.

```
(syntax-track-origin new-stx
                     orig-stx
                     id-stx)  → any
  new-stx : syntax?
  orig-stx : syntax?
  id-stx : syntax?
```

Adds properties to *new-stx* in the same way that macro expansion adds properties to a transformer result. In particular, it merges the properties of *orig-stx* into *new-stx*, first adding *id-stx* as an 'origin property, and it returns the property-extended syntax object. Use the syntax-track-origin procedure in a macro transformer that discards syntax (corresponding to *orig-stx* with a keyword *id-stx*) leaving some other syntax in its place (corresponding to *new-stx*).

For example, the expression

```
(or x y)
```

expands to

```
(let ([or-part x]) (if or-part or-part (or y)))
```

which, in turn, expands to

```
(let-values ([(or-part) x]) (if or-part or-part y))
```

The syntax object for the final expression will have an 'origin property whose value is (list (quote-syntax let) (quote-syntax or)).

11.8 Syntax Taints

§16.2.6 “Syntax Taints” in ***The Racket Guide*** introduces syntax taints.

The *tamper status* of a syntax object is either tainted, armed, or clean:

- A *tainted* identifier is rejected by the macro expander for use as either a binding or expression. If a syntax object is tainted, then any syntax object in the result of (syntax-e *stx*) is tainted, and datum->syntax with *stx* as its first argument produces a tainted syntax object.

Other derived operations, such as pattern matching in `syntax-case`, also taint syntax objects when extracting them from a tainted syntax object.

- An *armed* syntax object has a set of *dye packs*, which creates taints if the armed syntax object is used without first disarming the dye packs. In particular, if a syntax object is armed, `syntax-e`, `datum->syntax`, `quote-syntax`, and derived operations effectively treat the syntax object as tainted. The macro expander, in contrast, disarms dye packs before pulling apart syntax objects.

  Each dye pack, which is added to a syntax object with the `syntax-arm` function, is keyed by an inspector. A dye pack can be *disarm*ed using `syntax-disarm` with an inspector that is the same as or a superior of the dye pack's inspector.

- A *clean* syntax object has no immediate taints or dye packs, although it may contain syntax objects that are tainted or armed.

Taints cannot be removed, and attempting to arm a syntax object that is already tainted has no effect on the resulting syntax object.

The macro expander disarms any syntax object that it encounters in an expression position or as a module body. A syntax object is therefore disarmed when it is provided to a syntax transformer. The transformer's result, however, is *rearm*ed by copying to it any dye packs that were originally attached to the transformer's input. The rearming process obeys the following rules:

- If the result has a `'taint-mode` property (see §11.7 "Syntax Object Properties") that is `'opaque`, then dye packs are attached to the immediate syntax object.

- If the result has a `'taint-mode` property that is `'none`, then no dye pack is attached to the syntax object. The `'none` mode is rarely appropriate.

- If the result has a `'taint-mode` property that is `'transparent`, then the dye packs are propagated recursively to syntax object that corresponds to elements of the syntax object's datum as a list (or, more precisely, to the `car`s of the datum as reached by any number of `cdr`s), and the immediate syntax object loses its lexical context; If the immediate syntax object is already armed, then recursive propagation taints the elements. Recursive propagation uses syntax properties and shapes, as for the immediate rearming.

- If the result has a `'taint-mode` property that is `'transparent-binding`, then dye packs are attached in a way similar to `'transparent`, but further treating the syntax object corresponding to the second list element as having a `'transparent` value for the `'taint-mode` property if it does not already have a `'taint-mode` property value.

- If the result has no `'taint-mode` property value, but its datum is a pair, and if the syntax object corresponding to the `car` of the pair is an identifier bound to `begin`, `module`, or `#%plain-module-begin`, then dye packs are propagated as if the syntax object had the `'transparent` property value.

- If the result has no `'taint-mode` property value, but its datum is a pair, and if the syntax object corresponding to the `car` of the pair is an identifier bound to `define-values` or `define-syntaxes`, then dye packs are propagated as if the syntax object had the `'transparent-binding` property value.

For backward compatibility, a `'certify-mode` property is treated the same as a `'taint-mode` property if the former is not attached. To avoid accidental transfer of a `'taint-mode` or `'certify-mode` property value, the expander always removes any `'taint-mode` and `'certify-mode` property on a syntax object that is passed to a syntax transformer.

```
(syntax-tainted? stx) → boolean?
  stx : syntax?
```

Returns `#t` if *stx* is tainted, `#f` otherwise.

```
(syntax-arm stx [inspector use-mode?]) → syntax?
  stx : syntax?
  inspector : (or/c inspector? #f) = #f
  use-mode? : any/c = #f
```

Produces a syntax object like *stx*, but armed with a dye pack that is keyed by *inspector*.

A `#f` value for *inspector* is equivalent to an inspector that depends on the current dynamic context:

- when applying a syntax transformer is being applied, the declaration-time code inspector of the module in which a syntax transformer was bound;
- when a module is being visited, the module's declaration-time code inspector;
- `(current-code-inspector)`, otherwise.

If *use-mode?* is `#f`, then if *stx* is tainted or already armed with the key *inspector*, the result is *stx*.

If *use-mode?* is a true value, then a dye pack is not necessarily added directly to *stx*. Instead, the dye pack is pushed to interior syntax objects in the same way that the expander pushes armings into a syntax transformer's results when rearming (based on a `'taint-mode` syntax property or identifier bindings); see the expander's rearming rules for more information. To the degree that pushing dye packs into a syntax object must destructure *stx*, existing taints or dye packs can lead to tainted results rather than armed results.

```
(syntax-protect stx) → syntax?
  stx : syntax?
```

Equivalent to `(syntax-arm stx #f #f #t)`.

```
(syntax-disarm stx inspector) → syntax?
  stx : syntax?
  inspector : (or/c inspector? #f)
```

Produces a disarmed version of *stx*, removing any immediate dye packs that match *inspector*. An inspector matches when it is either the same as or a super-inspector of the dye pack's inspector. A `#f` value for *inspector* is replaced by a specific inspector in the same way as for `syntax-arm`.

```
(syntax-rearm stx from-stx [use-mode?]) → syntax?
  stx : syntax?
  from-stx : syntax?
  use-mode? : any/c = #f
```

Produces a rearmed or tainted version of *stx* by adding all immediate taints and dye packs of *from-stx*.

If *use-mode?* is a true value, *stx* is not necessarily tainted or armed directly. Instead, taints or dye packs are pushed to interior syntax objects in the same way as for `syntax-arm` or rearming by the expander.

```
(syntax-taint stx) → syntax?
  stx : syntax?
```

Returns tainted version of *stx*—equivalent to `(datum->syntax (syntax-arm stx) (syntax-e stx) stx stx)`—or *stx* if it is already tainted.

11.9 Expanding Top-Level Forms

```
(expand top-level-form) → syntax?
  top-level-form : any/c
```

Expands all non-primitive syntax in *top-level-form*, and returns a syntax object for the expanded form that contains only core forms, matching the grammar specified by §1.2.3.1 "Fully Expanded Programs".

Before *top-level-form* is expanded, its lexical context is enriched with `namespace-syntax-introduce`, just as for `eval`. Use `syntax->datum` to convert the returned syntax object into a printable datum.

Here's an example of using `expand` on a module:

```
(parameterize ([current-namespace (make-base-namespace)])
 (expand
  (datum->syntax
   #f
   '(module foo scheme
      (define a 3)
      (+ a 4)))))
```

Here's an example of using `expand` on a non-top-level form:

```
(define-namespace-anchor anchor)
(parameterize ([current-namespace
                (namespace-anchor->namespace anchor)])
 (expand
  (datum->syntax
   #f
   '(delay (+ 1 2)))))
```

(expand-syntax *stx*) → syntax?
  stx : syntax?

Like (expand *stx*), except that the argument must be a syntax object, and its lexical context is not enriched before expansion.

(expand-once *top-level-form*) → syntax?
  top-level-form : any/c

Partially expands *top-level-form* and returns a syntax object for the partially-expanded expression. Due to limitations in the expansion mechanism, some context information may be lost. In particular, calling `expand-once` on the result may produce a result that is different from expansion via `expand`.

Before *top-level-form* is expanded, its lexical context is enriched with `namespace-syntax-introduce`, as for `eval`.

(expand-syntax-once *stx*) → syntax?
  stx : syntax?

Like (expand-once *stx*), except that the argument must be a syntax object, and its lexical context is not enriched before expansion.

(expand-to-top-form *top-level-form*) → syntax?
  top-level-form : any/c

Partially expands *top-level-form* to reveal the outermost syntactic form. This partial expansion is mainly useful for detecting top-level uses of `begin`. Unlike the result of `expand-once`, expanding the result of `expand-to-top-form` with `expand` produces the same result as using `expand` on the original syntax.

Before `stx-or-sexpr` is expanded, its lexical context is enriched with `namespace-syntax-introduce`, as for `eval`.

```
(expand-syntax-to-top-form stx) → syntax?
  stx : syntax?
```

Like `(expand-to-top-form stx)`, except that the argument must be a syntax object, and its lexical context is not enriched before expansion.

11.9.1 Information on Expanded Modules

Information for an expanded `module` declaration is stored in a set of syntax properties (see §11.7 “Syntax Object Properties”) attached to the syntax object:

- `'module-direct-requires` — a list of module path indexes (or symbols) representing the modules explicitly imported into the module.
- `'module-direct-for-syntax-requires` — a list of module path indexes (or symbols) representing the modules explicitly for-syntax imported into the module.
- `'module-direct-for-template-requires` — a list of module path indexes (or symbols) representing the modules explicitly for-template imported into the module.
- `'module-variable-provides` — a list of provided items, where each item is one of the following:
  - `symbol` — represents a locally defined variable that is provided with its defined name.
  - `(cons provided-sym defined-sym)` — represents a locally defined variable that is provided with renaming; the first symbol is the exported name, and the second symbol is the defined name.
  - `(list* module-path-index provided-sym defined-sym)` — represents a re-exported and possibly re-named variable from the specified module; `module-path-index` is either a module path index or symbol (see §13.4.2 “Compiled Modules and References”), indicating the source module for the binding. The *provided-sym* is the external name for the re-export, and *defined-sym* is the originally defined name in the module specified by `module-path-index`.
- `'module-syntax-provides` — like `'module-variable-provides`, but for syntax exports instead of variable exports.

- `'module-indirect-provides` — a list of symbols for variables that are defined in the module but not exported; they may be exported indirectly through macro expansions. Definitions of macro-generated identifiers create uninterned symbols in this list.

11.10 File Inclusion

```
(require racket/include)
```

The bindings documented in this section are provided by the `racket/include` and `racket` libraries, but not `racket/base`.

```
(include path-spec)

path-spec = string
          | (file string)
          | (lib string ...+)
```

Inlines the syntax in the file designated by *path-spec* in place of the `include` expression.

A *path-spec* resembles a subset of the *mod-path* forms for `require`, but it specifies a file whose content need not be a module. That is, *string* refers to a file using a platform-independent relative path, `(file` *string*`)` refers to a file using platform-specific notation, and `(lib` *string* `...)` refers to a file within a collection.

If *path-spec* specifies a relative path, the path is resolved relative to the source for the `include` expression, if that source is a complete path string. If the source is not a complete path string, then *path-spec* is resolved relative to `(current-load-relative-directory)` if it is not `#f`, or relative to `(current-directory)` otherwise.

The included syntax is given the lexical context of the `include` expression, while the included syntax's source location refers to its actual source.

```
(include-at/relative-to context source path-spec)
```

Like `include`, except that the lexical context of *context* is used for the included syntax, and a relative *path-spec* is resolved with respect to the source of *source*. The *context* and *source* elements are otherwise discarded by expansion.

```
(include/reader path-spec reader-expr)
```

Like `include`, except that the procedure produced by the expression *reader-expr* is used to read the included file, instead of `read-syntax`.

The *reader-expr* is evaluated at expansion time in the transformer environment. Since it serves as a replacement for `read-syntax`, the expression's value should be a procedure that

consumes two inputs—a string representing the source and an input port—and produces a syntax object or `eof`. The procedure will be called repeatedly until it produces `eof`.

The syntax objects returned by the procedure should have source location information, but usually no lexical context; any lexical context in the syntax objects will be ignored.

```
(include-at/relative-to/reader context source path-spec reader-expr)
```

Combines `include-at/relative-to` and `include/reader`.

11.11 Syntax Utilities

```
(require racket/syntax)
```

The bindings documented in this section are provided by the `racket/syntax` library, not `racket/base` or `racket`.

11.11.1 Creating formatted identifiers

```
(format-id  lctx
            fmt
            v ...
           [#:source src
            #:props props
            #:cert ignored]) → identifier?
  lctx : (or/c syntax? #f)
  fmt : string?
  v : (or/c string? symbol? identifier? keyword? char? number?)
  src : (or/c syntax? #f) = #f
  props : (or/c syntax? #f) = #f
  ignored : (or/c syntax? #f) = #f
```

Like `format`, but produces an identifier using *lctx* for the lexical context, *src* for the source location, and *props* for the properties. An argument supplied with `#:cert` is ignored. (See `datum->syntax`.)

The format string must use only `~a` placeholders. Identifiers in the argument list are automatically converted to symbols.

Examples:

```
> (define-syntax (make-pred stx)
    (syntax-case stx ()
```

```
      [(make-pred name)
       (format-id #'name "~a?" (syntax-e #'name))]))

> (make-pred pair)
#<procedure:pair?>
> (make-pred none-such)
none-such?: undefined;
 cannot reference undefined identifier
> (define-syntax (better-make-pred stx)
    (syntax-case stx ()
      [(better-make-pred name)
       (format-id #'name #:source #'name
                  "~a?" (syntax-e #'name))]))

> (better-make-pred none-such)
none-such?: undefined;
 cannot reference undefined identifier
```

(Scribble doesn't show it, but the DrRacket pinpoints the location of the second error but not of the first.)

```
(format-symbol fmt v ...) → symbol?
  fmt : string?
  v : (or/c string? symbol? identifier? keyword? char? number?)
```

Like `format`, but produces a symbol. The format string must use only `~a` placeholders. Identifiers in the argument list are automatically converted to symbols.

Example:

```
> (format-symbol "make-~a" 'triple)
'make-triple
```

11.11.2 Pattern variables

```
(define/with-syntax pattern stx-expr)

  stx-expr : syntax?
```

Definition form of `with-syntax`. That is, it matches the syntax object result of `expr` against *pattern* and creates pattern variable definitions for the pattern variables of *pattern*.

Examples:

```
> (define/with-syntax (px ...) #'(a b c))
```

```
> (define/with-syntax (tmp ...) (generate-temporaries #'(px ...)))

> #'([tmp px] ...)
#<syntax:11:0 ((a7 a) (b8 b) (c9 c))>
> (define-pattern-variable name #'Alice)
define-pattern-variable: undefined;
 cannot reference undefined identifier
> #'(hello name)
#<syntax:13:0 (hello name)>
```

11.11.3 Error reporting

```
(current-syntax-context) → (or/c syntax? false/c)
(current-syntax-context stx) → void?
  stx : (or/c syntax? false/c)
```

The current contextual syntax object, defaulting to `#f`. It determines the special form name that prefixes syntax errors created by `wrong-syntax`.

```
(wrong-syntax stx format-string v ...) → any
  stx : syntax?
  format-string : string?
  v : any/c
```

Raises a syntax error using the result of `(current-syntax-context)` as the "major" syntax object and the provided *stx* as the specific syntax object. (The latter, *stx*, is usually the one highlighted by DrRacket.) The error message is constructed using the format string and arguments, and it is prefixed with the special form name as described under `current-syntax-context`.

Examples:

```
> (wrong-syntax #'here "expected ~s" 'there)
?: expected there
> (parameterize ([current-syntax-context #'(look over here)])
    (wrong-syntax #'here "expected ~s" 'there))
eval:15:0: look: expected there
  at: here
  in: (look over here)
```

A macro using `wrong-syntax` might set the syntax context at the very beginning of its transformation as follows:

```
(define-syntax (my-macro stx)
  (parameterize ([current-syntax-context stx])
    (syntax-case stx ()
      __)))
```

Then any calls to `wrong-syntax` during the macro's transformation will refer to `my-macro` (more precisely, the name that referred to `my-macro` where the macro was used, which may be different due to renaming, prefixing, etc).

11.11.4 Recording disappeared uses

```
(current-recorded-disappeared-uses)
 → (or/c (listof identifier?) false/c)
(current-recorded-disappeared-uses ids) → void?
  ids : (or/c (listof identifier?) false/c)
```

Parameter for tracking disappeared uses. Tracking is "enabled" when the parameter has a non-false value. This is done automatically by forms like `with-disappeared-uses`.

```
(with-disappeared-uses stx-expr)

  stx-expr : syntax?
```

Evaluates the *stx-expr*, catching identifiers looked up using `syntax-local-value/record`. Adds the caught identifiers to the `'disappeared-uses` syntax property of the resulting syntax object.

```
(syntax-local-value/record id predicate) → any/c
  id : identifier?
  predicate : (-> any/c boolean?)
```

Looks up *id* in the syntactic environment (as `syntax-local-value`). If the lookup succeeds and returns a value satisfying the predicate, the value is returned and *id* is recorded as a disappeared use. If the lookup fails or if the value does not satisfy the predicate, `#f` is returned and the identifier is not recorded as a disappeared use.

```
(record-disappeared-uses ids) → void?
  ids : (listof identifier?)
```

Add *ids* to `(current-recorded-disappeared-uses)`. If not used within the extent of a `with-disappeared-uses` form or similar, has no effect.

11.11.5 Miscellaneous utilities

```
(generate-temporary [name-base]) → identifier?
  name-base : any/c = 'g
```

Generates one fresh identifier. Singular form of `generate-temporaries`. If *name-base* is supplied, it is used as the basis for the identifier's name.

```
(internal-definition-context-apply intdef-ctx
                                   stx)       → syntax?
  intdef-ctx : internal-definition-context?
  stx : syntax?
```

Applies the renamings of *intdef-ctx* to *stx*.

```
(syntax-local-eval stx [intdef-ctx]) → any
  stx : syntax?
  intdef-ctx : (or/c internal-definition-context? #f) = #f
```

Evaluates *stx* as an expression in the current transformer environment (that is, at phase level 1), optionally extended with *intdef-ctx*.

Examples:

```
> (define-syntax (show-me stx)
    (syntax-case stx ()
      [(show-me expr)
       (begin
         (printf "at compile time produces ~s\n"
                 (syntax-local-eval #'expr))
         #'(printf "at run time produces ~s\n"
                   expr))]))

> (show-me (+ 2 5))
at compile time produces 7
at run time produces 7

> (define-for-syntax fruit 'apple)

> (define fruit 'pear)

> (show-me fruit)
at compile time produces apple
at run time produces pear
```

```
(with-syntax* ([pattern stx-expr] ...)
  body ...+)

  stx-expr : syntax?
```

Similar to `with-syntax`, but the pattern variables of each *pattern* are bound in the *stx-expr*s of subsequent clauses as well as the *body*s, and the *pattern*s need not bind distinct pattern variables; later bindings shadow earlier bindings.

Example:

```
> (with-syntax* ([(x y) (list #'val1 #'val2)]
                 [nest #'((x) (y))])
    #'nest)
#<syntax:21:0 ((val1) (val2))>
```

12 Input and Output

12.1 Ports

§8 “Input and Output” in ***The Racket Guide*** introduces Ports and I/O.

Ports produce and consume bytes. When a port is provided to a character-based operation, the port’s bytes are decoded; see §12.1.1 “Encodings and Locales”.

When a port corresponds to a file, network connection, or some other system resource, it must be explicitly closed via `close-input-port` or `close-output-port` (or indirectly via `custodian-shutdown-all`) to release low-level resources associated with the port. For any kind of port, after it is closed, attempting to read from or write to the port raises `exn:fail`.

The global variable `eof` is bound to the end-of-file value, and `eof-object?` returns `#t` only when applied to this value. Reading from a port produces an end-of-file result when the port has no more data, but some ports may also return end-of-file mid-stream. For example, a port connected to a Unix terminal returns an end-of-file when the user types control-D; if the user provides more input, the port returns additional bytes after the end-of-file.

Every port has a name, as reported by `object-name`. The name can be any value, and it is used mostly for error-reporting purposes. The `read-syntax` procedure uses the name of an input port as the default source location for the syntax objects that it produces.

12.1.1 Encodings and Locales

When a port is provided to a character-based operation, such as `read-char` or `read`, the port’s bytes are read and interpreted as a UTF-8 encoding of characters. Thus, reading a single character may require reading multiple bytes, and a procedure like `char-ready?` may need to peek several bytes into the stream to determine whether a character is available. In the case of a byte stream that does not correspond to a valid UTF-8 encoding, functions such as `read-char` may need to peek one byte ahead in the stream to discover that the stream is not a valid encoding.

When an input port produces a sequence of bytes that is not a valid UTF-8 encoding in a character-reading context, then bytes that constitute an invalid sequence are converted to the character `#\uFFFD`. Specifically, bytes 255 and 254 are always converted to `#\uFFFD`, bytes in the range 192 to 253 produce `#\uFFFD` when they are not followed by bytes that form a valid UTF-8 encoding, and bytes in the range 128 to 191 are converted to `#\uFFFD` when they are not part of a valid encoding that was started by a preceding byte in the range 192 to 253. To put it another way, when reading a sequence of bytes as characters, a minimal set of bytes are changed to the encoding of `#\uFFFD` so that the entire sequence of bytes is a valid UTF-8 encoding.

See §3.4 “Byte Strings” for procedures that facilitate conversions using UTF-8 or other en-

codings. See also reencode-input-port and reencode-output-port for obtaining a UTF-8-based port from one that uses a different encoding of characters.

A *locale* captures information about a user's language-specific interpretation of character sequences. In particular, a locale determines how strings are "alphabetized," how a lowercase character is converted to an uppercase character, and how strings are compared without regard to case. String operations such as string-ci=? are *not* sensitive to the current locale, but operations such as string-locale-ci=? (see §3.3 "Strings") produce results consistent with the current locale.

A locale also designates a particular encoding of code-point sequences into byte sequences. Racket generally ignores this aspect of the locale, with a few notable exceptions: command-line arguments passed to Racket as byte strings are converted to character strings using the locale's encoding; command-line strings passed as byte strings to other processes (through subprocess) are converted to byte strings using the locale's encoding; environment variables are converted to and from strings using the locale's encoding; filesystem paths are converted to and from strings (for display purposes) using the locale's encoding; and, finally, Racket provides functions such as string->bytes/locale to specifically invoke a locale-specific encoding.

A Unix user selects a locale by setting environment variables, such as LC_ALL. On Windows and Mac OS X, the operating system provides other mechanisms for setting the locale. Within Racket, the current locale can be changed by setting the current-locale parameter. The locale name within Racket is a string, and the available locale names depend on the platform and its configuration, but the "" locale means the current user's default locale; on Windows and Mac OS X, the encoding for "" is always UTF-8, and locale-sensitive operations use the operating system's native interface. (In particular, setting the LC_ALL and LC_CTYPE environment variables does not affect the locale "" on Mac OS X. Use getenv and current-locale to explicitly install the environment-specified locale, if desired.) Setting the current locale to #f makes locale-sensitive operations locale-insensitive, which means using the Unicode mapping for case operations and using UTF-8 for encoding.

```
(current-locale) → (or/c string? #f)
(current-locale locale) → void?
  locale : (or/c string? #f)
```

A parameter that determines the current locale for procedures such as string-locale-ci=?.

When locale sensitivity is disabled by setting the parameter to #f, strings are compared, etc., in a fully portable manner, which is the same as the standard procedures. Otherwise, strings are interpreted according to a locale setting (in the sense of the C library's setlocale). The "" locale is always an alias for the current machine's default locale, and it is the default. The "C" locale is also always available; setting the locale to "C" is the same as disabling locale sensitivity with #f only when string operations are restricted to the first 128 characters. Other locale names are platform-specific.

String or character printing with `write` is not affected by the parameter, and neither are symbol case or regular expressions (see §3.7 “Regular Expressions”).

12.1.2 Managing Ports

```
(input-port? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is an input port, `#f` otherwise.

```
(output-port? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is an output port, `#f` otherwise.

```
(port? v) → boolean?
  v : any/c
```

Returns `#t` if either `(input-port? v)` or `(output-port? v)` is `#t`, `#f` otherwise.

```
(close-input-port in) → void?
  in : input-port?
```

Closes the input port *in*. For some kinds of ports, closing the port releases lower-level resources, such as a file handle. If the port is already closed, `close-input-port` has no effect.

```
(close-output-port out) → void?
  out : output-port?
```

Closes the output port *out*. For some kinds of ports, closing the port releases lower-level resources, such as a file handle. Also, if the port is buffered, closing may first flush the port before closing it, and this flushing process can block. If the port is already closed, `close-output-port` has no effect.

```
(port-closed? port) → boolean?
  port : port?
```

Returns `#t` if the input or output port *port* is closed, `#f` otherwise.

```
(port-closed-evt port) → evt?
  port : port?
```

Return a synchronizable event that becomes ready when *port* is closed.

```
(current-input-port) → input-port?
(current-input-port in) → void?
  in : input-port?
```

A parameter that determines a default input port for many operations, such as `read`.

```
(current-output-port) → output-port?
(current-output-port out) → void?
  out : output-port?
```

A parameter that determines a default output port for many operations, such as `write`.

```
(current-error-port) → output-port?
(current-error-port out) → void?
  out : output-port?
```

A parameter that determines an output port that is typically used for errors and logging. For example, the default error display handler writes to this port.

```
(file-stream-port? port) → boolean?
  port : port?
```

Returns `#t` if the given port is a file-stream port (see §12.1.5 “File Ports”), `#f` otherwise.

```
(terminal-port? port) → boolean?
  port : port?
```

Returns `#t` if the given port is attached to an interactive terminal, `#f` otherwise.

```
eof : eof-object?
```

A value (distinct from all other values) that represents an end-of-file.

```
(eof-object? a) → boolean?
  a : any/c
```

Returns `#t` if `v` is `eof`, `#f` otherwise.

12.1.3 Port Buffers and Positions

Some ports—especially those that read from and write to files—are internally buffered:

- An input port is typically block-buffered by default, which means that on any read, the buffer is filled with immediately-available bytes to speed up future reads. Thus, if a file is modified between a pair of reads to the file, the second read can produce stale data. Calling `file-position` to set an input port's file position flushes its buffer.
- An output port is typically block-buffered by default, though a terminal output port is line-buffered, and the initial error output port is unbuffered. An output buffer is filled with a sequence of written bytes to be committed as a group, either when the buffer is full (in block mode), when a newline is written (in line mode), when the port is closed via `close-output-port`, or when a flush is explicitly requested via a procedure like `flush-output`.

If a port supports buffering, its buffer mode can be changed via `file-stream-buffer-mode` (even if the port is not a file-stream port).

For an input port, peeking always places peeked bytes into the port's buffer, even when the port's buffer mode is `'none`; furthermore, on some platforms, testing the port for input (via `char-ready?` or `sync`) may be implemented with a peek. If an input port's buffer mode is `'none`, then at most one byte is read for `read-bytes-avail!*`, `read-bytes-avail!`, `peek-bytes-avail!*`, or `peek-bytes-avail!`; if any bytes are buffered in the port (e.g., to satisfy a previous peek), the procedures may access multiple buffered bytes, but no further bytes are read.

In addition, the initial current output and error ports are automatically flushed when `read`, `read-line`, `read-bytes`, `read-string`, etc., are performed on the initial standard input port; more precisely, flushing is performed by the default port read handler (see `port-read-handler`).

```
(flush-output [out]) → void?
  out : output-port? = (current-output-port)
```

Forces all buffered data in the given output port to be physically written. Only file-stream ports, TCP ports, and custom ports (see §12.1.9 "Custom Ports") use buffers; when called on a port without a buffer, `flush-output` has no effect.

```
(file-stream-buffer-mode port) → (or/c 'none 'line 'block #f)
  port : port?
(file-stream-buffer-mode port mode) → void?
  port : port?
  mode : (or/c 'none 'line 'block)
```

Gets or sets the buffer mode for *port*, if possible. File-stream ports support setting the buffer mode, TCP ports (see §14.3 "Networking") support setting and getting the buffer mode, and custom ports (see §12.1.9 "Custom Ports") may support getting and setting buffer modes.

If *mode* is provided, it must be one of `'none`, `'line` (output only), or `'block`, and the port's buffering is set accordingly. If the port does not support setting the mode, the `exn:fail`

exception is raised.

If *mode* is not provided, the current mode is returned, or `#f` is returned if the mode cannot be determined. If `file-stream-port` is an input port and *mode* is `'line`, the `exn:fail:contract` exception is raised.

```
(file-position port) → exact-nonnegative-integer?
  port : port?
(file-position port pos) → void?
  port : port?
  pos : (or/c exact-nonnegative-integer? eof-object?)
```

Returns or sets the current read/write position of *port*.

Calling `file-position` without a position on a port other than a file-stream port or string port returns the number of bytes that have been read from that port if the position is known (see §12.1.4 "Counting Positions, Lines, and Columns"), otherwise the `exn:fail:filesystem` exception is raised.

For file-stream ports and string ports, the position-setting variant sets the read/write position to *pos* relative to the beginning of the file or (byte) string if *pos* is a number, or to the current end of the file or (byte) string if *pos* is `eof`. In position-setting mode, `file-position` raises the `exn:fail:contract` exception for port kinds other than file-stream ports and string ports. Furthermore, not all file-stream ports support setting the position; if `file-position` is called with a position argument on such a file-stream port, the `exn:fail:filesystem` exception is raised.

When `file-position` sets the position *pos* beyond the current size of an output file or (byte) string, the file/string is enlarged to size *pos* and the new region is filled with `0` bytes. If *pos* is beyond the end of an input file or (byte) string, then reading thereafter returns `eof` without changing the port's position.

When changing the file position for an output port, the port is first flushed if its buffer is not empty. Similarly, setting the position for an input port clears the port's buffer (even if the new position is the same as the old position). However, although input and output ports produced by `open-input-output-file` share the file position, setting the position via one port does not flush the other port's buffer.

12.1.4 Counting Positions, Lines, and Columns

By default, Racket keeps track of the *position* in a port as the number of bytes that have been read from or written to any port (independent of the read/write position, which is accessed or changed with `file-position`). Optionally, however, Racket can track the position in terms of characters (after UTF-8 decoding), instead of bytes, and it can track *line locations* and *column locations*; this optional tracking must be specifically enabled for a port via `port-`

count-lines! or the port-count-lines-enabled parameter. Position, line, and column locations for a port are used by read-syntax. Position and line locations are numbered from 1; column locations are numbered from 0.

When counting lines, Racket treats linefeed, return, and return-linefeed combinations as a line terminator and as a single position (on all platforms). Each tab advances the column count to one before the next multiple of 8. When a sequence of bytes in the range 128 to 253 forms a UTF-8 encoding of a character, the position/column is incremented once for each byte, and then decremented appropriately when a complete encoding sequence is discovered. See also §12.1 "Ports" for more information on UTF-8 decoding for ports.

A position is known for any port as long as its value can be expressed as a fixnum (which is more than enough tracking for realistic applications in, say, syntax-error reporting). If the position for a port exceeds the value of the largest fixnum, then the position for the port becomes unknown, and line and column tacking is disabled. Return-linefeed combinations are treated as a single character position only when line and column counting is enabled.

Custom ports can define their own counting functions, which are not subject to the rules above, except that the counting functions are invoked only when tracking is specifically enabled with port-count-lines!.

```
(port-count-lines! port) → void?
  port : port?
```

Turns on line and column counting for a port. Counting can be turned on at any time, though generally it is turned on before any data is read from or written to a port. At the point that line counting is turned on, port-next-location typically starts reporting as its last result (one more than) the number of characters read since line counting was enabled, instead of (one more than) bytes read since the port was opened.

When a port is created, if the value of the port-count-lines-enabled parameter is true, then line counting is automatically enabled for the port. Line counting cannot be disabled for a port after it is enabled.

```
(port-next-location port)
    (or/c exact-positive-integer? #f)
 → (or/c exact-nonnegative-integer? #f)
    (or/c exact-positive-integer? #f)
  port : port?
```

Returns three values: an integer or #f for the line number of the next read/written item, an integer or #f for the next item's column, and an integer or #f for the next item's position. The next column and position normally increase as bytes are read from or written to the port, but if line/character counting is enabled for *port*, the column and position results can decrease after reading or writing a byte that ends a UTF-8 encoding sequence.

If line counting is not enabled for a port, than the first two results are #f, and the last result is

one more than the number of bytes read so far. At the point when line counting is enabled, the first two results typically become non-#f, and last result starts reporting characters instead of bytes, typically starting from the point when line counting is enabled.

Even with line counting enabled, a port may return #f values if it somehow cannot keep track of lines, columns, or positions.

```
(set-port-next-location! port
                         line
                         column
                         position) → void?
  port : port?
  line : (or/c exact-positive-integer? #f)
  column : (or/c exact-nonnegative-integer? #f)
  position : (or/c exact-positive-integer? #f)
```

Sets the next line, column, and position for *port*. If line counting has not been enabled for *port* or if *port* is a custom port that defines its own counting function, then set-port-next-location! has no effect.

```
(port-count-lines-enabled) → boolean?
(port-count-lines-enabled on?) → void?
  on? : any/c
```

A parameter that determines whether line counting is enabled automatically for newly created ports. The default value is #f.

12.1.5 File Ports

A port created by open-input-file, open-output-file, subprocess, and related functions is a *file-stream port*. The initial input, output, and error ports in racket are also file-stream ports. The file-stream-port? predicate recognizes file-stream ports.

When an input or output file-stream port is created, it is placed into the management of the current custodian (see §13.7 “Custodians”).

```
(open-input-file path [#:mode mode-flag]) → input-port?
  path : path-string?
  mode-flag : (or/c 'binary 'text) = 'binary
```

Opens the file specified by *path* for input. The *mode-flag* argument specifies how the file’s bytes are translated on input:

- 'binary — bytes are returned from the port exactly as they are read from the file.

- `'text` — return and linefeed bytes (10 and 13) as read from the file are filtered by the port in a platform specific manner:
  - Unix and Mac OS X: no filtering occurs.
  - Windows: a return-linefeed combination from a file is returned by the port as a single linefeed; no filtering occurs for return bytes that are not followed by a linefeed, or for a linefeed that is not preceded by a return.

On Windows, `'text` mode works only with regular files; attempting to use `'text` with other kinds of files triggers an `exn:fail:filesystem` exception.

Otherwise, the file specified by *path* need not be a regular file. It might be a device that is connected through the filesystem, such as `"aux"` on Windows or `"/dev/null"` on Unix. In all cases, the port is buffered by default.

The port produced by `open-input-file` should be explicitly closed, either though `close-input-port` or indirectly via `custodian-shutdown-all`, to release the OS-level file handle. The input port will not be closed automatically if it is otherwise available for garbage collection (see §1.1.7 "Garbage Collection"); a will could be associated with an input port to close it more automatically (see §15.3 "Wills and Executors").

A path value that is the cleansed version of *path* is used as the name of the opened port.

Examples:

```
> (with-output-to-file some-file
    (lambda () (printf "hello world")))

(define in (open-input-file some-file))

> (read-string 11 in)
"hello world"
> (close-input-port in)
```

```
(open-output-file path
                  [#:mode mode-flag
                   #:exists exists-flag]) → output-port?
  path : path-string?
  mode-flag : (or/c 'binary 'text) = 'binary
                (or/c 'error 'append 'update 'can-update
  exists-flag :       'replace 'truncate
                      'must-truncate 'truncate/replace)
              = 'error
```

Opens the file specified by *path* for output. The *mode-flag* argument specifies how bytes written to the port are translated when written to the file:

- `'binary` — bytes are written to the file exactly as written to the port.
- `'text` — on Windows, a linefeed byte (10) written to the port is translated to a return-linefeed combination in the file; no filtering occurs for returns.

On Windows, `'text` mode works only with regular files; attempting to use `'text` with other kinds of files triggers an `exn:fail:filesystem` exception.

The *exists-flag* argument specifies how to handle/require files that already exist:

- `'error` — raise `exn:fail:filesystem` if the file exists.
- `'replace` — remove the old file, if it exists, and write a new one.
- `'truncate` — remove all old data, if the file exists.
- `'must-truncate` — remove all old data in an existing file; if the file does not exist, the `exn:fail:filesystem` exception is raised.
- `'truncate/replace` — try `'truncate`; if it fails (perhaps due to file permissions), try `'replace`.
- `'update` — open an existing file without truncating it; if the file does not exist, the `exn:fail:filesystem` exception is raised. Use `file-position` to change the current read/write position.
- `'can-update` — open an existing file without truncating it, or create the file if it does not exist.
- `'append` — append to the end of the file, whether it already exists or not; on Windows, `'append` is equivalent to `'update`, except that the file is not required to exist, and the file position is immediately set to the end of the file after opening it.

The file specified by *path* need not be a regular file. It might be a device that is connected through the filesystem, such as `"aux"` on Windows or `"/dev/null"` on Unix. The output port is block-buffered by default, unless the file corresponds to a terminal, in which case it is line-buffered by default.

The port produced by `open-output-file` should be explicitly closed, either though `close-output-port` or indirectly via `custodian-shutdown-all`, to release the OS-level file handle. The output port will not be closed automatically if it is otherwise available for garbage collection (see §1.1.7 “Garbage Collection”); a will could be associated with an output port to close it more automatically (see §15.3 “Wills and Executors”).

A path value that is the cleansed version of *path* is used as the name of the opened port.

Examples:

```
(define out (open-output-file some-file))

> (write "hello world" out)

> (close-output-port out)
```

```
(open-input-output-file path
                        [#:mode mode-flag
                         #:exists exists-flag])
 → input-port? output-port?
  path : path-string?
  mode-flag : (or/c 'binary 'text) = 'binary
  exists-flag :  (or/c 'error 'append 'update 'can-update
                       'replace 'truncate 'truncate/replace)
              = 'error
```

Like open-output-file, but producing two values: an input port and an output port. The two ports are connected in that they share the underlying file descriptor. This procedure is intended for use with special devices that can be opened by only one process, such as "COM1" in Windows. For regular files, sharing the file descriptor can be confusing. For example, using one port does not automatically flush the other port's buffer, and reading or writing in one port moves the file position (if any) for the other port. For regular files, use separate open-input-file and open-output-file calls to avoid confusion.

```
(call-with-input-file path
                      proc
                     [#:mode mode-flag]) → any
  path : path-string?
  proc : (input-port? . -> . any)
  mode-flag : (or/c 'binary 'text) = 'binary
```

Calls open-input-file with the *path* and *mode-flag* arguments, and passes the resulting port to *proc*. The result of *proc* is the result of the call-with-input-file call, but the newly opened port is closed when *proc* returns.

Examples:

```
> (with-output-to-file some-file
    (lambda () (printf "text in a file")))

> (call-with-input-file some-file
    (lambda (in) (read-string 14 in)))
"text in a file"
```

```
(call-with-output-file path
                       proc
                      [#:mode mode-flag
                       #:exists exists-flag]) → any
  path : path-string?
  proc : (output-port? . -> . any)
  mode-flag : (or/c 'binary 'text) = 'binary
  exists-flag : (or/c 'error 'append 'update
                      'replace 'truncate 'truncate/replace)
              = 'error
```

Analogous to `call-with-input-file`, but passing *path*, *mode-flag* and *exists-flag* to `open-output-file`.

Examples:

```
> (call-with-output-file some-file
    (lambda (out)
      (write 'hello out)))

> (call-with-input-file some-file
    (lambda (in)
      (read-string 5 in)))
"hello"
```

```
(call-with-input-file* path
                       proc
                      [#:mode mode-flag]) → any
  path : path-string?
  proc : (input-port? . -> . any)
  mode-flag : (or/c 'binary 'text) = 'binary
```

Like `call-with-input-file`, but the newly opened port is closed whenever control escapes the dynamic extent of the `call-with-input-file*` call, whether through *proc*'s return, a continuation application, or a prompt-based abort.

```
(call-with-output-file* path
                        proc
                       [#:mode mode-flag
                        #:exists exists-flag]) → any
  path : path-string?
  proc : (output-port? . -> . any)
  mode-flag : (or/c 'binary 'text) = 'binary
  exists-flag : (or/c 'error 'append 'update
                      'replace 'truncate 'truncate/replace)
              = 'error
```

Like call-with-output-file, but the newly opened port is closed whenever control escapes the dynamic extent of the call-with-output-file* call, whether through *proc*'s return, a continuation application, or a prompt-based abort.

```
(with-input-from-file path
                      thunk
                      [#:mode mode-flag]) → any
  path : path-string?
  thunk : (-> any)
  mode-flag : (or/c 'binary 'text) = 'binary
```

Like call-with-input-file*, but instead of passing the newly opened port to the given procedure argument, the port is installed as the current input port (see current-input-port) using parameterize around the call to *thunk*.

Examples:

```
> (with-output-to-file some-file
    (lambda () (printf "hello")))

> (with-input-from-file some-file
    (lambda () (read-string 5)))
"hello"
```

```
(with-output-to-file path
                     thunk
                     [#:mode mode-flag
                      #:exists exists-flag]) → any
  path : path-string?
  thunk : (-> any)
  mode-flag : (or/c 'binary 'text) = 'binary
  exists-flag : (or/c 'error 'append 'update
                      'replace 'truncate 'truncate/replace)
              = 'error
```

Like call-with-output-file*, but instead of passing the newly opened port to the given procedure argument, the port is installed as the current output port (see current-output-port) using parameterize around the call to *thunk*.

Examples:

```
> (with-output-to-file some-file
    (lambda () (printf "hello")))

> (with-input-from-file some-file
    (lambda () (read-string 5)))
```

```
"hello"
```

```
(port-try-file-lock? port mode) → boolean?
  port : file-stream-port?
  mode : (or/c 'shared 'exclusive)
```

Attempts to acquire a lock on the file using the current platform's facilities for file locking. Multiple processes can acquire a `'shared` lock on a file, but at most one process can hold an `'exclusive` lock, and `'shared` and `'exclusive` locks are mutually exclusive. When *mode* is `'shared`, then *port* must be an input port; when *mode* is `'exclusive`, then *port* must be an output port.

The result is `#t` if the requested lock is acquired, `#f` otherwise. When a lock is acquired, it is held until either it is released with `port-file-unlock` or the port is closed (perhaps because the process terminates).

Depending on the platform, locks may be merely advisory (i.e., locks affect only the ability of processes to acquire locks) or they may correspond to mandatory locks that prevent reads and writes to the locked file. Specifically, locks are mandatory on Windows and advisory on other platforms. Multiple tries for a `'shared` lock on a single port can succeed; on Unix and Mac OS X, a single `port-file-unlock` release the lock, while on other Windows, a `port-file-unlock` is needed for each successful `port-try-file-lock?`. On Unix and Mac OS X, multiple tries for a `'exclusive` lock can succeed and a single `port-file-unlock` releases the lock, while on Windows, a try for an `'exclusive` lock fails for a given port if the port already holds the lock.

A lock acquired for an input port from `open-input-output-file` can be released through `port-file-unlock` on the corresponding output port, and vice versa. If the output port from `open-input-output-file` holds an `'exclusive` lock, the corresponding input port can still acquire a `'shared` lock, even multiple times; on Windows, a `port-file-unlock` is needed for each successful lock try, while a single `port-file-unlock` balances the lock tries on Unix and Mac OS X. A `'shared` lock on an input port can be upgraded to an `'exclusive` lock through the corresponding output port on Unix and Mac OS X, in which case a single `port-file-unlock` (on either port) releases the lock, while such upgrades are not allowed on Windows.

Locking is normally supported only for file ports, and attempting to acquire a lock with other kinds of file-stream ports raises an `exn:fail:filesystem` exception.

```
(port-file-unlock port) → void?
  port : file-stream-port?
```

Releases a lock held by the current process on the file of *port*.

```
(port-file-identity port) → exact-positive-integer?
  port : file-stream-port?
```

Returns a number that represents the identity of the device and file read or written by *port*. For two ports whose open times overlap, the result of `port-file-identity` is the same for both ports if and only if the ports access the same device and file. For ports whose open times do not overlap, no guarantee can be provided for the port identities (even if the ports actually access the same file)—except as can be inferred through relationships with other ports. If *port* is closed, the `exn:fail` exception is raised. On Windows 95, 98, and Me, if *port* is connected to a pipe instead of a file, the `exn:fail:filesystem` exception is raised.

Examples:

```
(define file1 (open-output-file some-file))

(define file2 (open-output-file some-other-file))

> (port-file-identity file1)
807734784
> (port-file-identity file2)
808128000
> (close-output-port file1)

> (close-output-port file2)
```

12.1.6 String Ports

A *string port* reads or writes from a byte string. An input string port can be created from either a byte string or a string; in the latter case, the string is effectively converted to a byte string using `string->bytes/utf-8`. An output string port collects output into a byte string, but `get-output-string` conveniently converts the accumulated bytes to a string.

Input and output string ports do not need to be explicitly closed. The `file-position` procedure works for string ports in position-setting mode.

```
(open-input-bytes bstr [name]) → input-port?
  bstr : bytes?
  name : any/c = 'string
```

Creates an input string port that reads characters from *bstr* (see §3.4 “Byte Strings”). Modifying *bstr* afterward does not affect the byte stream produced by the port. The optional *name* argument is used as the name for the returned port.

```
(open-input-string str [name]) → input-port?
  str : string?
  name : any/c = 'string
```

Creates an input string port that reads bytes from the UTF-8 encoding (see §12.1.1 “Encodings and Locales”) of *str*. The optional *name* argument is used as the name for the returned port.

```
(open-output-bytes [name]) → output-port?
  name : any/c = 'string
```

Creates an output string port that accumulates the output into a byte string. The optional *name* argument is used as the name for the returned port.

```
(open-output-string [name]) → output-port?
  name : any/c = 'string
```

The same as open-output-bytes.

```
(get-output-bytes  out
                  [reset?
                   start-pos
                   end-pos]) → bytes?
  out : output-port?
  reset? : any/c = #f
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = #f
```

Returns the bytes accumulated in the string port *out* so far in a freshly allocated byte string (including any bytes written after the port’s current position, if any). The *out* port must be an output string port produced by open-output-bytes (or open-output-string) or a structure whose prop:output-port property refers to such an output port (transitively).

If *reset?* is true, then all bytes are removed from the port, and the port’s position is reset to 0; if *reset?* is #f, then all bytes remain in the port for further accumulation (so they are returned for later calls to get-output-bytes or get-output-string), and the port’s position is unchanged.

The *start-pos* and *end-pos* arguments specify the range of bytes in the port to return; supplying *start-pos* and *end-pos* is the same as using subbytes on the result of get-output-bytes, but supplying them to get-output-bytes can avoid an allocation. The *end-pos* argument can be #f, which corresponds to not passing a second argument to subbytes.

```
(get-output-string out) → string?
  out : output-port?
```

Returns (bytes->string/utf-8 (get-output-bytes *out*) #\?).

Examples:

```
> (define i (open-input-string "hello world"))

> (define o (open-output-string))

> (write (read i) o)

> (get-output-string o)
"hello"
```

12.1.7 Pipes

A Racket *pipe* is internal to Racket, and not related to OS-level pipes (which are file-stream ports) for communicating between different processes.

```
(make-pipe [limit input-name output-name])
 → input-port? output-port?
  limit : exact-positive-integer? = #f
  input-name : any/c = 'pipe
  output-name : any/c = 'pipe
```

Returns two port values: the first port is an input port and the second is an output port. Data written to the output port is read from the input port, with no intermediate buffering. Unlike some other kinds of ports, pipe ports do not need to be explicitly closed to be reclaimed by garbage collection.

If *limit* is `#f`, the new pipe holds an unlimited number of unread bytes (i.e., limited only by the available memory). If *limit* is a positive number, then the pipe will hold at most *limit* unread/unpeeked bytes; writing to the pipe's output port thereafter will block until a read or peek from the input port makes more space available. (Peeks effectively extend the port's capacity until the peeked bytes are read.)

The optional *input-name* and *output-name* are used as the names for the returned input and output ports, respectively.

```
(pipe-content-length pipe-port) → exact-nonnegative-integer?
  pipe-port : port?
```

Returns the number of bytes contained in a pipe, where *pipe-port* is either of the pipe's ports produced by `make-pipe`. The pipe's content length counts all bytes that have been written to the pipe and not yet read (though possibly peeked).

12.1.8 Structures as Ports

```
prop:input-port : struct-type-property?
```

```
prop:output-port : struct-type-property?
```

The prop:input-port and prop:output-port structure type properties identify structure types whose instances can serve as input and output ports, respectively.

Each property value can be either of the following:

- An input port (for prop:input-port) or output port (for prop:output-port): In this case, using the structure as port is equivalent to using the given input or output port.
- An exact, non-negative integer between 0 (inclusive) and the number of non-automatic fields in the structure type (exclusive, not counting supertype fields): The integer identifies a field in the structure, and the field must be designated as immutable. If the field contains an input port (for prop:input-port) or output port (for prop:output-port), the port is used. Otherwise, an empty string input port is used for prop:input-port, and a port that discards all data is used for prop:output-port.

Some procedures, such as file-position, work on both input and output ports. When given an instance of a structure type with both the prop:input-port and prop:output-port properties, the instance is used as an input port.

12.1.9 Custom Ports

The make-input-port and make-output-port procedures create *custom ports* with arbitrary control procedures (much like implementing a device driver). Custom ports are mainly useful to obtain fine control over the action of committing bytes as read or written.

```
(make-input-port  name
                  read-in
                  peek
                  close
                 [get-progress-evt
                  commit
                  get-location
                  count-lines!
                  init-position
                  buffer-mode])     → input-port?
  name : any/c
```

```
          (bytes?
           . -> . (or/c exact-nonnegative-integer?
read-in :                eof-object?
                         procedure?
                         evt?))
       (bytes? exact-nonnegative-integer? (or/c evt? #f)
               . -> . (or/c exact-nonnegative-integer?
                            eof-object?
peek :                      procedure?
                            evt?
                            #f))
close : (-> any)
get-progress-evt : (or/c (-> evt?) #f) = #f
          (or/c (exact-positive-integer? evt? evt? . -> . any)
commit :       #f)
       = #f
               (or/c
                (()
                 . ->* .
get-location :   ((or/c exact-positive-integer? #f)    = #f
                  (or/c exact-nonnegative-integer? #f)
                  (or/c exact-positive-integer? #f)))
                #f)
count-lines! : (-> any) = void
init-position : exact-positive-integer? = 1
              (or/c (case-> ((or/c 'block 'none) . -> . any)
buffer-mode :                (-> (or/c 'block 'none #f)))
                    #f)
            = #f
```

Creates an input port, which is immediately open for reading. If *close* procedure has no side effects, then the port need not be explicitly closed. See also `make-input-port/peek-to-read`.

The arguments implement the port as follows:

- *name* — the name for the input port.
- *read-in* — a procedure that takes a single argument: a mutable byte string to receive read bytes. The procedure's result is one of the following:
  - the number of bytes read, as an exact, non-negative integer;
  - `eof`;
  - a procedure of arity four (representing a "special" result, as discussed further below) and optionally of arity zero, but a procedure result is allowed only when *peek* is not `#f`;

– a pipe input port that supplies bytes to be used as long as the pipe has content (see pipe-content-length) or until *read-in* or *peek* is called again; or

– a synchronizable event (see §10.2.1 “Events”) other than a pipe input port or procedure of arity four; the event becomes ready when the read is complete (roughly): the event’s value can be one of the above three results or another event like itself; in the last case, a reading process loops with sync until it gets a non-event result.

The *read-in* procedure must not block indefinitely. If no bytes are immediately available for reading, the *read-in* must return 0 or an event, and preferably an event (to avoid busy waits). The *read-in* should not return 0 (or an event whose value is 0) when data is available in the port, otherwise polling the port will behave incorrectly. An event result from an event can also break polling.

If the result of a *read-in* call is not one of the above values, the exn:fail:contract exception is raised. If a returned integer is larger than the supplied byte string’s length, the exn:fail:contract exception is raised. If *peek* is #f and a procedure for a special result is returned, the exn:fail:contract exception is raised.

The *read-in* procedure can report an error by raising an exception, but only if no bytes are read. Similarly, no bytes should be read if eof, an event, or a procedure is returned. In other words, no bytes should be lost due to spurious exceptions or non-byte data.

A port’s reading procedure may be called in multiple threads simultaneously (if the port is accessible in multiple threads), and the port is responsible for its own internal synchronization. Note that improper implementation of such synchronization mechanisms might cause a non-blocking read procedure to block indefinitely.

If the result is a pipe input port, then previous *get-progress-evt* calls whose event is not yet ready must have been the pipe input port itself. Furthermore, *get-progress-evt* must continue to return the pipe as long as it contains data, or until the *read-in* or peek-in procedure is called again (instead of using the pipe, for whatever reason). If *read-in* or peek-in is called, any previously associated pipe (as returned by a previous call) is disassociated from the port and is not in use by any other thread as a result of the previous association.

If *peek*, *get-progress-evt*, and *commit* are all provided and non-#f, then the following is an acceptable implementation of *read-in*:

```
(lambda (bstr)
  (let* ([progress-evt (get-progress-evt)]
         [v (peek bstr 0 progress-evt)])
    (cond
     [(sync/timeout 0 progress-evt) 0] ; try again
     [(evt? v) (wrap-evt v (lambda (x) 0))] ; sync, try again
     [(and (number? v) (zero? v)) 0] ; try again
     [else
```

```
(if (commit (if (number? v) v 1)
                progress-evt
                always-evt)
    v      ; got a result
    0)]))) ; try again
```

An implementor may choose not to implement the *peek*, *get-progress-evt*, and *commit* procedures, however, and even an implementor who does supply them may provide a different *read-in* that uses a fast path for non-blocking reads.

- *peek* — either #f or a procedure that takes three arguments:
  - a mutable byte string to receive peeked bytes;
  - a non-negative number of bytes (or specials) to skip before peeking; and
  - either #f or a progress event produced by *get-progress-evt*.

  The results and conventions for *peek* are mostly the same as for *read-in*. The main difference is in the handling of the progress event, if it is not #f. If the given progress event becomes ready, the *peek* must abort any skip attempts and not peek any values. In particular, *peek* must not peek any values if the progress event is initially ready.

  Unlike *read-in*, *peek* should produce #f (or an event whose value is #f) if no bytes were peeked because the progress event became ready. Like *read-in*, a 0 result indicates that another attempt is likely to succeed, so 0 is inappropriate when the progress event is ready. Also like *read-in*, *peek* must not block indefinitely.

  The skip count provided to *peek* is a number of bytes (or specials) that must remain present in the port—in addition to the peek results—when the peek results are reported. If a progress event is supplied, then the peek is effectively canceled when another process reads data before the given number can be skipped. If a progress event is not supplied and data is read, then the peek must effectively restart with the original skip count.

  The system does not check that multiple peeks return consistent results, or that peeking and reading produce consistent results.

  If *peek* is #f, then peeking for the port is implemented automatically in terms of reads, but with several limitations. First, the automatic implementation is not thread-safe. Second, the automatic implementation cannot handle special results (non-byte and non-eof), so *read-in* cannot return a procedure for a special when *peek* is #f. Finally, the automatic peek implementation is incompatible with progress events, so if *peek* is #f, then progress-evt and *commit* must be #f. See also make-input-port/peek-to-read, which implements peeking in terms of *read-in* without these constraints.

- *close* — a procedure of zero arguments that is called to close the port. The port is not considered closed until the closing procedure returns. The port's procedures will never be used again via the port after it is closed. However, the closing procedure can be called simultaneously in multiple threads (if the port is accessible in multiple

threads), and it may be called during a call to the other procedures in another thread; in the latter case, any outstanding reads and peeks should be terminated with an error.

- *get-progress-evt* — either #f (the default), or a procedure that takes no arguments and returns an event. The event must become ready only after data is next read from the port or the port is closed. After the event becomes ready, it must remain so. See the description of *read-in* for information about the allowed results of this function when *read-in* returns a pipe input port. See also semaphore-peek-evt, which is sometimes useful for implementing *get-progress-evt*.

  If *get-progress-evt* is #f, then port-provides-progress-evts? applied to the port will produce #f, and the port will not be a valid argument to port-progress-evt.

  The result event will not be exposed directly by port-progress-evt. Instead, it will be wrapped in an event for which progress-evt? returns true.

- *commit* — either #f (the default), or a procedure that takes three arguments:

  – an exact, positive integer k_r;

  – a progress event produced by *get-progress-evt*;

  – an event, *done*, that is either a channel-put event, channel, semaphore, semaphore-peek event, always event, or never event.

  A *commit* corresponds to removing data from the stream that was previously peeked, but only if no other process removed data first. (The removed data does not need to be reported, because it has been peeked already.) More precisely, assuming that k_p bytes, specials, and mid-stream eofs have been previously peeked or skipped at the start of the port's stream, *commit* must satisfy the following constraints:

  – It must return only when the commit is complete or when the given progress event becomes ready.

  – It must commit only if k_p is positive.

  – If it commits, then it must do so with either k_r items or k_p items, whichever is smaller, and only if k_p is positive.

  – It must never choose *done* in a synchronization after the given progress event is ready, or after *done* has been synchronized once.

  – It must not treat any data as read from the port unless *done* is chosen in a synchronization.

  – It must not block indefinitely if *done* is ready; it must return soon after the read completes or soon after the given progress event is ready, whichever is first.

  – It can report an error by raising an exception, but only if no data has been committed. In other words, no data should be lost due to an exception, including a break exception.

  – It must return a true value if data has been committed, #f otherwise. When it returns a value, the given progress event must be ready (perhaps because data has just been committed).

  - It should return a byte string as a true result when line counting is enabled and *get-location* is #f (so that line counting is implemented the default way); the result byte string represents the data that was committed for the purposes of character and line counting. If any other true result is returned when a byte string is expected, it is treated like a byte string where each byte corresponds to a non-newline character.
  - It must raise an exception if no data (including eof) has been peeked from the beginning of the port's stream, or if it would have to block indefinitely to wait for the given progress event to become ready.

  A call to *commit* is parameterize-breaked to disable breaks.

- *get-location* — either #f (the default), or a procedure that takes no arguments and returns three values: the line number for the next item in the port's stream (a positive number or #f), the column number for the next item in the port's stream (a non-negative number or #f), and the position for the next item in the port's stream (a positive number or #f). See also §12.1.4 "Counting Positions, Lines, and Columns".

  This procedure is called to implement port-next-location, but only if line counting is enabled for the port via port-count-lines! (in which case *count-lines!* is called). The read and read-syntax procedures assume that reading a non-whitespace character increments the column and position by one.

- *count-lines!* — a procedure of no arguments that is called if and when line counting is enabled for the port. The default procedure is void.

- *init-position* — an exact, positive integer that determines the position of the port's first item, used when line counting is *not* enabled for the port. The default is 1.

- *buffer-mode* — either #f (the default) or a procedure that accepts zero or one arguments. If *buffer-mode* is #f, then the resulting port does not support a buffer-mode setting. Otherwise, the procedure is called with one symbol argument ('block or 'none) to set the buffer mode, and it is called with zero arguments to get the current buffer mode. In the latter case, the result must be 'block, 'none, or #f (unknown). See §12.1.3 "Port Buffers and Positions" for more information on buffer modes.

"Special" results: When *read-in* or *peek* (or an event produced by one of these) returns a procedure, the procedure is used to obtain a non-byte result. (This non-byte result is *not* intended to return a character or eof; in particular, read-char raises an exception if it encounters a special-result procedure, even if the procedure produces a byte.) A special-result procedure must accept four arguments, and it can optionally accept zero arguments:

- When the special read is triggered by read-syntax or read-syntax/recursive, the procedure is passed four arguments that represent a source location.

- When the special read is triggered by read, read-byte-or-special, read-char-or-special, peek-byte-or-special, or peek-char-or-special, the procedure

is passed no arguments if it accepts zero arguments, otherwise it is passed four arguments that are all `#f`.

The special-value procedure can return an arbitrary value, and it will be called zero or one times (not necessarily before further reads or peeks from the port). See §12.7.2 “Reader-Extension Procedures” for more details on the procedure’s result.

If *read-in* or *peek* returns a special procedure when called by any reading procedure other than `read`, `read-syntax`, `read-char-or-special`, `peek-char-or-special`, `read-byte-or-special`, or `peek-byte-or-special`, then the `exn:fail:contract` exception is raised.

Examples:

```
; A port with no input...
; Easy: (open-input-bytes #"")
; Hard:
> (define /dev/null-in
    (make-input-port 'null
                     (lambda (s) eof)
                     (lambda (skip s progress-evt) eof)
                     void
                     (lambda () never-evt)
                     (lambda (k progress-evt done-evt)
                       (error "no successful peeks!"))))

> (read-char /dev/null-in)
#<eof>
> (peek-char /dev/null-in)
#<eof>
> (read-byte-or-special /dev/null-in)
#<eof>
> (peek-byte-or-special /dev/null-in 100)
#<eof>
; A port that produces a stream of 1s:
> (define infinite-ones
    (make-input-port
     'ones
     (lambda (s)
       (bytes-set! s 0 (char->integer #\1)) 1)
     #f
     void))

> (read-string 5 infinite-ones)
"11111"
; But we can't peek ahead arbitrarily far, because the
```

```
; automatic peek must record the skipped bytes, so
; we'd run out of memory.
; An infinite stream of 1s with a specific peek procedure:
> (define infinite-ones
    (let ([one! (lambda (s)
                  (bytes-set! s 0 (char->integer #\1)) 1)])
      (make-input-port
       'ones
       one!
       (lambda (s skip progress-evt) (one! s))
       void)))

> (read-string 5 infinite-ones)
"11111"
; Now we can peek ahead arbitrarily far:
> (peek-string 5 (expt 2 5000) infinite-ones)
"11111"
; The port doesn't supply procedures to implement progress events:
> (port-provides-progress-evts? infinite-ones)
#f
> (port-progress-evt infinite-ones)
port-progress-evt: port does not provide progress evts
  port: #<input-port:ones>
; Non-byte port results:
> (define infinite-voids
    (make-input-port
     'voids
     (lambda (s) (lambda args 'void))
     (lambda (skip s evt) (lambda args 'void))
     void))

> (read-char infinite-voids)
read-char: non-character in an unsupported context
  port: #<input-port:voids>
> (read-char-or-special infinite-voids)
'void
; This port produces 0, 1, 2, 0, 1, 2, etc., but it is not
; thread-safe, because multiple threads might read and change n.
> (define mod3-cycle/one-thread
    (let* ([n 2]
           [mod! (lambda (s delta)
                   (bytes-set! s 0 (+ 48 (modulo (+ n delta) 3)))
                   1)])
      (make-input-port
       'mod3-cycle/not-thread-safe
       (lambda (s)
```

```
          (set! n (modulo (add1 n) 3))
          (mod! s 0))
        (lambda (s skip evt)
          (mod! s skip))
        void)))

> (read-string 5 mod3-cycle/one-thread)
"01201"
> (peek-string 5 (expt 2 5000) mod3-cycle/one-thread)
"20120"
; Same thing, but thread-safe and kill-safe, and with progress
; events. Only the server thread touches the stateful part
; directly. (See the output port examples for a simpler thread-
safe
; example, but this one is more general.)
> (define (make-mod3-cycle)
    (define read-req-ch (make-channel))
    (define peek-req-ch (make-channel))
    (define progress-req-ch (make-channel))
    (define commit-req-ch (make-channel))
    (define close-req-ch (make-channel))
    (define closed? #f)
    (define n 0)
    (define progress-sema #f)
    (define (mod! s delta)
      (bytes-set! s 0 (+ 48 (modulo (+ n delta) 3)))
      1)
    ; ----------------------------
    ; The server has a list of outstanding commit requests,
    ;  and it also must service each port operation (read,
    ;  progress-evt, etc.)
    (define (serve commit-reqs response-evts)
      (apply
       sync
       (handle-evt read-req-ch
                   (handle-read commit-reqs response-evts))
       (handle-evt progress-req-ch
                   (handle-progress commit-reqs response-evts))
       (handle-evt commit-req-ch
                   (add-commit commit-reqs response-evts))
       (handle-evt close-req-ch
                   (handle-close commit-reqs response-evts))
       (append
        (map (make-handle-response commit-reqs response-evts)
             response-evts)
        (map (make-handle-commit commit-reqs response-evts)
```

```
        commit-reqs))))
; Read/peek request: fill in the string and commit
(define ((handle-read commit-reqs response-evts) r)
  (let ([s (car r)]
        [skip (cadr r)]
        [ch (caddr r)]
        [nack (cadddr r)]
        [evt (car (cddddr r))]
        [peek? (cdr (cddddr r))])
    (let ([fail? (and evt
                      (sync/timeout 0 evt))])
      (unless (or closed? fail?)
        (mod! s skip)
        (unless peek?
          (commit! 1)))
      ; Add an event to respond:
      (serve commit-reqs
             (cons (choice-evt
                    nack
                    (channel-put-evt ch (if closed?
                                            0
                                            (if fail? #f 1))))
                   response-evts)))))
; Progress request: send a peek evt for the current
;  progress-sema
(define ((handle-progress commit-reqs response-evts) r)
  (let ([ch (car r)]
        [nack (cdr r)])
    (unless progress-sema
      (set! progress-sema (make-semaphore (if closed? 1 0))))
    ; Add an event to respond:
    (serve commit-reqs
           (cons (choice-evt
                  nack
                  (channel-put-evt
                   ch
                   (semaphore-peek-evt progress-sema)))
                 response-evts))))
; Commit request: add the request to the list
(define ((add-commit commit-reqs response-evts) r)
  (serve (cons r commit-reqs) response-evts))
; Commit handling: watch out for progress, in which case
;  the response is a commit failure; otherwise, try
;  to sync for a commit. In either event, remove the
;  request from the list
(define ((make-handle-commit commit-reqs response-evts) r)
```

```
      (let ([k (car r)]
            [progress-evt (cadr r)]
            [done-evt (caddr r)]
            [ch (cadddr r)]
            [nack (cddddr r)])
        ; Note: we don't check that k is <= the sum of
        ;  previous peeks, because the entire stream is actually
        ;  known, but we could send an exception in that case.
        (choice-evt
         (handle-evt progress-evt
                     (lambda (x)
                       (sync nack (channel-put-evt ch #f))
                       (serve (remq r commit-reqs) response-
evts)))
         ; Only create an event to satisfy done-evt if progress-
evt
         ;  isn't already ready.
         ; Afterward, if progress-evt becomes ready, then this
         ;  event-making function will be called again, because
         ;  the server controls all posts to progress-evt.
         (if (sync/timeout 0 progress-evt)
             never-evt
             (handle-evt done-evt
                         (lambda (v)
                           (commit! k)
                           (sync nack (channel-put-evt ch #t))
                           (serve (remq r commit-reqs)
                                  response-evts)))))))
    ; Response handling: as soon as the respondee listens,
    ;  remove the response
    (define ((make-handle-response commit-reqs response-evts) evt)
      (handle-evt evt
                  (lambda (x)
                    (serve commit-reqs
                           (remq evt response-evts)))))
    ; Close handling: post the progress sema, if any, and set
    ;   the closed? flag
    (define ((handle-close commit-reqs response-evts) r)
      (let ([ch (car r)]
            [nack (cdr r)])
        (set! closed? #t)
        (when progress-sema
          (semaphore-post progress-sema))
        (serve commit-reqs
               (cons (choice-evt nack
                                 (channel-put-evt ch (void)))
```

```
                response-evts))))
; Helper for reads and post-peek commits:
(define (commit! k)
  (when progress-sema
    (semaphore-post progress-sema)
    (set! progress-sema #f))
  (set! n (+ n k)))
; Start the server thread:
(define server-thread (thread (lambda () (serve null null))))
; ----------------------------
; Client-side helpers:
(define (req-evt f)
  (nack-guard-evt
   (lambda (nack)
     ; Be sure that the server thread is running:
     (thread-resume server-thread (current-thread))
     ; Create a channel to hold the reply:
     (let ([ch (make-channel)])
       (f ch nack)
       ch))))
(define (read-or-peek-evt s skip evt peek?)
  (req-evt (lambda (ch nack)
             (channel-put read-req-ch
                          (list* s skip ch nack evt peek?)))))
; Make the port:
(make-input-port 'mod3-cycle
                 ; Each handler for the port just sends
                 ;  a request to the server
                 (lambda (s) (read-or-peek-evt s 0 #f #f))
                 (lambda (s skip evt)
                   (read-or-peek-evt s skip evt #t))
                 (lambda () ; close
                   (sync (req-evt
                          (lambda (ch nack)
                            (channel-put progress-req-ch
                                         (list* ch nack))))))
                 (lambda () ; progress-evt
                   (sync (req-evt
                          (lambda (ch nack)
                            (channel-put progress-req-ch
                                         (list* ch nack))))))
                 (lambda (k progress-evt done-evt)  ; commit
                   (sync (req-evt
                          (lambda (ch nack)
                            (channel-put
                             commit-req-ch
```

```
                                (list* k progress-evt done-evt ch
                                       nack))))))))

> (define mod3-cycle (make-mod3-cycle))

> (let ([result1 #f]
        [result2 #f])
    (let ([t1 (thread
               (lambda ()
                 (set! result1 (read-string 5 mod3-cycle))))]
          [t2 (thread
               (lambda ()
                 (set! result2 (read-string 5 mod3-cycle))))])
      (thread-wait t1)
      (thread-wait t2)
      (string-append result1 "," result2)))
"02220,10101"
> (define s (make-bytes 1))

> (define progress-evt (port-progress-evt mod3-cycle))

> (peek-bytes-avail! s 0 progress-evt mod3-cycle)
1
> s
#"1"
> (port-commit-peeked 1 progress-evt (make-semaphore 1)
                      mod3-cycle)
#t
> (sync/timeout 0 progress-evt)
#<progress-evt>
> (peek-bytes-avail! s 0 progress-evt mod3-cycle)
0
> (port-commit-peeked 1 progress-evt (make-semaphore 1)
                      mod3-cycle)
#f
> (close-input-port mod3-cycle)
```

```
(make-output-port  name
                   evt
                   write-out
                   close
                  [write-out-special
                   get-write-evt
                   get-write-special-evt
                   get-location
                   count-lines!
                   init-position
                   buffer-mode])          → output-port?
  name : any/c
  evt : evt?
              (bytes? exact-nonnegative-integer?
                      exact-nonnegative-integer?
                      boolean?
                      boolean?
  write-out :         . -> .
                      (or/c exact-nonnegative-integer?
                            #f
                            evt?))
  close : (-> any)
                      (or/c (any/c boolean? boolean?
                                   . -> .
                                   (or/c any/c
  write-out-special :                    #f          = #f
                                         evt?))
                            #f)
                  (or/c
                   (bytes? exact-nonnegative-integer?
                           exact-nonnegative-integer?
  get-write-evt :          . -> .                     = #f
                           evt?)
                   #f)
                          (or/c
  get-write-special-evt :  (any/c . -> . evt?) = #f
                           #f)
                 (or/c
                  (()
                   . ->* .
  get-location :   ((or/c exact-positive-integer? #f)    = #f
                    (or/c exact-nonnegative-integer? #f)
                    (or/c exact-positive-integer? #f)))
                  #f)
  count-lines! : (-> any) = void
  init-position : exact-positive-integer? = 1
```

```
                (or/c (case->
                       ((or/c 'block 'line 'none) . -> . any)
buffer-mode :          (-> (or/c 'block 'line 'none #f)))
                      #f)
              = #f
```

Creates an output port, which is immediately open for writing. If *close* procedure has no side effects, then the port need not be explicitly closed. The port can buffer data within its *write-out* and *write-out-special* procedures.

- *name* — the name for the output port.
- *evt* — a synchronization event (see §10.2.1 “Events”; e.g., a semaphore or another port). The event is used in place of the output port when the port is supplied to synchronization procedures like `sync`. Thus, the event should be unblocked when the port is ready for writing at least one byte without blocking, or ready to make progress in flushing an internal buffer without blocking. The event must not unblock unless the port is ready for writing; otherwise, the guarantees of `sync` will be broken for the output port. Use `always-evt` if writes to the port always succeed without blocking.
- *write-out* — a procedure of five arguments:
  - an immutable byte string containing bytes to write;
  - a non-negative exact integer for a starting offset (inclusive) into the byte string;
  - a non-negative exact integer for an ending offset (exclusive) into the byte string;
  - a boolean; `#f` indicates that the port is allowed to keep the written bytes in a buffer, and that it is allowed to block indefinitely; `#t` indicates that the write should not block, and that the port should attempt to flush its buffer and completely write new bytes instead of buffering them;
  - a boolean; `#t` indicates that if the port blocks for a write, then it should enable breaks while blocking (e.g., using `sync/enable-break`); this argument is always `#f` if the fourth argument is `#t`.

  The procedure returns one of the following:

  - a non-negative exact integer representing the number of bytes written or buffered;
  - `#f` if no bytes could be written, perhaps because the internal buffer could not be completely flushed;
  - a pipe output port (when buffering is allowed and not when flushing) for buffering bytes as long as the pipe is not full and until *write-out* or *write-out-special* is called; or
  - a synchronizable event (see §10.2.1 “Events”) other than a pipe output port that acts like the result of `write-bytes-avail-evt` to complete the write.

Since *write-out* can produce an event, an acceptable implementation of *write-out* is to pass its first three arguments to the port's *get-write-evt*. Some port implementors, however, may choose not to provide *get-write-evt* (perhaps because writes cannot be made atomic), or may implement write-proc to enable a fast path for non-blocking writes or to enable buffering.

From a user's perspective, the difference between buffered and completely written data is (1) buffered data can be lost in the future due to a failed write, and (2) flush-output forces all buffered data to be completely written. Under no circumstances is buffering required.

If the start and end indices are the same, then the fourth argument to *write-out* will be #f, and the write request is actually a flush request for the port's buffer (if any), and the result should be 0 for a successful flush (or if there is no buffer).

The result should never be 0 if the start and end indices are different, otherwise the exn:fail:contract exception is raised. Similarly, the exn:fail:contract exception is raised if *write-out* returns a pipe output port when buffering is disallowed or when it is called for flushing. If a returned integer is larger than the supplied byte-string range, the exn:fail:contract exception is raised.

The #f result should be avoided, unless the next write attempt is likely to work. Otherwise, if data cannot be written, return an event instead.

An event returned by *write-out* can return #f or another event like itself, in contrast to events produced by write-bytes-avail-evt or *get-write-evt*. A writing process loops with sync until it obtains a non-event result.

The *write-out* procedure is always called with breaks disabled, independent of whether breaks were enabled when the write was requested by a client of the port. If breaks were enabled for a blocking operation, then the fifth argument to *write-out* will be #t, which indicates that *write-out* should re-enable breaks while blocking.

If the writing procedure raises an exception, due to write or commit operations, it must not have committed any bytes (though it may have committed previously buffered bytes).

A port's writing procedure may be called in multiple threads simultaneously (if the port is accessible in multiple threads). The port is responsible for its own internal synchronization. Note that improper implementation of such synchronization mechanisms might cause a non-blocking write procedure to block.

- *close* — a procedure of zero arguments that is called to close the port. The port is not considered closed until the closing procedure returns. The port's procedures will never be used again via the port after it is closed. However, the closing procedure can be called simultaneously in multiple threads (if the port is accessible in multiple threads), and it may be called during a call to the other procedures in another thread; in the latter case, any outstanding writes or flushes should be terminated immediately with an error.

- *write-out-special* — either #f (the default), or a procedure to handle write-special calls for the port. If #f, then the port does not support special output, and port-writes-special? will return #f when applied to the port.

If a procedure is supplied, it takes three arguments: the special value to write, a boolean that is `#f` if the procedure can buffer the special value and block indefinitely, and a boolean that is `#t` if the procedure should enable breaks while blocking. The result is one of the following:

- a non-event true value, which indicates that the special is written;
- `#f` if the special could not be written, perhaps because an internal buffer could not be completely flushed;
- a synchronizable event (see §10.2.1 "Events") that acts like the result of *get-write-special-evt* to complete the write.

Since *write-out-special* can return an event, passing the first argument to an implementation of *get-write-special-evt* is acceptable as a *write-out-special*.

As for *write-out*, the `#f` result is discouraged, since it can lead to busy waiting. Also as for *write-out*, an event produced by *write-out-special* is allowed to produce `#f` or another event like itself. The *write-out-special* procedure is always called with breaks disabled, independent of whether breaks were enabled when the write was requested by a client of the port.

- *get-write-evt* — either `#f` (the default) or a procedure of three arguments:
  - an immutable byte string containing bytes to write;
  - a non-negative exact integer for a starting offset (inclusive) into the byte string; and
  - a non-negative exact integer for an ending offset (exclusive) into the byte string.

  The result is a synchronizable event (see §10.2.1 "Events") to act as the result of `write-bytes-avail-evt` for the port (i.e., to complete a write or flush), which becomes available only as data is committed to the port's underlying device, and whose result is the number of bytes written.

  If *get-write-evt* is `#f`, then `port-writes-atomic?` will produce `#f` when applied to the port, and the port will not be a valid argument to procedures such as `write-bytes-avail-evt`. Otherwise, an event returned by *get-write-evt* must not cause data to be written to the port unless the event is chosen in a synchronization, and it must write to the port if the event is chosen (i.e., the write must appear atomic with respect to the synchronization).

  If the event's result integer is larger than the supplied byte-string range, the `exn:fail:contract` exception is raised by a wrapper on the event. If the start and end indices are the same (i.e., no bytes are to be written), then the event should produce `0` when the buffer is completely flushed. (If the port has no buffer, then it is effectively always flushed.)

  If the event raises an exception, due to write or commit operations, it must not have committed any new bytes (though it may have committed previously buffered bytes).

  Naturally, a port's events may be used in multiple threads simultaneously (if the port is accessible in multiple threads). The port is responsible for its own internal synchronization.

- *get-write-special-evt* — either #f (the default), or a procedure to handle write-special-evt calls for the port. This argument must be #f if either *write-out-special* or *get-write-evt* is #f, and it must be a procedure if both of those arguments are procedures.

  If it is a procedure, it takes one argument: the special value to write. The resulting event (with its constraints) is analogous to the result of *get-write-evt*.

  If the event raises an exception, due to write or commit operations, it must not have committed the special value (though it may have committed previously buffered bytes and values).

- *get-location* — either #f (the default), or a procedure that takes no arguments and returns three values: the line number for the next item written to the port's stream (a positive number or #f), the column number for the next item written to port's stream (a non-negative number or #f), and the position for the next item written to port's stream (a positive number or #f). See also §12.1.4 "Counting Positions, Lines, and Columns".

  This procedure is called to implement port-next-location for the port, but only if line counting is enabled for the port via port-count-lines! (in which case *count-lines!* is called).

- *count-lines!* — a procedure of no arguments that is called if and when line counting is enabled for the port. The default procedure is void.

- *init-position* — an exact, positive integer that determines the position of the port's first output item, used when line counting is *not* enabled for the port. The default is 1.

- *buffer-mode* — either #f (the default) or a procedure that accepts zero or one arguments. If *buffer-mode* is #f, then the resulting port does not support a buffer-mode setting. Otherwise, the procedure is called with one symbol argument ('block, 'line, or 'none) to set the buffer mode, and it is called with zero arguments to get the current buffer mode. In the latter case, the result must be 'block, 'line, 'none, or #f (unknown). See §12.1.3 "Port Buffers and Positions" for more information on buffer modes.

Examples:

```
; A port that writes anything to nowhere:
> (define /dev/null-out
    (make-output-port
     'null
     always-evt
     (lambda (s start end non-block? breakable?) (- end start))
     void
     (lambda (special non-block? breakable?) #t)
     (lambda (s start end) (wrap-evt
                            always-evt
                            (lambda (x)
```

```
                              (- end start))))
     (lambda (special) always-evt)))

> (display "hello" /dev/null-out)

> (write-bytes-avail #"hello" /dev/null-out)
5
> (write-special 'hello /dev/null-out)
#t
> (sync (write-bytes-avail-evt #"hello" /dev/null-out))
5
; A port that accumulates bytes as characters in a list,
;  but not in a thread-safe way:
> (define accum-list null)

> (define accumulator/not-thread-safe
    (make-output-port
     'accum/not-thread-safe
     always-evt
     (lambda (s start end non-block? breakable?)
       (set! accum-list
             (append accum-list
                     (map integer->char
                          (bytes->list (subbytes s start end)))))
       (- end start))
     void))

> (display "hello" accumulator/not-thread-safe)

> accum-list
'(#\h #\e #\l #\l #\o)
; Same as before, but with simple thread-safety:
> (define accum-list null)

> (define accumulator
    (let* ([lock (make-semaphore 1)]
           [lock-peek-evt (semaphore-peek-evt lock)])
      (make-output-port
       'accum
       lock-peek-evt
       (lambda (s start end non-block? breakable?)
         (if (semaphore-try-wait? lock)
             (begin
               (set! accum-list
                     (append accum-list
                             (map integer->char
```

```
                                  (bytes->list
                                   (subbytes s start end)))))
                (semaphore-post lock)
                (- end start))
              ; Cheap strategy: block until the list is unlocked,
              ;   then return 0, so we get called again
              (wrap-evt
               lock-peek
               (lambda (x) 0))))
        void)))

> (display "hello" accumulator)

> accum-list
'(#\h #\e #\l #\l #\o)
; A port that transforms data before sending it on
;  to another port. Atomic writes exploit the
;  underlying port's ability for atomic writes.
> (define (make-latin-1-capitalize port)
    (define (byte-upcase s start end)
      (list->bytes
       (map (lambda (b) (char->integer
                         (char-upcase
                          (integer->char b))))
            (bytes->list (subbytes s start end)))))
    (make-output-port
     'byte-upcase
     ; This port is ready when the original is ready:
     port
     ; Writing procedure:
     (lambda (s start end non-block? breakable?)
       (let ([s (byte-upcase s start end)])
         (if non-block?
             (write-bytes-avail* s port)
             (begin
               (display s port)
               (bytes-length s)))))
     ; Close procedure -- close original port:
     (lambda () (close-output-port port))
     #f
     ; Write event:
     (and (port-writes-atomic? port)
          (lambda (s start end)
            (write-bytes-avail-evt
             (byte-upcase s start end)
             port)))))
```

```
> (define orig-port (open-output-string))

> (define cap-port (make-latin-1-capitalize orig-port))

> (display "Hello" cap-port)

> (get-output-string orig-port)
"HELLO"
> (sync (write-bytes-avail-evt #"Bye" cap-port))
3
> (get-output-string orig-port)
"HELLOBYE"
```

12.1.10 More Port Constructors, Procedures, and Events

```
(require racket/port)
```

The bindings documented in this section are provided by the `racket/port` and `racket` libraries, but not `racket/base`.

Port String and List Conversions

```
(port->list [r in]) → (listof any/c)
  r : (input-port? . -> . any/c) = read
  in : input-port? = (current-input-port)
```

Returns a list whose elements are produced by calling *r* on *in* until it produces `eof`.

Examples:

```
> (define (read-number input-port)
    (define char (read-char input-port))
    (if (eof-object? char)
     char
     (string->number (string char))))

> (port->list read-number (open-input-string "12345"))
'(1 2 3 4 5)
```

```
(port->string [in]) → string?
  in : input-port? = (current-input-port)
```

Reads all characters from *in* and returns them as a string.

Example:

```
> (port->string (open-input-string "hello world"))
"hello world"
```

```
(port->bytes [in]) → bytes?
  in : input-port? = (current-input-port)
```

Reads all bytes from *in* and returns them as a byte string.

Example:

```
> (port->bytes (open-input-string "hello world"))
#"hello world"
```

```
(port->lines [in #:line-mode line-mode]) → (listof string?)
  in : input-port? = (current-input-port)
  line-mode : (or/c 'linefeed 'return 'return-linefeed 'any 'any-one)
            = 'any
```

Read all characters from *in*, breaking them into lines. The *line-mode* argument is the same as the second argument to `read-line`, but the default is `'any` instead of `'linefeed`.

Example:

```
> (port->lines
   (open-input-string "line 1\nline 2\n  line 3\nline 4"))
'("line 1" "line 2" "  line 3" "line 4")
```

```
(port->bytes-lines [in
                    #:line-mode line-mode]) → (listof bytes?)
  in : input-port? = (current-input-port)
  line-mode : (or/c 'linefeed 'return 'return-linefeed 'any 'any-one)
            = 'any
```

Like `port->lines`, but reading bytes and collecting them into lines like `read-bytes-line`.

Example:

```
> (port->bytes-lines
   (open-input-string "line 1\nline 2\n  line 3\nline 4"))
'(#"line 1" #"line 2" #"  line 3" #"line 4")
```

```
(display-lines  lst
                [out
                 #:separator separator]) → void?
  lst : list?
  out : output-port? = (current-output-port)
  separator : any/c = #"\n"
```

Uses `display` on each element of *lst* to *out*, adding *separator* after each element.

```
(call-with-output-string proc) → string?
  proc : (output-port? . -> . any)
```

Calls *proc* with an output port that accumulates all output into a string, and returns the string.

The port passed to *proc* is like the one created by `open-output-string`, except that it is wrapped via `dup-output-port`, so that *proc* cannot access the port's content using `get-output-string`. If control jumps back into *proc*, the port continues to accumulate new data, and `call-with-output-string` returns both the old data and newly accumulated data.

```
(call-with-output-bytes proc) → bytes?
  proc : (output-port? . -> . any)
```

Like `call-with-output-string`, but returns the accumulated result in a byte string instead of a string. Furthermore, the port's content is emptied when `call-with-output-bytes` returns, so that if control jumps back into *proc* and returns a second time, only the newly accumulated bytes are returned.

```
(with-output-to-string proc) → string?
  proc : (-> any)
```

Equivalent to

```
(call-with-output-string
 (lambda (p) (parameterize ([current-output-port p])
               (proc))))
```

```
(with-output-to-bytes proc) → bytes?
  proc : (-> any)
```

Equivalent to

```
(call-with-output-bytes
 (lambda (p) (parameterize ([current-output-port p])
               (proc))))
```

```
(call-with-input-string str proc) → any
  str : string?
  proc : (input-port? . -> . any)
```

Equivalent to `(proc (open-input-string str))`.

```
(call-with-input-bytes bstr proc) → any
  bstr : bytes?
  proc : (input-port? . -> . any)
```

Equivalent to `(proc (open-input-bytes bstr))`.

```
(with-input-from-string str proc) → any
  str : string?
  proc : (-> any)
```

Equivalent to

```
(parameterize ([current-input-port (open-input-string str)])
  (proc))
```

```
(with-input-from-bytes bstr proc) → any
  bstr : bytes?
  proc : (-> any)
```

Equivalent to

```
(parameterize ([current-input-port (open-input-bytes str)])
  (proc))
```

Creating Ports

```
(input-port-append close-at-eof? in ...) → input-port?
  close-at-eof? : any/c
  in : input-port?
```

Takes any number of input ports and returns an input port. Reading from the input port draws bytes (and special non-byte values) from the given input ports in order. If *close-at-eof?* is true, then each port is closed when an end-of-file is encountered from the port, or when the result input port is closed. Otherwise, data not read from the returned input port remains available for reading in its original input port.

See also `merge-input`, which interleaves data from multiple input ports as it becomes available.

```
(make-input-port/read-to-peek  name
                               read-in
                               fast-peek
                               close
                              [get-location
                               count-lines!
                               init-position
                               buffer-mode
                               buffering?
                               on-consumed]) → input-port?
  name : any/c
           (bytes?
            . -> . (or/c exact-nonnegative-integer?
  read-in :                eof-object?
                           procedure?
                           evt?))
             (or/c #f
                   (bytes? exact-nonnegative-integer?
                    (bytes? exact-nonnegative-integer?
                     . -> . (or/c exact-nonnegative-integer?
                                  eof-object?
                                  procedure?
  fast-peek :                     evt?
                                  #f))
                    . -> . (or/c exact-nonnegative-integer?
                                 eof-object?
                                 procedure?
                                 evt?
                                 #f)))
  close : (-> any)
                 (or/c
                  (->
                   (values
  get-location :    (or/c exact-positive-integer? #f)    = #f
                    (or/c exact-nonnegative-integer? #f)
                    (or/c exact-positive-integer? #f)))
                  #f)
  count-lines! : (-> any) = void
  init-position : exact-positive-integer? = 1
                (or/c (case-> ((or/c 'block 'none) . -> . any)
  buffer-mode :                (-> (or/c 'block 'none #f)))
                        #f)
              = #f
  buffering? : any/c = #f
```

```
                (or/c ((or/c exact-nonnegative-integer? eof-object?
                             procedure? evt?)
  on-consumed :
                       . -> . any)
                      #f)
              = #f
```

Similar to make-input-port, but if the given *read-in* returns an event, the event's value must be 0. The resulting port's peek operation is implemented automatically (in terms of *read-in*) in a way that can handle special non-byte values. The progress-event and commit operations are also implemented automatically. The resulting port is thread-safe, but not kill-safe (i.e., if a thread is terminated or suspended while using the port, the port may become damaged).

The *read-in*, *close*, *get-location*, *count-lines!*, *init-position*, and *buffer-mode* procedures are the same as for make-input-port.

The *fast-peek* argument can be either #f or a procedure of three arguments: a byte string to receive a peek, a skip count, and a procedure of two arguments. The *fast-peek* procedure can either implement the requested peek, or it can dispatch to its third argument to implement the peek. The *fast-peek* is not used when a peek request has an associated progress event.

The *buffering?* argument determines whether *read-in* can be called to read more characters than are immediately demanded by the user of the new port. If *buffer-mode* is not #f, then *buffering?* determines the initial buffer mode, and *buffering?* is enabled after a buffering change only if the new mode is 'block.

If *on-consumed* is not #f, it is called when data is read (or committed) from the port, as opposed to merely peeked. The argument to *on-consumed* is the result value of the port's reading procedure, so it can be an integer or any result from *read-in*.

```
(make-limited-input-port  in
                          limit
                         [close-orig?]) → input-port?
  in : input-port?
  limit : exact-nonnegative-integer?
  close-orig? : any/c = #t
```

Returns a port whose content is drawn from *in*, but where an end-of-file is reported after *limit* bytes (and non-byte special values) have been read. If *close-orig?* is true, then the original port is closed if the returned port is closed.

Bytes are consumed from *in* only when they are consumed from the returned port. In particular, peeking into the returned port peeks into the original port.

If *in* is used directly while the resulting port is also used, then the *limit* bytes provided by the port need not be contiguous parts of the original port's stream.

```
(make-pipe-with-specials [limit
                          in-name
                          out-name]) → input-port? output-port?
  limit : exact-nonnegative-integer? = #f
  in-name : any/c = 'pipe
  out-name : any/c = 'pipe
```

Returns two ports: an input port and an output port. The ports behave like those returned by `make-pipe`, except that the ports support non-byte values written with procedures such as `write-special` and read with procedures such as `get-byte-or-special`.

The *limit* argument determines the maximum capacity of the pipe in bytes, but this limit is disabled if special values are written to the pipe before *limit* is reached. The limit is re-enabled after the special value is read from the pipe.

The optional *in-name* and *out-name* arguments determine the names of the result ports.

```
(merge-input a-in b-in [buffer-limit]) → input-port?
  a-in : input-port?
  b-in : input-port?
  buffer-limit : (or/c exact-nonnegative-integer? #f) = 4096
```

Accepts two input ports and returns a new input port. The new port merges the data from two original ports, so data can be read from the new port whenever it is available from either of the two original ports. The data from the original ports are interleaved. When an end-of-file has been read from an original port, it no longer contributes characters to the new port. After an end-of-file has been read from both original ports, the new port returns end-of-file. Closing the merged port does not close the original ports.

The optional *buffer-limit* argument limits the number of bytes to be buffered from *a-in* and *b-in*, so that the merge process does not advance arbitrarily beyond the rate of consumption of the merged data. A `#f` value disables the limit. As for `make-pipe-with-specials`, *buffer-limit* does not apply when a special value is produced by one of the input ports before the limit is reached.

See also `input-port-append`, which concatenates input streams instead of interleaving them.

```
(open-output-nowhere [name special-ok?]) → output-port?
  name : any/c = 'nowhere
  special-ok? : any/c = #t
```

Creates and returns an output port that discards all output sent to it (without blocking). The *name* argument is used as the port's name. If the *special-ok?* argument is true, then the resulting port supports `write-special`, otherwise it does not.

```
(peeking-input-port  in
                    [name
                     skip
                     #:init-position init-position]) → input-port
  in : input-port?
  name : any/c = (object-name in)
  skip : exact-nonnegative-integer? = 0
  init-position : exact-positive-integer? = 1
```

Returns an input port whose content is determined by peeking into *in*. In other words, the resulting port contains an internal skip count, and each read of the port peeks into *in* with the internal skip count, and then increments the skip count according to the amount of data successfully peeked.

The optional *name* argument is the name of the resulting port. The *skip* argument is the port initial skip count, and it defaults to 0.

The resulting port's initial position (as reported by file-position) is (- *init-position* 1), no matter the position of *in*.

For example, when you read from a peeking port, you see the same answers as when you read from the original port.

Examples:

```
> (define an-original-port (open-input-string "123456789"))

> (define a-peeking-port (peeking-input-port an-original-port))

> (read-string 3 a-peeking-port)
"123"
> (read-string 3 an-original-port)
"123"
```

But beware, the read from the original port is invisible to the peeking port, which keeps its own separate internal counter, and thus interleaving reads on the two ports can produce confusing results. Continuing the example before, if we read three more characters from the peeking port, we end up skipping over the 456 in the port.

Example:

```
> (read-string 3 a-peeking-port)
"789"
```

```
(reencode-input-port  in
                     encoding
                     error-bytes
                    [close?
                     name
                     convert-newlines?
                     enc-error])        → input-port?
  in : input-port?
  encoding : string?
  error-bytes : (or/c #f bytes?)
  close? : any/c = #t
  name : any/c = (object-name in)
  convert-newlines? : any/c = #f
  enc-error : (string? input-port? . -> . any)
            = (lambda (msg port) (error ...))
```

Produces an input port that draws bytes from *in*, but converts the byte stream using `(bytes-open-converter encoding-str "UTF-8")`. In addition, if *convert-newlines?* is true, then decoded sequences that correspond to UTF-8 encodings of `"\r\n"`, `"\r\u0085"`, `"\r"`, `"\u0085"`, and `"\u2028"` are all converted to the UTF-8 encoding of `"\n"`.

If *error-bytes* is provided and not `#f`, then the given byte sequence is used in place of bytes from *in* that trigger conversion errors. Otherwise, if a conversion is encountered, *enc-error* is called, which must raise an exception.

If *close?* is true, then closing the result input port also closes *in*. The *name* argument is used as the name of the result input port.

In non-buffered mode, the resulting input port attempts to draw bytes from *in* only as needed to satisfy requests. Toward that end, the input port assumes that at least *n* bytes must be read to satisfy a request for *n* bytes. (This is true even if the port has already drawn some bytes, as long as those bytes form an incomplete encoding sequence.)

```
(reencode-output-port  out
                       encoding
                       error-bytes
                      [close?
                       name
                       newline-bytes
                       enc-error])    → output-port?
  out : output-port?
  encoding : string?
  error-bytes : (or/c #f bytes?)
  close? : any/c = #t
  name : any/c = (object-name out)
  newline-bytes : (or/c #f bytes?) = #f
```

```
enc-error : (string? output-port? . -> . any)
          = (lambda (msg port) (error ...))
```

Produces an output port that directs bytes to *out*, but converts its byte stream using (bytes-open-converter "UTF-8" encoding-str). In addition, if *newline-bytes* is not #f, then bytes written to the port that are the UTF-8 encoding of "\n" are first converted to *newline-bytes* (before applying the convert from UTF-8 to encoding-str).

If *error-bytes* is provided and not #f, then the given byte sequence is used in place of bytes that have been sent to the output port and that trigger conversion errors. Otherwise, *enc-error* is called, which must raise an exception.

If *close?* is true, then closing the result output port also closes *out*. The *name* argument is used as the name of the result output port.

The resulting port supports buffering, and the initial buffer mode is (or (file-stream-buffer-mode *out*) 'block). In 'block mode, the port's buffer is flushed only when it is full or a flush is requested explicitly. In 'line mode, the buffer is flushed whenever a newline or carriage-return byte is written to the port. In 'none mode, the port's buffer is flushed after every write. Implicit flushes for 'line or 'none leave bytes in the buffer when they are part of an incomplete encoding sequence.

The resulting output port does not support atomic writes. An explicit flush or special-write to the output port can hang if the most recently written bytes form an incomplete encoding sequence.

```
(dup-input-port in [close?]) → input-port?
  in : input-port?
  close? : any/c = #f
```

Returns an input port that draws directly from *in*. Closing the resulting port closes *in* only if *close?* is #t.

The new port is initialized with the port read handler of *in*, but setting the handler on the result port does not affect reading directly from *in*.

```
(dup-output-port out [close?]) → output-port?
  out : output-port?
  close? : any/c = #f
```

Returns an output port that propagates data directly to *out*. Closing the resulting port closes *out* only if *close?* is #t.

The new port is initialized with the port display handler and port write handler of *out*, but setting the handlers on the result port does not affect writing directly to *out*.

```
(relocate-input-port  in
                      line
                      column
                      position
                     [close?])  → input-port?
  in : input-port?
  line : (or/c exact-positive-integer? #f)
  column : (or/c exact-nonnegative-integer? #f)
  position : exact-positive-integer?
  close? : any/c = #t
```

Produces an input port that is equivalent to *in* except in how it reports location information. The resulting port's content starts with the remaining content of *in*, and it starts at the given line, column, and position. A `#f` for the line or column means that the line and column will always be reported as `#f`.

The *line* and *column* values are used only if line counting is enabled for *in* and for the resulting port, typically through `port-count-lines!`. The *column* value determines the column for the first line (i.e., the one numbered *line*), and later lines start at column `0`. The given *position* is used even if line counting is not enabled.

When line counting is on for the resulting port, reading from *in* instead of the resulting port increments location reports from the resulting port. Otherwise, the resulting port's position does not increment when data is read from *in*.

If *close?* is true, then closing the resulting port also closes *in*. If *close?* is `#f`, then closing the resulting port does not close *in*.

```
(relocate-output-port  out
                       line
                       column
                       position
                      [close?])  → output-port?
  out : output-port?
  line : (or/c exact-positive-integer? #f)
  column : (or/c exact-nonnegative-integer? #f)
  position : exact-positive-integer?
  close? : any/c = #t
```

Like `relocate-input-port`, but for output ports.

```
(transplant-input-port  in
                        get-location
                        init-pos
                       [close?
                        count-lines!])  → input-port?
```

```
  in : input-port?
                 (or/c
                  (->
                   (values
  get-location :    (or/c exact-positive-integer? #f)
                    (or/c exact-nonnegative-integer? #f)
                    (or/c exact-positive-integer? #f)))
                  #f)
  init-pos : exact-positive-integer?
  close? : any/c = #t
  count-lines! : (-> any) = void
```

Like `relocate-input-port`, except that arbitrary position information can be produced (when line counting is enabled) via *get-location*, which is used as for `make-input-port`. If *get-location* is `#f`, then the port counts lines in the usual way starting from *init-pos*, independent of locations reported by *in*.

If *count-lines!* is supplied, it is called when line counting is enabled for the resulting port. The default is `void`.

```
(transplant-output-port  in
                         get-location
                         init-pos
                        [close?
                         count-lines!]) → output-port?
  in : input-port?
                 (or/c
                  (->
                   (values
  get-location :    (or/c exact-positive-integer? #f)
                    (or/c exact-nonnegative-integer? #f)
                    (or/c exact-positive-integer? #f)))
                  #f)
  init-pos : exact-positive-integer?
  close? : any/c = #t
  count-lines! : (-> any) = void
```

Like `transplant-input-port`, but for output ports.

```
(filter-read-input-port  in
                         read-wrap
                         peek-wrap
                        [close?])   → input-port?
  in : input-port?
```

```
             (bytes? (or/c exact-nonnegative-integer?
                           eof-object?
                           procedure?
                           evt?)
  read-wrap :          . -> .
                       (or/c exact-nonnegative-integer?
                             eof-object?
                             procedure?
                             evt?))
             (bytes? exact-nonnegative-integer? (or/c evt? #f)
                     (or/c exact-nonnegative-integer?
                      eof-object?
                      procedure?
                      evt?
  peek-wrap :         #f)
              . -> . (or/c exact-nonnegative-integer?
                      eof-object?
                      procedure?
                      evt?
                      #f))
  close? : any/c = #t
```

Creates a port that draws from *in*, but each result from the port's read and peek procedures (in the sense of `make-input-port`) is filtered by *read-wrap* and *peek-wrap*. The filtering procedures each receive both the arguments and results of the read and peek procedures on *in* for each call.

If *close?* is true, then closing the resulting port also closes *in*.

```
(special-filter-input-port in proc [close?]) → input-port?
  in : input-port?
       (procedure? bytes? . -> . (or/c exact-nonnegative-integer?
                                       eof-object?
  proc :                               procedure?
                                       evt?))
  close? : any/c = #t
```

Produces an input port that is equivalent to *in*, except that when *in* produces a procedure to access a special value, *proc* is applied to the procedure to allow the special value to be replaced with an alternative. The *proc* is called with the special-value procedure and the byte string that was given to the port's read or peek function (see `make-input-port`), and the result is used as the read or peek function's result. The *proc* can modify the byte string to substitute a byte for the special value, but the byte string is guaranteed only to hold at least one byte.

If *close?* is true, then closing the resulting input port also closes *in*.

Port Events

```
(eof-evt in) → evt?
  in : input-port?
```

Returns a synchronizable event that is ready when *in* produces an `eof`. If *in* produces a mid-stream `eof`, the `eof` is consumed by the event only if the event is chosen in a synchronization.

```
(read-bytes-evt k in) → evt?
  k : exact-nonnegative-integer?
  in : input-port?
```

Returns a synchronizable event that is ready when *k* bytes can be read from *in*, or when an end-of-file is encountered in *in*. If *k* is `0`, then the event is ready immediately with `""`. For non-zero *k*, if no bytes are available before an end-of-file, the event's result is `eof`. Otherwise, the event's result is a byte string of up to *k* bytes, which contains as many bytes as are available (up to *k*) before an available end-of-file. (The result is a byte string on less than *k* bytes only when an end-of-file is encountered.)

Bytes are read from the port if and only if the event is chosen in a synchronization, and the returned bytes always represent contiguous bytes in the port's stream.

The event can be synchronized multiple times—event concurrently—and each synchronization corresponds to a distinct read request.

The *in* must support progress events, and it must not produce a special non-byte value during the read attempt.

```
(read-bytes!-evt bstr in progress-evt) → evt?
  bstr : (and/c bytes? (not/c immutable?))
  in : input-port?
  progress-evt : (or/c progress-evt? #f)
```

Like `read-bytes-evt`, except that the read bytes are placed into *bstr*, and the number of bytes to read corresponds to `(bytes-length bstr)`. The event's result is either `eof` or the number of read bytes.

The *bstr* may be mutated any time after the first synchronization attempt on the event and until either the event is selected, a non-`#f` *progress-evt* is ready, or the current custodian (at the time of synchronization) is shut down. Note that there is no time bound otherwise on when *bstr* might be mutated if the event is not selected by a synchronzation; nevertheless, multiple synchronization attempts can use the same result from `read-bytes!-evt` as long as there is no intervening read on *in* until one of the synchronization attempts selects the event.

```
(read-bytes-avail!-evt bstr in) → evt?
  bstr : (and/c bytes? (not/c immutable?))
  in : input-port?
```

Like read-bytes!-evt, except that the event reads only as many bytes as are immediately available, after at least one byte or one eof becomes available.

```
(read-string-evt k in) → evt?
  k : exact-nonnegative-integer?
  in : input-port?
```

Like read-bytes-evt, but for character strings instead of byte strings.

```
(read-string!-evt str in) → evt?
  str : (and/c string? (not/c immutable?))
  in : input-port?
```

Like read-bytes!-evt, but for a character string instead of a byte string.

```
(read-line-evt in mode) → evt?
  in : input-port?
  mode : (or/c 'linefeed 'return 'return-linefeed 'any 'any-one)
```

Returns a synchronizable event that is ready when a line of characters or end-of-file can be read from *in*. The meaning of *mode* is the same as for read-line. The event result is the read line of characters (not including the line separator).

A line is read from the port if and only if the event is chosen in a synchronization, and the returned line always represents contiguous bytes in the port's stream.

```
(read-bytes-line-evt in mode) → evt?
  in : input-port?
  mode : (or/c 'linefeed 'return 'return-linefeed 'any 'any-one)
```

Like read-line-evt, but returns a byte string instead of a string.

```
(peek-bytes-evt k skip progress-evt in) → evt?
  k : exact-nonnegative-integer?
  skip : exact-nonnegative-integer?
  progress-evt : (or/c progress-evt? #f)
  in : input-port?
(peek-bytes!-evt bstr skip progress-evt in) → evt?
  bstr : (and/c bytes? (not/c immutable?))
  skip : exact-nonnegative-integer?
```

```
  progress-evt : (or/c progress-evt? #f)
  in : input-port?
(peek-bytes-avail!-evt bstr
                       skip
                       progress-evt
                       in)          → evt?
  bstr : (and/c bytes? (not/c immutable?))
  skip : exact-nonnegative-integer?
  progress-evt : (or/c progress-evt? #f)
  in : input-port?
(peek-string-evt k skip progress-evt in) → evt?
  k : exact-nonnegative-integer?
  skip : exact-nonnegative-integer?
  progress-evt : (or/c progress-evt? #f)
  in : input-port?
(peek-string!-evt str skip progress-evt in) → evt?
  str : (and/c string? (not/c immutable?))
  skip : exact-nonnegative-integer?
  progress-evt : (or/c progress-evt? #f)
  in : input-port?
```

Like the `read-...-evt` functions, but for peeking. The *skip* argument indicates the number of bytes to skip, and *progress-evt* indicates an event that effectively cancels the peek (so that the event never becomes ready). The *progress-evt* argument can be `#f`, in which case the event is never canceled.

```
(regexp-match-evt pattern in) → any
  pattern : (or/c string? bytes? regexp? byte-regexp?)
  in : input-port?
```

Returns a synchronizable event that is ready when *pattern* matches the stream of bytes/characters from *in*; see also `regexp-match`. The event's value is the result of the match, in the same form as the result of `regexp-match`.

If *pattern* does not require a start-of-stream match, then bytes skipped to complete the match are read and discarded when the event is chosen in a synchronization.

Bytes are read from the port if and only if the event is chosen in a synchronization, and the returned match always represents contiguous bytes in the port's stream. If not-yet-available bytes from the port might contribute to the match, the event is not ready. Similarly, if *pattern* begins with a start-of-stream `^` and the *pattern* does not initially match, then the event cannot become ready until bytes have been read from the port.

The event can be synchronized multiple times—even concurrently—and each synchronization corresponds to a distinct match request.

The *in* port must support progress events. If *in* returns a special non-byte value during the match attempt, it is treated like eof.

Copying Streams

```
(convert-stream from-encoding
                in
                to-encoding
                out)          → void?
  from-encoding : string?
  in : input-port?
  to-encoding : string?
  out : output-port?
```

Reads data from *in*, converts it using (bytes-open-converter *from-encoding* *to-encoding*) and writes the converted bytes to *out*. The convert-stream procedure returns after reaching eof in *in*.

If opening the converter fails, the exn:fail exception is raised. Similarly, if a conversion error occurs at any point while reading from *in*, then exn:fail exception is raised.

```
(copy-port in out ...+) → void?
  in : input-port?
  out : output-port?
```

Reads data from *in* and writes it back out to *out*, returning when *in* produces eof. The copy is efficient, and it is without significant buffer delays (i.e., a byte that becomes available on *in* is immediately transferred to *out*, even if future reads on *in* must block). If *in* produces a special non-byte value, it is transferred to *out* using write-special.

This function is often called from a "background" thread to continuously pump data from one stream to another.

If multiple *out*s are provided, case data from *in* is written to every *out*. The different *out*s block output to each other, because each block of data read from *in* is written completely to one *out* before moving to the next *out*. The *out*s are written in the provided order, so non-blocking ports (e.g., file output ports) should be placed first in the argument list.

12.2 Byte and String Input

```
(read-char [in]) → (or/c character? eof-object?)
  in : input-port? = (current-input-port)
```

Reads a single character from *in*—which may involve reading several bytes to UTF-8-decode them into a character (see §12.1 “Ports”); a minimal number of bytes are read/peeked to perform the decoding. If no bytes are available before an end-of-file, then `eof` is returned.

```
(read-byte [in]) → (or/c byte? eof-object?)
  in : input-port? = (current-input-port)
```

Reads a single byte from *in*. If no bytes are available before an end-of-file, then `eof` is returned.

```
(read-line [in mode]) → (or/c string? eof-object?)
  in : input-port? = (current-input-port)
  mode : (or/c 'linefeed 'return 'return-linefeed 'any 'any-one)
       = 'linefeed
```

Returns a string containing the next line of bytes from *in*.

Characters are read from *in* until a line separator or an end-of-file is read. The line separator is not included in the result string (but it is removed from the port’s stream). If no characters are read before an end-of-file is encountered, `eof` is returned.

The *mode* argument determines the line separator(s). It must be one of the following symbols:

- `'linefeed` breaks lines on linefeed characters.
- `'return` breaks lines on return characters.
- `'return-linefeed` breaks lines on return-linefeed combinations. If a return character is not followed by a linefeed character, it is included in the result string; similarly, a linefeed that is not preceded by a return is included in the result string.
- `'any` breaks lines on any of a return character, linefeed character, or return-linefeed combination. If a return character is followed by a linefeed character, the two are treated as a combination.
- `'any-one` breaks lines on either a return or linefeed character, without recognizing return-linefeed combinations.

Return and linefeed characters are detected after the conversions that are automatically performed when reading a file in text mode. For example, reading a file in text mode on Windows automatically changes return-linefeed combinations to a linefeed. Thus, when a file is opened in text mode, `'linefeed` is usually the appropriate `read-line` mode.

```
(read-bytes-line [in mode]) → (or/c bytes? eof-object?)
  in : input-port? = (current-input-port)
  mode : (or/c 'linefeed 'return 'return-linefeed 'any 'any-one)
       = 'linefeed
```

Like read-line, but reads bytes and produces a byte string.

```
(read-string amt [in]) → (or/c string? eof-object?)
  amt : exact-nonnegative-integer?
  in : input-port? = (current-input-port)
```

To read an entire port as a string use port->string.

Returns a string containing the next amt characters from in.

If amt is 0, then the empty string is returned. Otherwise, if fewer than amt characters are available before an end-of-file is encountered, then the returned string will contain only those characters before the end-of-file; that is, the returned string's length will be less than amt. (A temporary string of size amt is allocated while reading the input, even if the size of the result is less than amt characters.) If no characters are available before an end-of-file, then eof is returned.

If an error occurs during reading, some characters may be lost; that is, if read-string successfully reads some characters before encountering an error, the characters are dropped.

```
(read-bytes amt [in]) → (or/c bytes? eof-object?)
  amt : exact-nonnegative-integer?
  in : input-port? = (current-input-port)
```

To read an entire port as bytes use port->bytes.

Like read-string, but reads bytes and produces a byte string.

```
(read-string! str [in start-pos end-pos])
 → (or/c exact-positive-integer? eof-object?)
  str : (and/c string? (not/c immutable?))
  in : input-port? = (current-input-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (string-length str)
```

Reads characters from in like read-string, but puts them into str starting from index start-pos (inclusive) up to end-pos (exclusive). Like substring, the exn:fail:contract exception is raised if start-pos or end-pos is out-of-range for str.

If the difference between start-pos and end-pos is 0, then 0 is returned and str is not modified. If no bytes are available before an end-of-file, then eof is returned. Otherwise, the return value is the number of characters read. If m characters are read and $m <$ end-pos - start-pos, then str is not modified at indices start-pos $+ m$ through end-pos.

```
(read-bytes! bstr [in start-pos end-pos])
 → (or/c exact-positive-integer? eof-object?)
  bstr : bytes?
  in : input-port? = (current-input-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like `read-string!`, but reads bytes, puts them into a byte string, and returns the number of bytes read.

```
(read-bytes-avail! bstr [in start-pos end-pos])
 → (or/c exact-positive-integer? eof-object? procedure?)
  bstr : bytes?
  in : input-port? = (current-input-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like `read-bytes!`, but returns without blocking after having read the immediately available bytes, and it may return a procedure for a "special" result. The `read-bytes-avail!` procedure blocks only if no bytes (or specials) are yet available. Also unlike `read-bytes!`, `read-bytes-avail!` never drops bytes; if `read-bytes-avail!` successfully reads some bytes and then encounters an error, it suppresses the error (treating it roughly like an end-of-file) and returns the read bytes. (The error will be triggered by future reads.) If an error is encountered before any bytes have been read, an exception is raised.

When *in* produces a special value, as described in §12.1.9 "Custom Ports", the result is a procedure of four arguments. The four arguments correspond to the location of the special value within the port, as described in §12.1.9 "Custom Ports". If the procedure is called more than once with valid arguments, the `exn:fail:contract` exception is raised. If `read-bytes-avail` returns a special-producing procedure, then it does not place characters in *bstr*. Similarly, `read-bytes-avail` places only as many bytes into *bstr* as are available before a special value in the port's stream.

```
(read-bytes-avail!* bstr
                    [in
                     start-pos
                     end-pos])
 → (or/c exact-nonnegative-integer? eof-object? procedure?)
  bstr : bytes?
  in : input-port? = (current-input-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like `read-bytes-avail!`, but returns `0` immediately if no bytes (or specials) are available for reading and the end-of-file is not reached.

```
(read-bytes-avail!/enable-break bstr
                                [in
                                 start-pos
                                 end-pos])
 → (or/c exact-positive-integer? eof-object? procedure?)
  bstr : bytes?
  in : input-port? = (current-input-port)
```

```
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like read-bytes-avail!, but breaks are enabled during the read (see also §9.6 “Breaks”). If breaking is disabled when read-bytes-avail!/enable-break is called, and if the exn:break exception is raised as a result of the call, then no bytes will have been read from *in*.

```
(peek-string amt skip-bytes-amt [in]) → (or/c string? eof-object?)
  amt : exact-nonnegative-integer?
  skip-bytes-amt : exact-nonnegative-integer?
  in : input-port? = (current-input-port)
```

Similar to read-string, except that the returned characters are preserved in the port for future reads. (More precisely, undecoded bytes are left for future reads.) The *skip-bytes-amt* argument indicates a number of bytes (*not* characters) in the input stream to skip before collecting characters to return; thus, in total, the next *skip-bytes-amt* bytes plus *amt* characters are inspected.

For most kinds of ports, inspecting *skip-bytes-amt* bytes and *amt* characters requires at least *skip-bytes-amt*+*amt* bytes of memory overhead associated with the port, at least until the bytes/characters are read. No such overhead is required when peeking into a string port (see §12.1.6 “String Ports”), a pipe port (see §12.1.7 “Pipes”), or a custom port with a specific peek procedure (depending on how the peek procedure is implemented; see §12.1.9 “Custom Ports”).

If a port produces eof mid-stream, peek skips beyond the eof always produce eof until the eof is read.

```
(peek-bytes amt skip-bytes-amt [in]) → (or/c bytes? eof-object?)
  amt : exact-nonnegative-integer?
  skip-bytes-amt : exact-nonnegative-integer?
  in : input-port? = (current-input-port)
```

Like peek-string, but peeks bytes and produces a byte string.

```
(peek-string!  str
               skip-bytes-amt
              [in
               start-pos
               end-pos])
 → (or/c exact-positive-integer? eof-object?)
  str : (and/c string? (not/c immutable?))
  skip-bytes-amt : exact-nonnegative-integer?
  in : input-port? = (current-input-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (string-length str)
```

Like read-string!, but for peeking, and with a skip-bytes-amt argument like peek-string.

```
(peek-bytes!  bstr
              skip-bytes-amt
             [in
              start-pos
              end-pos])
 → (or/c exact-positive-integer? eof-object?)
  bstr : (and/c bytes? (not/c immutable?))
  skip-bytes-amt : exact-nonnegative-integer?
  in : input-port? = (current-input-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like peek-string!, but peeks bytes, puts them into a byte string, and returns the number of bytes read.

```
(peek-bytes-avail!  bstr
                    skip-bytes-amt
                   [progress
                    in
                    start-pos
                    end-pos])
 → (or/c exact-nonnegative-integer? eof-object? procedure?)
  bstr : (and/c bytes? (not/c immutable?))
  skip-bytes-amt : exact-nonnegative-integer?
  progress : (or/c progress-evt? #f) = #f
  in : input-port? = (current-input-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like read-bytes-avail!, but for peeking, and with two extra arguments. The skip-bytes-amt argument is as in peek-bytes. The progress argument must be either #f or an event produced by port-progress-evt for in.

To peek, peek-bytes-avail! blocks until finding an end-of-file, at least one byte (or special) past the skipped bytes, or until a non-#f progress becomes ready. Furthermore, if progress is ready before bytes are peeked, no bytes are peeked or skipped, and progress may cut short the skipping process if it becomes available during the peek attempt.

The result of peek-bytes-avail! is 0 only in the case that progress becomes ready before bytes are peeked.

```
(peek-bytes-avail!*  bstr
                     skip-bytes-amt
                    [progress
                     in
                     start-pos
                     end-pos])
 → (or/c exact-nonnegative-integer? eof-object? procedure?)
  bstr : (and/c bytes? (not/c immutable?))
  skip-bytes-amt : exact-nonnegative-integer?
  progress : (or/c progress-evt? #f) = #f
  in : input-port? = (current-input-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like `read-bytes-avail!*`, but for peeking, and with *skip-bytes-amt* and *progress* arguments like `peek-bytes-avail!`. Since this procedure never blocks, it may return before even `skip-amt` bytes are available from the port.

```
(peek-bytes-avail!/enable-break  bstr
                                 skip-bytes-amt
                                [progress
                                 in
                                 start-pos
                                 end-pos])
 → (or/c exact-nonnegative-integer? eof-object? procedure?)
  bstr : (and/c bytes? (not/c immutable?))
  skip-bytes-amt : exact-nonnegative-integer?
  progress : (or/c progress-evt? #f) = #f
  in : input-port? = (current-input-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like `read-bytes-avail!/enable-break`, but for peeking, and with *skip-bytes-amt* and *progress* arguments like `peek-bytes-avail!`.

```
(read-char-or-special [in]) → (or/c character? eof-object? any/c)
  in : input-port? = (current-input-port)
```

Like `read-char`, but if the input port returns a non-byte value (through a value-generating procedure in a custom port; see §12.1.9 “Custom Ports” and §12.7.3 “Special Comments” for details), then the non-byte value is returned.

```
(read-byte-or-special [in]) → (or/c byte? eof-object? any/c)
  in : input-port? = (current-input-port)
```

Like `read-char-or-special`, but reads and returns a byte instead of a character.

```
(peek-char [in skip-bytes-amt]) → (or/c character? eof-object?)
  in : input-port? = (current-input-port)
  skip-bytes-amt : exact-nonnegative-integer? = 0
```

Like read-char, but peeks instead of reading, and skips *skip-bytes-amt* bytes (not characters) at the start of the port.

```
(peek-byte [in skip-bytes-amt]) → (or/c byte? eof-object?)
  in : input-port? = (current-input-port)
  skip-bytes-amt : exact-nonnegative-integer? = 0
```

Like peek-char, but reads and returns a byte instead of a character.

```
(peek-char-or-special [in skip-bytes-amt])
 → (or/c character? eof-object? any/c)
  in : input-port? = (current-input-port)
  skip-bytes-amt : exact-nonnegative-integer? = 0
```

Like peek-char, but if the input port returns a non-byte value after *skip-bytes-amt* byte positions, then it is returned.

```
(peek-byte-or-special [in
                       skip-bytes-amt
                       progress])
 → (or/c character? eof-object? any/c)
  in : input-port? = (current-input-port)
  skip-bytes-amt : exact-nonnegative-integer? = 0
  progress : (or/c progress-evt? #f) = #f
```

Like peek-char-or-special, but reads and returns a byte instead of a character, and it supports a *progress* argument like peek-bytes-avail!.

```
(port-progress-evt [in]) → progress-evt?
  in : (and/c input-port? port-provides-progress-evts?)
     = (current-input-port)
```

Returns a synchronizable event that becomes ready after any subsequent read from *in*, or after *in* is closed. After the event becomes ready, it remains ready..

```
(port-provides-progress-evts? in) → boolean
  in : input-port?
```

Returns #t if port-progress-evt can return an event for *in*. All built-in kinds of ports support progress events, but ports created with make-input-port (see §12.1.9 “Custom Ports”) may not.

```
(port-commit-peeked amt progress evt [in]) → boolean?
  amt : exact-nonnegative-integer?
  progress : progress-evt?
  evt : evt?
  in : input-port? = (current-input-port)
```

Attempts to commit as read the first *amt* previously peeked bytes, non-byte specials, and `eof`s from *in*, or the first `eof` or special value peeked from *in*. (Only mid-stream `eof`s can be committed. An `eof` when the port is exhausted does not correspond to data in the stream.)

The read commits only if *progress* does not become ready first (i.e., if no other process reads from *in* first), and only if *evt* is chosen by a `sync` within `port-commit-peeked` (in which case the event result is ignored); the *evt* must be either a channel-put event, channel, semaphore, semaphore-peek event, always event, or never event. Suspending the thread that calls `port-commit-peeked` may or may not prevent the commit from proceeding.

The result from `port-commit-peeked` is `#t` if data has been committed, and `#f` otherwise.

If no data has been peeked from *in* and *progress* is not ready, then `exn:fail:contract` exception is raised. If fewer than *amt* items have been peeked at the current start of *in*'s stream, then only the peeked items are committed as read. If *in*'s stream currently starts at an `eof` or a non-byte special value, then only the `eof` or special value is committed as read.

If *progress* is not a result of `port-progress-evt` applied to *in*, then `exn:fail:contract` exception is raised.

```
(byte-ready? [in]) → boolean?
  in : input-port? = (current-input-port)
```

Returns `#t` if `(read-byte in)` would not block (at the time that `byte-ready?` was called, at least). Equivalent to `(and (sync/timeout 0 in) #t)`.

```
(char-ready? [in]) → boolean?
  in : input-port? = (current-input-port)
```

Returns `#t` if `(read-char in)` would not block (at the time that `char-ready?` was called, at least). Depending on the initial bytes of the stream, multiple bytes may be needed to form a UTF-8 encoding.

```
(progress-evt? v) → boolean?
  v : any/c
(progress-evt? evt in) → boolean?
  evt : progress-evt?
  in : input-port?
```

With one argument, returns `#t` is *v* is a progress evt for some input port, `#f` otherwise.

With two arguments, returns `#t` if *evt* is a progress event for *in*, `#f` otherwise.

12.3 Byte and String Output

```
(write-char char [out]) → void?
  char : character?
  out : output-port? = (current-output-port)
```

Writes a single character to *out*; more precisely, the bytes that are the UTF-8 encoding of *char* are written to *out*.

```
(write-byte byte [out]) → void?
  byte : any/c
  out : output-port? = (current-output-port)
```

Writes a single byte to *out*.

```
(newline [out]) → void?
  out : output-port? = (current-output-port)
```

The same as `(write-char #\newline out)`.

```
(write-string str [out start-pos end-pos])
 → exact-nonnegative-integer?
  str : string?
  out : output-port? = (current-output-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (string-length str)
```

Writes characters to *out* from *str* starting from index *start-pos* (inclusive) up to *end-pos* (exclusive). Like `substring`, the `exn:fail:contract` exception is raised if *start-pos* or *end-pos* is out-of-range for *str*.

The result is the number of characters written to *out*, which is always `(- end-pos start-pos)`.

```
(write-bytes bstr [out start-pos end-pos])
 → exact-nonnegative-integer?
  bstr : bytes?
  out : output-port? = (current-output-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like write-string, but writes bytes instead of characters.

```
(write-bytes-avail  bstr
                   [out
                    start-pos
                    end-pos]) → exact-nonnegative-integer?
  bstr : bytes?
  out : output-port? = (current-output-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like write-bytes, but returns without blocking after writing as many bytes as it can immediately flush. It blocks only if no bytes can be flushed immediately. The result is the number of bytes written and flushed to *out*; if *start-pos* is the same as *end-pos*, then the result can be 0 (indicating a successful flush of any buffered data), otherwise the result is between 1 and (- *end-pos* *start-pos*), inclusive.

The write-bytes-avail procedure never drops bytes; if write-bytes-avail successfully writes some bytes and then encounters an error, it suppresses the error and returns the number of written bytes. (The error will be triggered by future writes.) If an error is encountered before any bytes have been written, an exception is raised.

```
(write-bytes-avail*  bstr
                    [out
                     start-pos
                     end-pos])
 → (or/c exact-nonnegative-integer? #f)
  bstr : bytes?
  out : output-port? = (current-output-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like write-bytes-avail, but never blocks, returns #f if the port contains buffered data that cannot be written immediately, and returns 0 if the port's internal buffer (if any) is flushed but no additional bytes can be written immediately.

```
(write-bytes-avail/enable-break  bstr
                                 [out
                                  start-pos
                                  end-pos])
 → exact-nonnegative-integer?
  bstr : bytes?
  out : output-port? = (current-output-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like write-bytes-avail, except that breaks are enabled during the write. The procedure provides a guarantee about the interaction of writing and breaks: if breaking is disabled when write-bytes-avail/enable-break is called, and if the exn:break exception is raised as a result of the call, then no bytes will have been written to *out*. See also §9.6 "Breaks".

```
(write-special v [out]) → boolean?
  v : any/c
  out : output-port? = (current-output-port)
```

Writes *v* directly to *out* if the port supports special writes, or raises exn:fail:contract if the port does not support special write. The result is always #t, indicating that the write succeeded.

```
(write-special-avail* v [out]) → boolean?
  v : any/c
  out : output-port? = (current-output-port)
```

Like write-special, but without blocking. If *v* cannot be written immediately, the result is #f without writing *v*, otherwise the result is #t and *v* is written.

```
(write-bytes-avail-evt  bstr
                       [out
                        start-pos
                        end-pos]) → evt?
  bstr : bytes?
  out : output-port? = (current-output-port)
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Similar to write-bytes-avail, but instead of writing bytes immediately, it returns a synchronizable event (see §10.2.1 "Events"). The *out* must support atomic writes, as indicated by port-writes-atomic?.

Synchronizing on the object starts a write from *bstr*, and the event becomes ready when bytes are written (unbuffered) to the port. If *start-pos* and *end-pos* are the same, then the synchronization result is 0 when the port's internal buffer (if any) is flushed, otherwise the result is a positive exact integer. If the event is not selected in a synchronization, then no bytes will have been written to *out*.

```
(write-special-evt v [out]) → evt?
  v : any/c
  out : output-port? = (current-output-port)
```

Similar to write-special, but instead of writing the special value immediately, it returns a synchronizable event (see §10.2.1 "Events"). The *out* must support atomic writes, as indicated by port-writes-atomic?.

Synchronizing on the object starts a write of the special value, and the event becomes ready when the value is written (unbuffered) to the port. If the event is not selected in a synchronization, then no value will have been written to *out*.

```
(port-writes-atomic? out) → boolean?
  out : output-port?
```

Returns #t if write-bytes-avail/enable-break can provide an exclusive-or guarantee (break or write, but not both) for *out*, and if the port can be used with procedures like write-bytes-avail-evt. Racket's file-stream ports, pipes, string ports, and TCP ports all support atomic writes; ports created with make-output-port (see §12.1.9 "Custom Ports") may support atomic writes.

```
(port-writes-special? out) → boolean?
  out : output-port?
```

Returns #t if procedures like write-special can write arbitrary values to the port. Racket's file-stream ports, pipes, string ports, and TCP ports all reject special values, but ports created with make-output-port (see §12.1.9 "Custom Ports") may support them.

12.4 Reading

```
(read [in]) → any
  in : input-port? = (current-input-port)
```

Reads and returns a single datum from *in*. If *in* has a handler associated to it via port-read-handler, then the handler is called. Otherwise, the default reader is used, as parameterized by the current-readtable parameter, as well as many other parameters.

See §1.3 "The Reader" for information on the default reader.

```
(read-syntax [source-name in]) → (or/c syntax? eof-object?)
  source-name : any/c = (object-name in)
  in : input-port? = (current-input-port)
```

Like read, but produces a syntax object with source-location information. The *source-name* is used as the source field of the syntax object; it can be an arbitrary value, but it should generally be a path for the source file.

See §1.3 "The Reader" for information on the default reader in read-syntax mode.

```
(read/recursive [in start readtable graph?]) → any
  in : input-port? = (current-input-port)
  start : (or/c char? #f) = #f
  readtable : (or/c readtable? #f) = (current-readtable)
  graph? : any/c = #t
```

Similar to calling `read`, but normally used during the dynamic extent of `read` within a reader-extension procedure (see §12.7.2 “Reader-Extension Procedures”). The main effect of using `read/recursive` instead of `read` is that graph-structure annotations (see §1.3.16 “Reading Graph Structure”) in the nested read are considered part of the overall read, at least when the *graph?* argument is true; since the result is wrapped in a placeholder, however, it is not directly inspectable.

If *start* is provided and not `#f`, it is effectively prefixed to the beginning of *in*’s stream for the read. (To prefix multiple characters, use `input-port-append`.)

The *readtable* argument is used for top-level parsing to satisfy the read request; recursive parsing within the read (e.g., to read the elements of a list) instead uses the current readtable as determined by the `current-readtable` parameter. A reader macro might call `read/recursive` with a character and readtable to effectively invoke the readtable’s behavior for the character. If *readtable* is `#f`, the default readtable is used for top-level parsing.

When *graph?* is `#f`, graph structure annotations in the read datum are local to the datum.

When called within the dynamic extent of `read`, the `read/recursive` procedure produces either an opaque placeholder value, a special-comment value, or an end-of-file. The result is a special-comment value (see §12.7.3 “Special Comments”) when the input stream’s first non-whitespace content parses as a comment. The result is end-of-file when `read/recursive` encounters an end-of-file. Otherwise, the result is a placeholder that protects graph references that are not yet resolved. When this placeholder is returned within an S-expression that is produced by any reader-extension procedure (see §12.7.2 “Reader-Extension Procedures”) for the same outermost `read`, it will be replaced with the actual read value before the outermost `read` returns.

See §12.7.1 “Readtables” for an extended example that uses `read/recursive`.

```
(read-syntax/recursive [source-name
                        in
                        start
                        readtable
                        graph?])     → any
  source-name : any/c = (object-name in)
  in : input-port? = (current-input-port)
  start : (or/c char? #f) = #f
  readtable : (or/c readtable? #f) = (current-readtable)
  graph? : any/c = #t
```

Analogous to calling `read/recursive`, but the resulting value encapsulates S-expression structure with source-location information. As with `read/recursive`, when `read-syntax/recursive` is used within the dynamic extent of `read-syntax`, the result from `read-syntax/recursive` is either a special-comment value, end-of-file, or opaque graph-structure placeholder (not a syntax object). The placeholder can be embedded in an S-

expression or syntax object returned by a reader macro, etc., and it will be replaced with the actual syntax object before the outermost `read-syntax` returns.

Using `read/recursive` within the dynamic extent of `read-syntax` does not allow graph structure for reading to be included in the outer `read-syntax` parsing, and neither does using `read-syntax/recursive` within the dynamic extent of `read`. In those cases, `read/recursive` and `read-syntax/recursive` produce results like `read` and `read-syntax`, except that a special-comment value is returned when the input stream starts with a comment (after whitespace).

See §12.7.1 "Readtables" for an extended example that uses `read-syntax/recursive`.

```
(read-language [in fail-thunk]) → (any/c any/c . -> . any)
  in : input-port? = (current-input-port)
  fail-thunk : (-> any) = (lambda () (error ...))
```

Reads from *in* in the same way as `read`, but stopping as soon as a reader language (or its absence) is determined.

A *reader language* is specified by `#lang` or `#!` (see §1.3.17 "Reading via an Extension") at the beginning of the input, though possibly after comment forms. The default readtable is used by `read-language` (instead of the value of `current-readtable`), and `#reader` forms (which might produce comments) are not allowed before `#lang` or `#!`.

See also §17.3.5 "Source-Handling Configuration" in ***The Racket Guide***.

When it finds a `#lang` or `#!` specification, instead of dispatching to a `read` or `read-syntax` function as `read` and `read-syntax` do, `read-language` dispatches to a `get-info` function (if any) exported by the same module. The result of the `get-info` function is the result of `read-language` if it is a function of two arguments; if `get-info` produces any other kind of result, the `exn:fail:contract` exception is raised.

The function produced by `get-info` reflects information about the expected syntax of the input stream. The first argument to the function serves as a key on such information; acceptable keys and the interpretation of results is up to external tools, such as DrRacket (see §20 "`drracket:module-language-tools`"). If no information is available for a given key, the result should be the second argument.

Examples:

```
> ((read-language (open-input-string "#lang algol60")) 'color-
lexer #f)
#<procedure:default-lexer>
> ((read-language (open-input-string "#lang algol60")) 'something-
else #f)
#f
```

The `get-info` function itself is applied to five arguments: the input port being read, the module path from which the `get-info` function was extracted, and the source line (positive

exact integer or #f), column (non-negative exact integer or #f), and position (positive exact integer or #f) of the start of the #lang or #! form. The get-info function may further read from the given input port to determine its result, but it should read no further than necessary. The get-info function should not read from the port after returning a function.

If *in* starts with a reader language specification but the relevant module does not export get-info (but perhaps does export read and read-syntax), then the result of read-language is #f.

If *in* has a #lang or #! specification, but parsing and resolving the specification raises an exception, the exception is propagated by read-language. Having at least #l or #! (after comments and whitespace) counts as starting a #lang or #! specification.

If *in* does not specify a reader language with #lang or #!, then *fail-thunk* is called. The default *fail-thunk* raises exn:fail:read or exn:fail:read:eof.

```
(read-case-sensitive) → boolean?
(read-case-sensitive on?) → void?
  on? : any/c
```

A parameter that controls parsing and printing of symbols. When this parameter's value is #f, the reader case-folds symbols (e.g., producing 'hi when the input is any one of hi, Hi, HI, or hI). The parameter also affects the way that write prints symbols containing uppercase characters; if the parameter's value is #f, then symbols are printed with uppercase characters quoted by a \ or |. The parameter's value is overridden by quoting \ or | vertical-bar quotes and the #cs and #ci prefixes; see §1.3.2 "Reading Symbols" for more information. While a module is loaded, the parameter is set to #t (see current-load).

```
(read-square-bracket-as-paren) → boolean?
(read-square-bracket-as-paren on?) → void?
  on? : any/c
```

A parameter that controls whether [and] are treated as parentheses. See §1.3.5 "Reading Pairs and Lists" for more information.

```
(read-curly-brace-as-paren) → boolean?
(read-curly-brace-as-paren on?) → void?
  on? : any/c
```

A parameter that controls whether { and } are treated as parentheses. See §1.3.5 "Reading Pairs and Lists" for more information.

```
(read-accept-box) → boolean?
(read-accept-box on?) → void?
  on? : any/c
```

A parameter that controls parsing #& input. See §1.3.12 "Reading Boxes" for more information.

```
(read-accept-compiled) → boolean?
(read-accept-compiled on?) → void?
  on? : any/c
```

A parameter that controls parsing `#~` compiled input. See §1.3 “The Reader” and `current-compile` for more information.

```
(read-accept-bar-quote) → boolean?
(read-accept-bar-quote on?) → void?
  on? : any/c
```

A parameter that controls parsing and printing of `|` in symbols. See §1.3.2 “Reading Symbols” and §1.4 “The Printer” for more information.

```
(read-accept-graph) → boolean?
(read-accept-graph on?) → void?
  on? : any/c
```

A parameter value that controls parsing input with sharing. See §1.3.16 “Reading Graph Structure” for more information.

```
(read-decimal-as-inexact) → boolean?
(read-decimal-as-inexact on?) → void?
  on? : any/c
```

A parameter that controls parsing input numbers with a decimal point or exponent (but no explicit exactness tag). See §1.3.3 “Reading Numbers” for more information.

```
(read-accept-dot) → boolean?
(read-accept-dot on?) → void?
  on? : any/c
```

A parameter that controls parsing input with a dot, which is normally used for literal cons cells. See §1.3.5 “Reading Pairs and Lists” for more information.

```
(read-accept-infix-dot) → boolean?
(read-accept-infix-dot on?) → void?
  on? : any/c
```

A parameter that controls parsing input with two dots to trigger infix conversion. See §1.3.5 “Reading Pairs and Lists” for more information.

```
(read-accept-quasiquote) → boolean?
(read-accept-quasiquote on?) → void?
  on? : any/c
```

A parameter that controls parsing input with ` or , which is normally used for `quasiquote`, `unquote`, and `unquote-splicing` abbreviations. See §1.3.7 “Reading Quotes” for more information.

```
(read-accept-reader) → boolean?
(read-accept-reader on?) → void?
  on? : any/c
```

A parameter that controls whether `#reader`, `#lang`, or `#!` are allowed for selecting a parser. See §1.3.17 “Reading via an Extension” for more information.

```
(read-accept-lang) → boolean?
(read-accept-lang on?) → void?
  on? : any/c
```

A parameter that (along with `read-accept-reader` controls whether `#lang` and `#!` are allowed for selecting a parser. See §1.3.17 “Reading via an Extension” for more information.

```
(current-reader-guard) → (any/c . -> . any)
(current-reader-guard proc) → void?
  proc : (any/c . -> . any)
```

A parameter whose value converts or rejects (by raising an exception) a module-path datum following `#reader`. See §1.3.17 “Reading via an Extension” for more information.

```
(current-readtable) → (or/c readtable? #f)
(current-readtable readtable) → void?
  readtable : (or/c readtable? #f)
```

A parameter whose value determines a readtable that adjusts the parsing of S-expression input, where `#f` implies the default behavior. See §12.7.1 “Readtables” for more information.

```
(read-on-demand-source) → (and/c path? complete-path?)
(read-on-demand-source path) → void?
  path : (and/c path? complete-path?)
```

A parameter that enables lazy parsing of compiled code, so that closure bodies and syntax objects are extracted (and validated) from marshaled compiled code on demand. Normally, this parameter is set by the default load handler when `load-on-demand-enabled` is `#t`.

Even when parsing is delayed, compiled code is loaded into memory. If the `PLT_DELAY_FROM_ZO` environment variable is set (to any value) on start-up, however, even loading from disk is delayed. If the file at path changes before the delayed code or syntax object is demanded, the read-on-demand most likely will encounter garbage, leading to an exception.

```
                          (case->
(port-read-handler in) →   (input-port? . -> . any)
                           (input-port?  any/c . -> . any))
  in : input-port?
(port-read-handler in proc) → void?
  in : input-port?
         (case->
  proc :  (input-port? . -> . any)
          (input-port? any/c . -> . any))
```

Gets or sets the *port read handler* for *in*. The handler called to read from the port when the built-in `read` or `read-syntax` procedure is applied to the port. (The port read handler is not used for `read/recursive` or `read-syntax/recursive`.)

A port read handler is applied to either one argument or two arguments:

- A single argument is supplied when the port is used with `read`; the argument is the port being read. The return value is the value that was read from the port (or end-of-file).
- Two arguments are supplied when the port is used with `read-syntax`; the first argument is the port being read, and the second argument is a value indicating the source. The return value is a syntax object that was read from the port (or end-of-file).

The default port read handler reads standard Racket expressions with Racket's built-in parser (see §1.3 "The Reader"). It handles a special result from a custom input port (see `make-custom-input-port`) by treating it as a single expression, except that special-comment values (see §12.7.3 "Special Comments") are treated as whitespace.

The default port read handler itself can be customized through a readtable; see §12.7.1 "Readtables" for more information.

12.5 Writing

```
(write datum [out]) → void?
  datum : any/c
  out : output-port? = (current-output-port)
```

Writes *datum* to *out*, normally in such a way that instances of core datatypes can be read back in. If *out* has a handler associated to it via `port-write-handler`, then the handler is called. Otherwise, the default printer is used (in `write` mode), as configured by various parameters.

See §1.4 “The Printer” for more information about the default printer. In particular, note that `write` may require memory proportional to the depth of the value being printed, due to the initial cycle check.

Examples:

```
> (write 'hi)
hi

> (write (lambda (n) n))
#<procedure>

> (define o (open-output-string))

> (write "hello" o)

> (get-output-string o)
"\"hello\""
```

```
(display datum [out]) → void?
  datum : any/c
  out : output-port? = (current-output-port)
```

Displays *datum* to *out*, similar to `write`, but usually in such a way that byte- and character-based datatypes are written as raw bytes or characters. If *out* has a handler associated to it via `port-display-handler`, then the handler is called. Otherwise, the default printer is used (in `display` mode), as configured by various parameters.

See §1.4 “The Printer” for more information about the default printer. In particular, note that `display` may require memory proportional to the depth of the value being printed, due to the initial cycle check.

```
(print datum [out quote-depth]) → void?
  datum : any/c
  out : output-port? = (current-output-port)
  quote-depth : (or/c 0 1) = 0
```

Writes *datum* to *out*, normally the same way as `write`. If *out* has a handler associated to it via `port-print-handler`, then the handler is called. Otherwise, the handler specified by `global-port-print-handler` is called; the default handler uses the default printer in `write` mode.

The optional *quote-depth* argument adjusts printing when the `print-as-expression` parameter is set to `#t`. In that case, *quote-depth* specifies the starting quote depth for printing *datum*.

The rationale for providing `print` is that `display` and `write` both have specific output conventions, and those conventions restrict the ways that an environment can change the

behavior of `display` and `write` procedures. No output conventions should be assumed for `print`, so that environments are free to modify the actual output generated by `print` in any way.

```
(displayln datum [out]) → void?
  datum : any/c
  out : output-port? = (current-output-port)
```

The same as `(display datum out)` followed by `(newline out)`, which is similar to `println` in Pascal or Java.

```
(fprintf out form v ...) → void?
  out : output-port?
  form : string?
  v : any/c
```

Prints formatted output to *out*, where *form* is a string that is printed directly, except for special formatting escapes:

- `~n` or `~%` prints a newline character (which is equivalent to `\n` in a literal format string)
- `~a` or `~A` displays the next argument among the *v*s
- `~s` or `~S` writes the next argument among the *v*s
- `~v` or `~V` prints the next argument among the *v*s
- `~.`⟨*c*⟩ where ⟨*c*⟩ is `a`, `A`, `s`, `S`, `v`, or `V`: truncates `display`, `write`, or `print` output to `(error-print-width)` characters, using `...` as the last three characters if the untruncated output would be longer
- `~e` or `~E` outputs the next argument among the *v*s using the current error value conversion handler (see `error-value->string-handler`) and current error printing width
- `~c` or `~C` `write-char`s the next argument in *v*s; if the next argument is not a character, the `exn:fail:contract` exception is raised
- `~b` or `~B` prints the next argument among the *v*s in binary; if the next argument is not an exact number, the `exn:fail:contract` exception is raised
- `~o` or `~O` prints the next argument among the *v*s in octal; if the next argument is not an exact number, the `exn:fail:contract` exception is raised
- `~x` or `~X` prints the next argument among the *v*s in hexadecimal; if the next argument is not an exact number, the `exn:fail:contract` exception is raised

- ~~ prints a tilde.
- ~⟨w⟩, where ⟨w⟩ is a whitespace character (see char-whitespace?), skips characters in *form* until a non-whitespace character is encountered or until a second end-of-line is encountered (whichever happens first). On all platforms, an end-of-line can be #\return, #\newline, or #\return followed immediately by #\newline.

The *form* string must not contain any ~ that is not one of the above escapes, otherwise the exn:fail:contract exception is raised. When the format string requires more *v*s than are supplied, the exn:fail:contract exception is raised. Similarly, when more *v*s are supplied than are used by the format string, the exn:fail:contract exception is raised.

Example:

```
> (fprintf (current-output-port)
           "~a as a string is ~s.\n"
           '(3 4)
           "(3 4)")
(3 4) as a string is "(3 4)".
```

```
(printf form v ...) → void?
  form : string?
  v : any/c
```

The same as (fprintf (current-output-port) *form v* ...).

```
(eprintf form v ...) → void?
  form : string?
  v : any/c
```

The same as (fprintf (current-error-port) *form v* ...).

```
(format form v ...) → string?
  form : string?
  v : any/c
```

Formats to a string. The result is the same as

```
(let ([o (open-output-string)])
  (fprintf o form v ...)
  (get-output-string o))
```

Example:

```
> (format "~a as a string is ~s.\n" '(3 4) "(3 4)")
"(3 4) as a string is \"(3 4)\".\n"
```

```
(print-pair-curly-braces) → boolean?
(print-pair-curly-braces on?) → void?
  on? : any/c
```

A parameter that controls pair printing. If the value is true, then pairs print using { and } instead of (and). The default is #f.

```
(print-mpair-curly-braces) → boolean?
(print-mpair-curly-braces on?) → void?
  on? : any/c
```

A parameter that controls pair printing. If the value is true, then mutable pairs print using { and } instead of (and). The default is #t.

```
(print-unreadable) → boolean?
(print-unreadable on?) → void?
  on? : any/c
```

A parameter that enables or disables printing of values that have no readable form (using the default reader), including structures that have a custom-write procedure (see prop:custom-write), but not including uninterned symbols and unreadable symbols (which print the same as interned symbols). If the parameter value is #f, an attempt to print an unreadable value raises exn:fail. The parameter value defaults to #t. See §1.4 “The Printer” for more information.

```
(print-graph) → boolean?
(print-graph on?) → void?
  on? : any/c
```

A parameter that controls printing data with sharing; defaults to #f. See §1.4 “The Printer” for more information.

```
(print-struct) → boolean?
(print-struct on?) → void?
  on? : any/c
```

A parameter that controls printing structure values in vector or prefab form; defaults to #t. See §1.4 “The Printer” for more information. This parameter has no effect on the printing of structures that have a custom-write procedure (see prop:custom-write).

```
(print-box) → boolean?
(print-box on?) → void?
  on? : any/c
```

A parameter that controls printing box values; defaults to `#t`. See §1.4.9 “Printing Boxes” for more information.

```
(print-vector-length) → boolean?
(print-vector-length on?) → void?
  on? : any/c
```

A parameter that controls printing vectors; defaults to `#f`. See §1.4.6 “Printing Vectors” for more information.

```
(print-hash-table) → boolean?
(print-hash-table on?) → void?
  on? : any/c
```

A parameter that controls printing hash tables; defaults to `#f`. See §1.4.8 “Printing Hash Tables” for more information.

```
(print-boolean-long-form) → boolean?
(print-boolean-long-form on?) → void?
  on? : any/c
```

A parameter that controls printing of booleans. When the parameter’s value is true, `#t` and `#f` print as `#true` and `#false`, otherwise they print as `#t` and `#f`. The default is `#f`.

```
(print-reader-abbreviations) → boolean?
(print-reader-abbreviations on?) → void?
  on? : any/c
```

A parameter that controls printing of two-element lists that start with `quote`, `'quasiquote`, `'unquote`, `'unquote-splicing`, `'syntax`, `'quasisyntax`, `'unsyntax`, or `'unsyntax-splicing`; defaults to `#f`. See §1.4.4 “Printing Pairs and Lists” for more information.

```
(print-as-expression) → boolean?
(print-as-expression on?) → void?
  on? : any/c
```

A parameter that controls printing in `print` mode (as opposed to `write` or `display`); defaults to `#t`. See §1.4 “The Printer” for more information.

```
(print-syntax-width)
 → (or/c +inf.0 0 (and/c exact-integer? (>/c 3)))
(print-syntax-width width) → void?
  width : (or/c +inf.0 0 (and/c exact-integer? (>/c 3)))
```

A parameter that controls printing of syntax objects. Up to *width* characters are used to show the datum form of a syntax object within `#<syntax`...`>` (after the syntax object’s source location, if any).

```
(current-write-relative-directory)
    (or/c (and/c path? complete-path?)
          (cons/c (and/c path? complete-path?)
 →                (and/c path? complete-path?))
          #f)
(current-write-relative-directory path) → void?
         (or/c (and/c path-string? complete-path?)
               (cons/c (and/c path-string? complete-path?)
  path :                (and/c path-string? complete-path?))
               #f)
```

A parameter that is used when writing compiled code (see §1.4.15 “Printing Compiled Code”) that contains pathname literals, including source-location pathnames for procedure names. When the parameter’s value is a *path*, paths that syntactically extend *path* are converted to relative paths; when the resulting compiled code is read, relative paths are converted back to complete paths using the current-load-relative-directory parameter (if it is not #f; otherwise, the path is left relative). When the parameter’s value is (cons *rel-to-path base-path*), then paths that syntactically extend *base-path* are converted as relative to *rel-to-path*; the *rel-to-path* must extend *base-path*, in which case 'up path elements (in the sense of build-path) may be used to make a path relative to *rel-to-path*.

```
(port-write-handler out) → (any/c output-port? . -> . any)
  out : output-port?
(port-write-handler out proc) → void?
  out : output-port?
  proc : (any/c output-port? . -> . any)
```

```
(port-display-handler out) → (any/c output-port? . -> . any)
  out : output-port?
(port-display-handler out proc) → void?
  out : output-port?
  proc : (any/c output-port? . -> . any)
```

```
(port-print-handler out)
 → ((any/c output-port?) ((or/c 0 1)) . ->* . any)
  out : output-port?
(port-print-handler out proc) → void?
  out : output-port?
  proc : (any/c output-port? . -> . any)
```

Gets or sets the *port write handler*, *port display handler*, or *port print handler* for out. This handler is called to output to the port when write, display, or print (respectively) is

applied to the port. Each handler must accept two arguments: the value to be printed and the destination port. The handler's return value is ignored.

A port print handler optionally accepts a third argument, which corresponds to the optional third argument to `print`; if a procedure given to `port-print-handler` does not accept a third argument, it is wrapped with a procedure that discards the optional third argument.

The default port display and write handlers print Racket expressions with Racket's built-in printer (see §1.4 "The Printer"). The default print handler calls the global port print handler (the value of the `global-port-print-handler` parameter); the default global port print handler is the same as the default write handler.

```
(global-port-print-handler)
 → (->* (any/c output-port?) ((or/c 0 1)) any)
(global-port-print-handler proc) → void?
  proc : (or/c (->* (any/c output-port?) ((or/c 0 1)) any)
               (any/c output-port? . -> . any))
```

A parameter that determines *global port print handler*, which is called by the default port print handler (see `port-print-handler`) to `print` values into a port. The default value uses the built-in printer (see §1.4 "The Printer") in `print` mode.

A global port print handler optionally accepts a third argument, which corresponds to the optional third argument to `print`. If a procedure given to `global-port-print-handler` does not accept a third argument, it is wrapped with a procedure that discards the optional third argument.

12.6 Pretty Printing

```
(require racket/pretty)
```

The bindings documented in this section are provided by the `racket/pretty` and `racket` libraries, but not `racket/base`.

```
(pretty-print v [port quote-depth]) → void?
  v : any/c
  port : output-port? = (current-output-port)
  quote-depth : (or/c 0 1) = 0
```

Pretty-prints the value *v* using the same printed form as the default `print` mode, but with newlines and whitespace inserted to avoid lines longer than `(pretty-print-columns)`, as controlled by `(pretty-print-current-style-table)`. The printed form ends in a newline, unless the `pretty-print-columns` parameter is set to `'infinity`. When *port* has line counting enabled (see §12.1.4 "Counting Positions, Lines, and Columns"), then

printing is sensitive to the column when printing starts—both for determining an initial line break and indenting subsequent lines.

In addition to the parameters defined in this section, `pretty-print` conforms to the `print-graph`, `print-struct`, `print-hash-table`, `print-vector-length`, `print-box`, and `print-as-expression` parameters.

The pretty printer detects structures that have the `prop:custom-write` property and calls the corresponding custom-write procedure. The custom-write procedure can check the parameter `pretty-printing` to cooperate with the pretty-printer. Recursive printing to the port automatically uses pretty printing, but if the structure has multiple recursively printed sub-expressions, a custom-write procedure may need to cooperate more to insert explicit newlines. Use `port-next-location` to determine the current output column, use `pretty-print-columns` to determine the target printing width, and use `pretty-print-newline` to insert a newline (so that the function in the `pretty-print-print-line` parameter can be called appropriately). Use `make-tentative-pretty-print-output-port` to obtain a port for tentative recursive prints (e.g., to check the length of the output).

```
(pretty-write v [port]) → void?
  v : any/c
  port : output-port? = (current-output-port)
```

Same as `pretty-print`, but `v` is printed like `write` instead of like `print`.

```
(pretty-display v [port]) → void?
  v : any/c
  port : output-port? = (current-output-port)
```

Same as `pretty-print`, but `v` is printed like `display` instead of like `print`.

```
(pretty-format v [columns]) → string?
  v : any/c
  columns : exact-nonnegative-integer? = (pretty-print-columns)
```

Like `pretty-print`, except that it returns a string containing the pretty-printed value, rather than sending the output to a port.

The optional argument *columns* argument is used to parameterize `pretty-print-columns`.

```
(pretty-print-handler v) → void?
  v : any/c
```

Pretty-prints `v` if `v` is not `#<void>`, or prints nothing if `v` is `#<void>`. Pass this procedure to `current-print` to install the pretty printer into the REPL run by `read-eval-print-loop`.

12.6.1 Basic Pretty-Print Options

```
(pretty-print-columns)
 → (or/c exact-positive-integer? 'infinity)
(pretty-print-columns width) → void?
  width : (or/c exact-positive-integer? 'infinity)
```

A parameter that determines the default width for pretty printing.

If the display width is `'infinity`, then pretty-printed output is never broken into lines, and a newline is not added to the end of the output.

```
(pretty-print-depth) → (or/c exact-nonnegative-integer? #f)
(pretty-print-depth depth) → void?
  depth : (or/c exact-nonnegative-integer? #f)
```

Parameter that controls the default depth for recursive pretty printing. Printing to *depth* means that elements nested more deeply than *depth* are replaced with "..."; in particular, a depth of `0` indicates that only simple values are printed. A depth of `#f` (the default) allows printing to arbitrary depths.

```
(pretty-print-exact-as-decimal) → boolean?
(pretty-print-exact-as-decimal as-decimal?) → void?
  as-decimal? : any/c
```

A parameter that determines how exact non-integers are printed. If the parameter's value is `#t`, then an exact non-integer with a decimal representation is printed as a decimal number instead of a fraction. The initial value is `#f`.

```
(pretty-print-.-symbol-without-bars) → boolean?
(pretty-print-.-symbol-without-bars on?) → void?
  on? : any/c
```

A parameter that controls the printing of the symbol whose print name is just a period. If set to a true value, then such a symbol is printed as only the period. If set to a false value, it is printed as a period with vertical bars surrounding it.

```
(pretty-print-show-inexactness) → boolean?
(pretty-print-show-inexactness show?) → void?
  show? : any/c
```

A parameter that determines how inexact numbers are printed. If the parameter's value is `#t`, then inexact numbers are always printed with a leading `#i`. The initial value is `#f`.

12.6.2 Per-Symbol Special Printing

```
(pretty-print-abbreviate-read-macros) → boolean?
(pretty-print-abbreviate-read-macros abbrev?) → void?
  abbrev? : any/c
```

A parameter that controls whether or not `quote`, `unquote`, `unquote-splicing`, etc., are abbreviated with `'`, `,`, `,@`, etc. By default, the abbreviations are enabled.

See also `pretty-print-remap-stylable`.

```
(pretty-print-style-table? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a style table for use with `pretty-print-current-style-table`, `#f` otherwise.

```
(pretty-print-current-style-table) → pretty-print-style-table?
(pretty-print-current-style-table style-table) → void?
  style-table : pretty-print-style-table?
```

A parameter that holds a table of style mappings. See `pretty-print-extend-style-table`.

```
(pretty-print-extend-style-table style-table
                                 symbol-list
                                 like-symbol-list)
 → pretty-print-style-table?
  style-table : pretty-print-style-table?
  symbol-list : (listof symbol?)
  like-symbol-list : (listof symbol?)
```

Creates a new style table by extending an existing *`style-table`*, so that the style mapping for each symbol of *`like-symbol-list`* in the original table is used for the corresponding symbol of *`symbol-list`* in the new table. The *`symbol-list`* and *`like-symbol-list`* lists must have the same length. The *`style-table`* argument can be `#f`, in which case the default mappings are used from the original table (see below).

The style mapping for a symbol controls the way that whitespace is inserted when printing a list that starts with the symbol. In the absence of any mapping, when a list is broken across multiple lines, each element of the list is printed on its own line, each with the same indentation.

The default style mapping includes mappings for the following symbols, so that the output follows popular code-formatting rules:

```
'lambda 'case-lambda
'define 'define-macro 'define-syntax
'let 'letrec 'let*
'let-syntax 'letrec-syntax
'let-values 'letrec-values 'let*-values
'let-syntaxes 'letrec-syntaxes
'begin 'begin0 'do
'if 'set! 'set!-values
'unless 'when
'cond 'case 'and 'or
'module
'syntax-rules 'syntax-case 'letrec-syntaxes+values
'import 'export 'link
'require 'require-for-syntax 'require-for-template 'provide
'public 'private 'override 'rename 'inherit 'field 'init
'shared 'send 'class 'instantiate 'make-object
```

```
(pretty-print-remap-stylable)
 → (any/c . -> . (or/c symbol? #f))
(pretty-print-remap-stylable proc) → void?
  proc : (any/c . -> . (or/c symbol? #f))
```

A parameter that controls remapping for styles and for the determination of the reader shorthands.

This procedure is called with each sub-expression that appears as the first element in a sequence. If it returns a symbol, the style table is used, as if that symbol were at the head of the sequence. If it returns `#f`, the style table is treated normally. Similarly, when determining whether to abbreviate reader macros, this parameter is consulted.

12.6.3 Line-Output Hook

```
(pretty-print-newline port width) → void?
  port : out-port?
  width : exact-nonnegative-integer?
```

Calls the procedure associated with the `pretty-print-print-line` parameter to print a newline to *port*, if *port* is the output port that is redirected to the original output port for printing, otherwise a plain newline is printed to *port*. The *width* argument should be the target column width, typically obtained from `pretty-print-columns`.

```
(pretty-print-print-line)
```

```
   ((or/c exact-nonnegative-integer? #f)
    output-port?
    exact-nonnegative-integer?
→   (or/c exact-nonnegative-integer? 'infinity)
    . -> .
    exact-nonnegative-integer?)
(pretty-print-print-line proc) → void?
        ((or/c exact-nonnegative-integer? #f)
         output-port?
         exact-nonnegative-integer?
 proc :  (or/c exact-nonnegative-integer? 'infinity)
         . -> .
         exact-nonnegative-integer?)
```

A parameter that determines a procedure for printing the newline separator between lines of a pretty-printed value. The procedure is called with four arguments: a new line number, an output port, the old line's length, and the number of destination columns. The return value from *proc* is the number of extra characters it printed at the beginning of the new line.

The *proc* procedure is called before any characters are printed with `0` as the line number and `0` as the old line length; *proc* is called after the last character of a value has been printed with `#f` as the line number and with the length of the last line. Whenever the pretty-printer starts a new line, *proc* is called with the new line's number (where the first new line is numbered `1`) and the just-finished line's length. The destination-columns argument to *proc* is always the total width of the destination printing area, or `'infinity` if pretty-printed values are not broken into lines.

The default *proc* procedure prints a newline whenever the line number is not `0` and the column count is not `'infinity`, always returning `0`. A custom *proc* procedure can be used to print extra text before each line of pretty-printed output; the number of characters printed before each line should be returned by *proc* so that the next line break can be chosen correctly.

The destination port supplied to *proc* is generally not the port supplied to `pretty-print` or `pretty-display` (or the current output port), but output to this port is ultimately redirected to the port supplied to `pretty-print` or `pretty-display`.

12.6.4 Value Output Hook

```
(pretty-print-size-hook)
    (any/c boolean? output-port?
→    . -> .
     (or/c #f exact-nonnegative-integer?))
(pretty-print-size-hook proc) → void?
```

```
        (any/c boolean? output-port?
 proc :  . -> .
         (or/c #f exact-nonnegative-integer?))
```

A parameter that determines a sizing hook for pretty-printing.

The sizing hook is applied to each value to be printed. If the hook returns `#f`, then printing is handled internally by the pretty-printer. Otherwise, the value should be an integer specifying the length of the printed value in characters; the print hook will be called to actually print the value (see `pretty-print-print-hook`).

The sizing hook receives three arguments. The first argument is the value to print. The second argument is a boolean: `#t` for printing like `display` and `#f` for printing like `write`. The third argument is the destination port; the port is the one supplied to `pretty-print` or `pretty-display` (or the current output port). The sizing hook may be applied to a single value multiple times during pretty-printing.

```
(pretty-print-print-hook)
 → (any/c boolean? output-port? . -> . void?)
(pretty-print-print-hook proc) → void?
  proc : (any/c boolean? output-port? . -> . void?)
```

A parameter that determines a print hook for pretty-printing. The print-hook procedure is applied to a value for printing when the sizing hook (see `pretty-print-size-hook`) returns an integer size for the value.

The print hook receives three arguments. The first argument is the value to print. The second argument is a boolean: `#t` for printing like `display` and `#f` for printing like `write`. The third argument is the destination port; this port is generally not the port supplied to `pretty-print` or `pretty-display` (or the current output port), but output to this port is ultimately redirected to the port supplied to `pretty-print` or `pretty-display`.

```
(pretty-print-pre-print-hook)
 → (any/c output-port? . -> . void)
(pretty-print-pre-print-hook proc) → void?
  proc : (any/c output-port? . -> . void)
```

A parameter that determines a hook procedure to be called just before an object is printed. The hook receives two arguments: the object and the output port. The port is the one supplied to `pretty-print` or `pretty-display` (or the current output port).

```
(pretty-print-post-print-hook)
 → (any/c output-port? . -> . void)
(pretty-print-post-print-hook proc) → void?
  proc : (any/c output-port? . -> . void)
```

A parameter that determines a hook procedure to be called just after an object is printed. The hook receives two arguments: the object and the output port. The port is the one supplied to `pretty-print` or `pretty-display` (or the current output port).

12.6.5 Additional Custom-Output Support

```
(pretty-printing) → boolean?
(pretty-printing on?) → void?
  on? : any/c
```

A parameter that is set to `#t` when the pretty printer calls a custom-write procedure (see `prop:custom-write`) for output in a mode that supports line breaks. When pretty printer calls a custom-write procedure merely to detect cycles or to try to print on a single line, it sets this parameter to `#f`.

```
(make-tentative-pretty-print-output-port out
                                         width
                                         overflow-thunk)
 → output-port?
  out : output-port?
  width : exact-nonnegative-integer?
  overflow-thunk : (-> any)
```

Produces an output port that is suitable for recursive pretty printing without actually producing output. Use such a port to tentatively print when proper output depends on the size of recursive prints. After printing, determine the size of the tentative output using `file-position`.

The *out* argument should be a pretty-printing port, such as the one supplied to a custom-write procedure when `pretty-printing` is set to true, or another tentative output port. The *width* argument should be a target column width, usually obtained from `pretty-print-columns`, possibly decremented to leave room for a terminator. The *overflow-thunk* procedure is called if more than *width* items are printed to the port or if a newline is printed to the port via `pretty-print-newline`; it can escape from the recursive print through a continuation as a shortcut, but *overflow-thunk* can also return, in which case it is called every time afterward that additional output is written to the port.

After tentative printing, either accept the result with `tentative-pretty-print-port-transfer` or reject it with `tentative-pretty-print-port-cancel`. Failure to accept or cancel properly interferes with graph-structure printing, calls to hook procedures, etc. Explicitly cancel the tentative print even when *overflow-thunk* escapes from a recursive print.

```
(tentative-pretty-print-port-transfer tentative-out
                                      orig-out)     → void?
```

```
  tentative-out : output-port?
  orig-out : output-port?
```

Causes the data written to *tentative-out* to be transferred as if written to *orig-out*. The *tentative-out* argument should be a port produced by `make-tentative-pretty-print-output-port`, and *orig-out* should be either a pretty-printing port (provided to a custom-write procedure) or another tentative output port.

```
(tentative-pretty-print-port-cancel tentative-out) → void?
  tentative-out : output-port?
```

Cancels the content of *tentative-out*, which was produced by `make-tentative-pretty-print-output-port`. The main effect of canceling is that graph-reference definitions are undone, so that a future print of a graph-referenced object includes the defining `#`⟨*n*⟩`=`.

12.7 Reader Extension

Racket's reader can be extended in three ways: through a reader-macro procedure in a readtable (see §12.7.1 "Readtables"), through a `#reader` form (see §1.3.17 "Reading via an Extension"), or through a custom-port byte reader that returns a "special" result procedure (see §12.1.9 "Custom Ports"). All three kinds of *reader extension procedures* accept similar arguments, and their results are treated in the same way by `read` and `read-syntax` (or, more precisely, by the default read handler; see `port-read-handler`).

12.7.1 Readtables

The dispatch table in §1.3.1 "Delimiters and Dispatch" corresponds to the default *readtable*. By creating a new readtable and installing it via the `current-readtable` parameter, the reader's behavior can be extended.

A readtable is consulted at specific times by the reader:

- when looking for the start of a datum;
- when determining how to parse a datum that starts with `#`;
- when looking for a delimiter to terminate a symbol or number;
- when looking for an opener (such as `(`), closer (such as `)`), or `.` after the first character parsed as a sequence for a pair, list, vector, or hash table; or
- when looking for an opener after `#`⟨*n*⟩ in a vector of specified length ⟨*n*⟩.

The readtable is ignored at other times. In particular, after parsing a character that is mapped to the default behavior of `;`, the readtable is ignored until the comment's terminating newline is discovered. Similarly, the readtable does not affect string parsing until a closing double-quote is found. Meanwhile, if a character is mapped to the default behavior of `(`, then it starts sequence that is closed by any character that is mapped to a closing parenthesis `)`. An apparent exception is that the default parsing of `|` quotes a symbol until a matching character is found, but the parser is simply using the character that started the quote; it does not consult the readtable.

For many contexts, `#f` identifies the default readtable. In particular, `#f` is the initial value for the `current-readtable` parameter, which causes the reader to behave as described in §1.3 "The Reader".

```
(readtable? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a readtable, `#f` otherwise.

```
(make-readtable readtable
                key
                mode
                action ...+) → readtable?
  readtable : readtable?
  key : (or/c character? #f)
        (or/c (or/c 'terminating-macro
                    'non-terminating-macro
  mode :                'dispatch-macro)
              character?)
         (or/c procedure?
  action :      readtable?)
```

Creates a new readtable that is like *readtable* (which can be `#f`), except that the reader's behavior is modified for each *key* according to the given *mode* and *action*. The `...+` for `make-readtable` applies to all three of *key*, *mode*, and *action*; in other words, the total number of arguments to `make-readtable` must be 1 modulo 3.

The possible combinations for *key*, *mode*, and *action* are as follows:

- *char* `'terminating-macro` *proc* — causes *char* to be parsed as a delimiter, and an unquoted/uncommented *char* in the input string triggers a call to the *reader macro* *proc*; the activity of *proc* is described further below. Conceptually, characters like `;`, `(`, and `)` are mapped to terminating reader macros in the default readtable.
- *char* `'non-terminating-macro` *proc* — like the `'terminating-macro` variant, but *char* is not treated as a delimiter, so it can be used in the middle of an identifier or number. Conceptually, `#` is mapped to a non-terminating macro in the default readtable.

- *char* 'dispatch-macro *proc* — like the 'non-terminating-macro variant, but for *char* only when it follows a # (or, more precisely, when the character follows one that has been mapped to the behavior of #hash in the default readtable).

- *char like-char readtable* — causes *char* to be parsed in the same way that *like-char* is parsed in *readtable*, where *readtable* can be #f to indicate the default readtable. Mapping a character to the same actions as | in the default reader means that the character starts quoting for symbols, and the same character terminates the quote; in contrast, mapping a character to the same action as a " means that the character starts a string, but the string is still terminated with a closing ". Finally, mapping a character to an action in the default readtable means that the character's behavior is sensitive to parameters that affect the original character; for example, mapping a character to the same action as a curly brace { in the default readtable means that the character is disallowed when the read-curly-brace-as-paren parameter is set to #f.

- #f 'non-terminating-macro *proc* — replaces the macro used to parse characters with no specific mapping: i.e., the characters (other than # or |) that can start a symbol or number with the default readtable.

If multiple 'dispatch-macro mappings are provided for a single *char*, all but the last one are ignored. Similarly, if multiple non-'dispatch-macro mappings are provided for a single *char*, all but the last one are ignored.

A reader macro *proc* must accept six arguments, and it can optionally accept two arguments. The first two arguments are always the character that triggered the reader macro and the input port for reading. When the reader macro is triggered by read-syntax (or read-syntax/recursive), the procedure is passed four additional arguments that represent a source location. When the reader macro is triggered by read (or read/recursive), the procedure is passed only two arguments if it accepts two arguments, otherwise it is passed six arguments where the last four are all #f. See §12.7.2 "Reader-Extension Procedures" for information on the procedure's results.

A reader macro normally reads characters from the given input port to produce a value to be used as the "reader macro-expansion" of the consumed characters. The reader macro might produce a special-comment value (see §12.7.3 "Special Comments") to cause the consumed character to be treated as whitespace, and it might use read/recursive or read-syntax/recursive.

```
(readtable-mapping readtable char)
   (or/c character?
         (or/c 'terminating-macro
→              'non-terminating-macro))
   (or/c #f procedure?)
   (or/c #f procedure?)
 readtable : readtable?
 char : character?
```

Produces information about the mappings in *readtable* for *char*. The result is three values:

- either a character (mapping to same behavior as the character in the default readtable), `'terminating-macro`, or `'non-terminating-macro`; this result reports the main (i.e., non-`'dispatch-macro`) mapping for `key`. When the result is a character, then `key` is mapped to the same behavior as the returned character in the default readtable.
- either `#f` or a reader-macro procedure; the result is a procedure when the first result is `'terminating-macro` or `'non-terminating-macro`.
- either `#f` or a reader-macro procedure; the result is a procedure when the character has a `'dispatch-macro` mapping in *readtable* to override the default dispatch behavior.

Note that reader-macro procedures for the default readtable are not directly accessible. To invoke default behaviors, use `read/recursive` or `read-syntax/recursive` with a character and the `#f` readtable.

Examples:

```
; Provides raise-read-error and raise-read-eof-error
> (require syntax/readerr)

> (define (skip-whitespace port)
    ; Skips whitespace characters, sensitive to the current
    ; readtable's definition of whitespace
    (let ([ch (peek-char port)])
      (unless (eof-object? ch)
        ; Consult current readtable:
        (let-values ([(like-ch/sym proc dispatch-proc)
                      (readtable-mapping (current-readtable) ch)])
          ; If like-ch/sym is whitespace, then ch is whitespace
          (when (and (char? like-ch/sym)
                     (char-whitespace? like-ch/sym))
            (read-char port)
            (skip-whitespace port))))))

> (define (skip-comments read-one port src)
    ; Recursive read, but skip comments and detect EOF
    (let loop ()
      (let ([v (read-one)])
        (cond
         [(special-comment? v) (loop)]
         [(eof-object? v)
          (let-values ([(l c p) (port-next-location port)])
```

```
            (raise-read-eof-error
             "unexpected EOF in tuple" src l c p 1))]
          [else v])))))

> (define (parse port read-one src)
    ; First, check for empty tuple
    (skip-whitespace port)
    (if (eq? #\> (peek-char port))
        null
        (let ([elem (read-one)])
          (if (special-comment? elem)
              ; Found a comment, so look for > again
              (parse port read-one src)
              ; Non-empty tuple:
              (cons elem
                    (parse-nonempty port read-one src))))))

> (define (parse-nonempty port read-one src)
    ; Need a comma or closer
    (skip-whitespace port)
    (case (peek-char port)
      [(#\>) (read-char port)
       ; Done
       null]
      [(#\,) (read-char port)
       ; Read next element and recur
       (cons (skip-comments read-one port src)
             (parse-nonempty port read-one src))]
      [else
       ; Either a comment or an error; grab location (in case
       ; of error) and read recursively to detect comments
       (let-values ([(l c p) (port-next-location port)]
                    [(v) (read-one)])
         (cond
          [(special-comment? v)
           ; It was a comment, so try again
           (parse-nonempty port read-one src)]
          [else
           ; Wasn't a comment, comma, or closer; error
           ((if (eof-object? v)
                raise-read-eof-error
                raise-read-error)
            "expected `,' or `>'" src l c p 1)]))]))

> (define (make-delims-table)
    ; Table to use for recursive reads to disallow delimiters
```

```
    ;  (except those in sub-expressions)
    (letrec ([misplaced-delimiter
              (case-lambda
               [(ch port) (unexpected-delimiter ch port #f #f #f #f)]
               [(ch port src line col pos)
                (raise-read-error
                 (format "misplaced ‘~a’ in tuple" ch)
                 src line col pos 1)])])
      (make-readtable (current-readtable)
                      #\, 'terminating-macro misplaced-delimiter
                      #\> 'terminating-macro misplaced-
delimiter)))

> (define (wrap l)
    `(make-tuple (list ,@l)))

> (define parse-open-tuple
    (case-lambda
     [(ch port)
      ; ‘read’ mode
      (wrap (parse port
                   (lambda ()
                     (read/recursive port #f
                                     (make-delims-table)))
                   (object-name port)))]
     [(ch port src line col pos)
      ; ‘read-syntax’ mode
      (datum->syntax
       #f
       (wrap (parse port
                    (lambda ()
                      (read-syntax/recursive src port #f
                                             (make-delims-table)))
                    src))
       (let-values ([(l c p) (port-next-location port)])
         (list src line col pos (and pos (- p pos))))))]))

> (define tuple-readtable
    (make-readtable #f #\< 'terminating-macro parse-open-tuple))

> (parameterize ([current-readtable tuple-readtable])
    (read (open-input-string "<1 , 2 , \"a\">")))
'(make-tuple (list 1 2 "a"))
> (parameterize ([current-readtable tuple-readtable])
    (read (open-input-string
           "< #|# 1 #|# , #|# 2 #|# , #|# \"a\" #|# >")))
```

```
'(make-tuple (list 1 2 "a"))
> (define tuple-readtable+
    (make-readtable tuple-readtable
                    #\* 'terminating-macro (lambda a
                                             (make-special-
comment #f))
                    #\_ #\space #f))

> (parameterize ([current-readtable tuple-readtable+])
    (read (open-input-string "< * 1 __,__  2 __,__ * \"a\" * >")))
'(make-tuple (list 1 2 "a"))
```

12.7.2 Reader-Extension Procedures

Calls to reader extension procedures can be triggered through `read`, `read/recursive`, or `read-syntax`. In addition, a special-read procedure can be triggered by calls to `read-char-or-special`, or by the context of `read-bytes-avail!`, `peek-bytes-avail!*`, `read-bytes-avail!`, and `peek-bytes-avail!*`.

Optional arities for reader-macro and special-result procedures allow them to distinguish reads via `read`, etc., from reads via `read-syntax`, etc. (where the source value is `#f` and no other location information is available).

When a reader-extension procedure is called in syntax-reading mode (via `read-syntax`, etc.), it should generally return a syntax object that has no lexical context (e.g., a syntax object created using `datum->syntax` with `#f` as the first argument and with the given location information as the third argument). Another possible result is a special-comment value (see §12.7.3 “Special Comments”). If the procedure’s result is not a syntax object and not a special-comment value, it is converted to one using `datum->syntax`.

When a reader-extension procedure is called in non-syntax-reading modes, it should generally not return a syntax object. If a syntax object is returned, it is converted to a plain value using `syntax->datum`.

In either context, when the result from a reader-extension procedure is a special-comment value (see §12.7.3 “Special Comments”), then `read`, `read-syntax`, etc. treat the value as a delimiting comment and otherwise ignore it.

Also, in either context, the result may be copied to prevent mutation to vectors or boxes before the read result is completed, and to support the construction of graphs with cycles. Mutable boxes, vectors, and prefab structures are copied, along with any pairs, boxes, vectors, prefab structures that lead to such mutable values, to placeholders produced by a recursive read (see `read/recursive`), or to references of a shared value. Graph structure (including cycles) is preserved in the copy.

12.7.3 Special Comments

```
(make-special-comment v) → special-comment?
  v : any/c
```

Creates a special-comment value that encapsulates `v`. The `read`, `read-syntax`, etc., procedures treat values constructed with `make-special-comment` as delimiting whitespace when returned by a reader-extension procedure (see §12.7.2 “Reader-Extension Procedures”).

```
(special-comment? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is the result of `make-special-comment`, `#f` otherwise.

```
(special-comment-value sc) → any
  sc : special-comment?
```

Returns the value encapsulated by the special-comment value `sc`. This value is never used directly by a reader, but it might be used by the context of a `read-char-or-special`, etc., call that detects a special comment.

12.8 Printer Extension

```
gen:custom-write : any/c
```

A generic interface (see §4.4 “Generic Interfaces”) that supplies a method, `write-proc` used by the default printer to `display`, `write`, or `print` instances of the structure type.

A `write-proc` method takes three arguments: the structure to be printed, the target port, and an argument that is `#t` for `write` mode, `#f` for `display` mode, or `0` or `1` indicating the current quoting depth for `print` mode. The procedure should print the value to the given port using `write`, `display`, `print`, `fprintf`, `write-special`, etc.

The port write handler, port display handler, and print handler are specially configured for a port given to a custom-write procedure. Printing to the port through `display`, `write`, or `print` prints a value recursively with sharing annotations. To avoid a recursive print (i.e., to print without regard to sharing with a value currently being printed), print instead to a string or pipe and transfer the result to the target port using `write-string` or `write-special`. To print recursively to a port other than the one given to the custom-write procedure, copy the given port’s write handler, display handler, and print handler to the other port.

The port given to a custom-write handler is not necessarily the actual target port. In particular, to detect cycles and sharing, the printer invokes a custom-write procedure with a port that records recursive prints, and does not retain any other output.

Recursive print operations may trigger an escape from the call to the custom-write procedure (e.g., for pretty-printing where a tentative print attempt overflows the line, or for printing error output of a limited width).

The following example definition of a `tuple` type includes custom-write procedures that print the tuple's list content using angle brackets in `write` and `print` mode and no brackets in `display` mode. Elements of the tuple are printed recursively, so that graph and cycle structure can be represented.

Examples:

```
(define (tuple-print tuple port mode)
  (when mode (write-string "<" port))
  (let ([l (tuple-ref tuple)]
        [recur (case mode
                 [(#t) write]
                 [(#f) display]
                 [else (lambda (p port) (print p port mode))])])
    (unless (zero? (vector-length l))
      (recur (vector-ref l 0) port)
      (for-each (lambda (e)
                  (write-string ", " port)
                  (recur e port))
                (cdr (vector->list l)))))
  (when mode (write-string ">" port)))

> (struct tuple (ref)
          #:methods gen:custom-write
          [(define write-proc tuple-print)])

> (display (tuple #(1 2 "a")))
1, 2, a

> (print (tuple #(1 2 "a")))
<1, 2, "a">

> (let ([t (tuple (vector 1 2 "a"))])
    (vector-set! (tuple-ref t) 0 t)
    (write t))
#0=<#0#, 2, "a">
```

`prop:custom-write` : `struct-type-property?`

A deprecated structure type property (see §4.3 "Structure Type Properties") that supplies a procedure that corresponds to `gen:custom-write`'s `write-proc`. `gen:custom-write` should be used instead.

```
(custom-write? v) → boolean?
  v : any/c
```

Returns `#t` if `v` has the `prop:custom-write` property, `#f` otherwise.

```
(custom-write-accessor v)
 → (custom-write? output-port? boolean? . -> . any)
  v : custom-write?
```

Returns the custom-write procedure associated with `v`.

```
prop:custom-print-quotable : struct-type-property?
custom-print-quotable? : struct-type-property?
custom-print-quotable-accessor : struct-type-property?
```

A property and associated predicate and accessor. The property value is one of `'self`, `'never`, `'maybe`, or `'always`. When a structure has this property in addition to a `prop:custom-write` property value, then the property value affects printing in `print` mode; see §1.4 "The Printer". When a value does not have the `prop:custom-print-quotable`, it is equivalent to having the `'self` property value, which is suitable both for self-quoting forms and printed forms that are unreadable.

12.9 Serialization

```
(require racket/serialize)
```

The bindings documented in this section are provided by the `racket/serialize` library, not `racket/base` or `racket`.

```
(serializable? v) → boolean?
  v : any/c
```

Returns `#t` if `v` appears to be serializable, without checking the content of compound values, and `#f` otherwise. See `serialize` for an enumeration of serializable values.

```
(serialize v) → any
  v : serializable?
```

Returns a value that encapsulates the value `v`. This value includes only readable values, so it can be written to a stream with `write` or `s-exp->fasl`, later read from a stream using `read` or `fasl->s-exp`, and then converted to a value like the original using `deserialize`. Serialization followed by deserialization produces a value with the same graph structure and mutability as the original value, but the serialized value is a plain tree (i.e., no sharing).

The following kinds of values are serializable:

- structures created through `serializable-struct` or `serializable-struct/versions`, or more generally structures with the `prop:serializable` property (see `prop:serializable` for more information);
- prefab structures;
- instances of classes defined with `define-serializable-class` or `define-serializable-class*`;
- booleans, numbers, characters, interned symbols, unreadable symbols, strings, byte strings, paths (for a specific convention), `#<void>`, and the empty list;
- pairs, mutable pairs, vectors, boxes, hash tables, and sets;
- `date`, `date*`, and `arity-at-least` structures; and
- module path index values.

Serialization succeeds for a compound value, such as a pair, only if all content of the value is serializable. If a value given to `serialize` is not completely serializable, the `exn:fail:contract` exception is raised.

If *v* contains a cycle (i.e., a collection of objects that are all reachable from each other), then *v* can be serialized only if the cycle includes a mutable value, where a prefab structure counts as mutable only if all of its fields are mutable.

The `serialize` and `deserialize` functions currently do not handle certain cyclic values that `read` and `write` can handle, such as `'#0=(#0#)`.

See `deserialize` for information on the format of serialized data.

```
(deserialize v) → any
  v : any/c
```

Given a value *v* that was produced by `serialize`, produces a value like the one given to `serialize`, including the same graph structure and mutability.

A serialized representation *v* is a list of six or seven elements:

- An optional list `'(1)`, `'(2)`, or `'(3)` that represents the version of the serialization format. If the first element of a representation is not a list, then the version is `0`. Version 1 adds support for mutable pairs, version 2 adds support for unreadable symbols, and version 3 adds support for `date*` structures.
- A non-negative exact integer *s-count* that represents the number of distinct structure types represented in the serialized data.
- A list *s-types* of length *s-count*, where each element represents a structure type. Each structure type is encoded as a pair. The `car` of the pair is `#f` for a structure whose deserialization information is defined at the top level, otherwise it is a quoted module path or a byte string (to be converted into a platform-specific path using `bytes->path`) for a module that exports the structure's deserialization information. The `cdr`

of the pair is the name of a binding (at the top level or exported from a module) for deserialization information, either a symbol or a string representing an unreadable symbol. These two are used with either `namespace-variable-binding` or `dynamic-require` to obtain deserialization information. See `make-deserialize-info` for more information on the binding's value. See also `deserialize-module-guard`.

- A non-negative exact integer, *`g-count`* that represents the number of graph points contained in the following list.

- A list *`graph`* of length *`g-count`*, where each element represents a serialized value to be referenced during the construction of other serialized values. Each list element is either a box or not:

  – A box represents a value that is part of a cycle, and for deserialization, it must be allocated with `#f` for each of its fields. The content of the box indicates the shape of the value:

    * a non-negative exact integer *`i`* for an instance of a structure type that is represented by the *`i`* th element of the *`s-types`* list;
    * `'c` for a pair, which fails on deserialization (since pairs are immutable; this case does not appear in output generated by `serialize`);
    * `'m` for a mutable pair;
    * `'b` for a box;
    * a pair whose `car` is `'v` and whose `cdr` is a non-negative exact integer *`s`* for a vector of length *`s`*;
    * a list whose first element is `'h` and whose remaining elements are symbols that determine the hash-table type:
      · `'equal` — `(make-hash)`
      · `'equal 'weak` — `(make-weak-hash)`
      · `'weak` — `(make-weak-hasheq)`
      · no symbols — `(make-hasheq)`
    * `'date*` for a `date*` structure, which fails on deserialization (since dates are immutable; this case does not appear in output generated by `serialize`);
    * `'date` for a `date` structure, which fails on deserialization (since dates are immutable; this case does not appear in output generated by `serialize`);
    * `'arity-at-least` for an `arity-at-least` structure, which fails on deserialization (since arity-at-least are immutable; this case does not appear in output generated by `serialize`); or
    * `'mpi` for a module path index, which fails on deserialization (since a module path index is immutable; this case does not appear in output generated by `serialize`).

    The `#f`-filled value will be updated with content specified by the fifth element of the serialization list *`v`*.

  – A non-box represents a *serial* value to be constructed immediately, and it is one of the following:

* a boolean, number, character, interned symbol, or empty list, representing itself.
* a string, representing an immutable string.
* a byte string, representing an immutable byte string.
* a pair whose `car` is `'?` and whose `cdr` is a non-negative exact integer *`i`*; it represents the value constructed for the *`i`*th element of *`graph`*, where *`i`* is less than the position of this element within *`graph`*.
* a pair whose `car` is a number *`i`*; it represents an instance of a structure type that is described by the *`i`*th element of the *`s-types`* list. The `cdr` of the pair is a list of serials representing arguments to be provided to the structure type's deserializer.
* a pair whose `car` is `'f`; it represents an instance of a prefab structure type. The `cadr` of the pair is a prefab structure type key, and the `cddr` is a list of serials representing the field values.
* a pair whose `car` is `'void`, representing `#<void>`.
* a pair whose `car` is `'su` and whose `cdr` is a character string; it represents an unreadable symbol.
* a pair whose `car` is `'u` and whose `cdr` is either a byte string or character string; it represents a mutable byte or character string.
* a pair whose `car` is `'p` and whose `cdr` is a byte string; it represents a path using the serializer's path convention (deprecated in favor of `'p+`).
* a pair whose `car` is `'p+`, whose `cadr` is a byte string, and whose `cddr` is one of the possible symbol results of `system-path-convention-type`; it represents a path using the specified convention.
* a pair whose `car` is `'c` and whose `cdr` is a pair of serials; it represents an immutable pair.
* a pair whose `car` is `'c!` and whose `cdr` is a pair of serials; it represents a pair (but formerly represented a mutable pair), and does not appear in output generated by `serialize`.
* a pair whose `car` is `'m` and whose `cdr` is a pair of serials; it represents a mutable pair.
* a pair whose `car` is `'v` and whose `cdr` is a list of serials; it represents an immutable vector.
* a pair whose `car` is `'v!` and whose `cdr` is a list of serials; it represents a mutable vector.
* a pair whose `car` is `'b` and whose `cdr` is a serial; it represents an immutable box.
* a pair whose `car` is `'b!` and whose `cdr` is a serial; it represents a mutable box.
* a pair whose `car` is `'h`, whose `cadr` is either `'!` or `'-` (mutable or immutable, respectively), whose `caddr` is a list of symbols (containing `'equal`, `'weak`, both, or neither) that determines the hash table type, and whose `cdddr` is a list of pairs, where the `car` of each pair is a serial for a hash-table key and the `cdr` is a serial for the corresponding value.

    * a pair whose `car` is `'date*` and whose `cdr` is a list of serials; it represents a `date*` structure.
    * a pair whose `car` is `'date` and whose `cdr` is a list of serials; it represents a `date` structure.
    * a pair whose `car` is `'arity-at-least` and whose `cdr` is a serial; it represents an `arity-at-least` structure.
    * a pair whose `car` is `'mpi` and whose `cdr` is a pair; it represents a module path index that joins the paired values.

- A list of pairs, where the `car` of each pair is a non-negative exact integer *i* and the `cdr` is a serial (as defined in the previous bullet). Each element represents an update to an *i*th element of *graph* that was specified as a box, and the serial describes how to construct a new value with the same shape as specified by the box. The content of this new value must be transferred into the value created for the box in *graph*.

- A final serial (as defined in the two bullets back) representing the result of `deserialize`.

The result of `deserialize` shares no mutable values with the argument to `deserialize`.

If a value provided to `serialize` is a simple tree (i.e., no sharing), then the fourth and fifth elements in the serialized representation will be empty.

```
(serialized=? v1 v2) → boolean?
  v1 : any/c
  v2 : any/c
```

Returns `#t` if *v1* and *v2* represent the same serialization information.

More precisely, it returns the same value that `(equal? (deserialize v1) (deserialize v2))` would return if

- all structure types whose deserializers are accessed with distinct module paths are actually distinct types;

- all structure types are transparent; and

- all structure instances contain only the constituent values recorded in each of *v1* and *v2*.

```
(deserialize-module-guard)
 → (module-path? symbol? . -> . void?)
(deserialize-module-guard guard) → void?
  guard : (module-path? symbol? . -> . void?)
```

A parameter whose value is called by `deserialize` before dynamically loading a module via `dynamic-require`. The two arguments provided to the procedure are the same as the arguments to be passed to `dynamic-require`. The procedure can raise an exception to disallow the `dynamic-require`.

```
(serializable-struct id maybe-super (field ...)
                     struct-option ...)
```

Like `struct`, but instances of the structure type are serializable with `serialize`. This form is allowed only at the top level or in a module's top level (so that deserialization information can be found later).

Serialization only supports cycles involving the created structure type when all fields are mutable (or when the cycle can be broken through some other mutable value).

In addition to the bindings generated by `struct`, `serializable-struct` binds `deserialize-info:`*id*`-v0` to deserialization information. Furthermore, in a module context, it automatically `provide`s this binding.

The `serializable-struct` form enables the construction of structure instances from places where *id* is not accessible, since deserialization must construct instances. Furthermore, `serializable-struct` provides limited access to field mutation, but only for instances generated through the deserialization information bound to `deserialize-info:`*id*`-v0`. See `make-deserialize-info` for more information.

The `-v0` suffix on the deserialization enables future versioning on the structure type through `serializable-struct/version`.

When a supertype is supplied as *maybe-super*, compile-time information bound to the supertype identifier must include all of the supertype's field accessors. If any field mutator is missing, the structure type will be treated as immutable for the purposes of marshaling (so cycles involving only instances of the structure type cannot be handled by the deserializer).

Examples:

```
> (serializable-struct point (x y))

> (point-x (deserialize (serialize (point 1 2))))
1
```

```
(define-serializable-struct id-maybe-super (field ...)
                            struct-option ...)
```

Like `serializable-struct`, but with the supertype syntax and default constructor name of `define-struct`.

```
(serializable-struct/versions id maybe-super vers (field ...)
                              (other-version-clause ...)
                              struct-option ...)
```

```
other-version-clause = (other-vers make-proc-expr
                                   cycle-make-proc-expr)
```

Like `serializable-struct`, but the generated deserializer binding is `deserialize-info:`*id*`-v`*vers*. In addition, `deserialize-info:`*id*`-v`*other-vers* is bound for each *other-vers*. The *vers* and each *other-vers* must be a literal, exact, nonnegative integer.

Each *make-proc-expr* should produce a procedure, and the procedure should accept as many argument as fields in the corresponding version of the structure type, and it produce an instance of *id*. Each *cycle-make-proc-expr* should produce a procedure of no arguments; this procedure should return two values: an instance `x` of *id* (typically with `#f` for all fields) and a procedure that accepts another instance of *id* and copies its field values into `x`.

Examples:

```
> (serializable-struct point (x y) #:mutable #:transparent)

> (define ps (serialize (point 1 2)))

> (deserialize ps)
(point 1 2)
> (define x (point 1 10))

> (set-point-x! x x)

> (define xs (serialize x))

> (deserialize xs)
#0=(point #0# 10)
> (serializable-struct/versions point 1 (x y z)
    ([0
      ; Constructor for simple v0 instances:
      (lambda (x y) (point x y 0))
      ; Constructor for v0 instance in a cycle:
      (lambda ()
        (let ([p0 (point #f #f 0)])
          (values
            p0
            (lambda (p)
              (set-point-x! p0 (point-x p))
              (set-point-y! p0 (point-y p))))))])
    #:mutable #:transparent)

> (deserialize (serialize (point 4 5 6)))
```

```
(point 4 5 6)
> (deserialize ps)
(point 1 2 0)
> (deserialize xs)
#0=(point #0# 10 0)
```

```
(define-serializable-struct/versions id-maybe-super vers (field ...)
                                     (other-version-clause ...)
                                     struct-option ...)
```

Like `serializable-struct/versions`, but with the supertype syntax and default constructor name of `define-struct`.

```
(make-deserialize-info make cycle-make) → any
  make : procedure?
  cycle-make : (-> (values any/c procedure?))
```

Produces a deserialization information record to be used by `deserialize`. This information is normally tied to a particular structure because the structure has a `prop:serializable` property value that points to a top-level variable or module-exported variable that is bound to deserialization information.

The *make* procedure should accept as many arguments as the structure's serializer put into a vector; normally, this is the number of fields in the structure. It should return an instance of the structure.

The *cycle-make* procedure should accept no arguments, and it should return two values: a structure instance `x` (with dummy field values) and an update procedure. The update procedure takes another structure instance generated by the *make*, and it transfers the field values of this instance into `x`.

```
prop:serializable : property?
```

This property identifies structures and structure types that are serializable. The property value should be constructed with `make-serialize-info`.

```
(make-serialize-info to-vector
                     deserialize-id
                     can-cycle?
                     dir)            → any
  to-vector : (any/c . -> . vector?)
  deserialize-id :  (or identifier?
                        symbol?
                        (cons/c symbol?
                                module-path-index?))
  can-cycle? : any/c
  dir : path-string?
```

Produces a value to be associated with a structure type through the `prop:serializable` property. This value is used by `serialize`.

The *to-vector* procedure should accept a structure instance and produce a vector for the instance's content.

The *deserialize-id* value indicates a binding for deserialize information, to either a module export or a top-level definition. It must be one of the following:

- If *deserialize-id* is an identifier, and if `(identifier-binding` *deserialize-id*`)` produces a list, then the third element is used for the exporting module, otherwise the top-level is assumed. In either case, `syntax-e` is used to obtain the name of an exported identifier or top-level definition.
- If *deserialize-id* is a symbol, it indicates a top-level variable that is named by the symbol.
- If *deserialize-id* is a pair, the `car` must be a symbol to name an exported identifier, and the `cdr` must be a module path index to specify the exporting module.

See `make-deserialize-info` and `deserialize` for more information.

The *can-cycle?* argument should be false if instances should not be serialized in such a way that deserialization requires creating a structure instance with dummy field values and then updating the instance later.

The *dir* argument should be a directory path that is used to resolve a module reference for the binding of *deserialize-id*. This directory path is used as a last resort when *deserialize-id* indicates a module that was loaded through a relative path with respect to the top level. Usually, it should be `(or (current-load-relative-directory) (current-directory))`.

12.10 Fast-Load Serialization

```
(require racket/fasl)
```

The bindings documented in this section are provided by the `racket/fasl` library, not `racket/base` or `racket`.

```
(s-exp->fasl v [out]) → (or/c (void) bytes?)
  v : any/c
  out : (or/c output-port? #f) = #f
(fasl->s-exp in) → any/c
  in : (or/c input-port? bytes?)
```

The s-exp->fasl function serializes v to a byte string, printing it directly to out if out is an output port or return the byte string otherwise. The fasl->s-exp function decodes a value from a byte string (supplied either directly or as an input port) that was encoded with s-exp->fasl.

The v argument must be a value that could be quoted as a literal, because s-exp->fasl essentially uses (compile `',v) to encode the value using Racket's built-in fast-load format for bytecode.

The byte-string encoding produced by s-exp->fasl is specific to a version of Racket. That is, the resulting byte string can be decoded back to v only using the same version with which it was encoded.

Examples:

```
> (define fasl (s-exp->fasl (list #("speed") 'racer #\!)))

> fasl
#"#~\0035.3T\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\1\0\0\0\1\0\0\0\0\36\0\0\0\237$\24q\23
> (fasl->s-exp fasl)
'(#("speed") racer #\!)
```

13 Reflection and Security

13.1 Namespaces

See §1.2.5 “Namespaces” for basic information on the namespace model.

A new namespace is created with procedures like `make-empty-namespace`, and `make-base-namespace`, which return a first-class namespace value. A namespace is used by setting the `current-namespace` parameter value, or by providing the namespace to procedures such as `eval` and `eval-syntax`.

```
(namespace? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a namespace value, `#f` otherwise.

```
(make-empty-namespace) → namespace?
```

Creates a new namespace that is empty, and whose module registry contains no mappings. The namespace’s base phase is the same as the base phase of the current namespace. Attach modules from an existing namespace to the new one with `namespace-attach-module`.

```
(make-base-empty-namespace) → namespace?
```

Creates a new empty namespace, but with `racket/base` attached. The namespace’s base phase is the same as the phase in which the `make-base-empty-namespace` function was created.

```
(make-base-namespace) → namespace?
```

Creates a new namespace with `racket/base` attached and `required` into the top-level environment. The namespace’s base phase is the same as the phase in which the `make-base-namespace` function was created.

```
(define-namespace-anchor id)
```

Binds *id* to a namespace anchor that can be used with `namespace-anchor->empty-namespace` and `namespace-anchor->namespace`.

This form can be used only in a top-level context or in a module-context.

```
(namespace-anchor? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a namespace-anchor value, `#f` otherwise.

```
(namespace-anchor->empty-namespace a) → namespace?
  a : namespace-anchor?
```

Returns an empty namespace that shares a module registry with the source of the anchor, and whose base phase the phase in which the anchor was created.

If the anchor is from a `define-namespace-anchor` form in a module context, then the source is the namespace in which the containing module is instantiated. If the anchor is from a `define-namespace-anchor` form in a top-level content, then the source is the namespace in which the anchor definition was evaluated.

```
(namespace-anchor->namespace a) → namespace?
  a : namespace-anchor?
```

Returns a namespace corresponding to the source of the anchor.

If the anchor is from a `define-namespace-anchor` form in a module context, then the result is a namespace for the module's body in the anchor's phase. The result is the same as a namespace obtained via `module->namespace`.

If the anchor is from a `define-namespace-anchor` form in a top-level content, then the result is the namespace in which the anchor definition was evaluated.

```
(current-namespace) → namespace?
(current-namespace n) → void?
  n : namespace?
```

A parameter that determines the current namespace.

```
(namespace-symbol->identifier sym) → identifier?
  sym : symbol?
```

Similar to `datum->syntax` restricted to symbols. The lexical information of the resulting identifier corresponds to the top-level environment of the current namespace; the identifier has no source location or properties.

```
(namespace-base-phase [namespace]) → exact-integer?
  namespace : namespace? = (current-namespace)
```

Returns the base phase of *namespace*.

```
(namespace-module-identifier [where]) → identifier?
  where : (or/c namespace? exact-integer? #f)
        = (current-namespace)
```

Returns an identifier whose binding is `module` in the base phase of *where* if it is a namespace, or in the *where* phase level otherwise.

The lexical information of the identifier includes bindings (in the same phase level) for all syntactic forms that appear in fully expanded code (see §1.2.3.1 "Fully Expanded Programs"), but using the name reported by the second element of `identifier-binding` for the binding; the lexical information may also include other bindings.

```
(namespace-variable-value sym
                          [use-mapping?
                           failure-thunk
                           namespace])   → any
  sym : symbol?
  use-mapping? : any/c = #t
  failure-thunk : (or/c (-> any) #f) = #f
  namespace : namespace? = (current-namespace)
```

Returns a value for *sym* in *namespace*, using *namespace*'s base phase. The returned value depends on *use-mapping?*:

- If *use-mapping?* is true (the default), and if *sym* maps to a top-level variable or an imported variable (see §1.2.5 "Namespaces"), then the result is the same as evaluating *sym* as an expression. If *sym* maps to syntax or imported syntax, then *failure-thunk* is called or the `exn:fail:syntax` exception is raised. If *sym* is mapped to an undefined variable or an uninitialized module variable, then *failure-thunk* is called of the `exn:fail:contract:variable` exception is raised.
- If *use-mapping?* is `#f`, the namespace's syntax and import mappings are ignored. Instead, the value of the top-level variable named *sym* in namespace is returned. If the variable is undefined, then *failure-thunk* is called or the `exn:fail:contract:variable` exception is raised.

If *failure-thunk* is not `#f`, `namespace-variable-value` calls *failure-thunk* to produce the return value in place of raising an `exn:fail:contract:variable` or `exn:fail:syntax` exception.

```
(namespace-set-variable-value! sym
                               v
                               [map?
                                namespace]) → void?
  sym : symbol?
  v : any/c
  map? : any/c = #f
  namespace : namespace? = (current-namespace)
```

Sets the value of *sym* in the top-level environment of *namespace* in the base phase, defining *sym* if it is not already defined.

If *map?* is supplied as true, then the namespace's identifier mapping is also adjusted (see §1.2.5 "Namespaces") in the phase level corresponding to the base phase, so that *sym* maps to the variable.

```
(namespace-undefine-variable! sym
                              [namespace]) → void?
  sym : symbol?
  namespace : namespace? = (current-namespace)
```

Removes the *sym* variable, if any, in the top-level environment of *namespace* in its base phase. The namespace's identifier mapping (see §1.2.5 "Namespaces") is unaffected.

```
(namespace-mapped-symbols [namespace]) → (listof symbol?)
  namespace : namespace? = (current-namespace)
```

Returns a list of all symbols that are mapped to variables, syntax, and imports in *namespace* for the phase level corresponding to the namespace's base phase.

```
(namespace-require quoted-raw-require-spec) → void?
  quoted-raw-require-spec : any/c
```

Performs the import corresponding to *quoted-raw-require-spec* in the top-level environment of the current namespace, like a top-level `#%require`. The *quoted-raw-require-spec* argument must be a datum that corresponds to a quoted *raw-require-spec* for `#%require`, which includes module paths.

Module paths in *quoted-raw-require-spec* are resolved with respect to `current-load-relative-directory` or `current-directory` (if the former is `#f`), even if the current namespace corresponds to a module body.

```
(namespace-require/copy quoted-raw-require-spec) → void?
  quoted-raw-require-spec : any/c
```

Like `namespace-require` for syntax exported from the module, but exported variables at the namespace's base phase are treated differently: the export's current value is copied to a top-level variable in the current namespace.

```
(namespace-require/constant quoted-raw-require-spec) → void?
  quoted-raw-require-spec : any/c
```

Like `namespace-require`, but for each exported variable at the namespace's base phase, the export's value is copied to a corresponding top-level variable that is made immutable. Despite setting the top-level variable, the corresponding identifier is bound as imported.

```
(namespace-require/expansion-time quoted-raw-require-spec)
 → void?
  quoted-raw-require-spec : any/c
```

Like namespace-require, but only the transformer part of the module is executed relative to the namespace's base phase; that is, the module is merely visited, and not instantiated (see §1.2.3.8 "Module Phases and Visits"). If the required module has not been instantiated before, the module's variables remain undefined.

```
(namespace-attach-module  src-namespace
                          modname
                         [dest-namespace]) → void?
  src-namespace : namespace?
  modname : module-path?
  dest-namespace : namespace? = (current-namespace)
```

Attaches the instantiated module named by *modname* in *src-namespace* (at its base phase) to the module registry of *dest-namespace*.

In addition to *modname*, every module that it imports (directly or indirectly) is also recorded in the current namespace's module registry, and instances at the same phase or lower are also attached to *dest-namespace* (while visits at the module's phase and instances at higher phases are not attached, nor even made available for on-demand visits). The inspector of the module invocation in *dest-namespace* is the same as inspector of the invocation in *src-namespace*.

If *modname* is not a symbol, the current module name resolver is called to resolve the path, but no module is loaded; the resolved form of *modname* is used as the module name in *dest-namespace*.

If *modname* refers to a submodule or a module with submodules, unless the module was loaded from bytecode (i.e., a ".zo" file) independently from submodules within the same top-level module, then declarations for all submodules within the module's top-level module are also attached to *dest-namespace*.

If *modname* does not refer to an instantiated module in *src-namespace*, or if the name of any module to be attached already has a different declaration or same-phase instance in *dest-namespace*, then the exn:fail:contract exception is raised.

If *src-namespace* and *dest-namespace* do not have the same base phase, then the exn:fail:contract exception is raised.

```
(namespace-attach-module-declaration  src-namespace
                                      modname
                                     [dest-namespace]) → void?
  src-namespace : namespace?
  modname : module-path?
  dest-namespace : namespace? = (current-namespace)
```

Like namespace-attach-module, but the module specified by *modname* need only be de-

clared (and not necessarily instantiated) in *src-namespace*, and the module is merely declared in *dest-namespace*.

```
(namespace-unprotect-module  inspector
                             modname
                            [namespace]) → void?
  inspector : inspector?
  modname : module-path?
  namespace : namespace? = (current-namespace)
```

Changes the inspector for the instance of the module referenced by *modname* in *namespace*'s module registry so that it is controlled by the current code inspector. The given *inspector* must currently control the invocation of the module in *namespace*'s module registry, otherwise the `exn:fail:contract` exception is raised. See also §13.10 "Code Inspectors".

```
(namespace-module-registry namespace) → any
  namespace : namespace?
```

Returns the module registry of the given namespace. This value is useful only for identification via `eq?`.

```
(module->namespace modname) → namespace?
  modname : module-path?
```

Returns a namespace that corresponds to the body of an instantiated module in the current namespace's module registry and in the current namespace's base phase, making the module at the available for on-demand visits at the namespace's base phase. The returned namespace has the same module registry as the current namespace. Modifying a binding in the namespace changes the binding seen in modules that require the namespace's module.

Module paths in a top-level `require` expression are resolved with respect to the namespace's module. New `provide` declarations are not allowed.

If the current code inspector does not control the invocation of the module in the current namespace's module registry, the `exn:fail:contract` exception is raised; see also §13.10 "Code Inspectors".

Bindings in the namespace cannot be modified if the `compile-enforce-module-constants` parameter was true when the module was declared, unless the module declaration itself included assignments to the binding via `set!`.

```
(namespace-syntax-introduce stx) → syntax-object?
  stx : syntax-object?
```

Returns a syntax object like *stx*, except that the current namespace's bindings are included in the syntax object's lexical information (see §1.2.2 "Syntax Objects"). The additional

context is overridden by any existing top-level bindings in the syntax object's lexical information, or by any existing or future module bindings in the lexical information.

```
(module-provide-protected? module-path-index
                           sym)              → boolean?
  module-path-index : (or/c symbol? module-path-index?)
  sym : symbol?
```

Returns `#f` if the module declaration for *module-path-index* defines *sym* and exports it unprotected, `#t` otherwise (which may mean that the symbol corresponds to an unexported definition, a protected export, or an identifier that is not defined at all within the module).

The *module-path-index* argument can be a symbol; see §13.4.2 "Compiled Modules and References" for more information on module path indices.

Typically, the arguments to `module-provide-protected?` correspond to the first two elements of a list produced by `identifier-binding`.

```
(variable-reference? v) → boolean?
  v : any/c
```

Return `#t` if *v* is a variable reference produced by `#%variable-reference`, `#f` otherwise.

```
(variable-reference-constant? varref) → boolean?
  varref : variable-reference?
```

Returns `#t` if the variable represented by *varref* will retain its current value (i.e., *varref* refers to a variable that cannot be further modified by `set!` or `define`), `#f` otherwise.

```
(variable-reference->empty-namespace varref) → namespace?
  varref : variable-reference?
```

Returns an empty namespace that shares module declarations and instances with the namespace in which *varref* is instantiated, and with the same phase as *varref*.

```
(variable-reference->namespace varref) → namespace?
  varref : variable-reference?
```

If *varref* refers to a module-level variable, then the result is a namespace for the module's body in the referenced variable's phase; the result is the same as a namespace obtained via `module->namespace`.

If *varref* refers to a top-level variable, then the result is the namespace in which the referenced variable is defined.

```
(variable-reference->resolved-module-path varref)
 → (or/c resolved-module-path? #f)
  varref : variable-reference?
```

If *varref* refers to a module-level variable, the result is a resolved module path naming the module.

If *varref* refers to a top-level variable, then the result is `#f`.

```
(variable-reference->module-path-index varref)
 → (or/c module-path-index? #f)
  varref : variable-reference?
```

If *varref* refers to a module-level variable, the result is a module path index naming the module.

If *varref* refers to a top-level variable, then the result is `#f`.

```
(variable-reference->module-source varref)
 → (or/c symbol? (and/c path? complete-path?) #f)
  varref : variable-reference?
```

If *varref* refers to a module-level variable, the result is a path or symbol naming the module's source (which is typically, but not always, the same as in the resolved module path). If the relevant module is a submodule, the result corresponds to the enclosing top-level module's source.

If *varref* refers to a top-level variable, then the result is `#f`.

```
(variable-reference->phase varref) → exact-nonnegative-integer?
  varref : variable-reference?
```

Returns the phase of the variable referenced by *varref*.

```
(variable-reference->module-base-phase varref) → exact-integer?
  varref : variable-reference?
```

Returns the phase in which the module is instantiated for the variable referenced by *varref*, or `0` if the variable for *varref* is not within a module.

For a variable with a module, the result is less than the result of `(variable-reference->phase varref)` by *n* when the variable is bound at phase level *n* within the module.

```
(variable-reference->module-declaration-inspector varref)
 → inspector?
  varref : variable-reference?
```

Returns the declaration inspector (see §13.10 “Code Inspectors”) for the module of *varref*, where *varref* must refer to an anonymous module variable as produced by `(#%variable-reference)`.

13.2 Evaluation and Compilation

```
(current-eval) → (any/c . -> . any)
(current-eval proc) → void?
  proc : (any/c . -> . any)
```

A parameter that determines the current *evaluation handler*. The evaluation handler is a procedure that takes a top-level form and evaluates it, returning the resulting values. The evaluation handler is called by `eval`, `eval-syntax`, the default load handler, and `read-eval-print-loop` to evaluate a top-level form. The handler should evaluate its argument in tail position.

The *top-level-form* provided to the handler can be a syntax object, a compiled form, a compiled form wrapped as a syntax object, or an arbitrary datum.

The default handler converts an arbitrary datum to a syntax object using `datum->syntax`, and then enriches its lexical information in the same way as `eval`. (If *top-level-form* is a syntax object, then its lexical information is not enriched.) The default evaluation handler partially expands the form to splice the body of top-level `begin` forms into the top level (see `expand-to-top-form`), and then individually compiles and evaluates each spliced form before continuing to expand, compile, and evaluate later forms.

```
(eval top-level-form [namespace]) → any
  top-level-form : any/c
  namespace : namespace? = (current-namespace)
```

See also §15.1.2 “Namespaces” in ***The Racket Guide***.

Calls the current evaluation handler to evaluate *top-level-form*. The evaluation handler is called in tail position with respect to the `eval` call, and `parameterized` to set `current-namespace` to *namespace*.

If *top-level-form* is a syntax object whose datum is not a compiled form, then its lexical information is enriched before it is sent to the evaluation handler:

- If *top-level-form* is a pair whose `car` is a symbol or identifier, and if applying `namespace-syntax-introduce` to the (`datum->syntax`-converted) identifier produces an identifier bound to `module` in a phase level that corresponds to *namespace*’s base phase, then only that identifier is enriched.
- For any other *top-level-form*, `namespace-syntax-introduce` is applied to the entire syntax object.

For interactive evaluation in the style of `read-eval-print-loop` and `load`, wrap each expression with `#%top-interaction`, which is normally bound to `#%top-interaction`, before passing it to `eval`.

```
(eval-syntax stx [namespace]) → any
  stx : syntax?
  namespace : namespace? = (current-namespace)
```

Like `eval`, except that *stx* must be a syntax object, and its lexical context is not enriched before it is passed to the evaluation handler.

```
(current-load)
    (path? (or/c #f
                 symbol?
                 (cons/c (or/c #f symbol?)
 →                       (non-empty-listof symbol?)))
           . -> .
           any)
(current-load proc) → void?
         (path? (or/c #f
                      symbol?
                      (cons/c (or/c #f symbol?)
  proc :                      (non-empty-listof symbol?)))
                . -> .
                any)
```

A parameter that determines the current *load handler* to load top-level forms from a file. The load handler is called by `load`, `load-relative`, `load/cd`, and the default compiled-load handler.

A load handler takes two arguments: a path (see §14.1 “Paths”) and an expected module name. The expected module name is a symbol or a list when the call is to load a module declaration in response to a `require` (in which case the file should contain a module declaration), or `#f` for any other load.

When loading a module from a stream that starts with a compiled module that contains submodules, the load handler should load only the requested module, where a symbol as the load handler’s indicates the root module and a list indicates a submodule whose path relative to the root module is given by the `cdr` of the list. The list starts with `#f` when a submodule should be loaded *only* if it can be loaded independently (i.e., from compiled form—never from source); if the submodule cannot be loaded independently, the load handler should return without loading from a file. When the expected module name is a list that starts with a symbol, the root module and any other submodules can be loaded from the given file, which might be from source, and the load handler still should not complain if the expected submodule is not found.

The default load handler reads forms from the file in `read-syntax` mode with line-counting enabled for the file port, unless the path has a `".zo"` suffix. It also `parameterizes` each read to set `read-accept-compiled`, `read-accept-reader`, and `read-accept-lang` to `#t`. In addition, if `load-on-demand-enabled` is `#t`, then `read-on-demand-source` is set to the cleansed, absolute form of `path` during the `read-syntax` call. After reading a single form, the form is passed to the current evaluation handler, wrapping the evaluation in a continuation prompt (see `call-with-continuation-prompt`) for the default continuation prompt tag with handler that propagates the abort to the continuation of the `load` call.

If the second argument to the load handler is a symbol, then:

- The `read-syntax` from the file is additionally `parameterized` as follows (to provide consistent reading of module source):

```
(current-readtable #f)
(read-case-sensitive #t)
(read-square-bracket-as-paren #t)
(read-curly-brace-as-paren #t)
(read-accept-box #t)
(read-accept-compiled #t)
(read-accept-bar-quote #t)
(read-accept-graph #t)
(read-decimal-as-inexact #t)
(read-accept-dot #t)
(read-accept-infix-dot #t)
(read-accept-quasiquote #t)
(read-accept-reader #t)
(read-accept-lang #t)
```

- If the read result is not a `module` form, or if a second `read-syntax` does not produce an end-of-file, then the `exn:fail` exception is raised without evaluating the form that was read from the file. (In previous versions, the module declaration was checked to match the name given as the second argument to the load handler, but this check is no longer performed.)
- The lexical information of the initial `module` identifier is enriched with a binding for `module`, so that the form corresponds to a module declaration independent of the current namespace's bindings.

If the second argument to the load handler is `#f`, then each expression read from the file is wrapped with `#%top-interaction`, which is normally bound to `#%top-interaction`, before passing it to the evaluation handler.

The return value from the default load handler is the value of the last form from the loaded file, or `#<void>` if the file contains no forms. If the given path is a relative path, then it is resolved using the value of `current-directory`.

```
(load file) → any
  file : path-string?
```

See also §15.1.2 “Namespaces” in ***The Racket Guide***.

Calls the current load handler in tail position. The call is `parameterized` to set `current-load-relative-directory` to the directory of *file*, which is resolved relative to the value of `current-directory`.

```
(load-relative file) → any
  file : path-string?
```

Like `load/use-compiled`, but when *file* is a relative path, it is resolved using the value of `current-load-relative-directory` instead of the value of `current-directory` if the former is not `#f`, otherwise `current-directory` is used.

```
(load/cd file) → any
  file : path-string?
```

Like `load`, but `load/cd` sets both `current-directory` and `current-load-relative-directory` before calling the load handler.

```
(current-load-extension)
 → (path? (or/c symbol? #f) . -> . any)
(current-load-extension proc) → void?
  proc : (path? (or/c symbol? #f) . -> . any)
```

A parameter that determines a *extension-load handler*, which is called by `load-extension` and the default compiled-load handler.

An extension-load handler takes the same arguments as a load handler, but the file should be a platform-specific *dynamic extension*, typically with the file suffix `".so"` (Unix), `".dll"` (Windows), or `".dylib"` (Mac OS X). The file is loaded using internal, OS-specific primitives. See *Inside: Racket C API* for more information on dynamic extensions.

```
(load-extension file) → any
  file : path-string?
```

Sets `current-load-relative-directory` like `load`, and calls the extension-load handler in tail position.

```
(load-relative-extension file) → any
  file : path-string?
```

Like `load-extension`, but resolves *file* using `current-load-relative-directory` like `load-relative`.

```
(current-load/use-compiled)
    (path? (or/c #f
                 symbol?
 →               (cons/c (or/c #f symbol?)
                         (non-empty-listof symbol?)))
           . -> . any)
(current-load/use-compiled proc) → void?
         (path? (or/c #f
                      symbol?
  proc :              (cons/c (or/c #f symbol?)
                              (non-empty-listof symbol?)))
                . -> . any)
```

A parameter that determines the current *compiled-load handler* to load from a file that may have a compiled form. The compiled-load handler is called by `load/use-compiled`.

The protocol for a compiled-load handler is the same as for the load handler (see `current-load`), except that a compiled-load handler is expected to set `current-load-relative-directory` itself. The default compiled-load handler, however, checks for a `".ss"` file when the given path ends with `".rkt"`, no `".rkt"` file exists, and when the handler's second argument is a symbol. In addition, the default compiled-load handler checks for `".zo"` (bytecode) files and `".so"` (native Unix), `".dll"` (native Windows), or `".dylib"` (native Mac OS X) files.

The check for a compiled file occurs whenever the given path *file* ends with any extension (e.g., `".rkt"` or `".scrbl"`), and the check consults the subdirectories indicated by the `use-compiled-file-paths` parameter relative to *file*. The subdirectories are checked in order. A `".zo"` version of the file (whose name is formed by passing *file* and `#".zo"` to `path-add-suffix`) is loaded if it exists directly in one of the indicated subdirectories, or a `".so"/".dll"/".dylib"` version of the file is loaded if it exists within a `"native"` subdirectory of a `use-compiled-file-paths` directory, in an even deeper subdirectory as named by `system-library-subpath`. A compiled file is loaded only if its modification date is not older than the date for *file*. If both `".zo"` and `".so"/".dll"/".dylib"` files are available, the `".so"/".dll"/".dylib"` file is used. If *file* ends with `".rkt"`, no such file exists, the handler's second argument is a symbol, and a `".ss"` file exists, then `".zo"` and `".so"/".dll"/".dylib"` files are used only with names based on *file* with its suffixed replaced by `".ss"`.

While a `".zo"`, `".so"`, `".dll"`, or `".dylib"` file is loaded, the current `load-relative` directory is set to the directory of the original *file*. If the file to be loaded has the suffix `".ss"` while the requested file has the suffix `".rkt"`, then the `current-module-declare-source` parameter is set to the full path of the loaded file, otherwise the `current-module-declare-source` parameter is set to `#f`.

If the original *file* is loaded or a `".zo"` variant is loaded, the load handler is called to load

the file. If any other kind of file is loaded, the extension-load handler is called.

When the default compiled-load handler loads a module from a bytecode (i.e., `".zo"`) file, the handler records the bytecode file path in the current namespace's module registry. More specifically, the handler records the path for the top-level module of the loaded module, which is an enclosing module if the loaded module is a submodule. Thereafter, loads via the default compiled-load handler for modules within the same top-level module use the recorded file, independent of the file that otherwise would be selected by the compiled-load handler (e.g., even if the `use-compiled-file-paths` parameter value changes). The default module name resolver transfers bytecode-file information when a module declaration is attached to a new namespace. This protocol supports independent but consistent loading of submodules from bytecode files.

```
(load/use-compiled file) → any
  file : path-string?
```

Calls the current compiled-load handler in tail position.

```
(current-load-relative-directory)
 → (or/c (and/c path-string? complete-path?) #f)
(current-load-relative-directory path) → void?
  path : (or/c (and/c path-string? complete-path?) #f)
```

A parameter that is set by `load`, `load-relative`, `load-extension`, `load-relative-extension`, and the default compiled-load handler, and used by `load-relative`, `load-relative-extension`, and the default compiled-load handler.

When a new path or string is provided as the parameter's value, it is immediately expanded (see §14.1 "Paths") and converted to a path. (The directory need not exist.)

```
(use-compiled-file-paths) → (listof path?)
(use-compiled-file-paths paths) → void?
  paths : (listof path-string?)
```

A list of relative paths, which defaults to `(list (string->path "compiled"))`. It is used by the compiled-load handler (see `current-load/use-compiled`).

```
(read-eval-print-loop) → any
```

Starts a new *REPL* using the current input, output, and error ports. The REPL wraps each expression to evaluate with `#%top-interaction`, which is normally bound to `#%top-interaction`, and it wraps each evaluation with a continuation prompt using the default continuation prompt tag and prompt handler (see `call-with-continuation-prompt`). The REPL also wraps the read and print operations with a prompt for the default tag whose handler ignores abort arguments and continues the loop. The `read-eval-print-loop` procedure does not return until `eof` is read, at which point it returns `#<void>`.

The read-eval-print-loop procedure can be configured through the current-prompt-read, current-eval, and current-print parameters.

```
(current-prompt-read) → (-> any)
(current-prompt-read proc) → void?
  proc : (-> any)
```

A parameter that determines a *prompt read handler*, which is a procedure that takes no arguments, displays a prompt string, and returns a top-level form to evaluate. The prompt read handler is called by read-eval-print-loop, and after printing a prompt, the handler typically should call the read interaction handler (as determined by the current-read-interaction parameter) with the port produced by the interaction port handler (as determined by the current-get-interaction-input-port parameter).

The default prompt read handler prints `> ` and returns the result of

```
(let ([in ((current-get-interaction-input-port))])
  ((current-read-interaction) (object-name in) in))
```

If the input and output ports are both terminals (in the sense of terminal-port?) and if the output port appears to be counting lines (because port-next-location returns a non-#f line and column), then the output port's line is incremented and its column is reset to 0 via set-port-next-location! before returning the read result.

```
(current-get-interaction-input-port) → (-> input-port?)
(current-get-interaction-input-port proc) → void?
  proc : (-> input-port?)
```

A parameter that determines the *interaction port handler*, which returns a port to use for read-eval-print-loop inputs.

The default interaction port handler returns the current input port. In addition, if that port is the initial current input port, the initial current output and error ports are flushed.

The racket/gui/base library adjusts this parameter's value by extending the current value. The extension wraps the result port so that GUI events can be handled when reading from the port blocks.

```
(current-read-interaction) → (any/c input-port? -> any)
(current-read-interaction proc) → void?
  proc : (any/c input-port? -> any)
```

A parameter that determines the current *read interaction handler*, which is procedure that takes an arbitrary value and an input port and returns an expression read from the input port.

The default read interaction handler accepts *src* and *in* and returns

```
(parameterize ([read-accept-reader #t]
               [read-accept-lang #f])
  (read-syntax src in))
```

```
(current-print) → (any/c -> any)
(current-print proc) → void?
  proc : (any/c -> any)
```

A parameter that determines the *print handler* that is called by `read-eval-print-loop` to print the result of an evaluation (and the result is ignored).

The default print handler `print`s the value to the current output port (as determined by the `current-output-port` parameter) and then outputs a newline, except that it prints nothing when the value is `#<void>`.

```
(current-compile)
 → (any/c boolean? . -> . compiled-expression?)
(current-compile proc) → void?
  proc : (any/c boolean? . -> . compiled-expression?)
```

A parameter that determines the current *compilation handler*. The compilation handler is a procedure that takes a top-level form and returns a compiled form; see §1.2.4 “Compilation” for more information on compilation.

The compilation handler is called by `compile`, and indirectly by the default evaluation handler and the default load handler.

The handler’s second argument is `#t` if the compiled form will be used only for immediate evaluation, or `#f` if the compiled form may be saved for later use; the default compilation handler is optimized for the special case of immediate evaluation.

When a compiled form is written to an output port, the written form starts with `#~`. See §1.4.15 “Printing Compiled Code” for more information.

For internal testing purposes, when the `PLT_VALIDATE_COMPILE` environment variable is set, the default compilation handler runs a bytecode validator on its own compilation results.

```
(compile top-level-form) → compiled-expression?
  top-level-form : any/c
```

Like `eval`, but calls the current compilation handler in tail position with *`top-level-form`*.

```
(compile-syntax stx) → compiled-expression?
  stx : syntax?
```

Like `eval-syntax`, but calls the current compilation handler in tail position with *`stx`*.

```
(compiled-expression? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a compiled form, `#f` otherwise.

```
(compile-enforce-module-constants) → boolean?
(compile-enforce-module-constants on?) → void?
  on? : any/c
```

A parameter that determines how a module declaration is compiled.

When constants are enforced, and when the macro-expanded body of a module contains no `set!` assignment to a particular variable defined within the module, then the variable is marked as constant when the definition is evaluated. Afterward, the variable's value cannot be assigned or undefined through `module->namespace`, and it cannot be defined by redeclaring the module.

Enforcing constants allows the compiler to inline some variable values, and it allows the native-code just-in-time compiler to generate code that skips certain run-time checks.

```
(compile-allow-set!-undefined) → boolean?
(compile-allow-set!-undefined allow?) → void?
  allow? : any/c
```

A parameter that determines how a `set!` expression is compiled when it mutates a global variable. If the value of this parameter is a true value, `set!` expressions for global variables are compiled so that the global variable is set even if it was not previously defined. Otherwise, `set!` expressions for global variables are compiled to raise the `exn:fail:contract:variable` exception if the global variable is not defined at the time the `set!` is performed. Note that this parameter is used when an expression is *compiled*, not when it is *evaluated*.

```
(compile-context-preservation-enabled) → boolean?
(compile-context-preservation-enabled on?) → void?
  on? : any/c
```

A parameter that determines whether compilation should avoid function-call inlining and other optimizations that may cause information to be lost from stack traces (as reported by `continuation-mark-set->context`). The default is `#f`, which allows such optimizations.

```
(eval-jit-enabled) → boolean?
(eval-jit-enabled on?) → void?
  on? : any/c
```

See also §18.2 “The Bytecode and Just-in-Time (JIT) Compilers” in ***The Racket Guide***.

A parameter that determines whether the native-code just-in-time compiler (*JIT*) is enabled for code (compiled or not) that is passed to the default evaluation handler. A true parameter value is effective only on platforms for which the JIT is supported.

The default is `#t`, unless the JIT is not supported by the current platform, unless it is disabled through the `-j`/`--no-jit` command-line flag to stand-alone Racket (or GRacket), and unless it is disabled through the `PLTNOMZJIT` environment variable (set to any value).

```
(load-on-demand-enabled) → boolean?
(load-on-demand-enabled on?) → void?
  on? : any/c
```

A parameter that determines whether the default load handler sets `read-on-demand-source`. See `current-load` for more information. The default is `#t`, unless it is disabled through the `-d`/`--no-delay` command-line flag.

13.3 The `racket/load` Language

```
#lang racket/load
```

The `racket/load` language supports evaluation where each top-level form in the module body is separately passed to `eval` in the same way as for `load`.

The namespace for evaluation shares the module registry with the `racket/load` module instance, but it has a separate top-level environment, and it is initialized with the bindings of `racket`. A single namespace is created for each instance of the `racket/load` module (i.e., multiple modules using the `racket/load` language share a namespace). The `racket/load` library exports only `#%module-begin` and `#%top-interaction` forms that effectively swap in the evaluation namespace and call `eval`.

For example, the body of a module using `racket/load` can include `module` forms, so that running the following module prints `5`:

```
#lang racket/load

(module m racket/base
  (provide x)
  (define x 5))

(module n racket/base
  (require 'm)
  (display x))

(require 'n)
```

Definitions in a module using `racket/load` are evaluated in the current namespace, which means that `load` and `eval` can see the definitions. For example, running the following module prints `6`:

```
#lang racket/load

(define x 6)
(display (eval 'x))
```

Since all forms within a `racket/load` module are evaluated in the top level, bindings cannot be exported from the module using `provide`. Similarly, since evaluation of the module-body forms is inherently dynamic, compilation of the module provides essentially no benefit. For these reasons, use `racket/load` for interactive exploration of top-level forms only, and not for constructing larger programs.

13.4 Module Names and Loading

13.4.1 Resolving Module Names

The name of a declared module is represented by a *resolved module path*, which encapsulates either a symbol or a complete filesystem path (see §14.1 "Paths"). A symbol normally refers to a predefined module or module declared through reflective evaluation (e.g., `eval`). A filesystem path normally refers to a module declaration that was loaded on demand via `require` or other forms.

A *module path* is a datum that matches the grammar for *module-path* for `require`. A module path is relative to another module.

```
(resolved-module-path? v) → boolean?
  v : any/c
```

Returns `#f` if *v* is a resolved module path, `#f` otherwise.

```
(make-resolved-module-path path) → resolved-module-path?
        (or/c symbol?
              (and/c path? complete-path?)
  path :      (cons/c (or/c symbol?
                            (and/c path? complete-path?))
                      (non-empty-listof symbol?)))
```

Returns a resolved module path that encapsulates *path*, where a list *path* corresponds to a submodule path. If *path* is a path or starts with a path, the path normally should be cleansed (see `cleanse-path`) and simplified (see `simplify-path`).

A resolved module path is interned. That is, if two resolved module path values encapsulate paths that are `equal?`, then the resolved module path values are `eq?`.

```
(resolved-module-path-name module-path)
    (or/c symbol?
          (and/c path? complete-path?)
 →        (cons/c (or/c symbol?
                        (and/c path? complete-path?))
                  (non-empty-listof symbol?)))
  module-path : resolved-module-path?
```

Returns the path or symbol encapsulated by a resolved module path. A list result corresponds to a submodule path.

```
(module-path? v) → boolean?
  v : any/c
```

Returns `#t` if `v` corresponds to a datum that matches the grammar for *module-path* for `require`, `#f` otherwise. Note that a path (in the sense of `path?`) is a module path.

```
(current-module-name-resolver)
    (case->
     (resolved-module-path? (or/c #f namespace?) . -> . any)
     (module-path?
 →    (or/c #f resolved-module-path?)
      (or/c #f syntax?)
      boolean?
      . -> .
      resolved-module-path?))
(current-module-name-resolver proc) → void?
        (case->
         (resolved-module-path? (or/c #f namespace?) . -> . any)
         (module-path?
  proc :  (or/c #f resolved-module-path?)
          (or/c #f syntax?)
          boolean?
          . -> .
          resolved-module-path?))
```

A parameter that determines the current *module name resolver*, which manages the conversion from other kinds of module references to a resolved module path. For example, when the expander encounters `(require module-path)` where *module-path* is not an identifier, then the expander passes `'module-path` to the module name resolver to obtain a symbol or resolved module path. When such a `require` appears within a module, the *module path resolver* is also given the name of the enclosing module, so that a relative reference can be converted to an absolute symbol or resolved module path.

A module name resolver takes two and four arguments:

- When given two arguments, the first is a name for a module that is now declared in the current namespace, and the second is optionally a namespace from which the declaration was copied. The module name resolver's result in this case is ignored.

  The current module name resolver is called with two arguments by `namespace-attach-module` or `namespace-attach-module-declaration` to notify the resolver that a module declaration was attached to the current namespace (and should not be loaded in the future for the namespace's module registry). Evaluation of a module declaration also calls the current module name resolver with two arguments, where the first is the declared module and the second is `#f`. No other Racket operation invokes the module name resolver with two arguments, but other tools (such as DrRacket) might call this resolver in this mode to avoid redundant module loads.

- When given four arguments, the first is a module path, equivalent to a quoted *`module-path`* for `require`. The second is name for the source module, if any, to which the path is relative; if the second argument is `#f`, the module path is relative to `(or (current-load-relative-directory) (current-directory))`. The third argument is a syntax object that can be used for error reporting, if it is not `#f`. If the last argument is `#t`, then the module declaration should be loaded (if it is not already), otherwise the module path should be simply resolved to a name. The result is the resolved name.

For the second case, the standard module name resolver keeps a table per module registry containing loaded module name. If a resolved module path is not in the table, and `#f` is not provided as the fourth argument to the module name resolver, then the name is put into the table and the corresponding file is loaded with a variant of `load/use-compiled` that passes the expected module name to the compiled-load handler.

While loading a file, the default module name resolver sets the `current-module-declare-name` parameter to the resolved module name (while the compiled-load handler sets `current-module-declare-source`). Also, the default module name resolver records in a private continuation mark the module being loaded, and it checks whether such a mark already exists; if such a continuation mark does exist in the current continuation, then the `exn:fail` exception is raised with a message about a dependency cycle.

The default module name resolver cooperates with the default compiled-load handler: on a module-attach notification, bytecode-file information recorded by the compiled-load handler for the source namespace's module registry is transferred to the target namespace's module registry.

Module loading is suppressed (i.e., `#f` is supplied as a fourth argument to the module name resolver) when resolving module paths in syntax objects (see §1.2.2 "Syntax Objects"). When a syntax object is manipulated, the current namespace might not match the original namespace for the syntax object, and the module should not necessarily be loaded in the current namespace.

```
(current-module-declare-name)
 → (or/c resolved-module-path? #f)
(current-module-declare-name name) → void?
  name : (or/c resolved-module-path? #f)
```

A parameter that determines a module name that is used when evaluating a `module` declaration (when the parameter value is not `#f`). In that case, the *id* from the `module` declaration is ignored, and the parameter's value is used as the name of the declared module.

When declaring submodules, `current-module-declare-name` determines the name used for the submodule's root module, while its submodule path relative to the root module is unaffected.

```
(current-module-declare-source)
 → (or/c symbol? (and/c path? complete-path?) #f)
(current-module-declare-source src) → void?
  src : (or/c symbol? (and/c path? complete-path?) #f)
```

A parameter that determines source information to be associated with a module when evaluating a `module` declaration. Source information is used in error messages and reflected by `variable-reference->module-source`. When the parameter value is `#f`, the module's name (as determined by `current-module-declare-name`) is used as the source name instead of the parameter value.

13.4.2 Compiled Modules and References

While expanding a `module` declaration, the expander resolves module paths for imports to load module declarations as necessary and to determine imported bindings, but the compiled form of a `module` declaration preserves the original module path. Consequently, a compiled module can be moved to another filesystem, where the module name resolver can resolve inter-module references among compiled code.

When a module reference is extracted from compiled form (see `module-compiled-imports`) or from syntax objects in macro expansion (see §11.2 "Syntax Object Content"), the module reference is reported in the form of a *module path index*. A module path index is a semi-interned (multiple references to the same relative module tend to use the same module path index value, but not always) opaque value that encodes a module path (see `module-path?`) and either a resolved module path or another module path index to which it is relative.

A module path index that uses both `#f` for its path and base module path index represents "self"—i.e., the module declaration that was the source of the module path index—and such a module path index can be used as the root for a chain of module path indexes at compile

time. For example, when extracting information about an identifier's binding within a module, if the identifier is bound by a definition within the same module, the identifier's source module is reported using the "self" module path index. If the identifier is instead defined in a module that is imported via a module path (as opposed to a literal module name), then the identifier's source module will be reported using a module path index that contains the `required` module path and the "self" module path index. A "self" module path index has a submodule path when the module that it refers to is a submodule.

A module path index has state. When it is *resolved* to a resolved module path, then the resolved module path is stored with the module path index. In particular, when a module is loaded, its root module path index is resolved to match the module's declaration-time name. This resolved path is forgotten, however, in identifiers that the module contributes to the compiled and marshaled form of other modules. The transient nature of resolved names allows the module code to be loaded with a different resolved name than the name when it was compiled.

```
(module-path-index? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a module path index, `#f` otherwise.

```
(module-path-index-resolve mpi) → resolved-module-path?
  mpi : module-path-index?
```

Returns a resolved module path for the resolved module name, computing the resolved name (and storing it in *mpi*) if it has not been computed before.

Resolving a module path index uses the current module name resolver (see `current-module-name-resolver`). Depending on the kind of module paths encapsulated by *mpi*, the computed resolved name can depend on the value of `current-load-relative-directory` or `current-directory`.

```
(module-path-index-split mpi)
  → (or/c module-path? #f)
    (or/c module-path-index? resolved-module-path? #f)
  mpi : module-path-index?
```

Returns two values: a module path, and a base path—either a module path index, resolved module path, or `#f`—to which the first path is relative.

A `#f` second result means that the path is relative to an unspecified directory (i.e., its resolution depends on the value of `current-load-relative-directory` and/or `current-directory`).

A `#f` for the first result implies a `#f` for the second result, and means that *mpi* represents "self" (see above). Such a module path index may have a non-`#f` submodule path as reported by `module-path-index-submodule`.

```
(module-path-index-submodule mpi)
 → (or/c #f (non-empty-listof symbol?))
  mpi : module-path-index?
```

Returns a non-empty list of symbols if *mpi* is a "self" (see above) module path index that refers to a submodule. The result is always `#f` if either result of `(module-path-index-split mpi)` is non-`#f`.

```
(module-path-index-join path base [submod]) → module-path-index?
  path : (or/c module-path? #f)
  base : (or/c module-path-index? resolved-module-path? #f)
  submod : (or/c #f (non-empty-listof symbol?)) = #f
```

Combines *path*, *base*, and *submod* to create a new module path index. The *path* argument can `#f` only if *base* is also `#f`. The *submod* argument can be a list only when *path* and *base* are both `#f`.

```
(compiled-module-expression? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a compiled `module` declaration, `#f` otherwise. See also `current-compile`.

```
(module-compiled-name compiled-module-code)
 → (or/c symbol? (cons/c symbol? (non-empty-listof symbol?)))
  compiled-module-code : compiled-module-expression?
(module-compiled-name compiled-module-code
                      name)
 → compiled-module-expression?
  compiled-module-code : compiled-module-expression?
  name : (or/c symbol? (cons/c symbol? (non-empty-listof symbol?)))
```

Takes a module declaration in compiled form and either gets the module's declared name (when *name* is not provided) or returns a revised module declaration with the given *name*.

The name is a symbol for a top-level module, and it list of symbols for a submodule, where a list reflects the submodule path to the module starting with the top-level module's declared name.

```
(module-compiled-submodules compiled-module-code
                            non-star?)
 → (listof compiled-module-expression?)
  compiled-module-code : compiled-module-expression?
  non-star? : any/c
```

```
(module-compiled-submodules compiled-module-code
                            non-star?
                            submodules)
 → compiled-module-expression?
  compiled-module-code : compiled-module-expression?
  non-star? : any/c
  submodules : (listof compiled-module-expression?)
```

Takes a module declaration in compiled form and either gets the module's submodules (when *submodules* is not provided) or returns a revised module declaration with the given *submodules*. The `pre-module?` argument determines whether the result or new submodule list corresponds to `module` declarations (when *non-star?* is true) or `module*` declarations (when *non-star?* is `#f`).

```
(module-compiled-imports compiled-module-code)
 → (listof (cons/c (or/c exact-integer? #f)
                   (listof module-path-index?)))
  compiled-module-code : compiled-module-expression?
```

Takes a module declaration in compiled form and returns an association list mapping phase level shifts (where `#f` corresponds to a shift into the label phase level) to module references for the module's explicit imports.

```
(module-compiled-exports compiled-module-code)
 → (listof (cons/c (or/c exact-integer? #f) list?))
   (listof (cons/c (or/c exact-integer? #f) list?))
  compiled-module-code : compiled-module-expression?
```

Returns two association lists mapping phase level values (where `#f` corresponds to the label phase level) to exports at the corresponding phase. The first association list is for exported variables, and the second is for exported syntax. Beware however, that value bindings re-exported though a rename transformer are in the syntax list instead of the value list.

Each associated list, which is represented by `list?` in the result contracts above, more precisely matches the contract

```
(listof (list/c symbol?
                (listof
                 (or/c module-path-index?
                       (list/c module-path-index?
                               (or/c exact-integer? #f)
                               symbol?
                               (or/c exact-integer? #f))))))
```

For each element of the list, the leading symbol is the name of the export.

The second part—the list of module path index values, etc.—describes the origin of the exported identifier. If the origin list is `null`, then the exported identifier is defined in the module. If the exported identifier is re-exported, instead, then the origin list provides information on the import that was re-exported. The origin list has more than one element if the binding was imported multiple times from (possibly) different sources.

For each origin, a module path index by itself means that the binding was imported with a phase level shift of `0` (i.e., a plain `require` without `for-meta`, `for-syntax`, etc.), and imported identifier has the same name as the re-exported name. An origin represented with a list indicates explicitly the import, the import phase level shift (where `#f` corresponds to a `for-label` import), the import name of the re-exported binding, and the phase level of the import.}

```
(module-compiled-language-info compiled-module-code)
 → (or/c #f (vector/c module-path? symbol? any/c))
  compiled-module-code : compiled-module-expression?
```

See also §17.3.6 "Module-Handling Configuration" in ***The Racket Guide***.

Returns information intended to reflect the "language" of the module's implementation as originally attached to the syntax of the module's declaration though the `'module-language` syntax property. See also `module`.

If no information is available for the module, the result is `#f`. Otherwise, the result is `(vector mp name val)` such that `((dynamic-require mp name) val)` should return function that takes two arguments. The function's arguments are a key for reflected information and a default value. Acceptable keys and the interpretation of results is up to external tools, such as DrRacket. If no information is available for a given key, the result should be the given default value.

See also `module->language-info` and `racket/language-info`.

13.4.3 Dynamic Module Access

```
(dynamic-require mod provided [fail-thunk]) → any
       (or/c module-path?
  mod :       resolved-module-path?
              module-path-index?)
  provided : (or/c symbol? #f 0 void?)
  fail-thunk : (-> any) = (lambda () ....)
```

Dynamically instantiates the module specified by *mod* in the current namespace's registry at the namespace's base phase, if it is not yet instantiated. The current module name resolver may load a module declaration to resolve *mod* (see `current-module-name-resolver`); the path is resolved relative to `current-load-relative-directory` and/or `current-directory`.

If *provided* is #f, then the result is #<void>, and the module is not visited (see §1.2.3.8 “Module Phases and Visits”) or even made available (for on-demand visits) in phases above the base phase.

When *provided* is a symbol, the value of the module’s export with the given name is returned, and still the module is not visited or made available in higher phases. If the module exports *provided* as syntax, then a use of the binding is expanded and evaluated in a fresh namespace to which the module is attached, which means that the module is visited in the fresh namespace. If the module has no such exported variable or syntax, then *fail-thunk* is called; the default *fail-thunk* raises exn:fail:contract. If the variable named by *provided* is exported protected (see §13.10 “Code Inspectors”), then the exn:fail:contract exception is raised.

If *provided* is 0, then the module is instantiated but not visited, the same as when *provided* is #f. With 0, however, the module is made available in higher phases.

If *provided* is #<void>, then the module is visited but not instantiated (see §1.2.3.8 “Module Phases and Visits”), and the result is #<void>.

```
(dynamic-require-for-syntax  mod
                             provided
                            [fail-thunk]) → any
  mod : module-path?
  provided : (or/c symbol? #f)
  fail-thunk : (-> any) = (lambda () ....)
```

Like dynamic-require, but in a phase that is 1 more than the namespace’s base phase.

```
(module-declared? mod [load?]) → boolean?
        (or/c module-path? module-path-index?
  mod :       resolved-module-path?)
  load? : any/c = #f
```

Returns #t if the module indicated by *mod* is declared (but not necessarily instantiated or visited) in the current namespace, #f otherwise.

If *load?* is #t and *mod* is not a resolved module path, the module is loaded in the process of resolving *mod* (as for dynamic-require and other functions). Checking for the declaration of a submodule does not trigger an exception if the submodule cannot be loaded because it does not exist, either within a root module that does exist or because the root module does not exist.

```
(module->language-info mod [load?])
 → (or/c #f (vector/c module-path? symbol? any/c))
        (or/c module-path? module-path-index?
  mod :       resolved-module-path?)
  load? : any/c = #f
```

Returns information intended to reflect the "language" of the implementation of *mod*. If *mod* is a resolved module path or *load?* is #f, the module named by *mod* must be declared (but not necessarily instantiated or visited) in the current namespace; otherwise, *mod* may be loaded (as for dynamic-require and other functions). The information returned by module->language-info is the same as would have been returned by module-compiled-language-info applied to the module's implementation as compiled code.

```
(module->imports mod)
  → (listof (cons/c (or/c exact-integer? #f)
                    (listof module-path-index?)))
  mod : (or/c module-path? module-path-index?
              resolved-module-path?)
```

Like module-compiled-imports, but produces the imports of *mod*, which must be declared (but not necessarily instantiated or visited) in the current namespace.

```
(module->exports mod)
  → (listof (cons/c (or/c exact-integer? #f) list?))
    (listof (cons/c (or/c exact-integer? #f) list?))
  mod : (or/c module-path? resolved-module-path?)
```

Like module-compiled-exports, but produces the exports of *mod*, which must be declared (but not necessarily instantiated or visited) in the current namespace.

```
(module-predefined? mod) → boolean?
  mod : (or/c module-path? resolved-module-path?)
```

Reports whether *mod* refers to a module that is predefined for the running Racket instance. Predefined modules always have a symbolic resolved module path, and they may be predefined always or specifically within a particular executable (such as one created by `raco exe` or create-embedding-executable).

13.5 Impersonators and Chaperones

An *impersonator* is a wrapper for a value where the wrapper redirects some of the value's operations. Impersonators apply only to procedures, structures for which an accessor or mutator is available, structure types, hash tables, vectors, boxes, and prompt tags. An impersonator is equal? to the original value, but not eq? to the original value.

A *chaperone* is a kind of impersonator whose refinement of a value's operation is restricted to side effects (including, in particular, raising an exception) or chaperoning values supplied to or produced by the operation. For example, a vector chaperone can redirect vector-ref to raise an exception if the accessed vector slot contains a string, or it can cause the result of vector-ref to be a chaperoned variant of the value that is in the accessed vector slot, but it

cannot redirect vector-ref to produce a value that is arbitrarily different from the value in the vector slot.

A non-chaperone impersonator, in contrast, can refine an operation to swap one value for any other. An impersonator cannot be applied to an immutable value or refine the access to an immutable field in an instance of a structure type, since arbitrary redirection of an operation amounts to mutation of the impersonated value.

Beware that each of the following operations can be redirected to an arbitrary procedure through an impersonator on the operation's argument—assuming that the operation is available to the creator of the impersonator:

- a structure-field accessor
- a structure-field mutator
- a structure type property accessor
- application of a procedure
- unbox
- set-box!
- vector-ref
- vector-set!
- hash-ref
- hash-set
- hash-set!
- hash-remove
- hash-remove!
- call-with-continuation-prompt
- abort-current-continuation

Derived operations, such as printing a value, can be redirected through impersonators due to their use of accessor functions. The equal?, equal-hash-code, and equal-secondary-hash-code operations, in contrast, may bypass impersonators (but they are not obliged to).

In addition to redirecting operations that work on a value, a impersonator can include *impersonator properties* for an impersonated value. An impersonator property is similar to a structure type property, but it applies to impersonators instead of structure types and their instances.

```
(impersonator? v) → boolean?
  v : any/c
```

Returns #t if v is an impersonator, #f otherwise.

Programs and libraries generally should avoid impersonator? and treat impersonators the same as non-impersonator values. In rare cases, impersonator? may be needed to guard against redirection by an impersonator of an operation to an arbitrary procedure.

```
(chaperone? v) → boolean?
  v : any/c
```

Returns #t if v is a chaperone, #f otherwise.

Programs and libraries generally should avoid chaperone? for the same reason that they should avoid impersonator?.

```
(impersonator-of? v1 v2) → boolean?
  v1 : any/c
  v2 : any/c
```

Indicates whether v1 can be considered equivalent modulo impersonators to v2.

For values that include no impersonators, v1 and v2 can be considered impersonators of each other if they are equal?.

Otherwise, all impersonators of v2 must be intact in v1, in the sense that parts of v2 must be derived from v1 through one of the impersonator constructors (e.g., impersonate-procedure or chaperone-procedure).

See also prop:impersonator-of.

```
(chaperone-of? v1 v2) → boolean?
  v1 : any/c
  v2 : any/c
```

Indicates whether v1 can be considered equivalent modulo chaperones to v2.

For values that include no chaperones, v1 and v2 can be considered chaperones of each other if they are equal?, except that the mutability of vectors and boxes with v1 and v2 must be the same.

Otherwise, all chaperones of v2 must be intact in v1, in the sense that parts of v2 must be derived from v1 through one of the chaperone constructors (e.g., chaperone-procedure).

13.5.1 Impersonator Constructors

```
(impersonate-procedure proc
                       wrapper-proc
                       prop
                       prop-val ...
                       ...)
 → (and/c procedure? impersonator?)
  proc : procedure?
  wrapper-proc : procedure?
  prop : impersonator-property?
  prop-val : any
```

Returns an impersonator procedure that has the same arity, name, and other attributes as *proc*. When the impersonator procedure is applied, the arguments are first passed to *wrapper-proc*, and then the results from *wrapper-proc* are passed to *proc*. The *wrapper-proc* can also supply a procedure that processes the results of *proc*.

The arity of *wrapper-proc* must include the arity of *proc*. The allowed keyword arguments of *wrapper-proc* must be a superset of the allowed keywords of *proc*. The required keyword arguments of *wrapper-proc* must be a subset of the required keywords of *proc*.

For applications without keywords, the result of *wrapper-proc* must be either the same number of values as supplied to it or one more than the number of supplied values, where an extra result is supplied before the others. The additional result, if any, must be a procedure that accepts as many results as produced by *proc*; it must return the same number of results. If *wrapper-proc* returns the same number of values as it is given (i.e., it does not return a procedure to impersonator *proc*'s result), then *proc* is called in tail position with respect to the call to the impersonator.

For applications that include keyword arguments, *wrapper-proc* must return an additional value before any other values but after the result-impersonating procedure (if any). The additional value must be a list of replacements for the keyword arguments that were supplied to the impersonator (i.e., not counting optional arguments that were not supplied). The arguments must be ordered according to the sorted order of the supplied arguments' keywords.

Pairs of *prop* and *prop-val* (the number of arguments to `procedure-impersonator` must be even) add impersonator properties or override impersonator-property values of *proc*.

If any *prop* is `impersonator-prop:application-mark` and if the associated *prop-val* is a pair, then the call to *proc* is wrapped with `with-continuation-mark` using `(car` *prop-val*`)` as the mark key and `(cdr` *prop-val*`)` as the mark value. In addition, if `continuation-mark-set-first` with `(car` *prop-val*`)` produces a value for the immediate continuation frame of the call to the impersonated procedure, the value is also installed as an immediate value for `(car` *prop-val*`)` as a mark during the call to *wrapper-proc*

(which allows tail-calls of impersonators with respect to wrapping impersonators to be detected within *wrapper-proc*).

```
(impersonate-struct v
                    orig-proc
                    redirect-proc ...
                    ...
                    prop
                    prop-val ...
                    ...)                → any/c
  v : any/c
              (or/c struct-accessor-procedure?
  orig-proc :       struct-mutator-procedure?
                    struct-type-property-accessor-procedure?)
  redirect-proc : procedure?
  prop : impersonator-property?
  prop-val : any
```

Returns an impersonator of *v*, which redirects certain operations on the impersonated value. The *orig-proc*s indicate the operations to redirect, and the corresponding *redirect-proc*s supply the redirections.

The protocol for a *redirect-proc* depends on the corresponding *orig-proc*:

- A structure-field accessor: *redirect-proc* must accept two arguments, *v* and the value *field-v* that *orig-proc* produces for *v*; it must return a replacement for *field-v*. The corresponding field must not be immutable, and either the field's structure type must be accessible via the current inspector or one of the other *orig-proc*s must be a structure-field mutator for the same field.
- A structure-field mutator: *redirect-proc* must accept two arguments, *v* and the value *field-v* supplied to the mutator; it must return a replacement for *field-v* to be propagated to *orig-proc* and *v*.
- A property accessor: *redirect-proc* uses the same protocol as for a structure-field accessor. The accessor's property must have been created with `'can-impersonate` as the second argument to `make-struct-type-property`.

Pairs of *prop* and *prop-val* (the number of arguments to `impersonate-struct` must be odd) add impersonator properties or override impersonator-property values of *v*.

Each *orig-proc* must indicate a distinct operation. If no *orig-proc*s are supplied, then no *prop*s must be supplied, and *v* is returned unimpersonated.

```
(impersonate-vector vec
                    ref-proc
                    set-proc
                    prop
                    prop-val ...
                    ...)
 → (and/c vector? impersonator?)
  vec : (and/c vector? (not/c immutable?))
  ref-proc : (vector? exact-nonnegative-integer? any/c . -> . any/c)
  set-proc : (vector? exact-nonnegative-integer? any/c . -> . any/c)
  prop : impersonator-property?
  prop-val : any
```

Returns an impersonator of *vec*, which redirects the `vector-ref` and `vector-set!` operations.

The *ref-proc* must accept *vec*, an index passed to `vector-ref`, and the value that `vector-ref` on *vec* produces for the given index; it must produce a replacement for the value, which is the result of `vector-ref` on the impersonator.

The *set-proc* must accept *vec*, an index passed to `vector-set!`, and the value passed to `vector-set!`; it must produce a replacement for the value, which is used with `vector-set!` on the original *vec* to install the value.

Pairs of *prop* and *prop-val* (the number of arguments to `impersonate-vector` must be odd) add impersonator properties or override impersonator-property values of *vec*.

```
(impersonate-box box
                 unbox-proc
                 set-proc
                 prop
                 prop-val ...
                 ...)          → (and/c box? impersonator?)
  box : (and/c box? (not/c immutable?))
  unbox-proc : (box? any/c . -> . any/c)
  set-proc : (box? any/c . -> . any/c)
  prop : impersonator-property?
  prop-val : any
```

Returns an impersonator of *box*, which redirects the `unbox` and `set-box!` operations.

The *unbox-proc* must accept *box* and the value that `unbox` produces on *box*; it must produce a replacement value, which is the result of `unbox` on the impersonator.

The *set-proc* must accept *box* and the value passed to `set-box!`; it must produce a replacement value, which is used with `set-box!` on the original *box* to install the value.

Pairs of *prop* and *prop-val* (the number of arguments to impersonate-box must be odd) add impersonator properties or override impersonator-property values of *box*.

```
(impersonate-hash hash
                  ref-proc
                  set-proc
                  remove-proc
                  key-proc
                  prop
                  prop-val ...
                  ...)          → (and/c hash? impersonator?)
  hash : (and/c hash? (not/c immutable?))
             (hash? any/c . -> . (values
  ref-proc :                       any/c
                                   (hash? any/c any/c . -> . any/c)))
  set-proc : (hash? any/c any/c . -> . (values any/c any/c))
  remove-proc : (hash? any/c . -> . any/c)
  key-proc : (hash? any/c . -> . any/c)
  prop : impersonator-property?
  prop-val : any
```

Returns an impersonator of *hash*, which redirects the hash-ref, hash-set! or hash-set (as applicable), and hash-remove or hash-remove! (as applicable) operations. When hash-set or hash-remove is used on an impersonator of a hash table, the result is an impersonator with the same redirecting procedures. In addition, operations like hash-iterate-key or hash-map, which extract keys from the table, use *key-proc* to filter keys extracted from the table. Operations like hash-iterate-value or hash-iterate-map implicitly use hash-ref and therefore redirect through *ref-proc*.

The *ref-proc* must accept *hash* and a key passed to hash-ref. It must return a replacement key as well as a procedure. The returned procedure is called only if the returned key is found in *hash* via hash-ref, in which case the procedure is called with *hash*, the previously returned key, and the found value. The returned procedure must itself return a replacement for the found value.

The *set-proc* must accept *hash*, a key passed to hash-set! or hash-set, and the value passed to hash-set! or hash-set; it must produce two values: a replacement for the key and a replacement for the value. The returned key and value are used with hash-set! or hash-set on the original *hash* to install the value.

The *remove-proc* must accept *hash* and a key passed to hash-remove! or hash-remove; it must produce the a replacement for the key, which is used with hash-remove! or hash-remove on the original *hash* to remove any mapping using the (impersonator-replaced) key.

The *key-proc* must accept *hash* and a key that has been extracted from *hash* (by hash-iterate-key or other operations that use hash-iterate-key internally); it must produce a replacement for the key, which is then reported as a key extracted from the table.

The hash-iterate-value, hash-map, or hash-for-each functions use a combination of hash-iterate-key and hash-ref. If a key produced by *key-proc* does not yield a value through hash-ref, then the exn:fail:contract exception is raised.

Pairs of *prop* and *prop-val* (the number of arguments to impersonate-hash must be odd) add impersonator properties or override impersonator-property values of *hash*.

```
(impersonate-prompt-tag prompt-tag
                        handle-proc
                        abort-proc
                        prop
                        prop-val ...
                        ...)
 → (and/c continuation-prompt-tag? impersonator?)
  prompt-tag : continuation-prompt-tag?
  handle-proc : procedure?
  abort-proc : procedure?
  prop : impersonator-property?
  prop-val : any
```

Returns an impersonator of *prompt-tag*, which redirects the call-with-continuation-prompt and abort-current-continuation operations.

The *handle-proc* must accept the values that the handler of a continuation prompt would take and it must produce replacement values, which will be passed to the handler.

The *abort-proc* must accept the values passed to abort-current-continuation; it must produce replacement values, which are aborted to the appropriate prompt.

Pairs of *prop* and *prop-val* (the number of arguments to impersonate-prompt-tag must be odd) add impersonator properties or override impersonator-property values of *prompt-tag*.

Examples:

```
> (define tag
    (impersonate-prompt-tag
     (make-continuation-prompt-tag)
     (lambda (n) (* n 2))
     (lambda (n) (+ n 1))))

> (call-with-continuation-prompt
    (lambda ()
      (abort-current-continuation tag 5))
    tag
    (lambda (n) n))
12
```

```
(impersonate-continuation-mark-key key
                                   get-proc
                                   set-proc
                                   prop
                                   prop-val ...
                                   ...)
 → (and/c continuation-mark? impersonator?)
  key : continuation-mark-key?
  get-proc : procedure?
  set-proc : procedure?
  prop : impersonator-property?
  prop-val : any
```

Returns an impersonator of *key*, which redirects `with-continuation-mark` and continuation mark accessors such as `continuation-mark-set->list`.

The *get-proc* must accept the value attached to a continuation mark and it must produce a replacement value, which will be returned by the continuation mark accessor.

The *set-proc* must accept a value passed to `with-continuation-mark`; it must produce a replacement value, which is attached to the continuation frame.

Pairs of *prop* and *prop-val* (the number of arguments to `impersonate-prompt-tag` must be odd) add impersonator properties or override impersonator-property values of *key*.

Examples:

```
> (define mark-key
    (impersonate-continuation-mark-key
     (make-continuation-mark-key)
     (lambda (l) (map char-upcase l))
     (lambda (s) (string->list s))))

> (with-continuation-mark mark-key "quiche"
    (continuation-mark-set-first
     (current-continuation-marks)
     mark-key))
'(#\Q #\U #\I #\C #\H #\E)
```

```
prop:impersonator-of : struct-type-property?
```

A structure type property (see §4.3 “Structure Type Properties”) that supplies a procedure for extracting an impersonated value from a structure that represents an impersonator. The property is used for `impersonator-of?` as well as `equal?`.

The property value must be a procedure of one argument, which is a structure whose structure type has the property. The result can be `#f` to indicate the structure does not represent an

impersonator, otherwise the result is a value for which the original structure is an impersonator (so the original structure is an `impersonator-of?` and `equal?` to the result value). The result value must have the same `prop:impersonator-of` and `prop:equal+hash` property values as the original structure, and the property values must be inherited from the same structure type (which ensures some consistency between `impersonator-of?` and `equal?`).

13.5.2 Chaperone Constructors

```
(chaperone-procedure proc
                     wrapper-proc
                     prop
                     prop-val ...
                     ...)
 → (and/c procedure? chaperone?)
  proc : procedure?
  wrapper-proc : procedure?
  prop : impersonator-property?
  prop-val : any
```

Like `impersonate-procedure`, but for each value supplied to *wrapper-proc*, the corresponding result must be the same or a chaperone of (in the sense of `chaperone-of?`) the supplied value. The additional result, if any, that precedes the chaperoned values must be a procedure that accepts as many results as produced by *proc*; it must return the same number of results, each of which is the same or a chaperone of the corresponding original result.

For applications that include keyword arguments, *wrapper-proc* must return an additional value before any other values but after the result-chaperoning procedure (if any). The additional value must be a list of chaperones of the keyword arguments that were supplied to the chaperone procedure (i.e., not counting optional arguments that were not supplied). The arguments must be ordered according to the sorted order of the supplied arguments' keywords.

```
(chaperone-struct v
                  orig-proc
                  redirect-proc ...
                  ...
                  prop
                  prop-val ...
                  ...)              → any/c
  v : any/c
              (or/c struct-accessor-procedure?
                    struct-mutator-procedure?
  orig-proc :       struct-type-property-accessor-procedure?
                    (one-of/c struct-info))
  redirect-proc : procedure?
```

```
  prop : impersonator-property?
  prop-val : any
```

Like impersonate-struct, but with the following refinements:

- With a structure-field accessor as *orig-proc*, *redirect-proc* must accept two arguments, *v* and the value *field-v* that *orig-proc* produces for *v*; it must return a chaperone of *field-v*. The corresponding field may be immutable.
- With structure-field mutator as *orig-proc*, *redirect-proc* must accept two arguments, *v* and the value *field-v* supplied to the mutator; it must return a chaperone of *field-v* to be propagated to *orig-proc* and *v*.
- A property accessor can be supplied as *orig-proc*, and the property need not have been created with 'can-impersonate. The corresponding *redirect-proc* uses the same protocol as for a structure-field accessor.
- With struct-info as *orig-proc*, the corresponding *redirect-proc* must accept two values, which are the results of struct-info on *v*; it must return each values or a chaperone of each value. The *redirect-proc* is not called if struct-info would return #f as its first argument. An *orig-proc* can be struct-info only if some other *orig-proc* is supplied.

```
(chaperone-vector vec
                  ref-proc
                  set-proc
                  prop
                  prop-val ...
                  ...)          → (and/c vector? chaperone?)
  vec : vector?
  ref-proc : (vector? exact-nonnegative-integer? any/c . -> . any/c)
  set-proc : (vector? exact-nonnegative-integer? any/c . -> . any/c)
  prop : impersonator-property?
  prop-val : any
```

Like impersonate-vector, but with support for immutable vectors. The *ref-proc* procedure must produce the same value or a chaperone of the original value, and *set-proc* must produce the value that is given or a chaperone of the value. The *set-proc* will not be used if *vec* is immutable.

```
(chaperone-box box
               unbox-proc
               set-proc
               prop
               prop-val ...
               ...)          → (and/c box? chaperone?)
```

```
  box : box?
  unbox-proc : (box? any/c . -> . any/c)
  set-proc : (box? any/c . -> . any/c)
  prop : impersonator-property?
  prop-val : any
```

Like impersonate-box, but with support for immutable boxes. The *unbox-proc* procedure must produce the same value or a chaperone of the original value, and *set-proc* must produce the same value or a chaperone of the value that it is given. The *set-proc* will not be used if *box* is immutable.

```
(chaperone-hash hash
                ref-proc
                set-proc
                remove-proc
                key-proc
                prop
                prop-val ...
                ...)          → (and/c hash? chaperone?)
  hash : hash?
  ref-proc :  (hash? any/c . -> . (values
                                   any/c
                                   (hash? any/c any/c . -> . any/c)))
  set-proc : (hash? any/c any/c . -> . (values any/c any/c))
  remove-proc : (hash? any/c . -> . any/c)
  key-proc : (hash? any/c . -> . any/c)
  prop : impersonator-property?
  prop-val : any
```

Like impersonate-hash, but with constraints on the given functions and support for immutable hashes. The *ref-proc* procedure must return a found value or a chaperone of the value. The *set-proc* procedure must produce two values: the key that it is given or a chaperone of the key and the value that it is given or a chaperone of the value. The *remove-proc* and *key-proc* procedures must produce the given key or a chaperone of the key.

```
(chaperone-struct-type struct-type
                       struct-info-proc
                       make-constructor-proc
                       guard-proc
                       prop
                       prop-val ...
                       ...)
 → (and/c struct-type? chaperone?)
  struct-type : struct-type?
  struct-info-proc : procedure?
  make-constructor-proc : (procedure? . -> . procedure?)
```

```
  guard-proc : procedure?
  prop : impersonator-property?
  prop-val : any
```

Returns a chaperoned value like *struct-type*, but with struct-type-info and struct-type-make-constructor operations on the chaperoned structure type redirected. In addition, when a new structure type is created as a subtype of the chaperoned structure type, *guard-proc* is interposed as an extra guard on creation of instances of the subtype.

The *struct-info-proc* must accept 8 arguments—the result of struct-type-info on *struct-type*. It must return 8 values, where each is the same or a chaperone of the corresponding argument. The 8 values are used as the results of struct-type-info for the chaperoned structure type.

The *make-constructor-proc* must accept a single procedure argument, which is a constructor produced by struct-type-make-constructor on *struct-type*. It must return the same or a chaperone of the procedure, which is used as the result of struct-type-make-constructor on the chaperoned structure type.

The *guard-proc* must accept as many argument as a constructor for *struct-type*; it must return the same number of arguments, each the same or a chaperone of the corresponding argument. The *guard-proc* is added as a constructor guard when a subtype is created of the chaperoned structure type.

Pairs of *prop* and *prop-val* (the number of arguments to chaperone-struct-type must be even) add impersonator properties or override impersonator-property values of *struct-type*.

```
(chaperone-evt evt proc prop prop-val ... ...)
 → (and/c evt? chaperone?)
  evt : evt?
  proc : (evt? . -> . (values evt? (any/c . -> . any/c)))
  prop : impersonator-property?
  prop-val : any
```

Returns a chaperoned value like *evt*, but with *proc* as an event generator when the result is synchronized with functions like sync.

The *proc* generator is called on synchronization, much like the procedure passed to guard-evt, except that *proc* is given *evt*. The *proc* must return two values: a synchronizable event that is a chaperone of *evt*, and a procedure that is used to check the event's result if it is chosen in a selection. The latter procedure accepts the result of *evt*, and it must return a chaperone of that value.

Pairs of *prop* and *prop-val* (the number of arguments to chaperone-evt must be even) add impersonator properties or override impersonator-property values of *evt*.

```
(chaperone-prompt-tag prompt-tag
                      handle-proc
                      abort-proc
                      prop
                      prop-val ...
                      ...)
 → (and/c continuation-prompt-tag? chaperone?)
  prompt-tag : continuation-prompt-tag?
  handle-proc : procedure?
  abort-proc : procedure?
  prop : impersonator-property?
  prop-val : any
```

Like `impersonate-prompt-tag`, but produces a chaperoned value. The *handle-proc* procedure must produce the same values or chaperones of the original values, and *abort-proc* must produce the same values or chaperones of the values that it is given.

Examples:

```
> (define bad-chaperone
    (chaperone-prompt-tag
     (make-continuation-prompt-tag)
     (lambda (n) (* n 2))
     (lambda (n) (+ n 1))))

> (call-with-continuation-prompt
    (lambda ()
      (abort-current-continuation bad-chaperone 5))
    bad-chaperone
    (lambda (n) n))
abort-current-continuation: non-chaperone result;
 received a value that is not a chaperone of the original
value
  original: 5
  received: 6
> (define good-chaperone
    (chaperone-prompt-tag
     (make-continuation-prompt-tag)
     (lambda (n) (if (even? n) n (error "not even")))
     (lambda (n) (if (even? n) n (error "not even")))))

> (call-with-continuation-prompt
    (lambda ()
      (abort-current-continuation good-chaperone 2))
    good-chaperone
```

```
    (lambda (n) n))
2
```

```
(chaperone-continuation-mark-key key
                                 get-proc
                                 set-proc
                                 prop
                                 prop-val ...
                                 ...)
 → (and/c continuation-mark-key? chaperone?)
  key : continuation-mark-key?
  get-proc : procedure?
  set-proc : procedure?
  prop : impersonator-property?
  prop-val : any
```

Like `impersonate-continuation-mark-key`, but produces a chaperoned value. The *get-proc* procedure must produce the same value or a chaperone of the original value, and *set-proc* must produce the same value or a chaperone of the value that it is given.

Examples:

```
> (define bad-chaperone
    (chaperone-continuation-mark-key
     (make-continuation-mark-key)
     (lambda (l) (map char-upcase l))
     string->list))

> (with-continuation-mark bad-chaperone "timballo"
    (continuation-mark-set-first
     (current-continuation-marks)
     bad-chaperone))
with-continuation-mark: non-chaperone result;
 received a value that is not a chaperone of the original
value
  original: "timballo"
  received: '(#\t #\i #\m #\b #\a #\l #\l #\o)
> (define (checker s)
    (if (> (string-length s) 5)
        s
        (error "expected string of length at least 5")))

> (define good-chaperone
    (chaperone-continuation-mark-key
     (make-continuation-mark-key)
     checker
     checker))
```

```
> (with-continuation-mark good-chaperone "zabaione"
    (continuation-mark-set-first
     (current-continuation-marks)
     good-chaperone))
"zabaione"
```

13.5.3 Impersonator Properties

```
                                        impersonator-property?
(make-impersonator-property name) → (-> any/c boolean?)
                                        (-> impersonator? any)
  name : symbol?
```

Creates a new impersonator property and returns three values:

- an *impersonator property descriptor*, for use with `impersonate-procedure`, `chaperone-procedure`, and other impersonator constructors;
- an *impersonator property predicate* procedure, which takes an arbitrary value and returns `#t` if the value is an impersonator with a value for the property, `#f` otherwise;
- an *impersonator property accessor* procedure, which returns the value associated with an impersonator for the property; if a value given to the accessor is not an impersonator or does not have a value for the property (i.e. if the corresponding impersonator property predicate returns `#f`), the `exn:fail:contract` exception is raised.

```
(impersonator-property? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a impersonator property descriptor value, `#f` otherwise.

```
(impersonator-property-accessor-procedure? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is an accessor procedure produced by `make-impersonator-property`, `#f` otherwise.

```
impersonator-prop:application-mark : impersonator-property?
```

An impersonator property that is recognized by `impersonate-procedure` and `chaperone-procedure`.

13.6 Security Guards

```
(security-guard? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a security guard value as created by `make-security-guard`, `#f` otherwise.

A *security guard* provides a set of access-checking procedures to be called when a thread initiates access of a file, directory, or network connection through a primitive procedure. For example, when a thread calls `open-input-file`, the thread's current security guard is consulted to check whether the thread is allowed read access to the file. If access is granted, the thread receives a port that it may use indefinitely, regardless of changes to the security guard (although the port's custodian could shut down the port; see §13.7 "Custodians").

A thread's current security guard is determined by the `current-security-guard` parameter. Every security guard has a parent, and a parent's access procedures are called whenever a child's access procedures are called. Thus, a thread cannot increase its own access arbitrarily by installing a new guard. The initial security guard enforces no access restrictions other than those enforced by the host platform.

```
(make-security-guard parent
                     file-guard
                     network-guard
                     [link-guard])  → security-guard?
  parent : security-guard?
  file-guard : (symbol?
                (or/c path? #f)
                (listof symbol?)
                . -> . any)
  network-guard : (symbol?
                   (or/c (and/c string? immutable?) #f)
                   (or/c (integer-in 1 65535) #f)
                   (or/c 'server 'client)
                   . -> . any)
  link-guard : (or/c (symbol? path? path? . -> . any) #f) = #f
```

Creates a new security guard as child of *parent*.

The *file-guard* procedure must accept three arguments:

- a symbol for the primitive procedure that triggered the access check, which is useful for raising an exception to deny access.
- a path (see §14.1 "Paths") or `#f` for pathless queries, such as `(current-directory)`, `(filesystem-root-list)`, and `(find-system-path symbol)`. A path provided

to *file-guard* is not expanded or otherwise normalized before checking access; it may be a relative path, for example.

- a list containing one or more of the following symbols:
  - 'read — read a file or directory
  - 'write — modify or create a file or directory
  - 'execute — execute a file
  - 'delete — delete a file or directory
  - 'exists — determine whether a file or directory exists, or that a path string is well-formed

  The 'exists symbol is never combined with other symbols in the last argument to *file-guard*, but any other combination is possible. When the second argument to *file-guard* is #f, the last argument always contains only 'exists.

The *network-guard* procedure must accept four arguments:

- a symbol for the primitive operation that triggered the access check, which is useful for raising an exception to deny access.
- an immutable string representing the target hostname for a client connection or the accepting hostname for a listening server; #f for a listening server or UDP socket that accepts connections at all of the host's address; or #f an unbound UDP socket.
- an exact integer between 1 and 65535 (inclusive) representing the port number, or #f for an unbound UDP socket. In the case of a client connection, the port number is the target port on the server. For a listening server, the port number is the local port number.
- a symbol, either 'client or 'server, indicating whether the check is for the creation of a client connection or a listening server. The opening of an unbound UDP socket is identified as a 'client connection; explicitly binding the socket is identified as a 'server action.

The *link-guard* argument can be #f or a procedure of three arguments:

- a symbol for the primitive procedure that triggered the access check, which is useful for raising an exception to deny access.
- a complete path (see §14.1 "Paths") representing the file to create as link.
- a path representing the content of the link, which may be relative the second-argument path; this path is not expanded or otherwise normalized before checking access.

If *link-guard* is #f, then a default procedure is used that always raises exn:fail.

The return value of *file-guard*, *network-guard*, or *link-guard* is ignored. To deny access, the procedure must raise an exception or otherwise escape from the context of the primitive call. If the procedure returns, the parent's corresponding procedure is called on the same inputs, and so on up the chain of security guards.

The *file-guard*, *network-guard*, and *link-guard* procedures are invoked in the thread that called the access-checked primitive. Breaks may or may not be enabled (see §9.6 "Breaks"). Full continuation jumps are blocked going into or out of the *file-guard* or *network-guard* call (see §1.1.12 "Prompts, Delimited Continuations, and Barriers").

```
(current-security-guard) → security-guard?
(current-security-guard guard) → void?
  guard : security-guard?
```

A parameter that determines the current security guard that controls access to the filesystem and network.

13.7 Custodians

See §1.1.16 "Custodians" for basic information on the Racket custodian model.

```
(custodian? v) → boolean?
  v : any/c
```

Returns #t if *v* is a custodian value, #f otherwise.

```
(make-custodian [cust]) → custodian?
  cust : custodian? = (current-custodian)
```

Creates a new custodian that is subordinate to *cust*. When *cust* is directed (via custodian-shutdown-all) to shut down all of its managed values, the new subordinate custodian is automatically directed to shut down its managed values as well.

```
(custodian-shutdown-all cust) → void?
  cust : custodian?
```

In GRacket, eventspaces managed by *cust* are also shut down.

Closes all file-stream ports, TCP ports, TCP listeners, and UDP sockets that are managed by *cust* (and its subordinates), and empties all custodian boxes associated with *cust* (and its subordinates). It also removes *cust* (and its subordinates) as managers of all threads; when a thread has no managers, it is killed (or suspended; see thread/suspend-to-kill) If the current thread is to be killed, all other shut-down actions take place before killing the thread.

```
(current-custodian) → custodian?
(current-custodian cust) → void?
  cust : custodian?
```

In GRacket, custodians also manage eventspaces.

A parameter that determines a custodian that assumes responsibility for newly created threads, file-stream ports, TCP ports, TCP listeners, UDP sockets, and byte converters.

```
(custodian-managed-list cust super) → list?
  cust : custodian?
  super : custodian?
```

Returns a list of immediately managed objects (not including custodian boxes) and subordinate custodians for *cust*, where *cust* is itself subordinate to *super* (directly or indirectly). If *cust* is not strictly subordinate to *super*, the `exn:fail:contract` exception is raised.

```
(custodian-memory-accounting-available?) → boolean?
```

Memory accounting is normally available in Racket 3m, which is the main variant of Racket, and not normally available in Racket CGC.

Returns `#t` if Racket is compiled with support for per-custodian memory accounting, `#f` otherwise.

```
(custodian-require-memory limit-cust
                          need-amt
                          stop-cust) → void?
  limit-cust : custodian?
  need-amt : exact-nonnegative-integer?
  stop-cust : custodian?
```

Registers a required-memory check if Racket is compiled with support for per-custodian memory accounting, otherwise the `exn:fail:unsupported` exception is raised.

If a check is registered, and if Racket later reaches a state after garbage collection (see §1.1.7 “Garbage Collection”) where allocating *need-amt* bytes charged to *limit-cust* would fail or trigger some shutdown, then *stop-cust* is shut down.

```
(custodian-limit-memory  limit-cust
                         limit-amt
                        [stop-cust]) → void?
  limit-cust : custodian?
  limit-amt : exact-nonnegative-integer?
  stop-cust : custodian? = limit-cust
```

Registers a limited-memory check if Racket is compiled with support for per-custodian memory accounting, otherwise the `exn:fail:unsupported` exception is raised.

If a check is registered, and if Racket later reaches a state after garbage collection (see §1.1.7 “Garbage Collection”) where *limit-cust* owns more than *limit-amt* bytes, then *stop-cust* is shut down.

A custodian's limit is checked only after a garbage collection, except that it may also be checked during certain large allocations that are individually larger than the custodian's limit. A single garbage collection may shut down multiple custodians, even if shutting down only one of the custodians would have reduced memory use for other custodians.

For reliable shutdown, *limit-amt* for `custodian-limit-memory` must be much lower than the total amount of memory available (minus the size of memory that is potentially used and not charged to *limit-cust*). Moreover, if individual allocations that are initially charged to *limit-cust* can be arbitrarily large, then *stop-cust* must be the same as *limit-cust*, so that excessively large immediate allocations can be rejected with an `exn:fail:out-of-memory` exception.

```
(make-custodian-box cust v) → custodian-box?
  cust : custodian?
  v : any/c
```

Returns a custodian box that contains *v* as long as *cust* has not been shut down.

A custodian box is a synchronizable event for use with functions like `sync`. The custodian box becomes ready when its custodian is shut down.

```
(custodian-box? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a custodian box produced by `make-custodian-box`, `#f` otherwise.

```
(custodian-box-value cb) → any
  cb : custodian-box?
```

Returns the value in the given custodian box, or `#f` if the value has been removed.

13.8 Thread Groups

A *thread group* is a collection of threads and other thread groups that have equal claim to the CPU. By nesting thread groups and by creating certain threads within certain groups, a programmer can control the amount of CPU allocated to a set of threads. Every thread belongs to a thread group, which is determined by the `current-thread-group` parameter when the thread is created. Thread groups and custodians (see §13.7 “Custodians”) are independent.

The root thread group receives all of the CPU that the operating system gives Racket. Every thread or nested group in a particular thread group receives equal allocation of the CPU (a portion of the group's access), although a thread may relinquish part of its allocation by sleeping or synchronizing with other processes.

```
(make-thread-group [group]) → thread-group?
  group : thread-group? = (current-thread-group)
```

Creates a new thread group that belongs to *group*.

```
(thread-group? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a thread group value, `#f` otherwise.

```
(current-thread-group) → thread-group?
(current-thread-group group) → void?
  group : thread-group?
```

A parameter that determines the thread group for newly created threads.

13.9 Structure Inspectors

An *inspector* provides access to structure fields and structure type information without the normal field accessors and mutators. (Inspectors are also used to control access to module bindings; see §13.10 "Code Inspectors".) Inspectors are primarily intended for use by debuggers.

When a structure type is created, an inspector can be supplied. The given inspector is not the one that will control the new structure type; instead, the given inspector's parent will control the type. By using the parent of the given inspector, the structure type remains opaque to "peer" code that cannot access the parent inspector.

The `current-inspector` parameter determines a default inspector argument for new structure types. An alternate inspector can be provided though the `#:inspector` option of the `define-struct` form (see §4.1 "Defining Structure Types: `struct`"), or through an optional `inspector` argument to `make-struct-type`.

```
(inspector? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is an inspector, `#f` otherwise.

```
(make-inspector [inspector]) → inspector?
  inspector : inspector? = (current-inspector)
```

Returns a new inspector that is a subinspector of *inspector*. Any structure type controlled by the new inspector is also controlled by its ancestor inspectors, but no other inspectors.

```
(make-sibling-inspector [inspector]) → inspector?
  inspector : inspector? = (current-inspector)
```

Returns a new inspector that is a subinspector of the same inspector as *inspector*. That is, *inspector* and the result inspector control mutually disjoint sets of structure types.

```
(current-inspector) → inspector?
(current-inspector insp) → void?
  insp : inspector?
```

A parameter that determines the default inspector for newly created structure types.

```
(struct-info v) → (or/c struct-type? #f) boolean?
  v : any/c
```

Returns two values:

- *struct-type*: a structure type descriptor or `#f`; the result is a structure type descriptor of the most specific type for which *v* is an instance, and for which the current inspector has control, or the result is `#f` if the current inspector does not control any structure type for which the `struct` is an instance.
- *skipped?*: `#f` if the first result corresponds to the most specific structure type of *v*, `#t` otherwise.

```
(struct-type-info struct-type)
    symbol?
    exact-nonnegative-integer?
    exact-nonnegative-integer?
    struct-accessor-procedure?
→
    struct-mutator-procedure?
    (listof exact-nonnegative-integer?)
    (or/c struct-type? #f)
    boolean?
  struct-type : struct-type?
```

Returns eight values that provide information about the structure type descriptor *struct-type*, assuming that the type is controlled by the current inspector:

- *name*: the structure type's name as a symbol;
- *init-field-cnt*: the number of fields defined by the structure type provided to the constructor procedure (not counting fields created by its ancestor types);

- *auto-field-cnt*: the number of fields defined by the structure type without a counterpart in the constructor procedure (not counting fields created by its ancestor types);
- *accessor-proc*: an accessor procedure for the structure type, like the one returned by make-struct-type;
- *mutator-proc*: a mutator procedure for the structure type, like the one returned by make-struct-type;
- *immutable-k-list*: an immutable list of exact non-negative integers that correspond to immutable fields for the structure type;
- *super-type*: a structure type descriptor for the most specific ancestor of the type that is controlled by the current inspector, or #f if no ancestor is controlled by the current inspector;
- *skipped?*: #f if the seventh result is the most specific ancestor type or if the type has no supertype, #t otherwise.

If the type for *struct-type* is not controlled by the current inspector, the exn:fail:contract exception is raised.

```
(struct-type-make-constructor struct-type)
 → struct-constructor-procedure?
  struct-type : struct-type?
```

Returns a constructor procedure to create instances of the type for *struct-type*. If the type for *struct-type* is not controlled by the current inspector, the exn:fail:contract exception is raised.

```
(struct-type-make-predicate struct-type) → any
  struct-type : any/c
```

Returns a predicate procedure to recognize instances of the type for *struct-type*. If the type for *struct-type* is not controlled by the current inspector, the exn:fail:contract exception is raised.

```
(object-name v) → any
  v : any/c
```

Returns a value for the name of *v* if *v* has a name, #f otherwise. The argument *v* can be any value, but only (some) procedures, structures, structure types, structure type properties, regexp values, and ports have names. See also §1.2.6 “Inferred Value Names”.

The name (if any) of a procedure is always a symbol. The procedure-rename function creates a procedure with a specific name.

The name of a structure, structure type, structure type property is always a symbol. If a structure is a procedure as implemented by one of its fields (i.e., the `prop:procedure` property value for the structure's type is an integer), then its name is the implementing procedure's name; otherwise, its name matches the name of the structure type that it instantiates.

The name of a regexp value is a string or byte string. Passing the string or byte string to `regexp`, `byte-regexp`, `pregexp`, or `byte-pregexp` (depending on the kind of regexp whose name was extracted) produces a value that matches the same inputs.

The name of a port can be any value, but many tools use a path or string name as the port's for (to report source locations, for example).

13.10 Code Inspectors

In the same way that inspectors control access to structure fields (see §13.9 "Structure Inspectors"), inspectors also control access to module bindings. The default inspector for module bindings is determined by the `current-code-inspector` parameter, instead of the `current-inspector` parameter.

When a `module` declaration is evaluated, the value of the `current-code-inspector` parameter is associated with the module declaration. When the module is invoked via `require` or `dynamic-require`, a sub-inspector of the module's declaration-time inspector is created, and this sub-inspector is associated with the module invocation. Any inspector that controls the sub-inspector (including the declaration-time inspector and its superior) controls the module invocation. In particular, if the value of `current-code-inspector` never changes, then no control is lost for any module invocation, since the module's invocation is associated with a sub-inspector of `current-code-inspector`.

When an inspector that controls a module invocation is installed `current-code-inspector`, it enables the following `module->namespace` on the module, and it enables access to the module's protected exports (i.e., those identifiers exported from the module with `protect-out`) via `dynamic-require`.

When a `module` form is expanded or a namespace is created, the value of `current-code-inspector` is associated with the module or namespace's top-level lexical information. Syntax objects with that lexical information gain access to the protected and unexported bindings of any module that the inspector controls. In the case of a `module`, the inspector sticks with such syntax objects even the syntax object is used in the expansion of code in a less powerful context; furthermore, if the syntax object is an identifier that is compiled as a variable reference, the inspector sticks with the variable reference even if it appears in a module form that is evaluated (i.e., declared) with a weaker inspector. When a syntax object or variable reference is within compiled code that is printed (see §1.4.15 "Printing Compiled Code"), the associated inspector is not preserved.

When compiled code in printed form is `read` back in, no inspectors are associated with

the code. When the code is evaluated, the instantiated syntax-object literals and module-variable references acquire value of current-code-inspector as their inspector.

When a module instantiation is attached to multiple namespaces, each with its own module registry, the inspector for the module invocation can be registry-specific. The invocation inspector in a particular module registry can be changed via namespace-unprotect-module (but changing the inspector requires control over the old one).

```
(current-code-inspector) → inspector?
(current-code-inspector insp) → void?
  insp : inspector?
```

A parameter that determines an inspector to control access to module bindings and redefinitions.

13.11 Sandboxed Evaluation

```
(require racket/sandbox)
```

The bindings documented in this section are provided by the racket/sandbox library, not racket/base or racket.

The racket/sandbox module provides utilities for creating "sandboxed" evaluators, which are configured in a particular way and can have restricted resources (memory and time), filesystem and network access, and much more. Sandboxed evaluators can be configured through numerous parameters — and the defaults are set for the common use case where sandboxes are very limited.

```
(make-evaluator  language
                 input-program ...
                [#:requires requires
                 #:allow-for-require allow-for-require
                 #:allow-for-load allow-for-load
                 #:allow-read allow-read])
 → (any/c . -> . any)
             (or/c module-path?
  language :       (list/c 'special symbol?)
                   (cons/c 'begin list?))
  input-program : any/c
             (listof (or/c module-path? path-string?
  requires :                 (cons/c 'for-syntax (listof module-path?))))
           = null
  allow-for-require : (listof (or/c module-path? path?)) = null
  allow-for-load : (listof path-string?) = null
  allow-read : (listof (or/c module-path? path-string?)) = null
```

```
(make-module-evaluator module-decl
                       [#:language lang
                        #:allow-for-require allow-for-require
                        #:allow-for-load allow-for-load
                        #:allow-read allow-read])
 → (any/c . -> . any)
  module-decl : (or/c syntax? pair? path? input-port? string? bytes?)
  lang : (or/c #f module-path?) = #f
  allow-for-require : (listof (or/c module-path? path?)) = null
  allow-for-load : (listof path-string?) = null
  allow-read : (listof (or/c module-path? path-string?)) = null
```

The make-evaluator function creates an evaluator with a *language* and *requires* specification, and starts evaluating the given *input-program*s. The make-module-evaluator function creates an evaluator that works in the context of a given module. The result in either case is a function for further evaluation.

The returned evaluator operates in an isolated and limited environment. In particular, filesystem access is restricted, which may interfere with using modules from the filesystem. See below for information on the *allow-for-require*, *allow-for-load*, and *allow-read* arguments. When *language* is a module path or when *requires* is provided, the indicated modules are implicitly included in the *allow-for-require* list. (For backward compatibility, non-module-path? path strings are allowed in *requires*; they are implicitly converted to paths before addition to *allow-for-require*.)

Each *input-program* or *module-decl* argument provides a program in one of the following forms:

- an input port used to read the program;
- a string or a byte string holding the complete input;
- a path that names a file holding the input; or
- an S-expression or a syntax object, which is evaluated as with eval (see also get-uncovered-expressions).

In the first three cases above, the program is read using sandbox-reader, with line-counting enabled for sensible error messages, and with 'program as the source (used for testing coverage). In the last case, the input is expected to be the complete program, and is converted to a syntax object (using 'program as the source), unless it already is a syntax object.

The returned evaluator function accepts additional expressions (each time it is called) in essentially the same form: a string or byte string holding a sequence of expressions, a path for a file holding expressions, an S-expression, or a syntax object. If the evaluator receives

an eof value, it is terminated and raises errors thereafter. See also kill-evaluator, which terminates the evaluator without raising an exception.

For make-evaluator, multiple *input-program*s are effectively concatenated to form a single program. The way that the *input-program*s are evaluated depends on the *language* argument:

- The *language* argument can be a module path (i.e., a datum that matches the grammar for *module-path* of require).

  In this case, the *input-program*s are automatically wrapped in a module, and the resulting evaluator works within the resulting module's namespace.

- The *language* argument can be a list starting with 'special, which indicates a built-in language with special input configuration. The possible values are '(special r5rs) or a value indicating a teaching language: '(special beginner), '(special beginner-abbr), '(special intermediate), '(special intermediate-lambda), or '(special advanced).

  In this case, the *input-program*s are automatically wrapped in a module, and the resulting evaluator works within the resulting module's namespace. In addition, certain parameters (such as such as read-accept-infix-dot) are set to customize reading programs from strings and ports.

  This option is provided mainly for older test systems. Using make-module-evaluator with input starting with #lang is generally better.

- Finally, *language* can be a list whose first element is 'begin.

  In this case, a new namespace is created using sandbox-namespace-specs, which by default creates a new namespace using sandbox-make-namespace (which, in turn, uses make-base-namespace or make-gui-namespace depending on sandbox-gui-available and gui-available?).

  In the new namespace, *language* is evaluated as an expression to further initialize the namespace.

The *requires* list adds additional imports to the module or namespace for the *input-program*s, even in the case that require is not made available through the *language*.

The following examples illustrate the difference between an evaluator that puts the program in a module and one that merely initializes a top-level namespace:

```
> (define base-module-eval
    ; a module cannot have free variables...
    (make-evaluator 'racket/base '(define (f) later)))
program:1:0: later: unbound identifier in module
  in: later
```

```
> (define base-module-eval
    (make-evaluator 'racket/base '(define (f) later)
                                 '(define later 5)))

> (base-module-eval '(f))
5
> (define base-top-eval
    ; non-module code can have free variables:
    (make-evaluator '(begin) '(define (f) later)))

> (base-top-eval '(+ 1 2))
3
> (base-top-eval '(define later 5))

> (base-top-eval '(f))
5
```

The `make-module-evaluator` function is essentially a restriction of `make-evaluator`, where the program must be a module, and all imports are part of the program. In some cases it is useful to restrict the program to be a module using a specific module in its language position — use the optional *`lang`* argument to specify such a restriction (the default, `#f`, means no restriction is enforced). When the program is specified as a path, then the path is implicitly added to the *`allow-for-load`* list.

```
(define base-module-eval2
  ; equivalent to base-module-eval:
  (make-module-evaluator '(module m racket/base
                            (define (f) later)
                            (define later 5))))
```

The `make-module-evaluator` function can be convenient for testing module files: pass in a path value for the file name, and you get back an evaluator in the module's context which you can use with your favorite test facility.

In all cases, the evaluator operates in an isolated and limited environment:

- It uses a new custodian and namespace. When `gui-available?` and `sandbox-gui-available` produce true, it is also runs in its own eventspace.
- The evaluator works under the `sandbox-security-guard`, which restricts file system and network access.
- The evaluator is contained in a memory-restricted environment, and each evaluation is wrapped in a `call-with-limits` (when memory accounting is available); see also `sandbox-memory-limit`, `sandbox-eval-limits` and `set-eval-limits`.

Note that these limits apply to the creation of the sandbox environment too — so, for example, if the memory that is required to create the sandbox is higher than the limit, then `make-evaluator` will fail with a memory limit exception.

The *`allow-for-require`* and *`allow-for-load`* arguments adjust filesystem permissions to extend the set of files that are usable by the evaluator. The *`allow-for-require`* argument lists modules that can be `required` along with their imports (transitively). The *`allow-for-load`* argument lists files that can be `load`ed. (The precise permissions needed for `require` versus `load` can differ.) The *`allow-read`* argument is for backward compatibility, only; each `module-path?` element of *`allow-read`* is effectively moved to *`allow-for-require`*, while other elements are moved to `all-for-load`.

The sandboxed environment is well isolated, and the evaluator function essentially sends it an expression and waits for a result. This form of communication makes it impossible to have nested (or concurrent) calls to a single evaluator. Usually this is not a problem, but in some cases you can get the evaluator function available inside the sandboxed code, for example:

```
(let ([e (make-evaluator 'racket/base)])
  (e (,e 1)))
```

An error will be signaled in such cases.

If the value of `sandbox-propagate-exceptions` is true (the default) when the sandbox is created, then exceptions (both syntax and run-time) are propagated as usual to the caller of the evaluation function (i.e., catch them with `with-handlers`). If the value of `sandbox-propagate-exceptions` is `#f` when the sandbox is created, then uncaught exceptions in a sandbox evaluation cause the error to be printed to the sandbox's error port, and the caller of the evaluation receives `#<void>`.

Finally, the fact that a sandboxed evaluator accept syntax objects makes it usable as the value for `current-eval`, which means that you can easily start a sandboxed read-eval-print-loop. For example, here is a quick implementation of a networked REPL:

```
(define e (make-evaluator 'racket/base))
(let-values ([(i o) (tcp-accept (tcp-listen 9999))])
  (parameterize ([current-input-port  i]
                 [current-output-port o]
                 [current-error-port  o]
                 [current-eval e])
    (read-eval-print-loop)
    (fprintf o "\nBye...\n")
    (close-output-port o)))
```

Note that in this code it is only the REPL interactions that are going over the network connection; using I/O operations inside the REPL will still use the usual sandbox parameters

(defaulting to no I/O). In addition, the code works only from an existing toplevel REPL — specifically, `read-eval-print-loop` reads a syntax value and gives it the lexical context of the current namespace. Here is a variation that uses the networked ports for user I/O, and works when used from a module (by using a new namespace):

```
(let-values ([(i o) (tcp-accept (tcp-listen 9999))])
  (parameterize ([current-input-port   i]
                 [current-output-port  o]
                 [current-error-port   o]
                 [sandbox-input        i]
                 [sandbox-output       o]
                 [sandbox-error-output o]
                 [current-namespace (make-empty-namespace)])
    (parameterize ([current-eval
                    (make-evaluator 'racket/base)])
      (read-eval-print-loop))
    (fprintf o "\nBye...\n")
    (close-output-port o)))
```

```
(exn:fail:sandbox-terminated? v) → boolean?
  v : any/c
(exn:fail:sandbox-terminated-reason exn) → symbol/c
  exn : exn:fail:sandbox-terminated?
```

A predicate and accessor for exceptions that are raised when a sandbox is terminated. Once a sandbox raises such an exception, it will continue to raise it on further evaluation attempts.

`call-with-limits`. The `resource` field holds a symbol, either `'time` or `'memory`.

13.11.1 Customizing Evaluators

The sandboxed evaluators that `make-evaluator` creates can be customized via many parameters. Most of the configuration parameters affect newly created evaluators; changing them has no effect on already-running evaluators.

The default configuration options are set for a very restricted sandboxed environment — one that is safe to make publicly available. Further customizations might be needed in case more privileges are needed, or if you want tighter restrictions. Another useful approach for customizing an evaluator is to begin with a relatively unrestricted configuration and add the desired restrictions. This approach is made possible by the `call-with-trusted-sandbox-configuration` function.

```
(call-with-trusted-sandbox-configuration thunk) → any
  thunk : (-> any)
```

Invokes the *thunk* in a context where sandbox configuration parameters are set for minimal restrictions. More specifically, there are no memory or time limits, and the existing existing inspectors, security guard, exit handler, and logger are used. (Note that the I/O ports settings are not included.)

```
(sandbox-init-hook) → (-> any)
(sandbox-init-hook thunk) → void?
  thunk : (-> any)
```

A parameter that determines a thunk to be called for initializing a new evaluator. The hook is called just before the program is evaluated in a newly-created evaluator context. It can be used to setup environment parameters related to reading, writing, evaluation, and so on. Certain languages (`'(special r5rs)` and the teaching languages) have initializations specific to the language; the hook is used after that initialization, so it can override settings.

```
(sandbox-reader) → (any/c . -> . any)
(sandbox-reader proc) → void?
  proc : (any/c . -> . any)
```

A parameter that specifies a function that reads all expressions from `(current-input-port)`. The function is used to read program source for an evaluator when a string, byte string, or port is supplied. The reader function receives a value to be used as input source (i.e., the first argument to `read-syntax`), and it should return a list of syntax objects. The default reader calls `read-syntax`, accumulating results in a list until it receives `eof`.

Note that the reader function is usually called as is, but when it is used to read the program input for `make-module-evaluator`, `read-accept-lang` and `read-accept-reader` are set to `#t`.

```
                    (or/c #f
                          string? bytes?
(sandbox-input) →         input-port?
                          'pipe
                          (-> input-port?))
(sandbox-input in) → void?
       (or/c #f
             string? bytes?
  in :       input-port?
             'pipe
             (-> input-port?))
```

A parameter that determines the initial `current-input-port` setting for a newly created evaluator. It defaults to `#f`, which creates an empty port. The following other values are allowed:

- a string or byte string, which is converted to a port using `open-input-string` or `open-input-bytes`;

- an input port;
- the symbol `'pipe`, which triggers the creation of a pipe, where `put-input` can return the output end of the pipe or write directly to it;
- a thunk, which is called to obtain a port (e.g., using `current-input-port` means that the evaluator input is the same as the calling context's input).

```
                    (or/c #f
                          output-port?
(sandbox-output) →        'pipe
                          'bytes
                          'string
                          (-> output-port?))
(sandbox-output in) → void?
       (or/c #f
             output-port?
  in :       'pipe
             'bytes
             'string
             (-> output-port?))
```

A parameter that determines the initial `current-output-port` setting for a newly created evaluator. It defaults to `#f`, which creates a port that discards all data. The following other values are allowed:

- an output port, which is used as-is;
- the symbol `'bytes`, which causes `get-output` to return the complete output as a byte string;
- the symbol `'string`, which is similar to `'bytes`, but makes `get-output` produce a string;
- the symbol `'pipe`, which triggers the creation of a pipe, where `get-output` returns the input end of the pipe;
- a thunk, which is called to obtain a port (e.g., using `current-output-port` means that the evaluator output is not diverted).

```
                           (or/c #f
                                 output-port?
(sandbox-error-output) →         'pipe
                                 'bytes
                                 'string
                                 (-> output-port?))
```

```
(sandbox-error-output in) → void?
      (or/c #f
            output-port?
            'pipe
 in :       'bytes
            'string
            (-> output-port?))
```

Like sandbox-output, but for the initial current-error-port value. An evaluator's error output is set after its output, so using current-output-port (the parameter itself, not its value) for this parameter value means that the error port is the same as the evaluator's initial output port.

The default is (lambda () (dup-output-port (current-error-port))), which means that the error output of the generated evaluator goes to the calling context's error port.

```
(sandbox-coverage-enabled) → boolean?
(sandbox-coverage-enabled enabled?) → void?
  enabled? : any/c
```

A parameter that controls whether syntactic coverage information is collected by sandbox evaluators. Use get-uncovered-expressions to retrieve coverage information.

```
(sandbox-propagate-breaks) → boolean?
(sandbox-propagate-breaks propagate?) → void?
  propagate? : any/c
```

When both this boolean parameter and (break-enabled) are true, breaking while an evaluator is running propagates the break signal to the sandboxed context. This makes the sandboxed evaluator break, typically, but beware that sandboxed evaluation can capture and avoid the breaks (so if safe execution of code is your goal, make sure you use it with a time limit). Also, beware that a break may be received after the evaluator's result, in which case the evaluation result is lost. Finally, beware that a break may be propagated after an evaluator has produced a result, so that the break is visible on the next interaction with the evaluator (or the break is lost if the evaluator is not used further). The default is #t.

```
(sandbox-propagate-exceptions) → boolean?
(sandbox-propagate-exceptions propagate?) → void?
  propagate? : any/c
```

A parameter that controls how uncaught exceptions during a sandbox evaluation are treated. When the parameter value is #t, then the exception is propagated to the caller of sandbox. When the parameter value is #f, the exception message is printed to the sandbox's error port, and the caller of the sandbox receives #<void> for the evaluation. The default is #t.

```
(sandbox-namespace-specs) → (cons/c (-> namespace?)
                                    (listof module-path?))
(sandbox-namespace-specs spec) → void?
  spec : (cons/c (-> namespace?)
                 (listof module-path?))
```

A parameter that holds a list of values that specify how to create a namespace for evaluation in `make-evaluator` or `make-module-evaluator`. The first item in the list is a thunk that creates the namespace, and the rest are module paths for modules to be attached to the created namespace using `namespace-attach-module`.

The default is `(list sandbox-make-namespace)`.

The module paths are needed for sharing module instantiations between the sandbox and the caller. For example, sandbox code that returns `posn` values (from the `lang/posn` module) will not be recognized as such by your own code by default, since the sandbox will have its own instance of `lang/posn` and thus its own struct type for `posn`s. To be able to use such values, include `'lang/posn` in the list of module paths.

When testing code that uses a teaching language, the following piece of code can be helpful:

```
(sandbox-namespace-specs
 (let ([specs (sandbox-namespace-specs)])
   `(,(car specs)
     ,@(cdr specs)
     lang/posn
     ,@(if (gui-available?) '(mrlib/cache-image-snip) '()))))
```

```
(sandbox-make-namespace) → namespace?
```

Calls `make-gui-namespace` when `(sandbox-gui-available)` produces true, `make-base-namespace` otherwise.

```
(sandbox-gui-available) → boolean?
(sandbox-gui-available avail?) → void?
  avail? : any/c
```

Determines whether the `racket/gui` module can be used when a sandbox evaluator is created. If `gui-available?` produces `#f` during the creation of a sandbox evaluator, this parameter is forced to `#f` during initialization of the sandbox. The default value of the parameter is `#t`.

Various aspects of the library change when the GUI library is available, such as using a new eventspace for each evaluator.

```
(sandbox-override-collection-paths) → (listof path-string?)
(sandbox-override-collection-paths paths) → void?
  paths : (listof path-string?)
```

A parameter that determines a list of collection directories to prefix `current-library-collection-paths` in an evaluator. This parameter is useful for cases when you want to test code using an alternate, test-friendly version of a collection, for example, testing code that uses a GUI (like the `htdp/world` teachpack) can be done using a fake library that provides the same interface but no actual interaction. The default is `null`.

```
(sandbox-security-guard)
 → (or/c security-guard? (-> security-guard?))
(sandbox-security-guard guard) → void?
  guard : (or/c security-guard? (-> security-guard?))
```

A parameter that determines the initial `(current-security-guard)` for sandboxed evaluations. It can be either a security guard, or a function to construct one. The default is a function that restricts the access of the current security guard by forbidding all filesystem I/O except for specifications in `sandbox-path-permissions`, and it uses `sandbox-network-guard` for network connections.

```
(sandbox-path-permissions)
    (listof (list/c (or/c 'execute 'write 'delete
 →                        'read-bytecode 'read 'exists)
                    (or/c byte-regexp? bytes? string? path?)))
(sandbox-path-permissions perms) → void?
          (listof (list/c (or/c 'execute 'write 'delete
  perms :                       'read-bytecode 'read 'exists)
                          (or/c byte-regexp? bytes? string? path?)))
```

A parameter that configures the behavior of the default sandbox security guard by listing paths and access modes that are allowed for them. The contents of this parameter is a list of specifications, each is an access mode and a byte-regexp for paths that are granted this access.

The access mode symbol is one of: `'execute`, `'write`, `'delete`, `'read`, or `'exists`. These symbols are in decreasing order: each implies access for the following modes too (e.g., `'read` allows reading or checking for existence).

The path regexp is used to identify paths that are granted access. It can also be given as a path (or a string or a byte string), which is (made into a complete path, cleansed, simplified, and then) converted to a regexp that allows the path and sub-directories; e.g., `"/foo/bar"` applies to `"/foo/bar/baz"`.

An additional mode symbol, `'read-bytecode`, is not part of the linear order of these modes. Specifying this mode is similar to specifying `'read`, but it is not implied by any other mode.

(For example, even if you specify `'write` for a certain path, you need to also specify `'read-bytecode` to grant this permission.) The sandbox usually works in the context of a lower code inspector (see `sandbox-make-code-inspector`) which prevents loading of untrusted bytecode files — the sandbox is set-up to allow loading bytecode from files that are specified with `'read-bytecode`. This specification is given by default to the Racket collection hierarchy (including user-specific libraries) and to libraries that are explicitly specified in an `#:allow-read` argument. (Note that this applies for loading bytecode files only, under a lower code inspector it is still impossible to use protected module bindings (see §13.10 "Code Inspectors").)

The default value is null, but when an evaluator is created, it is augmented by `'read-bytecode` permissions that make it possible to use collection libraries (including `sandbox-override-collection-paths`). See `make-evaluator` for more information.

```
(sandbox-network-guard)
    (symbol?
     (or/c (and/c string? immutable?) #f)
 →   (or/c (integer-in 1 65535) #f)
     (or/c 'server 'client)
     . -> . any)
(sandbox-network-guard proc) → void?
         (symbol?
          (or/c (and/c string? immutable?) #f)
  proc :  (or/c (integer-in 1 65535) #f)
          (or/c 'server 'client)
          . -> . any)
```

A parameter that specifies a procedure to be used (as is) by the default `sandbox-security-guard`. The default forbids all network connection.

```
(sandbox-exit-handler) → (any/c . -> . any)
(sandbox-exit-handler handler) → void?
  handler : (any/c . -> . any)
```

A parameter that determines the initial `(exit-handler)` for sandboxed evaluations. The default kills the evaluator with an appropriate error message (see `exn:fail:sandbox-terminated-reason`).

```
(sandbox-memory-limit) → (or/c nonnegative-number? #f)
(sandbox-memory-limit limit) → void?
  limit : (or/c nonnegative-number? #f)
```

A parameter that determines the total memory limit on the sandbox in megabytes (it can hold a rational or a floating point number). When this limit is exceeded, the sandbox is terminated. This value is used when the sandbox is created and the limit cannot be changed

afterwards. It defaults to 30mb. See sandbox-eval-limits for per-evaluation limits and a description of how the two limits work together.

Note that (when memory accounting is enabled) memory is attributed to the highest custodian that refers to it. This means that if you inspect a value that sandboxed evaluation returns outside of the sandbox, your own custodian will be charged for it. To ensure that it is charged back to the sandbox, you should remove references to such values when the code is done inspecting it.

This policy has an impact on how the sandbox memory limit interacts with the per-expression limit specified by sandbox-eval-limits: values that are reachable from the sandbox, as well as from the interaction will count against the sandbox limit. For example, in the last interaction of this code,

```
(define e (make-evaluator 'racket/base))
(e '(define a 1))
(e '(for ([i (in-range 20)]) (set! a (cons (make-
bytes 500000) a))))
```

the memory blocks are allocated within the interaction limit, but since they're chained to the defined variable, they're also reachable from the sandbox — so they will count against the sandbox memory limit but not against the interaction limit (more precisely, no more than one block counts against the interaction limit).

```
(sandbox-eval-limits)
    (or/c (list/c (or/c nonnegative-number? #f)
 →                (or/c nonnegative-number? #f))
          #f)
(sandbox-eval-limits limits) → void?
           (or/c (list/c (or/c nonnegative-number? #f)
  limits :                (or/c nonnegative-number? #f))
                 #f)
```

A parameter that determines the default limits on *each* use of a make-evaluator function, including the initial evaluation of the input program. Its value should be a list of two numbers; where the first is a timeout value in seconds, and the second is a memory limit in megabytes (note that they don't have to be integers). Either one can be #f for disabling the corresponding limit; alternately, the parameter can be set to #f to disable all per-evaluation limits (useful in case more limit kinds are available in future versions). The default is (list 30 20).

Note that these limits apply to the creation of the sandbox environment too — even (make-evaluator 'racket/base) can fail if the limits are strict enough. For example,

```
(parameterize ([sandbox-eval-limits '(0.25 5)])
```

```
(make-evaluator 'racket/base '(sleep 2)))
```

will throw an error instead of creating an evaluator. Therefore, to avoid surprises you need to catch errors that happen when the sandbox is created.

When limits are set, `call-with-limits` (see below) is wrapped around each use of the evaluator, so consuming too much time or memory results in an exception. Change the limits of a running evaluator using `set-eval-limits`.

A custodian's limit is checked only after a garbage collection, except that it may also be checked during certain large allocations that are individually larger than the custodian's limit.

The memory limit that is specified by this parameter applies to each individual evaluation, but not to the whole sandbox — that limit is specified via `sandbox-memory-limit`. When the global limit is exceeded, the sandbox is terminated, but when the per-evaluation limit is exceeded the `exn:fail:resource` exception is raised. For example, say that you evaluate an expression like

```
(for ([i (in-range 1000)])
  (set! a (cons (make-bytes 1000000) a))
  (collect-garbage))
```

then, assuming sufficiently small limits,

- if a global limit is set but no per-evaluation limit, the sandbox will eventually be terminated and no further evaluations possible;
- if there is a per-evaluation limit, but no global limit, the evaluation will abort with an error and it can be used again — specifically, `a` will still hold a number of blocks, and you can evaluate the same expression again which will add more blocks to it;
- if both limits are set, with the global one larger than the per-evaluation limit, then the evaluation will abort and you will be able to repeat it, but doing so several times will eventually terminate the sandbox (this will be indicated by the error message, and by the `evaluator-alive?` predicate).

```
(sandbox-eval-handlers)
  → (list/c (or/c #f ((-> any) . -> . any))
            (or/c #f ((-> any) . -> . any)))
(sandbox-eval-handlers handlers) → void?
  handlers : (list/c (or/c #f ((-> any) . -> . any))
                     (or/c #f ((-> any) . -> . any)))
```

A parameter that determines two (optional) handlers that wrap sandboxed evaluations. The first one is used when evaluating the initial program when the sandbox is being set-up, and the second is used for each interaction. Each of these handlers should expect a thunk as an argument, and they should execute these thunks — possibly imposing further restrictions.

The default values are `#f` and `call-with-custodian-shutdown`, meaning no additional restrictions on initial sandbox code (e.g., it can start background threads), and a custodian-shutdown around each interaction that follows. Another useful function for this is `call-with-killing-threads` which kills all threads, but leaves other resources intact.

```
(sandbox-make-inspector) → (-> inspector?)
(sandbox-make-inspector make) → void?
  make : (-> inspector?)
```

A parameter that determines the (nullary) procedure that is used to create the inspector for sandboxed evaluation. The procedure is called when initializing an evaluator. The default parameter value is `(lambda () (make-inspector (current-inspector)))`.

```
(sandbox-make-code-inspector) → (-> inspector?)
(sandbox-make-code-inspector make) → void?
  make : (-> inspector?)
```

A parameter that determines the (nullary) procedure that is used to create the code inspector for sandboxed evaluation. The procedure is called when initializing an evaluator. The default parameter value is `(lambda () (make-inspector (current-code-inspector)))`.

The `current-load/use-compiled` handler is setup to allow loading of bytecode files under the original code inspector when `sandbox-path-permissions` allows it through a `'read-bytecode` mode symbol, which makes loading libraries possible.

```
(sandbox-make-logger) → (-> logger?)
(sandbox-make-logger make) → void?
  make : (-> logger?)
```

A parameter that determines the procedure used to create the logger for sandboxed evaluation. The procedure is called when initializing an evaluator, and the default parameter value is `current-logger`. This means that it is not creating a new logger (this might change in the future).

13.11.2 Interacting with Evaluators

The following functions are used to interact with a sandboxed evaluator in addition to using it to evaluate code.

```
(evaluator-alive? evaluator) → boolean?
  evaluator : (any/c . -> . any)
```

Determines whether the evaluator is still alive.

```
(kill-evaluator evaluator) → void?
  evaluator : (any/c . -> . any)
```

Releases the resources that are held by *evaluator* by shutting down the evaluator's custodian. Attempting to use an evaluator after killing raises an exception, and attempts to kill a dead evaluator are ignored.

Killing an evaluator is similar to sending an `eof` value to the evaluator, except that an `eof` value will raise an error immediately.

```
(break-evaluator evaluator) → void?
  evaluator : (any/c . -> . any)
```

Sends a break to the running evaluator. The effect of this is as if Ctrl-C was typed when the evaluator is currently executing, which propagates the break to the evaluator's context.

```
(get-user-custodian evaluator) → void?
  evaluator : (any/c . -> . any)
```

Retrieves the *evaluator*'s toplevel custodian. This returns a value that is different from `(evaluator '(current-custodian))` or `call-in-sandbox-context evaluator current-custodian` — each sandbox interaction is wrapped in its own custodian, which is what these would return.

(One use for this custodian is with `current-memory-use`, where the per-interaction sub-custodians will not be charged with the memory for the whole sandbox.)

```
(set-eval-limits evaluator secs mb) → void?
  evaluator : (any/c . -> . any)
  secs : (or/c exact-nonnegative-integer? #f)
  mb : (or/c exact-nonnegative-integer? #f)
```

Changes the per-expression limits that *evaluator* uses to `sec` seconds and *mb* megabytes (either one can be `#f`, indicating no limit).

This procedure should be used to modify an existing evaluator limits, because changing the `sandbox-eval-limits` parameter does not affect existing evaluators. See also `call-with-limits`.

```
(set-eval-handler evaluator handler) → void?
  evaluator : (any/c . -> . any)
  handler : (or/c #f ((-> any) . -> . any))
```

Changes the per-expression handler that the *evaluator* uses around each interaction. A `#f` value means no handler is used.

This procedure should be used to modify an existing evaluator handler, because changing the `sandbox-eval-handlers` parameter does not affect existing evaluators. See also `call-with-custodian-shutdown` and `call-with-killing-threads` for two useful handlers that are provided.

```
(call-with-custodian-shutdown thunk) → any
  thunk : (-> any)
(call-with-killing-threads thunk) → any
  thunk : (-> any)
```

These functions are useful for use as an evaluation handler. `call-with-custodian-shutdown` will execute the *thunk* in a fresh custodian, then shutdown that custodian, making sure that *thunk* could not have left behind any resources. `call-with-killing-threads` is similar, except that it kills threads that were left, but leaves other resources as is.

```
(put-input evaluator) → output-port?
  evaluator : (any/c . -> . any)
(put-input evaluator i/o) → void?
  evaluator : (any/c . -> . any)
  i/o : (or/c bytes? string? eof-object?)
```

If `(sandbox-input)` is `'pipe` when an evaluator is created, then this procedure can be used to retrieve the output port end of the pipe (when used with no arguments), or to add a string or a byte string into the pipe. It can also be used with `eof`, which closes the pipe.

```
(get-output evaluator) → (or/c #f input-port? bytes? string?)
  evaluator : (any/c . -> . any)
(get-error-output evaluator)
 → (or/c #f input-port? bytes? string?)
  evaluator : (any/c . -> . any)
```

Returns the output or error-output of the *evaluator*, in a way that depends on the setting of `(sandbox-output)` or `(sandbox-error-output)` when the evaluator was created:

- if it was `'pipe`, then `get-output` returns the input port end of the created pipe;
- if it was `'bytes` or `'string`, then the result is the accumulated output, and the output port is reset so each call returns a different piece of the evaluator's output (note that any allocations of such output are still subject to the sandbox memory limit);
- otherwise, it returns `#f`.

```
(get-uncovered-expressions evaluator
                           [prog?
                            src])      → (listof syntax?)
```

```
evaluator : (any/c . -> . any)
prog? : any/c = #t
src : any/c = default-src
```

Retrieves uncovered expression from an evaluator, as longs as the sandbox-coverage-enabled parameter had a true value when the evaluator was created. Otherwise, an exception is raised to indicate that no coverage information is available.

The *prog?* argument specifies whether to obtain expressions that were uncovered after only the original input program was evaluated (#t) or after all later uses of the evaluator (#f). Using #t retrieves a list that is saved after the input program is evaluated, and before the evaluator is used, so the result is always the same.

A #t value of *prog?* is useful for testing student programs to find out whether a submission has sufficient test coverage built in. A #f value is useful for writing test suites for a program to ensure that your tests cover the whole code.

The second optional argument, *src*, specifies that the result should be filtered to hold only syntax objects whose source matches *src*. The default is the source that was used in the program code, if there was one. Note that 'program is used as the source value if the input program was given as S-expressions or as a string (and in these cases it will be the default for filtering). If given #f, the result is the unfiltered list of expressions.

The resulting list of syntax objects has at most one expression for each position and span. Thus, the contents may be unreliable, but the position information is reliable (i.e., it always indicates source code that would be painted red in DrRacket when coverage information is used).

Note that if the input program is a sequence of syntax values, either make sure that they have 'program as the source field, or use the *src* argument. Using a sequence of S-expressions (not syntax objects) for an input program leads to unreliable coverage results, since each expression may be assigned a single source location.

```
(call-in-sandbox-context evaluator
                         thunk
                        [unrestricted?]) → any
  evaluator : (any/c . -> . any)
  thunk : (-> any)
  unrestricted? : boolean? = #f
```

Calls the given *thunk* in the context of a sandboxed evaluator. The call is performed under the resource limits and evaluation handler that are used for evaluating expressions, unless *unrestricted?* is specified as true.

This process is usually similar to (*evaluator* (list *thunk*)), except that it relies on the common meaning of list expressions as function application (which is not true in all

languages), and it relies on eval allowing non-S-expression input. In addition, you can avoid some of the sandboxed restrictions by using your own permissions, for example,

```
(let ([guard (current-security-guard)])
  (call-in-sandbox-context
    ev
    (lambda ()
      (parameterize ([current-security-guard guard])
        ; can access anything you want here
        (delete-file "/some/file")))))
```

13.11.3 Miscellaneous

```
gui? : boolean?
```

For backward compatibility, only: the result of gui-available? at the time that racket/sandbox was instantiated.

The value of gui? is no longer used by racket/sandbox itself. Instead, gui-available? and sandbox-gui-available are checked at the time that a sandbox evaluator is created.

```
(call-with-limits secs mb thunk) → any
  secs : (or/c exact-nonnegative-integer? #f)
  mb : (or/c exact-nonnegative-integer? #f)
  thunk : (-> any)
```

Executes the given *thunk* with memory and time restrictions: if execution consumes more than *mb* megabytes or more than sec seconds, then the computation is aborted and the exn:fail:resource exception is raised. Otherwise the result of the thunk is returned as usual (a value, multiple values, or an exception). Each of the two limits can be #f to indicate the absence of a limit. See also custodian-limit-memory for information on memory limits.

Sandboxed evaluators use call-with-limits, according to the sandbox-eval-limits setting and uses of set-eval-limits: each expression evaluation is protected from time-outs and memory problems. Use call-with-limits directly only to limit a whole testing session, instead of each expression.

```
(with-limits sec-expr mb-expr body ...)
```

A macro version of call-with-limits.

```
(exn:fail:resource? v) → boolean?
  v : any/c
(exn:fail:resource-resource exn) → (or/c 'time 'memory)
  exn : exn:fail:resource?
```

A predicate and accessor for exceptions that are raised by `call-with-limits`. The `resource` field holds a symbol, either `'time` or `'memory`.

14 Operating System

14.1 Paths

When a Racket procedure takes a filesystem path as an argument, the path can be provided either as a string or as an instance of the *path* datatype. If a string is provided, it is converted to a path using `string->path`. A Racket procedure that generates a filesystem path always generates a path value.

By default, paths are created and manipulated for the current platform, but procedures that merely manipulate paths (without using the filesystem) can manipulate paths using conventions for other supported platforms. The `bytes->path` procedure accepts an optional argument that indicates the platform for the path, either `'unix` or `'windows`. For other functions, such as `build-path` or `simplify-path`, the behavior is sensitive to the kind of path that is supplied. Unless otherwise specified, a procedure that requires a path accepts only paths for the current platform.

Two path values are `equal?` when they are use the same convention type and when their byte-string representations are `equal?`. A path string (or byte string) cannot be empty, and it cannot contain a nul character or byte. When an empty string or a string containing nul is provided as a path to any procedure except `absolute-path?`, `relative-path?`, or `complete-path?`, the `exn:fail:contract` exception is raised.

Most Racket primitives that accept paths first *cleanse* the path before using it. Procedures that build paths or merely check the form of a path do not cleanse paths, with the exceptions of `cleanse-path`, `expand-user-path`, and `simplify-path`. For more information about path cleansing and other platform-specific details, see §14.1.3 “Unix and Mac OS X Paths” for Unix and Mac OS X paths and §14.1.4 “Windows Path Conventions” for Windows paths.

14.1.1 Manipulating Paths

```
(path? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a path value for the current platform (not a string, and not a path for a different platform), `#f` otherwise.

```
(path-string? v) → boolean?
  v : any/c
```

Return `#t` if `v` is either a path value for the current platform or a non-empty string without nul characters, `#f` otherwise.

```
(path-for-some-system? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a path value for some platform (not a string), `#f` otherwise.

```
(string->path str) → path?
  str : string?
```

Produces a path whose byte-string name is `(string->bytes/locale string (char->integer #\?))`.

Beware that the current locale might not encode every string, in which case `string->path` can produce the same path for different *str*s. See also `string->path-element`, which should be used instead of `string->path` when a string represents a single path element.

See also `string->some-system-path`.

```
(bytes->path bstr [type]) → path?
  bstr : bytes?
  type : (or/c 'unix 'windows) = (system-path-convention-type)
```

Produces a path (for some platform) whose byte-string name is *bstr*. The optional *type* specifies the convention to use for the path.

For converting relative path elements from literals, use instead `bytes->path-element`, which applies a suitable encoding for individual elements.

```
(path->string path) → string?
  path : path?
```

Produces a string that represents *path* by decoding *path*'s byte-string name using the current locale's encoding; `?` is used in the result string where encoding fails, and if the encoding result is the empty string, then the result is `"?"`.

The resulting string is suitable for displaying to a user, string-ordering comparisons, etc., but it is not suitable for re-creating a path (possibly modified) via `string->path`, since decoding and re-encoding the path's byte string may lose information.

Furthermore, for display and sorting based on individual path elements (such as pathless file names), use `path-element->string`, instead, to avoid special encodings use to represent some relative paths. See §14.1.4 "Windows Path Conventions" for specific information about the conversion of Windows paths.

See also `some-system-path->string`.

```
(path->bytes path) → bytes?
  path : path-for-some-system?
```

Produces *path*'s byte string representation. No information is lost in this translation, so that (bytes->path (path->bytes *path*) (path-convention-type *path*)) always produces a path that is equal? to *path*. The *path* argument can be a path for any platform.

Conversion to and from byte values is useful for marshaling and unmarshaling paths, but manipulating the byte form of a path is generally a mistake. In particular, the byte string may start with a \\?\REL encoding for Windows paths. Instead of path->bytes, use split-path and path-element->bytes to manipulate individual path elements.

```
(string->path-element str) → path?
  str : string?
```

Like string->path, except that *str* corresponds to a single relative element in a path, and it is encoded as necessary to convert it to a path. See §14.1.3 "Unix and Mac OS X Paths" for more information on the conversion for Unix and Mac OS X paths, and see §14.1.4 "Windows Path Conventions" for more information on the conversion for Windows paths.

If *str* does not correspond to any path element (e.g., it is an absolute path, or it can be split), or if it corresponds to an up-directory or same-directory indicator on Unix and Mac OS X, then exn:fail:contract exception is raised.

As for path->string, information can be lost from *str* in the locale-specific conversion to a path.

```
(bytes->path-element bstr [type]) → path-for-some-system?
  bstr : bytes?
  type : (or/c 'unix 'windows) = (system-path-convention-type)
```

Like bytes->path, except that *bstr* corresponds to a single relative element in a path. In terms of conversions and restrictions on *bstr*, bytes->path-element is like string->path-element.

The bytes->path-element procedure is generally the best choice for reconstructing a path based on another path (where the other path is deconstructed with split-path and path-element->bytes) when ASCII-level manipulation of path elements is necessary.

```
(path-element->string path) → string?
  path : path-element?
```

Like path->string, except that trailing path separators are removed (as by split-path). On Windows, any \\?\REL encoding prefix is also removed; see §14.1.4 "Windows Path Conventions" for more information on Windows paths.

The *path* argument must be such that split-path applied to *path* would return 'relative as its first result and a path as its second result, otherwise the exn:fail:contract exception is raised.

The path-element->string procedure is generally the best choice for presenting a pathless file or directory name to a user.

```
(path-element->bytes path) → bytes?
  path : path-element?
```

Like path->bytes, except that any encoding prefix is removed, etc., as for path-element->string.

For any reasonable locale, consecutive ASCII characters in the printed form of *path* are mapped to consecutive byte values that match each character's code-point value, and a leading or trailing ASCII character is mapped to a leading or trailing byte, respectively. The *path* argument can be a path for any platform.

The path-element->bytes procedure is generally the right choice (in combination with split-path) for extracting the content of a path to manipulate it at the ASCII level (then reassembling the result with bytes->path-element and build-path).

```
(path-convention-type path) → (or/c 'unix 'windows)
  path : path-for-some-system?
```

Accepts a path value (not a string) and returns its convention type.

```
(system-path-convention-type) → (or/c 'unix 'windows)
```

Returns the path convention type of the current platform: 'unix for Unix and Mac OS X, 'windows for Windows.

```
(build-path base sub ...) → path-for-some-system?
  base : (or/c path-string? path-for-some-system? 'up 'same)
       (or/c (and/c (or/c path-string? path-for-some-system?)
                    (not/c complete-path?))
             (or/c 'up 'same))
  sub :
```

Creates a path given a base path and any number of sub-path extensions. If *base* is an absolute path, the result is an absolute path, otherwise the result is a relative path.

The *base* and each *sub* must be either a relative path, the symbol 'up (indicating the relative parent directory), or the symbol 'same (indicating the relative current directory). For Windows paths, if *base* is a drive specification (with or without a trailing slash) the first *sub* can be an absolute (driveless) path. For all platforms, the last *sub* can be a filename.

The *base* and *sub* arguments can be paths for any platform. The platform for the resulting path is inferred from the *base* and *sub* arguments, where string arguments imply a path for the current platform. If different arguments are for different platforms, the exn:fail:contract exception is raised. If no argument implies a platform (i.e., all are 'up or 'same), the generated path is for the current platform.

Each *sub* and *base* can optionally end in a directory separator. If the last *sub* ends in a separator, it is included in the resulting path.

If *base* or *sub* is an illegal path string (because it is empty or contains a nul character), the exn:fail:contract exception is raised.

The build-path procedure builds a path *without* checking the validity of the path or accessing the filesystem.

See §14.1.3 "Unix and Mac OS X Paths" for more information on the construction of Unix and Mac OS X paths, and see §14.1.4 "Windows Path Conventions" for more information on the construction of Windows paths.

The following examples assume that the current directory is "/home/joeuser" for Unix examples and "C:\Joe's Files" for Windows examples.

```
(define p1 (build-path (current-directory) "src" "racket"))
 ; Unix: p1 is "/home/joeuser/src/racket"
 ; Windows: p1 is "C:\\Joe's Files\\src\\racket"
(define p2 (build-path 'up 'up "docs" "Racket"))
 ; Unix: p2 is "../../docs/Racket"
 ; Windows: p2 is "..\\..\\docs\\Racket"
(build-path p2 p1)
 ; Unix and Windows: raises exn:fail:contract; p1 is absolute
(build-path p1 p2)
 ; Unix: is "/home/joeuser/src/racket/../../docs/Racket"
 ; Windows: is "C:\\Joe's Files\\src\\racket\\..\\..\\docs\\Racket"
```

```
(build-path/convention-type type
                            base
                            sub ...) → path-for-some-system?
  type : (or/c 'unix 'windows)
  base : (or/c path-string? path-for-some-system? 'up 'same)
       (or/c (and/c (or/c path-string? path-for-some-system?)
  sub :              (not/c complete-path?))
             (or/c 'up 'same))
```

Like build-path, except a path convention type is specified explicitly.

```
(absolute-path? path) → boolean?
  path : (or/c path-string? path-for-some-system?)
```

Returns `#t` if *path* is an absolute path, `#f` otherwise. The *path* argument can be a path for any platform. If *path* is not a legal path string (e.g., it contains a nul character), `#f` is returned. This procedure does not access the filesystem.

```
(relative-path? path) → boolean?
  path : (or/c path-string? path-for-some-system?)
```

Returns `#t` if *path* is a relative path, `#f` otherwise. The *path* argument can be a path for any platform. If *path* is not a legal path string (e.g., it contains a nul character), `#f` is returned. This procedure does not access the filesystem.

```
(complete-path? path) → boolean?
  path : (or/c path-string? path-for-some-system?)
```

Returns `#t` if *path* is a completely determined path (*not* relative to a directory or drive), `#f` otherwise. The *path* argument can be a path for any platform. Note that for Windows paths, an absolute path can omit the drive specification, in which case the path is neither relative nor complete. If *path* is not a legal path string (e.g., it contains a nul character), `#f` is returned.

This procedure does not access the filesystem.

```
(path->complete-path path [base]) → path-for-some-system?
  path : (or/c path-string? path-for-some-system?)
  base : (or/c path-string? path-for-some-system?)
       = (current-directory)
```

Returns *path* as a complete path. If *path* is already a complete path, it is returned as the result. Otherwise, *path* is resolved with respect to the complete path *base*. If *base* is not a complete path, the `exn:fail:contract` exception is raised.

The *path* and *base* arguments can be paths for any platform; if they are for different platforms, the `exn:fail:contract` exception is raised.

This procedure does not access the filesystem.

```
(path->directory-path path) → path-for-some-system?
  path : (or/c path-string? path-for-some-system?)
```

Returns *path* if *path* syntactically refers to a directory and ends in a separator, otherwise it returns an extended version of *path* that specifies a directory and ends with a separator. For example, on Unix and Mac OS X, the path `"x/y/"` syntactically refers to a directory and ends in a separator, but `"x/y"` would be extended to `"x/y/"`, and `"x/.."` would be extended to `"x/../"`. The *path* argument can be a path for any platform, and the result will be for the same platform.

This procedure does not access the filesystem.

```
(resolve-path path) → path?
  path : path-string?
```

Cleanses *path* and returns a path that references the same file or directory as *path*. On Unix and Mac OS X, if *path* is a soft link to another path, then the referenced path is returned (this may be a relative path with respect to the directory owning *path*), otherwise *path* is returned (after expansion).

```
(cleanse-path path) → path-for-some-system?
  path : (or/c path-string? path-for-some-system?)
```

Cleanses *path* (as described at the beginning of this chapter) without consulting the filesystem.

```
(expand-user-path path) → path?
  path : path-string?
```

Cleanses *path*. In addition, on Unix and Mac OS X, a leading ~ is treated as user's home directory and expanded; the username follows the ~ (before a / or the end of the path), where ~ by itself indicates the home directory of the current user.

```
(simplify-path path [use-filesystem?]) → path-for-some-system?
  path : (or/c path-string? path-for-some-system?)
  use-filesystem? : boolean? = #t
```

Eliminates redundant path separators (except for a single trailing separator), up-directory .., and same-directory . indicators in *path*, and changes / separators to \ separators in Windows paths, such that the result accesses the same file or directory (if it exists) as *path*.

In general, the pathname is normalized as much as possible—without consulting the filesystem if *use-filesystem?* is #f, and (on Windows) without changing the case of letters within the path. If *path* syntactically refers to a directory, the result ends with a directory separator.

When *path* is simplified and *use-filesystem?* is true (the default), a complete path is returned; if *path* is relative, it is resolved with respect to the current directory, and up-directory indicators are removed taking into account soft links (so that the resulting path refers to the same directory as before).

When *use-filesystem?* is #f, up-directory indicators are removed by deleting a preceding path element, and the result can be a relative path with up-directory indicators remaining at the beginning of the path; up-directory indicators are dropped when they refer to the parent of a root directory. Similarly, the result can be the same as (build-path 'same) (but with a trailing separator) if eliminating up-directory indicators leaves only same-directory indicators.

The *path* argument can be a path for any platform when *use-filesystem?* is #f, and the resulting path is for the same platform.

The filesystem might be accessed when *use-filesystem?* is true, but the source or simplified path might be a non-existent path. If *path* cannot be simplified due to a cycle of links, the exn:fail:filesystem exception is raised (but a successfully simplified path may still involve a cycle of links if the cycle did not inhibit the simplification).

See §14.1.3 "Unix and Mac OS X Paths" for more information on simplifying Unix and Mac OS X paths, and see §14.1.4 "Windows Path Conventions" for more information on simplifying Windows paths.

```
(normal-case-path path) → path-for-some-system?
  path : (or/c path-string? path-for-some-system?)
```

Returns *path* with "normalized" case letters. For Unix and Mac OS X paths, this procedure always returns the input path, because filesystems for these platforms can be case-sensitive. For Windows paths, if *path* does not start \\?\, the resulting string uses only lowercase letters, based on the current locale. In addition, for Windows paths when the path does not start \\?\, all /s are converted to \s, and trailing spaces and .s are removed.

The *path* argument can be a path for any platform, but beware that local-sensitive decoding and conversion of the path may be different on the current platform than for the path's platform.

This procedure does not access the filesystem.

```
                     (or/c path-for-some-system? 'relative #f)
(split-path path) → (or/c path-for-some-system? 'up 'same)
                     boolean?
  path : (or/c path-string? path-for-some-system?)
```

Deconstructs *path* into a smaller path and an immediate directory or file name. Three values are returned:

- base is either
  - a path,
  - 'relative if *path* is an immediate relative directory or filename, or
  - #f if *path* is a root directory.
- name is either
  - a directory-name path,
  - a filename,

  - 'up if the last part of *path* specifies the parent directory of the preceding path (e.g., .. on Unix), or
  - 'same if the last part of *path* specifies the same directory as the preceding path (e.g., . on Unix).

- must-be-dir? is #t if *path* explicitly specifies a directory (e.g., with a trailing separator), #f otherwise. Note that must-be-dir? does not specify whether name is actually a directory or not, but whether *path* syntactically specifies a directory.

Compared to *path*, redundant separators (if any) are removed in the result base and name. If base is #f, then name cannot be 'up or 'same. The *path* argument can be a path for any platform, and resulting paths for the same platform.

This procedure does not access the filesystem.

See §14.1.3 "Unix and Mac OS X Paths" for more information on splitting Unix and Mac OS X paths, and see §14.1.4 "Windows Path Conventions" for more information on splitting Windows paths.

```
(path-replace-suffix path suffix) → path-for-some-system?
  path : (or/c path-string? path-for-some-system?)
  suffix : (or/c string? bytes?)
```

Returns a path that is the same as *path*, except that the suffix for the last element of the path is changed to *suffix*. If the last element of *path* has no suffix, then *suffix* is added to the path. A suffix is defined as a . followed by any number of non-. characters/bytes at the end of the path element, as long as the path element is not ".." or ".". The *path* argument can be a path for any platform, and the result is for the same platform. If *path* represents a root, the exn:fail:contract exception is raised.

```
(path-add-suffix path suffix) → path-for-some-system?
  path : (or/c path-string? path-for-some-system?)
  suffix : (or/c string? bytes?)
```

Similar to path-replace-suffix, but any existing suffix on *path* is preserved by replacing every . in the last path element with _, and then the *suffix* is added to the end.

14.1.2 More Path Utilities

```
(require racket/path)
```

The bindings documented in this section are provided by the racket/path and racket libraries, but not racket/base.

```
(explode-path path)
 → (listof (or/c path-for-some-system? 'up 'same))
  path : (or/c path-string? path-for-some-system?)
```

Returns the list of path elements that constitute *path*. If *path* is simplified in the sense of simple-form-path, then the result is always a list of paths, and the first element of the list is a root.

```
(file-name-from-path path) → (or/c path-for-some-system? #f)
  path : (or/c path-string? path-for-some-system?)
```

Returns the last element of *path*. If *path* is syntactically a directory path (see split-path), then the result is #f.

```
(filename-extension path) → (or/c bytes? #f)
  path : (or/c path-string? path-for-some-system?)
```

Returns a byte string that is the extension part of the filename in *path* without the . separator. If *path* is syntactically a directory (see split-path) or if the path has no extension, #f is returned.

```
(find-relative-path  base
                     path
                    [#:more-than-root? more-than-root?])
 → path-for-some-system?
  base : (or/c path-string? path-for-some-system?)
  path : (or/c path-string?  path-for-some-system?)
  more-than-root? : any/c = #f
```

Finds a relative pathname with respect to *base* that names the same file or directory as *path*. Both *base* and *path* must be simplified in the sense of simple-form-path. If *path* shares no subpath in common with *base*, *path* is returned.

If *more-than-root?* is true, if *base* and *path* share only a Unix root in common, and if neither *base* nor *path* is just a root path, then *path* is returned.

```
(normalize-path path [wrt]) → path?
  path : path-string?
  wrt :  (and/c path-string? complete-path?)
      = (current-directory)
```

Returns a complete version of *path* by making the path complete, expanding the complete path, and resolving all soft links (which requires consulting the filesystem). If *path* is relative, then *wrt* is used as the base path.

For most purposes, simple-form-path is the preferred mechanism to normalize a path, because it works for paths that include non-existent directory components, and it avoids unnecessarily expanding soft links.

Letter case is *not* normalized by `normalize-path`. For this and other reasons, such as whether the path is syntactically a directory, the result of `normalize-path` is not suitable for comparisons that determine whether two paths refer to the same file or directory (i.e., the comparison may produce false negatives).

An error is signaled by `normalize-path` if the input path contains an embedded path for a non-existent directory, or if an infinite cycle of soft links is detected.

```
(path-element? path) → boolean?
  path : any/c
```

Returns `#t` if *path* is a *path element*: a path value for some platform (see `path-for-some-system?`) such that `split-path` applied to *path* would return `'relative` as its first result and a path as its second result. Otherwise, the result is `#f`.

```
(path-only path) → (or/c #f path-for-some-system?)
  path : (or/c path-string? path-for-some-system?)
```

If *path* is a filename, the file's path is returned. If *path* is syntactically a directory, *path* is returned (as a path, if it was a string). If *path* has no directory part `#f` is returned.

```
(simple-form-path path) → path?
  path : path-string?
```

Returns `(simplify-path (path->complete-path path))`, which ensures that the result is a complete path containing no up- or same-directory indicators.

```
(some-system-path->string path) → string?
  path : path-for-some-system?
```

Converts *path* to a string using a UTF-8 encoding of the path's bytes.

Use this function when working with paths for a different system (whose encoding of pathnames might be unrelated to the current locale's encoding) and when starting and ending with strings.

```
(string->some-system-path str kind) → path-for-some-system?
  str : string?
  kind : (or/c 'unix 'windows)
```

Converts *str* to a *kind* path using a UTF-8 encoding of the path's bytes.

Use this function when working with paths for a different system (whose encoding of pathnames might be unrelated to the current locale's encoding) and when starting and ending with strings.

14.1.3 Unix and Mac OS X Paths

In Unix and Mac OS X paths, a / separates elements of the path, . as a path element always means the directory indicated by preceding path, and .. as a path element always means the parent of the directory indicated by the preceding path. A leading ~ in a path is not treated specially, but expand-user-path can be used to convert a leading ~ element to a user-specific directory. No other character or byte has a special meaning within a path. Multiple adjacent / are equivalent to a single / (i.e., they act as a single path separator).

A path root is always /. A path starting with / is an absolute, complete path, and a path starting with any other character is a relative path.

Any pathname that ends with a / syntactically refers to a directory, as does any path whose last element is . or ..

A Unix and Mac OS X path is cleansed by replacing multiple adjacent /s with a single /.

For (bytes->path-element *bstr*), bstr must not contain any /, otherwise the exn:fail:contract exception is raised. The result of (path-element->bytes *path*) or (path-element->string *path*) is always the same as the result of (path->bytes *path*) and (path->string *path*). Since that is not the case for other platforms, however, path-element->bytes and path-element->string should be used when converting individual path elements.

On Mac OS X, Finder aliases are zero-length files.

14.1.4 Windows Path Conventions

In general, a Windows pathname consists of an optional drive specifier and a drive-specific path. A Windows path can be *absolute* but still relative to the current drive; such paths start with a / or \ separator and are not UNC paths or paths that start with \\?\.

A path that starts with a drive specification is *complete*. Roughly, a drive specification is either a Roman letter followed by a colon, a UNC path of the form \\⟨*machine*⟩\⟨*volume*⟩, or a \\?\ form followed by something other than REL\⟨*element*⟩, or RED\⟨*element*⟩. (Variants of \\?\ paths are described further below.)

Racket fails to implement the usual Windows path syntax in one way. Outside of Racket, a pathname "C:rant.txt" can be a drive-specific relative path. That is, it names a file "rant.txt" on drive "C:", but the complete path to the file is determined by the current working directory for drive "C:". Racket does not support drive-specific working directories (only a working directory across all drives, as reflected by the current-directory parameter). Consequently, Racket implicitly converts a path like "C:rant.txt" into "C:\rant.txt".

- *Racket-specific:* Whenever a path starts with a drive specifier ⟨*letter*⟩`:` that is not followed by a `/` or `\`, a `\` is inserted as the path is cleansed.

Otherwise, Racket follows standard Windows path conventions, but also adds `\\?\REL` and `\\?\RED` conventions to deal with paths inexpressible in the standard convention, plus conventions to deal with excessive `\`s in `\\?\` paths.

In the following, ⟨*letter*⟩ stands for a Roman letter (case does not matter), ⟨*machine*⟩ stands for any sequence of characters that does not include `\` or `/` and is not `?`, ⟨*volume*⟩ stands for any sequence of characters that does not include `\` or `/` , and ⟨*element*⟩ stands for any sequence of characters that does not include `\`.

- Trailing spaces and `.` in a path element are ignored when the element is the last one in the path, unless the path starts with `\\?\` or the element consists of only spaces and `.`s.

- The following special "files", which access devices, exist in all directories, case-insensitively, and with all possible endings after a period or colon, except in path-names that start with `\\?\`: `"NUL"`, `"CON"`, `"PRN"`, `"AUX"`, `"COM1"`, `"COM2"`, `"COM3"`, `"COM4"`, `"COM5"`, `"COM6"`, `"COM7"`, `"COM8"`, `"COM9"`, `"LPT1"`, `"LPT2"`, `"LPT3"`, `"LPT4"`, `"LPT5"`, `"LPT6"`, `"LPT7"`, `"LPT8"`, `"LPT9"`.

- Except for `\\?\` paths, `/`s are equivalent to `\`s. Except for `\\?\` paths and the start of UNC paths, multiple adjacent `/`s and `\`s count as a single `\`. In a path that starts `\\?\` paths, elements can be separated by either a single or double `\`.

- A directory can be accessed with or without a trailing separator. In the case of a non-`\\?\` path, the trailing separator can be any number of `/`s and `\`s; in the case of a `\\?\` path, a trailing separator must be a single `\`, except that two `\`s can follow `\\?\`⟨*letter*⟩`:`.

- Except for `\\?\` paths, a single `.` as a path element means "the current directory," and a `..` as a path element means "the parent directory." Up-directory path elements (i.e., `..`) immediately after a drive are ignored.

- A pathname that starts `\\`⟨*machine*⟩`\`⟨*volume*⟩ (where a `/` can replace any `\`) is a UNC path, and the starting `\\`⟨*machine*⟩`\`⟨*volume*⟩ counts as the drive specifier.

- Normally, a path element cannot contain any of the following characters:

  `< > : " / \ |`

  Except for `\`, path elements containing these characters can be accessed using a `\\?\` path (assuming that the underlying filesystem allows the characters).

- In a pathname that starts `\\?\`⟨*letter*⟩`:\`, the `\\?\`⟨*letter*⟩`:\` prefix counts as the path's drive, as long as the path does not both contain non-drive elements and end with two

consecutive \s, and as long as the path contains no sequence of three or more \s. Two \s can appear in place of the \ before ⟨*letter*⟩. A / cannot be used in place of a \ (but /s can be used in element names, though the result typically does not name an actual directory or file).

- In a pathname that starts \\?\UNC\⟨*machine*⟩\⟨*volume*⟩, the \\?\UNC\⟨*machine*⟩\⟨*volume*⟩ prefix counts as the path's drive, as long as the path does not end with two consecutive \s, and as long as the path contains no sequence of three or more \s. Two \s can appear in place of the \ before UNC, the \s after UNC, and/or the \s after⟨*machine*⟩. The letters in the UNC part can be uppercase or lowercase, and / cannot be used in place of \s (but / can be used in element names).

- *Racket-specific:* A pathname that starts \\?\REL\⟨*element*⟩ or \\?\REL\\⟨*element*⟩ is a relative path, as long as the path does not end with two consecutive \s, and as long as the path contains no sequence of three or more \s. This Racket-specific path form supports relative paths with elements that are not normally expressible in Windows paths (e.g., a final element that ends in a space). The REL part must be exactly the three uppercase letters, and /s cannot be used in place of \s. If the path starts \\?\REL\.. then for as long as the path continues with repetitions of \.., each element counts as an up-directory element; a single \ must be used to separate the up-directory elements. As soon as a second \ is used to separate the elements, or as soon as a non-.. element is encountered, the remaining elements are all literals (never up-directory elements). When a \\?\REL path value is converted to a string (or when the path value is written or displayed), the string does not contain the starting \\?\REL or the immediately following \s; converting a path value to a byte string preserves the \\?\REL prefix.

- *Racket-specific:* A pathname that starts \\?\RED\⟨*element*⟩ or \\?\RED\\⟨*element*⟩ is a drive-relative path, as long as the path does not end with two consecutive \s, and as long as the path contains no sequence of three or more \s. This Racket-specific path form supports drive-relative paths (i.e., absolute given a drive) with elements that are not normally expressible in Windows paths. The RED part must be exactly the three uppercase letters, and /s cannot be used in place of \s. Unlike \\?\REL paths, a .. element is always a literal path element. When a \\?\RED path value is converted to a string (or when the path value is written or displayed), the string does not contain the starting \\?\RED and it contains a single starting \; converting a path value to a byte string preserves the \\?\RED prefix.

Three additional Racket-specific rules provide meanings to character sequences that are otherwise ill-formed as Windows paths:

- *Racket-specific:* In a pathname of the form \\?\⟨*any*⟩\\ where ⟨*any*⟩ is any non-empty sequence of characters other than ⟨*letter*⟩: or \⟨*letter*⟩:, the entire path counts as the path's (non-existent) drive.

- *Racket-specific:* In a pathname of the form \\?\⟨*any*⟩\\\⟨*elements*⟩, where ⟨*any*⟩ is any non-empty sequence of characters and ⟨*elements*⟩ is any sequence that does not

start with a \, does not end with two \s, and does not contain a sequence of three \s, then \\?\⟨*any*⟩\\ counts as the path's (non-existent) drive.

- *Racket-specific:* In a pathname that starts \\?\ and does not match any of the patterns from the preceding bullets, \\?\ counts as the path's (non-existent) drive.

Outside of Racket, except for \\?\ paths, pathnames are typically limited to 259 characters. Racket internally converts pathnames to \\?\ form as needed to avoid this limit. The operating system cannot access files through \\?\ paths that are longer than 32,000 characters or so.

Where the above descriptions says "character," substitute "byte" for interpreting byte strings as paths. The encoding of Windows paths into bytes preserves ASCII characters, and all special characters mentioned above are ASCII, so all of the rules are the same.

Beware that the \ path separator is an escape character in Racket strings. Thus, the path \\?\REL\..\\.. as a string must be written "\\\\?\\REL\\..\\\\..".

A path that ends with a directory separator syntactically refers to a directory. In addition, a path syntactically refers to a directory if its last element is a same-directory or up-directory indicator (not quoted by a \\?\ form), or if it refers to a root.

Windows paths are cleansed as follows: In paths that start \\?\, redundant \s are removed, an extra \ is added in a \\?\REL if an extra one is not already present to separate up-directory indicators from literal path elements, and an extra \ is similarly added after \\?\RED if an extra one is not already present. When \\?\ acts as the root and the path contains, to additional /s (which might otherwise be redundant) are included after the root. For other paths, multiple /s are converted to single /s (except at the beginning of a shared folder name), a / is inserted after the colon in a drive specification if it is missing.

For (bytes->path-element *bstr*), /s, colons, trailing dots, trailing whitespace, and special device names (e.g., "aux") in *bstr* are encoded as a literal part of the path element by using a \\?\REL prefix. The bstr argument must not contain a \, otherwise the exn:fail:contract exception is raised.

For (path-element->bytes *path*) or (path-element->string *path*), if the byte-string form of *path* starts with a \\?\REL, the prefix is not included in the result.

For (build-path *base-path sub-path* ...), trailing spaces and periods are removed from the last element of *base-path* and all but the last *sub-path* (unless the element consists of only spaces and periods), except for those that start with \\?\. If *base-path* starts \\?\, then after each non-\\?\REL\ and non-\\?\RED\ *sub-path* is added, all /s in the addition are converted to \s, multiple consecutive \s are converted to a single \, added . elements are removed, and added .. elements are removed along with the preceding element; these conversions are not performed on the original *base-path* part of the result or on any \\?\REL\ or \\?\RED\ or *sub-path*. If a \\?\REL\ or \\?\RED\ *sub-path* is added to a non-\\?\ *base-path*, the *base-path* (with any additions up to the \\?\REL\

or \\?\RED\ *sub-path*) is simplified and converted to a \\?\ path. In other cases, a \ may be added or removed before combining paths to avoid changing the root meaning of the path (e.g., combining //x and y produces /x/y, because //x/y would be a UNC path instead of a drive-relative path).

For (simplify-path *path use-filesystem?*), *path* is expanded, and if *path* does not start with \\?\, trailing spaces and periods are removed, a / is inserted after the colon in a drive specification if it is missing, and a \ is inserted after \\?\ as a root if there are elements and no extra \ already. Otherwise, if no indicators or redundant separators are in *path*, then *path* is returned.

For (split-path *path*) producing *base*, *name*, and *must-be-dir?*, splitting a path that does not start with \\?\ can produce parts that start with \\?\. For example, splitting C:/x~/aux/ produces \\?\C:\x~\ and \\?\REL\\aux; the \\?\ is needed in these cases to preserve a trailing space after x and to avoid referring to the AUX device instead of an "aux" file.

14.2 Filesystem

14.2.1 Locating Paths

(find-system-path *kind*) → path?
  kind : symbol?

Returns a machine-specific path for a standard type of path specified by *kind*, which must be one of the following:

- 'home-dir — the current user's home directory.

  On Unix and Mac OS X, this directory is determined by expanding the path "~", which is expanded by first checking for a HOME environment variable. If none is defined, the USER and LOGNAME environment variables are consulted (in that order) to find a user name, and then system files are consulted to locate the user's home directory.

  On Windows, the user's home directory is the user-specific profile directory as determined by the Windows registry. If the registry cannot provide a directory for some reason, the value of the USERPROFILE environment variable is used instead, as long as it refers to a directory that exists. If USERPROFILE also fails, the directory is the one specified by the HOMEDRIVE and HOMEPATH environment variables. If those environment variables are not defined, or if the indicated directory still does not exist, the directory containing the current executable is used as the home directory.

- 'pref-dir — the standard directory for storing the current user's preferences. On Unix, the directory is ".racket" in the user's home directory. On Windows, it is

"Racket" in the user's application-data folder as specified by the Windows registry; the application-data folder is usually "Application Data" in the user's profile directory. On Mac OS X, it is "Library/Preferences" in the user's home directory. This directory might not exist.

- 'pref-file — a file that contains a symbol-keyed association list of preference values. The file's directory path always matches the result returned for 'pref-dir. The file name is "racket-prefs.rktd" on Unix and Windows, and it is "org.racket-lang.prefs.rktd" on Mac OS X. The file's directory might not exist. See also get-preference.

- 'temp-dir — the standard directory for storing temporary files. On Unix and Mac OS X, this is the directory specified by the TMPDIR environment variable, if it is defined, otherwise it is the first path that exists among "/var/tmp", "/usr/tmp", and "/tmp". On Windows, the result is the directory specified by the TMP or TEMP environment variable, if it is defined, otherwise it is the current directory.

- 'init-dir — the directory containing the initialization file used by the Racket executable. It is the same as the current user's home directory.

- 'init-file — the file loaded at start-up by the Racket executable. The directory part of the path is the same path as returned for 'init-dir. The file name is platform-specific:

  - Unix and Mac OS X: ".racketrc"
  - Windows: "racketrc.rktl"

- 'links-file — the user-specific collection links file for specifying the location of library collections. This file is specified by the PLTLINKSFILE environment variable, and it can be overridden by the --links or -C command-line flag. If no environment variable or flag is specified, or if the value is not a legal path name, then this file defaults to "links.rktd" in the directory reported by (find-system-path 'addon-dir).

- 'addon-dir — a directory for installing user-specific Racket extensions. This directory is specified by the PLTADDONDIR environment variable, and it can be overridden by the --addon or -A command-line flag. If no environment variable or flag is specified, or if the value is not a legal path name, then this directory defaults to "Library/Racket" in the user's home directory on Mac OS X and 'pref-dir otherwise. This directory might not exist.

- 'doc-dir — the standard directory for storing the current user's documents. On Unix, it's the same as 'home-dir. On Mac OS X, it's the "Documents" directory in the user's home directory. On Windows, it is the user's documents folder as specified by the Windows registry; the documents folder is usually "My Documents" in the user's home directory.

- 'desk-dir — the directory for the current user's desktop. On Unix, it's the same as 'home-dir. On Windows, it is the user's desktop folder as specified by the Windows

registry; the documents folder is usually "Desktop" in the user's home directory. On Mac OS X, it is the desktop directory, which is specifically "~/Desktop" on Mac OS X.

- 'sys-dir — the directory containing the operating system for Windows. On Unix and Mac OS X, the result is "/".
- 'exec-file — the path of the Racket executable as provided by the operating system for the current invocation. For some operating systems, the path can be relative.

For GRacket, the executable path is the name of a GRacket executable.

- 'run-file — the path of the current executable; this may be different from result for 'exec-file because an alternate path was provided through a --name or -N command-line flag to the Racket (or GRacket) executable, or because an embedding executable installed an alternate path. In particular a "launcher" script created by make-racket-launcher sets this path to the script's path.
- 'collects-dir — a path to the main collection of libraries (see §17.2 "Libraries and Collections"). If this path is relative, then it is relative to the executable as reported by (find-system-path 'exec-file)—though the latter could be a soft-link or relative to the user's executable search path, so that the two results should be combined with find-executable-path. The 'collects-dir path is normally embedded in the Racket executable, but it can be overridden by the --collects or -X command-line flag.
- 'orig-dir — the current directory at start-up, which can be useful in converting a relative-path result from (find-system-path 'exec-file) or (find-system-path 'run-file) to a complete path.

```
(path-list-string->path-list str
                             default-path-list)
 → (listof path?)
  str : (or/c string? bytes?)
  default-path-list : (listof path?)
```

Parses a string or byte string containing a list of paths, and returns a list of path strings. On Unix and Mac OS X, paths in a path list are separated by a :; on Windows, paths are separated by a ;, and all "s in the string are discarded. Whenever the path list contains an empty path, the list *default-path-list* is spliced into the returned list of paths. Parts of *str* that do not form a valid path are not included in the returned list.

```
(find-executable-path  program-sub
                      [related-sub
                       deepest?])   → (or/c path? #f)
  program-sub : path-string?
  related-sub : (or/c path-string? #f) = #f
  deepest? : any/c = #f
```

Finds a path for the executable *program-sub*, returning #f if the path cannot be found.

If *related-sub* is not #f, then it must be a relative path string, and the path found for *program-sub* must be such that the file or directory *related-sub* exists in the same directory as the executable. The result is then the full path for the found *related-sub*, instead of the path for the executable.

This procedure is used by the Racket executable to find the standard library collection directory (see §17.2 "Libraries and Collections"). In this case, program is the name used to start Racket and related is "collects". The *related-sub* argument is used because, on Unix and Mac OS X, *program-sub* may involve a sequence of soft links; in this case, *related-sub* determines which link in the chain is relevant.

If *related-sub* is not #f, then when find-executable-path does not find a *program-sub* that is a link to another file path, the search can continue with the destination of the link. Further links are inspected until *related-sub* is found or the end of the chain of links is reached. If *deepest?* is #f (the default), then the result corresponds to the first path in a chain of links for which *related-sub* is found (and further links are not actually explored); otherwise, the result corresponds to the last link in the chain for which *related-sub* is found.

If *program-sub* is a pathless name, find-executable-path gets the value of the PATH environment variable; if this environment variable is defined, find-executable-path tries each path in PATH as a prefix for *program-sub* using the search algorithm described above for path-containing *program-sub*s. If the PATH environment variable is not defined, *program-sub* is prefixed with the current directory and used in the search algorithm above. (On Windows, the current directory is always implicitly the first item in PATH, so find-executable-path checks the current directory first on Windows.)

14.2.2 Files

```
(file-exists? path) → boolean?
  path : path-string?
```

Returns #t if a file (not a directory) *path* exists, #f otherwise.

On Windows, file-exists? reports #t for all variations of the special filenames (e.g., "LPT1", "x:/baddir/LPT1").

```
(link-exists? path) → boolean?
  path : path-string?
```

Returns #t if a link *path* exists (Unix and Mac OS X), #f otherwise.

The predicates file-exists? or directory-exists? work on the final destination of a

link or series of links, while `link-exists?` only follows links to resolve the base part of *path* (i.e., everything except the last name in the path).

This procedure never raises the `exn:fail:filesystem` exception.

```
(delete-file path) → void?
  path : path-string?
```

Deletes the file with path *path* if it exists, otherwise the `exn:fail:filesystem` exception is raised. If *path* is a link, the link is deleted rather than the destination of the link.

```
(rename-file-or-directory old
                          new
                         [exists-ok?]) → void?
  old : path-string?
  new : path-string?
  exists-ok? : any/c = #f
```

Renames the file or directory with path *old*—if it exists—to the path *new*. If the file or directory is not renamed successfully, the `exn:fail:filesystem` exception is raised.

This procedure can be used to move a file/directory to a different directory (on the same disk) as well as rename a file/directory within a directory. Unless *exists-ok?* is provided as a true value, *new* cannot refer to an existing file or directory. Even if *exists-ok?* is true, *new* cannot refer to an existing file when *old* is a directory, and vice versa.

If *new* exists and is replaced, the replacement is atomic on Unix and Mac OS X, but it is not guaranteed to be atomic on Windows. Furthermore, if *new* exists and is opened by any process for reading or writing, then attempting to replace it will typically fail on Windows.

If *old* is a link, the link is renamed rather than the destination of the link, and it counts as a file for replacing any existing *new*.

```
(file-or-directory-modify-seconds path
                                  [secs-n
                                   fail-thunk]) → any
  path : path-string?
  secs-n : (or/c exact-integer? #f) = #f
  fail-thunk : (-> any)
             = (lambda () (raise (make-exn:fail:filesystem ....)))
```

Returns the file or directory's last modification date in seconds since midnight UTC, January 1, 1970 (see also §14.6 "Time") when *secs-n* is not provided or is `#f`.

For FAT filesystems on Windows, directories do not have modification dates. Therefore, the creation date is returned for a directory, but the modification date is returned for a file.

If *secs-n* is provided and not #f, the access and modification times of *path* are set to the given time.

On error (e.g., if no such file exists), *fail-thunk* is called, and the default *fail-thunk* raises exn:fail:filesystem.

```
(file-or-directory-permissions path [mode])
 → (listof (or/c 'read 'write 'execute))
  path : path-string?
  mode : #f = #f
(file-or-directory-permissions path mode) → (integer-in 0 65535)
  path : path-string?
  mode : 'bits
(file-or-directory-permissions path mode) → void
  path : path-string?
  mode : (integer-in 0 65535)
```

When given one argument or #f as the second argument, returns a list containing 'read, 'write, and/or 'execute to indicate permission the given file or directory path by the current user and group. On Unix and Mac OS X, permissions are checked for the current effective user instead of the real user.

If 'bits is supplied as the second argument, the result is a platform-specific integer encoding of the file or directory properties (mostly permissions), and the result is independent of the current user and group. The lowest nine bits of the encoding are somewhat portable, reflecting permissions for the file or directory's owner, members of the file or directory's group, or other users:

- #o100 : owner has read permission
- #o200 : owner has write permission
- #o400 : owner has execute permission
- #o010 : group has read permission
- #o020 : group has write permission
- #o040 : group has execute permission
- #o001 : others have read permission
- #o002 : others have write permission
- #o004 : others have execute permission

See also user-read-bit, etc. On Windows, permissions from all three (owner, group, and others) are always the same, and read and execute permission are always available. On Unix and Mac OS X, higher bits have a platform-specific meaning.

If an integer is supplied as the second argument, its is used as an encoding of properties (mostly permissions) to install for the file.

In all modes, the exn:fail:filesystem exception is raised on error (e.g., if no such file exists).

```
(file-or-directory-identity path [as-link?])
 → exact-positive-integer?
  path : path-string?
  as-link? : any/c = #f
```

Returns a number that represents the identity of *path* in terms of the device and file or directory that it accesses. This function can be used to check whether two paths correspond to the same filesystem entity under the assumption that the path's entity selection does not change.

If *as-link?* is a true value, then if *path* refers to a filesystem link, the identity of the link is returned instead of the identity of the referenced file or directory (if any).

```
(file-size path) → exact-nonnegative-integer?
  path : path-string?
```

Returns the (logical) size of the specified file in bytes. On Mac OS X, this size excludes the resource-fork size. On error (e.g., if no such file exists), the exn:fail:filesystem exception is raised.

```
(copy-file src dest [exists-ok?]) → void?
  src : path-string?
  dest : path-string?
  exists-ok? : any/c = #f
```

Creates the file *dest* as a copy of *src*, if *dest* does not already exist. If *dest* already exists and *exists-ok?* is #f, the copy fails with exn:fail:filesystem:exists? exception is raised; otherwise, if *dest* exists, its content is replaced with the content of *src*. File permissions are transferred from *src* to *dest*; on Windows, the modification time of *src* is also transferred to *dest*. If *src* refers to a link, the target of the link is copied, rather than the link itself; if *dest* refers to a link and *exists-ok?* is true, the target of the link is updated.

```
(make-file-or-directory-link to path) → void?
  to : path-string?
  path : path-string?
```

Creates a link *path* to *to* on Unix and Mac OS X. The creation will fail if *path* already exists. The *to* need not refer to an existing file or directory, and *to* is not expanded before writing the link. If the link is not created successfully,the exn:fail:filesystem exception is raised. On Windows, the exn:fail:unsupported exception is raised always.

14.2.3 Directories

See also: `rename-file-or-directory`, `file-or-directory-modify-seconds`, `file-or-directory-permissions`.

```
(current-directory) → path-string?
(current-directory path) → void?
  path : path-string?
```

A parameter that determines the current directory for resolving relative paths.

When the parameter procedure is called to set the current directory, the path argument is cleansed using `cleanse-path`, simplified using `simplify-path`, and then converted to a directory path with `path->directory-path`; cleansing and simplification raise an exception if the path is ill-formed. Thus, the current value of `current-directory` is always a cleansed, simplified, complete, directory path.

The path is not checked for existence when the parameter is set.

```
(current-drive) → path?
```

Returns the current drive name Windows. For other platforms, the `exn:fail:unsupported` exception is raised. The current drive is always the drive of the current directory.

```
(directory-exists? path) → boolean?
  path : path-string?
```

Returns `#t` if *path* refers to a directory, `#f` otherwise.

```
(make-directory path) → void?
  path : path-string?
```

Creates a new directory with the path *path*. If the directory is not created successfully, the `exn:fail:filesystem` exception is raised.

```
(delete-directory path) → void?
  path : path-string?
```

Deletes an existing directory with the path *path*. If the directory is not deleted successfully, the `exn:fail:filesystem` exception is raised.

```
(directory-list [path #:build? build?]) → (listof path?)
  path : path-string? = (current-directory)
  build? : any/c = #f
```

See also the in-directory sequence constructor.

Returns a list of all files and directories in the directory specified by *path*. If *build?* is #f, the resulting paths are all path elements; otherwise, the individual results are combined with *path* using build-path. On Windows, an element of the result list may start with \\?\REL\\.

(filesystem-root-list) → (listof path?)

Returns a list of all current root directories. Obtaining this list can be particularly slow on Windows.

14.2.4 Declaring Paths Needed at Run Time

```
(require racket/runtime-path)
```

The bindings documented in this section are provided by the racket/runtime-path library, not racket/base or racket.

The racket/runtime-path library provides forms for accessing files and directories at run time using a path that are usually relative to an enclosing source file. Unlike using collection-path, define-runtime-path exposes each run-time path to tools like the executable and distribution creators, so that files and directories needed at run time are carried along in a distribution.

In addition to the bindings described below, racket/runtime-path provides #%datum in phase level 1, since string constants are often used as compile-time expressions with define-runtime-path.

(define-runtime-path *id* *expr*)

Uses *expr* as both a compile-time (i.e., phase 1) expression and a run-time (i.e., phase 0) expression. In either context, *expr* should produce a path, a string that represents a path, a list of the form (list 'lib *str* ...+), or a list of the form (list 'so *str*).

For run time, *id* is bound to a path that is based on the result of *expr*. The path is normally computed by taking a relative path result from *expr* and adding it to a path for the enclosing file (which is computed as described below). However, tools like the executable creator can also arrange (by colluding with racket/runtime-path) to have a different base path substituted in a generated executable. If *expr* produces an absolute path, it is normally returned directly, but again may be replaced by an executable creator. In all cases, the executable creator preserves the relative locations of all paths. When *expr* produces a relative or absolute path, then the path bound to *id* is always an absolute path.

If *expr* produces a list of the form (list 'lib *str* ...+), the value bound to *id* is an absolute path. The path refers to a collection-based file similar to using the value as a module

path.

If *expr* produces a list of the form (list 'so *str*), the value bound to *id* can be either *str* or an absolute path; it is an absolute path when adding the platform-specific shared-library extension — as produced by (system-type 'so-suffix) — and then searching in the Racket-specific shared-object library directories (as determined by get-lib-search-dirs) locates the path. In this way, shared-object libraries that are installed specifically for Racket get carried along in distributions.

If *expr* produces a list of the form (list 'module *module-path var-ref*), the value bound to *id* is a module path index, where *module-path* is treated as relative (if it is relative) to the module that is the home of the variable reference *var-ref*, where *var-ref* can be #f if *module-path* is absolute. In an executable, the corresponding module is carried along, including all of its dependencies.

For compile-time, the *expr* result is used by an executable creator—but not the result when the containing module is compiled. Instead, *expr* is preserved in the module as a compile-time expression (in the sense of begin-for-syntax). Later, at the time that an executable is created, the compile-time portion of the module is executed (again), and the result of *expr* is the file to be included with the executable. The reason for the extra compile-time execution is that the result of *expr* might be platform-dependent, so the result should not be stored in the (platform-independent) bytecode form of the module; the platform at executable-creation time, however, is the same as at run time for the executable. Note that *expr* is still evaluated at run-time; consequently, avoid procedures like collection-path, which depends on the source installation, and instead use relative paths and forms like (list 'lib *str* ...+).

If a path is needed only on some platforms and not on others, use define-runtime-path-list with an *expr* that produces an empty list on platforms where the path is not needed.

The enclosing path for a define-runtime-path is determined as follows from the define-runtime-path syntactic form:

- If the form has a source module according to syntax-source-module, then the source location is determined by preserving the original expression as a syntax object, extracting its source module path at run time (again using syntax-source-module), and then resolving the resulting module path index.
- If the expression has no source module, the syntax-source location associated with the form is used, if is a string or path.
- If no source module is available, and syntax-source produces no path, then current-load-relative-directory is used if it is not #f. Finally, current-directory is used if all else fails.

In the latter two cases, the path is normally preserved in (platform-specific) byte form. If it is is within the result of find-collects-dir, however, it the path is recorded relative to (find-collects-dir), and it is reconstructed using (find-collects-dir) at run time.

Examples:

```
; Access a file "data.txt" at run-time that is originally
; located in the same directory as the module source file:
(define-runtime-path data-file "data.txt")
(define (read-data)
  (with-input-from-file data-file
    (lambda ()
      (read-bytes (file-size data-file)))))

; Load a platform-specific shared object (using ffi-lib)
; that is located in a platform-specific sub-directory of the
; module's source directory:
(define-runtime-path libfit-path
  (build-path "compiled" "native" (system-library-subpath #f)
              (path-replace-suffix "libfit"
                                   (system-type 'so-suffix))))
(define libfit (ffi-lib libfit-path))

; Load a platform-specific shared object that might be installed
; as part of the operating system, or might be installed
; specifically for Racket:
(define-runtime-path libssl-so
  (case (system-type)
    [(windows) '(so "ssleay32")]
    [else '(so "libssl")]))
(define libssl (ffi-lib libssl-so))
```

(define-runtime-paths (*id* ...) *expr*)

Like `define-runtime-path`, but declares and binds multiple paths at once. The *expr* should produce as many values as *id*s.

(define-runtime-path-list *id* *expr*)

Like `define-runtime-path`, but *expr* should produce a list of paths.

(define-runtime-module-path-index *id* *module-path-expr*)

Similar to `define-runtime-path`, but *id* is bound to a module path index that encapsulates the result of *module-path-expr* relative to the enclosing module.

Use `define-runtime-module-path` to bind a module path that is passed to a reflective function like `dynamic-require` while also creating a module dependency for building and distributing executables.

```
(define-runtime-module-path id module-path)
```

Similar to `define-runtime-path`, but *id* is bound to a resolved module path. The resolved module path for *id* corresponds to *module-path* (with the same syntax as a module path for `require`), which can be relative to the enclosing module.

The `define-runtime-module-path-index` form is usually preferred, because it creates a weaker link to the referenced module. Unlike `define-runtime-module-path-index`, the `define-runtime-module-path` form creates a `for-label` dependency from an enclosing module to *module-path*. Since the dependency is merely `for-label`, *module-path* is not instantiated or visited when the enclosing module is instantiated or visited (unless such a dependency is created by other `requires`), but the code for the referenced module is loaded when the enclosing module is loaded.

```
(runtime-paths module-path)
```

This form is mainly for use by tools such as executable builders. It expands to a quoted list containing the run-time paths declared by *module-path*, returning the compile-time results of the declaration `exprs`, except that paths are converted to byte strings. The enclosing module must require (directly or indirectly) the module specified by *module-path*, which is an unquoted module path. The resulting list does *not* include module paths bound through `define-runtime-module-path`.

14.2.5 More File and Directory Utilities

```
(require racket/file)
```

The bindings documented in this section are provided by the `racket/file` and `racket` libraries, but not `racket/base`.

```
(file->string path [#:mode mode-flag]) → string?
  path : path-string?
  mode-flag : (or/c 'binary 'text) = 'binary
```

Reads all characters from *path* and returns them as a string. The *mode-flag* argument is the same as for `open-input-file`.

```
(file->bytes path [#:mode mode-flag]) → bytes?
  path : path-string?
  mode-flag : (or/c 'binary 'text) = 'binary
```

Reads all characters from *path* and returns them as a byte string. The *mode-flag* argument is the same as for `open-input-file`.

```
(file->value path [#:mode mode-flag]) → any
  path : path-string?
  mode-flag : (or/c 'binary 'text) = 'binary
```

Reads a single S-expression from *path* using `read`. The *mode-flag* argument is the same as for `open-input-file`.

```
(file->list path [proc #:mode mode-flag]) → (listof any/c)
  path : path-string?
  proc : (input-port? . -> . any/c) = read
  mode-flag : (or/c 'binary 'text) = 'binary
```

Repeatedly calls *proc* to consume the contents of *path*, until `eof` is produced. The *mode-flag* argument is the same as for `open-input-file`.

```
(file->lines  path
              [#:mode mode-flag
               #:line-mode line-mode]) → (listof string?)
  path : path-string?
  mode-flag : (or/c 'binary 'text) = 'binary
  line-mode :  (or/c 'linefeed 'return 'return-linefeed 'any 'any-one)
            = 'any
```

Read all characters from *path*, breaking them into lines. The *line-mode* argument is the same as the second argument to `read-line`, but the default is `'any` instead of `'linefeed`. The *mode-flag* argument is the same as for `open-input-file`.

```
(file->bytes-lines  path
                    [#:mode mode-flag
                     #:line-mode line-mode]) → (listof bytes?)
  path : path-string?
  mode-flag : (or/c 'binary 'text) = 'binary
  line-mode :  (or/c 'linefeed 'return 'return-linefeed 'any 'any-one)
            = 'any
```

Like `file->lines`, but reading bytes and collecting them into lines like `read-bytes-line`.

```
(display-to-file  v
                  path
                  [#:mode mode-flag
                   #:exists exists-flag]) → void?
  v : any/c
  path : path-string?
```

```
  mode-flag : (or/c 'binary 'text) = 'binary
  exists-flag : (or/c 'error 'append 'update
                      'replace 'truncate 'truncate/replace)
              = 'error
```

Uses `display` to print *v* to *path*. The *mode-flag* and *exists-flag* arguments are the same as for `open-output-file`.

```
(write-to-file  v
                path
               [#:mode mode-flag
                #:exists exists-flag]) → void?
  v : any/c
  path : path-string?
  mode-flag : (or/c 'binary 'text) = 'binary
  exists-flag : (or/c 'error 'append 'update
                      'replace 'truncate 'truncate/replace)
              = 'error
```

Like `display-to-file`, but using `write` instead of `display`.

```
(display-lines-to-file  lst
                        path
                       [#:separator separator
                        #:mode mode-flag
                        #:exists exists-flag]) → void?
  lst : list?
  path : path-string?
  separator : any/c = #"\n"
  mode-flag : (or/c 'binary 'text) = 'binary
  exists-flag : (or/c 'error 'append 'update
                      'replace 'truncate 'truncate/replace)
              = 'error
```

Displays each element of *lst* to *path*, adding *separator* after each element. The *mode-flag* and *exists-flag* arguments are the same as for `open-output-file`.

```
(copy-directory/files src dest) → void?
  src : path-string?
  dest : path-string?
```

Copies the file or directory *src* to *dest*, raising `exn:fail:filesystem` if the file or directory cannot be copied, possibly because *dest* exists already. If *src* is a directory, the copy applies recursively to the directory's content. If a source is a link, the target of the link is copied rather than the link itself.

```
(delete-directory/files path) → void?
  path : path-string?
```

Deletes the file or directory specified by *path*, raising `exn:fail:filesystem` if the file or directory cannot be deleted. If *path* is a directory, then `delete-directory/files` is first applied to each file and directory in *path* before the directory is deleted.

```
(find-files predicate [start-path]) → (listof path?)
  predicate : (path? . -> . any/c)
  start-path : (or/c path-string? #f) = #f
```

Traverses the filesystem starting at *start-path* and creates a list of all files and directories for which *predicate* returns true. If *start-path* is `#f`, then the traversal starts from `(current-directory)`. In the resulting list, each directory precedes its content.

The *predicate* procedure is called with a single argument for each file or directory. If *start-path* is `#f`, the argument is a pathname string that is relative to the current directory. Otherwise, it is a path building on *start-path*. Consequently, supplying `(current-directory)` for *start-path* is different from supplying `#f`, because *predicate* receives complete paths in the former case and relative paths in the latter. Another difference is that *predicate* is not called for the current directory when *start-path* is `#f`.

The `find-files` traversal follows soft links. To avoid following links, use the more general `fold-files` procedure.

If *start-path* does not refer to an existing file or directory, then *predicate* will be called exactly once with *start-path* as the argument.

The `find-files` procedure raises and exception if it encounters a directory for which `directory-list` fails.

```
(pathlist-closure path-list) → (listof path?)
  path-list : (listof path-string?)
```

Given a list of paths, either absolute or relative to the current directory, returns a list such that

- if a nested path is given, all of its ancestors are also included in the result (but the same ancestor is not added twice);
- if a path refers to directory, all of its descendants are also included in the result;
- ancestor directories appear before their descendants in the result list.

```
(fold-files  proc
             init-val
            [start-path
             follow-links?]) → any
        (or/c (path? (or/c 'file 'dir 'link) any/c
                . -> . any/c)
  proc :
              (path? (or/c 'file 'dir 'link) any/c
                . -> . (values any/c any/c)))
  init-val : any/c
  start-path : (or/c path-string? #f) = #f
  follow-links? : any/c = #t
```

Traverses the filesystem starting at *start-path*, calling *proc* on each discovered file, directory, and link. If *start-path* is `#f`, then the traversal starts from `(current-directory)`.

The *proc* procedure is called with three arguments for each file, directory, or link:

- If *start-path* is `#f`, the first argument is a pathname string that is relative to the current directory. Otherwise, the first argument is a pathname that starts with *start-path*. Consequently, supplying `(current-directory)` for *start-path* is different from supplying `#f`, because *proc* receives complete paths in the former case and relative paths in the latter. Another difference is that *proc* is not called for the current directory when *start-path* is `#f`.

- The second argument is a symbol, either `'file`, `'dir`, or `'link`. The second argument can be `'link` when *follow-links?* is `#f`, in which case the filesystem traversal does not follow links. If *follow-links?* is `#t`, then *proc* will only get a `'link` as a second argument when it encounters a dangling symbolic link (one that does not resolve to an existing file or directory).

- The third argument is the accumulated result. For the first call to *proc*, the third argument is *init-val*. For the second call to *proc* (if any), the third argument is the result from the first call, and so on. The result of the last call to *proc* is the result of `fold-files`.

The *proc* argument is used in an analogous way to the procedure argument of `foldl`, where its result is used as the new accumulated result. There is an exception for the case of a directory (when the second argument is `'dir`): in this case the procedure may return two values, the second indicating whether the recursive scan should include the given directory or not. If it returns a single value, the directory is scanned. In the cases of files or links (when the second argument is `'file` or `'link`), a second value is permitted but ignored.

If the *start-path* is provided but no such path exists, or if paths disappear during the scan, then an exception is raised.

```
(make-directory* path) → void?
  path : path-string?
```

Creates directory specified by *path*, creating intermediate directories as necessary.

```
(make-temporary-file [template
                      copy-from-filename
                      directory])         → path?
  template : string? = "rkttmp~a"
  copy-from-filename : (or/c path-string? #f 'directory) = #f
  directory : (or/c path-string? #f) = #f
```

Creates a new temporary file and returns a pathname string for the file. Instead of merely generating a fresh file name, the file is actually created; this prevents other threads or processes from picking the same temporary name.

The *template* argument must be a format string suitable for use with `format` and one additional string argument (where the string contains only digits). If the resulting string is a relative path, it is combined with the result of `(find-system-path 'temp-dir)`, unless *directory* is provided and non-`#f`, in which case the file name generated from *template* is combined with *directory* to obtain a full path.

The *template* argument's default is only the string `"rkttmp~a"` when there is no source location information for the callsite of `make-temporary-file` (or if `make-temporary-file` is used in a higher-order position). If there is such information, then the template string is based on the source location.

If *copy-from-filename* is provided as path, the temporary file is created as a copy of the named file (using `copy-file`). If *copy-from-filename* is `#f`, the temporary file is created as empty. If *copy-from-filename* is `'directory`, then the temporary "file" is created as a directory.

When a temporary file is created, it is not opened for reading or writing when the pathname is returned. The client program calling `make-temporary-file` is expected to open the file with the desired access and flags (probably using the `'truncate` flag; see `open-output-file`) and to delete it when it is no longer needed.

```
(get-preference  name
                [failure-thunk
                 flush-mode
                 filename
                 #:use-lock? use-lock?
                 #:timeout-lock-there timeout-lock-there
                 #:lock-there lock-there])                → any
  name : symbol?
```

```
failure-thunk : (-> any) = (lambda () #f)
flush-mode : any/c = 'timestamp
filename : (or/c string-path? #f) = #f
use-lock? : any/c = #t
timeout-lock-there : (or/c (path? . -> . any) #f) = #f
lock-there : (or/c (path? . -> . any) #f)
               (make-handle-get-preference-locked
           =    0.01 name failure-thunk flush-mode filename
                #:lock-there timeout-lock-there)
```

Extracts a preference value from the file designated by `(find-system-path 'pref-file)`, or by *filename* if it is provided and is not `#f`. In the former case, if the preference file doesn't exist, `get-preferences` attempts to read an old preferences file, and then a `"racket-prefs.rktd"` file in the `"defaults"` collection, instead. If none of those files exists, the preference set is empty.

The preference file should contain a list of symbol–value lists written with the default parameter settings. Keys starting with `racket:`, `mzscheme:`, `mred:`, and `plt:` in any letter case are reserved for use by Racket implementors. If the preference file does not contain a list of symbol–value lists, an error is logged via `log-error` and *failure-thunk* is called.

The result of `get-preference` is the value associated with *name* if it exists in the association list, or the result of calling *failure-thunk* otherwise.

Preference settings are cached (weakly) across calls to `get-preference`, using `(path->complete-path` *filename*`)` as a cache key. If *flush-mode* is provided as `#f`, the cache is used instead of re-consulting the preferences file. If *flush-mode* is provided as `'timestamp` (the default), then the cache is used only if the file has a timestamp that is the same as the last time the file was read. Otherwise, the file is re-consulted.

On platforms for which `preferences-lock-file-mode` returns `'file-lock` and when *use-lock?* is true, preference-file reading is guarded by a lock; multiple readers can share the lock, but writers take the lock exclusively. If the preferences file cannot be read because the lock is unavailable, *lock-there* is called on the path of the lock file; if *lock-there* is `#f`, an exception is raised. The default *lock-there* handler retries about 5 times (with increasing delays between each attempt) before trying *timeout-lock-there*, and the default *timeout-lock-there* triggers an exception.

See also `put-preferences`. For a more elaborate preference system, see `preferences:get`.

Old preferences files: When a *filename* is not provided and the file indicated by `(find-system-path 'pref-file)` does not exist, the following paths are checked for compatibility with old versions of Racket:

- Windows: `(build-path (find-system-path 'pref-dir) 'up "PLT`

Scheme" "plt-prefs.ss")

- Mac OS X: (build-path (find-system-path 'pref-dir) "org.plt-scheme.prefs.ss")
- Unix: (expand-user-path "~/.plt-scheme/plt-prefs.ss")

```
(put-preferences names
                 vals
                [locked-proc
                 filename])  → void?
  names : (listof symbol?)
  vals : list?
  locked-proc : (or/c #f (path? . -> . any)) = #f
  filename : (or/c #f path-string?) = #f
```

Installs a set of preference values and writes all current values to the preference file designated by (find-system-path 'pref-file), or *filename* if it is supplied and not #f.

The *names* argument supplies the preference names, and *vals* must have the same length as *names*. Each element of *vals* must be an instance of a built-in data type whose write output is readable (i.e., the print-unreadable parameter is set to #f while writing preferences).

Current preference values are read from the preference file before updating, and a write lock is held starting before the file read, and lasting until after the preferences file is updated. The lock is implemented by the existence of a file in the same directory as the preference file; see preferences-lock-file-mode for more information. If the directory of the preferences file does not already exist, it is created.

If the write lock is already held, then *locked-proc* is called with a single argument: the path of the lock file. The default *locked-proc* (used when the *locked-proc* argument is #f) reports an error; an alternative thunk might wait a while and try again, or give the user the choice to delete the lock file (in case a previous update attempt encountered disaster and locks are implemented by the presence of the lock file).

If *filename* is #f or not supplied, and the preference file does not already exist, then values read from the "defaults" collection (if any) are written for preferences that are not mentioned in *names*.

```
(preferences-lock-file-mode) → (or/c 'exists 'file-lock)
```

Reports the way that the lock file is used to implement preference-file locking on the current platform.

The 'exists mode is currently used on all platforms except Windows. In 'exists mode, the existence of the lock file indicates that a write lock is held, and readers need no lock (because the preferences file is atomically updated via rename-file-or-directory).

The 'file-lock mode is currently used on Windows. In 'file-lock mode, shared and exclusive locks (in the sense of port-try-file-lock?) on the lock file reflect reader and writer locks on the preference-file content. (The preference file itself is not locked, because a lock would interfere with replacing the file via rename-file-or-directory.)

```
(make-handle-get-preference-locked  delay
                                    name
                                   [failure-thunk
                                    flush-mode
                                    filename
                                    #:lock-there lock-there
                                    #:max-delay max-delay])
 → (path-string? . -> . any)
  delay : real?
  name : symbol?
  failure-thunk : (-> any) = (lambda () #f)
  flush-mode : any/c = 'timestamp
  filename : (or/c path-string? #f) = #f
  lock-there : (or/c (path? . -> . any) #f) = #f
  max-delay : real? = 0.2
```

Creates a procedure suitable for use as the #:lock-there argument to get-preference, where the *name*, *failure-thunk*, *flush-mode*, and *filename* are all passed on to get-preference by the result procedure to retry the preferences lookup.

Before calling get-preference, the result procedure uses (sleep *delay*) to pause. Then, if (* 2 *delay*) is less than *max-delay*, the result procedure calls make-handle-get-preference-locked to generate a new retry procedure to pass to get-preference, but with a *delay* of (* 2 *delay*). If (* 2 *delay*) is not less than *max-delay*, then get-preference is called with the given *lock-there*, instead.

```
(call-with-file-lock/timeout  filename
                              kind
                              thunk
                              failure-thunk
                             [#:lock-file lock-file
                              #:delay delay
                              #:max-delay max-delay]) → any
  filename : (or/c path-string? #f)
  kind : (or/c 'shared 'exclusive)
  thunk : (-> any)
  failure-thunk : (-> any)
  lock-file : (or/c #f path-string?) = #f
  delay : (and/c real? (not/c negative?)) = 0.01
  max-delay : (and/c real? (not/c negative?)) = 0.2
```

Obtains a lock for the filename *lock-file* and then calls *thunk*. The *filename* argument specifies a file path prefix that is used only to generate the lock filename when *lock-file* is #f. Specifically, when *lock-file* is #f, then call-with-file-lock/timeout uses make-lock-file-name to build the lock filename. If the lock file does not yet exist, it is created; beware that the lock file is *not* deleted by call-with-file-lock/timeout.

When *thunk* returns, call-with-file-lock/timeout releases the lock, returning the result of *thunk*. The call-with-file-lock/timeout function will retry after *delay* seconds and continue retrying with exponential backoff until delay reaches *max-delay*. If call-with-file-lock/timeout fails to obtain the lock, *failure-thunk* is called in tail position. The *kind* argument specifies whether the lock is 'shared or 'exclusive in the sense of port-try-file-lock?.

Examples:

```
> (call-with-file-lock/timeout filename 'exclusive
    (lambda () (printf "File is locked\n"))
    (lambda () (printf "Failed to obtain lock for file\n")))
File is locked

> (call-with-file-lock/timeout #f 'exclusive
    (lambda ()
      (call-with-file-lock/timeout filename 'shared
        (lambda () (printf "Shouldn't get here\n"))
        (lambda () (printf "Failed to obtain lock for file\n"))))
    (lambda () (printf "Shouldn't get here either\n"))
    #:lock-file (make-lock-file-name filename))
Failed to obtain lock for file
```

```
(make-lock-file-name path) → path?
  path : (or path-string? path-for-some-system?)
(make-lock-file-name dir name) → path?
  dir : (or path-string? path-for-some-system?)
  name : path-element?
```

Creates a lock filename by prepending "_LOCK" on Windows or ".LOCK" on other platforms to the file portion of the path.

Example:

```
> (make-lock-file-name "/home/george/project/important-file")
#<path:/home/george/project/.LOCKimportant-file>
```

```
user-read-bit : #o400
user-write-bit : #o200
user-execute-bit : #o100
```

```
group-read-bit : #o040
group-write-bit : #o020
group-execute-bit : #o010
other-read-bit : #o004
other-write-bit : #o002
other-execute-bit : #o001
```

Constants that are useful with `file-or-directory-permissions` and bitwise operations such as `bitwise-ior`, and `bitwise-and`.

14.3 Networking

14.3.1 TCP

```
(require racket/tcp)
```

The bindings documented in this section are provided by the `racket/tcp` and `racket` libraries, but not `racket/base`.

For information about TCP in general, see *TCP/IP Illustrated, Volume 1* by W. Richard Stevens.

```
(tcp-listen  port-no
             [max-allow-wait
              reuse?
              hostname])      → tcp-listener?
  port-no : (and/c exact-nonnegative-integer?
                   (integer-in 0 65535))
  max-allow-wait : exact-nonnegative-integer? = 4
  reuse? : any/c = #f
  hostname : (or/c string? #f) = #f
```

Creates a "listening" server on the local machine at the port number specified by *port-no*. If *port-no* is 0 the socket binds to an ephemeral port, which can be determined by calling `tcp-addresses`. The *max-allow-wait* argument determines the maximum number of client connections that can be waiting for acceptance. (When *max-allow-wait* clients are waiting acceptance, no new client connections can be made.)

If the *reuse?* argument is true, then `tcp-listen` will create a listener even if the port is involved in a `TIME_WAIT` state. Such a use of *reuse?* defeats certain guarantees of the TCP protocol; see Stevens's book for details. Furthermore, on many modern platforms, a true value for *reuse?* overrides `TIME_WAIT` only if the listener was previously created with a true value for *reuse?*.

If *hostname* is #f (the default), then the listener accepts connections to all of the listening machine's addresses. Otherwise, the listener accepts connections only at the interface(s) associated with the given hostname. For example, providing "127.0.0.1" as *hostname* creates a listener that accepts only connections to "127.0.0.1" (the loopback interface) from the local machine.

(Racket implements a listener with multiple sockets, if necessary, to accommodate multiple addresses with different protocol families. On Linux, if *hostname* maps to both IPv4 and IPv6 addresses, then the behavior depends on whether IPv6 is supported and IPv6 sockets can be configured to listen to only IPv6 connections: if IPv6 is not supported or IPv6 sockets are not configurable, then the IPv6 addresses are ignored; otherwise, each IPv6 listener accepts only IPv6 connections.)

The return value of tcp-listen is a *TCP listener*. This value can be used in future calls to tcp-accept, tcp-accept-ready?, and tcp-close. Each new TCP listener value is placed into the management of the current custodian (see §13.7 "Custodians").

If the server cannot be started by tcp-listen, the exn:fail:network exception is raised.

```
(tcp-connect  hostname
              port-no
             [local-hostname
              local-port-no]) → input-port? output-port?
  hostname : string?
  port-no :  (and/c exact-nonnegative-integer?
                    (integer-in 1 65535))
  local-hostname : (or/c string? #f) = #f
                   (or/c (and/c exact-nonnegative-integer?
  local-port-no :                   (integer-in 1 65535))      = #f
                        #f)
```

Attempts to connect as a client to a listening server. The *hostname* argument is the server host's Internet address name, and *port-no* is the port number where the server is listening.

(If *hostname* is associated with multiple addresses, they are tried one at a time until a connection succeeds. The name "localhost" generally specifies the local machine.)

The optional *local-hostname* and *local-port-no* specify the client's address and port. If both are #f (the default), the client's address and port are selected automatically. If *local-hostname* is not #f, then *local-port-no* must be non-#f. If *local-port-no* is non-#f and *local-hostname* is #f, then the given port is used but the address is selected automatically.

Two values are returned by tcp-connect: an input port and an output port. Data can be received from the server through the input port and sent to the server through the output port. If the server is a Racket program, it can obtain ports to communicate to the client with tcp-accept. These ports are placed into the management of the current custodian (see §13.7

“Custodians”).

Initially, the returned input port is block-buffered, and the returned output port is block-buffered. Change the buffer mode using `file-stream-buffer-mode`.

Both of the returned ports must be closed to terminate the TCP connection. When both ports are still open, closing the output port with `close-output-port` sends a TCP close to the server (which is seen as an end-of-file if the server reads the connection through a port). In contrast, `tcp-abandon-port` (see below) closes the output port, but does not send a TCP close until the input port is also closed.

Note that the TCP protocol does not support a state where one end is willing to send but not read, nor does it include an automatic message when one end of a connection is fully closed. Instead, the other end of a connection discovers that one end is fully closed only as a response to sending data; in particular, some number of writes on the still-open end may appear to succeed, though writes will eventually produce an error.

If a connection cannot be established by `tcp-connect`, the `exn:fail:network` exception is raised.

```
(tcp-connect/enable-break  hostname
                           port-no
                          [local-hostname]
                           local-port-no)
 → input-port? output-port?
  hostname : string?
  port-no :          (and/c exact-nonnegative-integer?
                            (integer-in 1 65535))
  local-hostname : (or/c string? #f) = #f
                   (or/c (and/c exact-nonnegative-integer?
  local-port-no :                   (integer-in 1 65535))
                        #f)
```

Like `tcp-connect`, but breaking is enabled (see §9.6 “Breaks”) while trying to connect. If breaking is disabled when `tcp-connect/enable-break` is called, then either ports are returned or the `exn:break` exception is raised, but not both.

```
(tcp-accept listener) → input-port? output-port?
  listener : tcp-listener?
```

Accepts a client connection for the server associated with *listener*. If no client connection is waiting on the listening port, the call to `tcp-accept` will block. (See also `tcp-accept-ready?`.)

Two values are returned by `tcp-accept`: an input port and an output port. Data can be received from the client through the input port and sent to the client through the output port.

These ports are placed into the management of the current custodian (see §13.7 “Custodians”).

In terms of buffering and connection states, the ports act the same as ports from tcp-connect.

If a connection cannot be accepted by tcp-accept, or if the listener has been closed, the exn:fail:network exception is raised.

```
(tcp-accept/enable-break listener) → input-port? output-port?
  listener : tcp-listener?
```

Like tcp-accept, but breaking is enabled (see §9.6 “Breaks”) while trying to accept a connection. If breaking is disabled when tcp-accept/enable-break is called, then either ports are returned or the exn:break exception is raised, but not both.

```
(tcp-accept-ready? listener) → boolean?
  listener : tcp-listener?
```

Tests whether an unaccepted client has connected to the server associated with *listener*. If a client is waiting, the return value is #t, otherwise it is #f. A client is accepted with the tcp-accept procedure, which returns ports for communicating with the client and removes the client from the list of unaccepted clients.

If the listener has been closed, the exn:fail:network exception is raised.

```
(tcp-close listener) → void?
  listener : tcp-listener?
```

Shuts down the server associated with *listener*. All unaccepted clients receive an end-of-file from the server; connections to accepted clients are unaffected.

If the listener has already been closed, the exn:fail:network exception is raised.

The listener’s port number may not become immediately available for new listeners (with the default reuse? argument of tcp-listen). For further information, see Stevens’s explanation of the TIME_WAIT TCP state.

```
(tcp-listener? v) → boolean?
  v : any/c
```

Returns #t if *v* is a TCP listener created by tcp-listen, #f otherwise.

```
(tcp-accept-evt listener) → evt?
  listener : tcp-listener?
```

Returns a synchronizable event (see §10.2.1 “Events”) that is in a blocking state when `tcp-accept` on *`listener`* would block. If the event is chosen in a synchronization, the result is a list of two items, which correspond to the two results of `tcp-accept`. (If the event is not chosen, no connections are accepted.) The ports are placed into the management of the custodian that is the current custodian (see §13.7 “Custodians”) at the time that `tcp-accept-evt` is called.

```
(tcp-abandon-port tcp-port) → void?
  tcp-port : tcp-port?
```

Like `close-output-port` or `close-input-port` (depending on whether *`tcp-port`* is an input or output port), but if *`tcp-port`* is an output port and its associated input port is not yet closed, then the other end of the TCP connection does not receive a TCP close message until the input port is also closed.

The TCP protocol does not include a “no longer reading” state on connections, so `tcp-abandon-port` is equivalent to `close-input-port` on input TCP ports.

```
(tcp-addresses tcp-port [port-numbers?])
    (or/c (values string? string?)
 →        (values string? (integer-in 1 65535)
                  string? (integer-in 0 65535)))
  tcp-port : (or/c tcp-port? tcp-listener?)
  port-numbers? : any/c = #f
```

Returns two strings when *`port-numbers?`* is `#f` (the default). The first string is the Internet address for the local machine a viewed by the given TCP port’s connection or for the TCP listener. (For most machines, the answer corresponds to the current machine’s only Internet address, but when a machine serves multiple addresses, the result is connection-specific or listener-specific.) If a listener is given and it has no specific host, the first string result is `"0.0.0.0"`. The second string is the Internet address for the other end of the connection, or always `"0.0.0.0"` for a listener.

If *`port-numbers?`* is true, then four results are returned: a string for the local machine’s address, an exact integer between `1` and `65535` for the local machine’s port number, a string for the remote machine’s address, and an exact integer between `1` and `65535` for the remote machine’s port number or `0` for a listener.

If the given port has been closed, the `exn:fail:network` exception is raised.

```
(tcp-port? v) → boolean?
  v : any/c
```

Returns `#t` if *`v`* is a *TCP port*—which is a port returned by `tcp-accept`, `tcp-connect`, `tcp-accept/enable-break`, or `tcp-connect/enable-break`—`#f` otherwise.

14.3.2 UDP

```
(require racket/udp)
```

The bindings documented in this section are provided by the `racket/udp` and `racket` libraries, but not `racket/base`.

For information about UDP in general, see *TCP/IP Illustrated, Volume 1* by W. Richard Stevens.

```
(udp-open-socket [family-hostname
                  family-port-no]) → udp?
  family-hostname : (or/c string? #f) = #f
  family-port-no : (or/c string? #f) = #f
```

Creates and returns a *UDP socket* to send and receive datagrams (broadcasting is allowed). Initially, the socket is not bound or connected to any address or port.

If *family-hostname* or *family-port-no* is not `#f`, then the socket's protocol family is determined from these arguments. The socket is *not* bound to the hostname or port number. For example, the arguments might be the hostname and port to which messages will be sent through the socket, which ensures that the socket's protocol family is consistent with the destination. Alternately, the arguments might be the same as for a future call to `udp-bind!`, which ensures that the socket's protocol family is consistent with the binding. If neither *family-hostname* nor *family-port-no* is non-`#f`, then the socket's protocol family is IPv4.

```
(udp-bind! udp-socket
           hostname-string
           port-no)        → void?
  udp-socket : udp?
  hostname-string : (or/c string? #f)
  port-no : (and/c exact-nonnegative-integer?
                   (integer-in 0 65535))
```

Binds an unbound *udp-socket* to the local port number *port-no*. If *port-no* is 0 the *udp-socket* is bound to an ephemeral port, which can be determined by calling `udp-addresses`.

If *hostname-string* is `#f`, then the socket accepts connections to all of the listening machine's IP addresses at *port-no*. Otherwise, the socket accepts connections only at the IP address associated with the given name. For example, providing `"127.0.0.1"` as *hostname-string* typically creates a listener that accepts only connections to `"127.0.0.1"` from the local machine.

A socket cannot receive datagrams until it is bound to a local address and port. If a socket is not bound before it is used with a sending procedure `udp-send`, `udp-send-to`, etc., the

sending procedure binds the socket to a random local port. Similarly, if an event from udp-send-evt or udp-send-to-evt is chosen for a synchronization (see §10.2.1 “Events”), the socket is bound; if the event is not chosen, the socket may or may not become bound.

The binding of a bound socket cannot be changed, with one exception: on some systems, if the socket is bound automatically when sending, if the socket is disconnected via udp-connect!, and if the socket is later used again in a send, then the later send may change the socket's automatic binding.

If *udp-socket* is already bound or closed, the exn:fail:network exception is raised.

```
(udp-connect! udp-socket
              hostname-string
              port-no)        → void?
  udp-socket : udp?
  hostname-string : (or/c string? #f)
            (or/c (and/c exact-nonnegative-integer?
  port-no :              (integer-in 1 65535))
                  #f)
```

Connects the socket to the indicated remote address and port if *hostname-string* is a string and *port-no* is an exact integer.

If *hostname-string* is #f, then *port-no* also must be #f, and the port is disconnected (if connected). If one of *hostname-string* or *port-no* is #f and the other is not, the exn:fail:contract exception is raised.

A connected socket can be used with udp-send (not udp-send-to), and it accepts datagrams only from the connected address and port. A socket need not be connected to receive datagrams. A socket can be connected, re-connected, and disconnected any number of times.

If *udp-socket* is closed, the exn:fail:network exception is raised.

```
(udp-send-to  udp-socket
              hostname
              port-no
              bstr
             [start-pos
              end-pos])   → void
  udp-socket : udp?
  hostname : string?
            (and/c exact-nonnegative-integer?
  port-no :        (integer-in 1 65535))
  bstr : bytes?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Sends (subbytes bytes start-pos end-pos) as a datagram from the unconnected udp-socket to the socket at the remote machine hostname-address on the port port-no. The udp-socket need not be bound or connected; if it is not bound, udp-send-to binds it to a random local port. If the socket's outgoing datagram queue is too full to support the send, udp-send-to blocks until the datagram can be queued.

If start-pos is greater than the length of bstr, or if end-pos is less than start-pos or greater than the length of bstr, the exn:fail:contract exception is raised.

If udp-socket is closed or connected, the exn:fail:network exception is raised.

```
(udp-send udp-socket bstr [start-pos end-pos]) → void
  udp-socket : udp?
  bstr : bytes?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like udp-send-to, except that udp-socket must be connected, and the datagram goes to the connection target. If udp-socket is closed or unconnected, the exn:fail:network exception is raised.

```
(udp-send-to* udp-socket
              hostname
              port-no
              bstr
             [start-pos
              end-pos])  → boolean?
  udp-socket : udp?
  hostname : string?
  port-no : (and/c exact-nonnegative-integer?
                   (integer-in 1 65535))
  bstr : bytes?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like udp-send-to, but never blocks; if the socket's outgoing queue is too full to support the send, #f is returned, otherwise the datagram is queued and the result is #t.

```
(udp-send* udp-socket bstr [start-pos end-pos]) → boolean?
  udp-socket : udp?
  bstr : bytes?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like udp-send, except that (like udp-send-to) it never blocks and returns #f or #t.

```
(udp-send-to/enable-break  udp-socket
                           hostname
                           port-no
                           bstr
                          [start-pos
                           end-pos])   → void
  udp-socket : udp?
  hostname : string?
  port-no : (and/c exact-nonnegative-integer?
                   (integer-in 1 65535))
  bstr : bytes?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like udp-send-to, but breaking is enabled (see §9.6 “Breaks”) while trying to send the datagram. If breaking is disabled when udp-send-to/enable-break is called, then either the datagram is sent or the exn:break exception is raised, but not both.

```
(udp-send/enable-break  udp-socket
                        bstr
                       [start-pos
                        end-pos])   → void
  udp-socket : udp?
  bstr : bytes?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like udp-send, except that breaks are enabled like udp-send-to/enable-break.

```
(udp-receive!  udp-socket
               bstr
              [start-pos
                                  exact-nonnegative-integer?
               end-pos])   → string?
                                  (integer-in 1 65535)
  udp-socket : udp?
  bstr : (and/c bytes? (not immutable?))
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Accepts up to *end-pos-start-pos* bytes of *udp-socket*’s next incoming datagram into *bstr*, writing the datagram bytes starting at position *start-pos* within *bstr*. The *udp-socket* must be bound to a local address and port (but need not be connected). If no incoming datagram is immediately available, udp-receive! blocks until one is available.

Three values are returned: the number of received bytes (between 0 and *end-pos*-*start-pos*, a hostname string indicating the source address of the datagram, and an integer indicating the source port of the datagram. If the received datagram is longer than *end-pos*-*start-pos* bytes, the remainder is discarded.

If *start-pos* is greater than the length of *bstr*, or if *end-pos* is less than *start-pos* or greater than the length of *bstr*, the exn:fail:contract exception is raised.

```
(udp-receive!* udp-socket
               bstr
              [start-pos
               end-pos])
   (or/c exact-nonnegative-integer? #f)
 → (or/c string? #f)
   (or/c (integer-in 1 65535) #f)
  udp-socket : udp?
  bstr : (and/c bytes? (not immutable?))
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like udp-receive!, except that it never blocks. If no datagram is available, the three result values are all #f.

```
(udp-receive!/enable-break udp-socket
                           bstr
                          [start-pos
                           end-pos])
   exact-nonnegative-integer?
 → string?
   (integer-in 1 65535)
  udp-socket : udp?
  bstr : (and/c bytes? (not immutable?))
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Like udp-receive!, but breaking is enabled (see §9.6 “Breaks”) while trying to receive the datagram. If breaking is disabled when udp-receive!/enable-break is called, then either a datagram is received or the exn:break exception is raised, but not both.

```
(udp-close udp-socket) → void?
  udp-socket : udp?
```

Closes *udp-socket*, discarding unreceived datagrams. If the socket is already closed, the exn:fail:network exception is raised.

```
(udp? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a socket created by `udp-open-socket`, `#f` otherwise.

```
(udp-bound? udp-socket) → boolean?
  udp-socket : udp?
```

Returns `#t` if *udp-socket* is bound to a local address and port, `#f` otherwise.

```
(udp-connected? udp-socket) → boolean?
  udp-socket : udp?
```

Returns `#t` if *udp-socket* is connected to a remote address and port, `#f` otherwise.

```
(udp-send-ready-evt udp-socket) → evt?
  udp-socket : udp?
```

Returns a synchronizable event (see §10.2.1 “Events”) that is in a blocking state when `udp-send-to` on *udp-socket* would block.

```
(udp-receive-ready-evt udp-socket) → evt?
  udp-socket : udp?
```

Returns a synchronizable event (see §10.2.1 “Events”) that is in a blocking state when `udp-receive!` on *udp-socket* would block.

```
(udp-send-to-evt  udp-socket
                  hostname
                  port-no
                  bstr
                 [start-pos
                  end-pos])   → evt?
  udp-socket : udp?
  hostname : string?
  port-no :  (and/c exact-nonnegative-integer?
                    (integer-in 1 65535))
  bstr : bytes?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Returns a synchronizable event. The event is in a blocking state when `udp-send-to` on *udp-socket* would block. Otherwise, if the event is chosen in a synchronization, data is sent as for `(udp-send-to udp-socket hostname-address port-no bstr start-pos end-pos)`, and the synchronization result is `#<void>`. (No bytes are sent if the event is not chosen.)

```
(udp-send-evt  udp-socket
               bstr
              [start-pos
               end-pos])   → evt?
  udp-socket : udp?
  bstr : bytes?
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Returns a synchronizable event. The event is in a blocking state when `udp-send` on *udp-socket* would block. Otherwise, if the event is chosen in a synchronization, data is sent as for `(udp-send-to` *udp-socket bstr start-pos end-pos*`)`, and the synchronization result is `#<void>`. (No bytes are sent if the event is not chosen.) If *udp-socket* is closed or unconnected, the `exn:fail:network` exception is raised during a synchronization attempt.

```
(udp-receive!-evt  udp-socket
                   bstr
                  [start-pos
                   end-pos])   → evt?
  udp-socket : udp?
  bstr : (and/c bytes? (not immutable?))
  start-pos : exact-nonnegative-integer? = 0
  end-pos : exact-nonnegative-integer? = (bytes-length bstr)
```

Returns a synchronizable event. The event is in a blocking state when `udp-receive` on *udp-socket* would block. Otherwise, if the event is chosen in a synchronization, data is received into *bstr* as for `(udp-receive!` *udp-socket bytes start-pos end-pos*`)`, and the synchronization result is a list of three values, corresponding to the three results from `udp-receive!`. (No bytes are received and the *bstr* content is not modified if the event is not chosen.)

```
(udp-addresses udp-port [port-numbers?])
    (or/c (values string? string?)
 →        (values string? (integer-in 0 65535)
                  string? (integer-in 0 65535)))
  udp-port : udp?
  port-numbers? : any/c = #f
```

Returns two strings when *port-numbers?* is `#f` (the default). The first string is the Internet address for the local machine a viewed by the given UDP socket's connection. (For most machines, the answer corresponds to the current machine's only Internet address, but when a machine serves multiple addresses, the result is connection-specific.) The second string is the Internet address for the other end of the connection.

If *port-numbers?* is true, then four results are returned: a string for the local machine's address, an exact integer between `1` and `65535` for the local machine's port number or `0` if the

socket is unbound, a string for the remote machine's address, and an exact integer between 1 and 65535 for the remote machine's port number or 0 if the socket is unconnected.

If the given port has been closed, the exn:fail:network exception is raised.

14.4 Processes

```
(subprocess stdout
            stdin
            stderr
            command
            arg ...)
  subprocess?
→ (or/c (and/c input-port? file-stream-port?) #f)
  (or/c (and/c output-port? file-stream-port?) #f)
  (or/c (and/c input-port? file-stream-port?) #f)
 stdout : (or/c (and/c output-port? file-stream-port?) #f)
 stdin : (or/c (and/c input-port? file-stream-port?) #f)
 stderr : (or/c (and/c output-port? file-stream-port?) #f 'stdout)
 command : path-string?
 arg : (or/c path? string-no-nuls? bytes-no-nuls?)
(subprocess stdout
            stdin
            stderr
            command
            exact
            arg)
  subprocess?
→ (or/c (and/c input-port? file-stream-port?) #f)
  (or/c (and/c output-port? file-stream-port?) #f)
  (or/c (and/c input-port? file-stream-port?) #f)
 stdout : (or/c (and/c output-port? file-stream-port?) #f)
 stdin : (or/c (and/c input-port? file-stream-port?) #f)
 stderr : (or/c (and/c output-port? file-stream-port?) #f)
 command : path-string?
 exact : 'exact
 arg : string?
```

Creates a new process in the underlying operating system to execute *command* asynchronously. See also system and process from racket/system.

The *command* argument is a path to a program executable, and the *arg*s are command-line arguments for the program. See find-executable-path for locating an executable based on the PATH environment variable. On Unix and Mac OS X, command-line arguments are

passed as byte strings, and string *arg*s are converted using the current locale's encoding (see §12.1.1 "Encodings and Locales"). On Windows, command-line arguments are passed as strings, and bytes strings are converted using UTF-8.

On Windows, the first *arg* can be replaced with `'exact`, which triggers a Windows-specific behavior: the sole *arg* is used exactly as the command-line for the subprocess. Otherwise, on Windows, a command-line string is constructed from *command* and *arg* so that a typical Windows console application can parse it back to an array of arguments. If `'exact` is provided on a non-Windows platform, the `exn:fail:contract` exception is raised.

For information on the Windows command-line conventions, search for "command line parsing" at `http://msdn.microsoft.com/`.

When provided as a port, *stdout* is used for the launched process's standard output, *stdin* is used for the process's standard input, and *stderr* is used for the process's standard error. All provided ports must be file-stream ports. Any of the ports can be `#f`, in which case a system pipe is created and returned by `subprocess`. The *stderr* argument can be `'stdout`, in which case the same file-stream port or system pipe that is supplied as standard output is also used for standard error. For each port or `'stdout` that is provided, no pipe is created and the corresponding returned value is `#f`.

The `subprocess` procedure returns four values:

- a subprocess value representing the created process;
- an input port piped from the process's standard output, or `#f` if `stdout-output-port` was a port;
- an output port piped to the process standard input, or `#f` if `stdin-input-port` was a port;
- an input port piped from the process's standard error, or `#f` if `stderr-output-port` was a port or `'stdout`.

Important: All ports returned from `subprocess` must be explicitly closed, usually with `close-input-port` or `close-output-port`.

The returned ports are file-stream ports (see §12.1.5 "File Ports"), and they are placed into the management of the current custodian (see §13.7 "Custodians"). The `exn:fail` exception is raised when a low-level error prevents the spawning of a process or the creation of operating system pipes for process communication.

If the `subprocess-group-enabled` parameter's value is true, then the new process is created as a new OS-level process group. In that case, `subprocess-kill` attempts to terminate all processes within the group, which may include additional processes created by the subprocess. See `subprocess-kill` for details, and see `subprocess-group-enabled` for additional caveats.

The `current-subprocess-custodian-mode` parameter determines whether the subprocess itself is registered with the current custodian so that a custodian shutdown calls `subprocess-kill` for the subprocess.

```
(subprocess-wait subproc) → void?
  subproc : subprocess?
```

Blocks until the process represented by *subproc* terminates. The *subproc* value also can be used with `sync` and `sync/timeout`.

```
(subprocess-status subproc) → (or/c 'running
                                    exact-nonnegative-integer?)
  subproc : subprocess?
```

Returns `'running` if the process represented by *subproc* is still running, or its exit code otherwise. The exit code is an exact integer, and `0` typically indicates success. If the process terminated due to a fault or signal, the exit code is non-zero.

```
(subprocess-kill subproc force?) → void?
  subproc : subprocess?
  force? : any/c
```

Terminates the subprocess represented by *subproc*. The precise action depends on whether *force?* is true, whether the process was created in its own group by setting the `subprocess-group-enabled` parameter to a true value, and the current platform:

- *force?* is true, not a group, all platforms: Terminates the process if the process still running.
- *force?* is false, not a group, on Unix or Mac OS X: Sends the process an interrupt signal instead of a kill signal.
- *force?* is false, not a group, on Windows: No action is taken.
- *force?* is true, a group, on Unix or Mac OS X: Terminates all processes in the group, but only if `subprocess-status` has never produced a non-`'running` result for the subprocess and only if functions like `subprocess-wait` and `sync` have not detected the subprocess's completion. Otherwise, no action is taken (because the immediate process is known to have terminated while the continued existence of the group is unknown).
- *force?* is true, a group, on Windows: Terminates the process if the process still running.
- *force?* is false, a group, on Unix or Mac OS X: The same as when *force?* is `#t`, but when the group is sent a signal, it is an interrupt signal instead of a kill signal.
- *force?* is false, a group, on Windows: All processes in the group receive a CTRL-BREAK signal (independent of whether the immediate subprocess has terminated).

If an error occurs during termination, the exn:fail exception is raised.

```
(subprocess-pid subproc) → exact-nonnegative-integer?
  subproc : subprocess?
```

Returns the operating system's numerical ID (if any) for the process represented by *subproc*. The result is valid only as long as the process is running.

```
(subprocess? v) → boolean?
  v : any/c
```

Returns #t if *v* is a subprocess value, #f otherwise.

```
(current-subprocess-custodian-mode)
 → (or/c #f 'kill 'interrupt)
(current-subprocess-custodian-mode mode) → void?
  mode : (or/c #f 'kill 'interrupt)
```

A parameter that determines whether a subprocess (as created by subprocess or wrappers like process) is registered with the current custodian. If the parameter value is #f, then the subprocess is not registered with the custodian—although any created ports are registered. If the parameter value is 'kill or 'interrupt, then the subprocess is shut down through subprocess-kill, where 'kill supplies a #t value for the *force?* argument and 'interrupt supplies a #f value. The shutdown may occur either before or after ports created for the subprocess are closed.

Custodian-triggered shutdown is limited by details of process handling in the host system. For example, process and system may create an intermediate shell process to run a program, in which case custodian-based termination shuts down the shell process and probably not the process started by the shell. See also subprocess-kill. Process groups (see subprocess-group-enabled) can address some limitations, but not all of them.

```
(subprocess-group-enabled) → boolean?
(subprocess-group-enabled on?) → void?
  on? : any/c
```

A parameter that determines whether a subprocess is created as a new process group. See subprocess-kill for more information.

Beware that creating a group may interfere with the job control in an interactive shell, since job control is based on process groups.

```
(shell-execute verb
               target
               parameters
               dir
               show-mode) → #f
```

```
verb : (or/c string? #f)
target : string?
parameters : string?
dir : path-string?
show-mode : symbol?
```

Performs the action specified by *verb* on *target* in Windows. For platforms other than Windows, the `exn:fail:unsupported` exception is raised.

For example,

```
(shell-execute #f "http://racket-lang.org" ""
               (current-directory) 'sw_shownormal)
```

Opens the Racket home page in a browser window.

The *verb* can be `#f`, in which case the operating system will use a default verb. Common verbs include `"open"`, `"edit"`, `"find"`, `"explore"`, and `"print"`.

The *target* is the target for the action, usually a filename path. The file could be executable, or it could be a file with a recognized extension that can be handled by an installed application.

The *parameters* argument is passed on to the system to perform the action. For example, in the case of opening an executable, the *parameters* is used as the command line (after the executable name).

The *dir* is used as the current directory when performing the action.

The *show-mode* sets the display mode for a Window affected by the action. It must be one of the following symbols; the description of each symbol's meaning is taken from the Windows API documentation.

- `'sw_hide` or `'SW_HIDE` — Hides the window and activates another window.
- `'sw_maximize` or `'SW_MAXIMIZE` — Maximizes the window.
- `'sw_minimize` or `'SW_MINIMIZE` — Minimizes the window and activates the next top-level window in the z-order.
- `'sw_restore` or `'SW_RESTORE` — Activates and displays the window. If the window is minimized or maximized, Windows restores it to its original size and position.
- `'sw_show` or `'SW_SHOW` — Activates the window and displays it in its current size and position.
- `'sw_showdefault` or `'SW_SHOWDEFAULT` — Uses a default.

- `'sw_showmaximized` or `'SW_SHOWMAXIMIZED` — Activates the window and displays it as a maximized window.
- `'sw_showminimized` or `'SW_SHOWMINIMIZED` — Activates the window and displays it as a minimized window.
- `'sw_showminnoactive` or `'SW_SHOWMINNOACTIVE` — Displays the window as a minimized window. The active window remains active.
- `'sw_showna` or `'SW_SHOWNA` — Displays the window in its current state. The active window remains active.
- `'sw_shownoactivate` or `'SW_SHOWNOACTIVATE` — Displays a window in its most recent size and position. The active window remains active.
- `'sw_shownormal` or `'SW_SHOWNORMAL` — Activates and displays a window. If the window is minimized or maximized, Windows restores it to its original size and position.

If the action fails, the `exn:fail` exception is raised. If the action succeeds, the result is `#f`.

In future versions of Racket, the result may be a subprocess value if the operating system did returns a process handle (but if a subprocess value is returned, its process ID will be `0` instead of the real process ID).

14.4.1 Simple Subprocesses

```
(require racket/system)
```

The bindings documented in this section are provided by the `racket/system` and `racket` libraries, but not `racket/base`.

```
(system command) → boolean?
  command : (or/c string-no-nuls? bytes-no-nuls?)
```

Executes a Unix, Mac OS X, or Windows shell command synchronously (i.e., the call to `system` does not return until the subprocess has ended). The *command* argument is a string or byte string containing no nul characters. If the command succeeds, the return value is `#t`, `#f` otherwise.

See also `current-subprocess-custodian-mode` and `subprocess-group-enabled`, which affect the subprocess used to implement `system`.

The resulting process writes to `(current-output-port)`, reads from `(current-input-port)`, and logs errors to `(current-error-port)`. To gather the process's non-error output to a string, for example, use `with-output-to-string`, which sets `current-output-port` while calling the given function:

```
(with-output-to-string (lambda () (system "date")))
```

```
(system* command arg ...) → boolean?
  command : path-string?
  arg : (or/c path? string-no-nuls? bytes-no-nuls?)
(system* command exact arg) → boolean?
  command : path-string?
  exact : 'exact
  arg : string?
```

Like system, except that *command* is a filename that is executed directly (instead of through a shell command; see find-executable-path for locating an executable based on the PATH environment variable), and the *arg*s are the arguments. The executed file is passed the specified string arguments (which must contain no nul characters).

On Windows, the first argument after *command* can be 'exact, and the final *arg* is a complete command line. See subprocess for details.

```
(system/exit-code command) → byte?
  command : (or/c string-no-nuls? bytes-no-nuls?)
```

Like system, except that the result is the exit code returned by the subprocess. A 0 result normally indicates success.

```
(system*/exit-code command arg ...) → byte?
  command : path-string?
  arg : (or/c path? string-no-nuls? bytes-no-nuls?)
(system*/exit-code command exact arg) → byte?
  command : path-string?
  exact : 'exact
  arg : string?
```

Like system*, but returns the exit code like system/exit-code.

```
(process command)
    (list input-port?
          output-port?
 →        exact-nonnegative-integer?
          input-port?
          ((or/c 'status 'wait 'interrupt 'kill) . -> . any))
  command : (or/c string-no-nuls? bytes-no-nuls?)
```

Executes a shell command asynchronously (using sh on Unix and Mac OS X, cmd on Windows). The result is a list of five values:

- an input port piped from the subprocess's standard output,
- an output port piped to the subprocess standard input,
- the system process id of the subprocess,
- an input port piped from the subprocess's standard error, and
- a procedure of one argument, either `'status`, `'wait`, `'interrupt`, `'exit-code` or `'kill`:
  - `'status` returns the status of the subprocess as one of `'running`, `'done-ok`, or `'done-error`.
  - `'exit-code` returns the integer exit code of the subprocess or `#f` if it is still running.
  - `'wait` blocks execution in the current thread until the subprocess has completed.
  - `'interrupt` sends the subprocess an interrupt signal on Unix and Mac OS X, and takes no action on Windows. The result is `#<void>`.
  - `'kill` terminates the subprocess and returns `#<void>`. Note that the immediate process created by `process` is a shell process that may run another program; terminating the shell process may not terminate processes that the shell starts, particularly on Windows.

On Unix and Mac OS X, if *command* runs a single program, then `sh` typically runs the program in such a way that it replaces `sh` in the same process. For reliable and precise control over process creation, however, use `process*`.

Important: All three ports returned from `process` must be explicitly closed with `close-input-port` or `close-output-port`.

See also `current-subprocess-custodian-mode` and `subprocess-group-enabled`, which affect the subprocess used to implement `process`. In particular, the `'interrupt` and `'kill` process-control messages are implemented via `subprocess-kill`, so they can affect a process group instead of a single process.

```
(process* command arg ...) → list?
  command : path-string?
  arg : (or/c path? string-no-nuls? bytes-no-nuls?)
(process* command exact arg) → list?
  command : path-string?
  exact : 'exact
  arg : string?
```

Like `process`, except that *command* is a filename that is executed directly like `system*`, and the *arg*s are the arguments. On Windows, as for `system*`, the first *arg* can be replaced with `'exact`.

```
(process/ports out in error-out command) → list?
  out : (or/c #f output-port?)
  in : (or/c #f input-port?)
  error-out : (or/c #f output-port? 'stdout)
  command : (or/c path? string-no-nuls? bytes-no-nuls?)
```

Like process, except that *out* is used for the process's standard output, *in* is used for the process's standard input, and *error-out* is used for the process's standard error. Any of the ports can be #f, in which case a system pipe is created and returned, as in process. If *error-out* is 'stdout, then standard error is redirected to standard output. For each port or 'stdout that is provided, no pipe is created, and the corresponding value in the returned list is #f.

```
(process*/ports out
                in
                error-out
                command
                arg ...)  → list?
  out : (or/c #f output-port?)
  in : (or/c #f input-port?)
  error-out : (or/c #f output-port? 'stdout)
  command : path-string?
  arg : (or/c path? string-no-nuls? bytes-no-nuls?)
(process*/ports out
                in
                error-out
                command
                exact
                arg)      → list?
  out : (or/c #f output-port?)
  in : (or/c #f input-port?)
  error-out : (or/c #f output-port? 'stdout)
  command : path-string?
  exact : 'exact
  arg : string?
```

Like process*, but with the port handling of process/ports.

The contracts of system and related functions may signal a contract error with references to the following functions.

```
(string-no-nuls? x) → boolean?
  x : any/c
```

Ensures that *x* is a string and does not contain "\u0000".

```
(bytes-no-nuls? x) → boolean?
  x : any/c
```

Ensures that *x* is a byte-string and does not contain #"\0".

14.5 Logging

A *logger* accepts events that contain information to be logged for interested parties. A *log receiver* represents an interested party that receives logged events asynchronously. Each event has a level of importance, and a log receiver subscribes to logging events at a certain level of importance and higher. The levels, in decreasing order of importance, are `'fatal`, `'error`, `'warning`, `'info`, and `'debug`.

To help organize logged events, loggers can be named and hierarchical. Every event reported to a logger is also propagated to its parent (if any), but the event message is prefixed with the name (if any) of the logger to which is was originally reported. A logger is not required to have a parent or name.

On start-up, Racket creates an initial logger that is used to record events from the core run-time system. For example, an `'info` event is reported for each garbage collection (see §1.1.7 "Garbage Collection"). For this initial logger, two log receivers are also created: one that writes events to the process's original error output port, and one that writes events to the system log. The level of written events in each case is system-specific, and the default can be changed through command-line flags (see §17.1.4 "Command Line") or through environment variables:

- If the `PLTSTDERR` environment variable is defined and is not overridden by a command-line flag, it determines the level of the log receiver that propagates events to the original error port. The environment variable's value should be `"none"`, `"fatal"`, `"error"`, `"warning"`, `"info"`, or `"debug"`.

  The default is `"error"`.

- If the `PLTSYSLOG` environment variable is defined and is not overridden by a command-line flag, it determines the level of the log receiver that propagates events to the system log. The possible values are the same as for `PLTSTDERR`.

  The default is `"none"` for Unix or `"error"` for Windows and Mac OS X.

The `current-logger` parameter determines the *current logger* that is used by forms such as `log-warning`. On start-up, the initial value of this parameter is the initial logger. The run-time system sometimes uses the current logger to report events. For example, the byte-code compiler sometimes reports `'warning` events when it detects an expression that would produce a run-time error if evaluated.

14.5.1 Creating Loggers

```
(logger? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a logger, `#f` otherwise.

```
(make-logger [name parent]) → logger?
  name : (or/c symbol? #f) = #f
  parent : (or/c logger? #f) = #f
```

Creates a new logger with an optional name and parent.

```
(logger-name logger) → (or/c symbol? #f)
  logger : logger?
```

Reports *logger*'s name, if any.

```
(current-logger) → logger?
(current-logger logger) → void?
  logger : logger?
```

A parameter that determines the current logger.

14.5.2 Logging Events

```
(log-message logger level message data) → void?
  logger : logger?
  level : (or/c 'fatal 'error 'warning 'info 'debug)
  message : string?
  data : any/c
```

Reports an event to *logger*, which in turn distributes the information to any log receivers attached to *logger* or its ancestors that are interested in events at *level* or higher.

If *logger* has a name, then *message* is prefixed with the logger's name followed by `": "` before it is sent to receivers.

```
(log-level? logger level) → boolean?
  logger : logger?
  level : (or/c 'fatal 'error 'warning 'info 'debug)
```

Reports whether any log receiver attached to *logger* or one of its ancestors is interested in *level* events (or potentially lower). Use this function to avoid work generating an event for `log-message` if no receiver is interested in the information; this shortcut is built into `log-fatal`, `log-error`, `log-warning`, `log-info`, and `log-debug`, however, so it should not be used with those forms.

The result of this function can change if a garbage collection determines that a log receiver is no longer accessible (and therefore that any event information it receives will never become accessible).

```
(log-max-level logger)
 → (or/c #f 'fatal 'error 'warning 'info 'debug)
  logger : logger?
```

Similar to `log-level?`, but reports the maximum level of logging for which `log-level?` on *logger* returns `#t`. The result is `#f` if `log-level?` with *logger* currently returns `#f` for all levels.

```
(log-fatal string-expr)
(log-error string-expr)
(log-warning string-expr)
(log-info string-expr)
(log-debug string-expr)
```

Log an event with the current logger, evaluating *string-expr* only if the logger has receivers that are interested in the event. In addition, the current continuation's continuation marks are sent to the logger with the message string.

For each `log-`*level*,

```
(log-level string-expr)
```

is equivalent to

```
(let ([l (current-logger)])
  (when (log-level? l 'level)
    (log-message l 'level string-expr
                 (current-continuation-marks))))
```

14.5.3 Receiving Logged Events

```
(log-receiver? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a log receiver, `#f` otherwise.

```
(make-log-receiver logger level) → log-receiver?
  logger : logger?
  level : (or/c 'fatal 'error 'warning 'info 'debug)
```

Creates a log receiver to receive events of importance *level* and higher as reported to *logger* and its descendants.

A log receiver is a synchronizable event. It becomes ready as an synchronizable event when a logging event is received, so use `sync` to receive an logged event. The log receiver's synchronization value is a vector containing three values: the level of the event as a symbol, an immutable string for the event message, and an arbitrary value that was supplied as the last argument to `log-message` when the event was logged.

14.6 Time

```
(current-seconds) → exact-integer?
```

Returns the current time in seconds since midnight UTC, January 1, 1970.

```
(current-inexact-milliseconds) → real?
```

Returns the current time in milliseconds since midnight UTC, January 1, 1970. The result may contain fractions of a millisecond.

Example:

```
> (current-inexact-milliseconds)
1289513737015.418
```

In this example, `1289513737015` is in milliseconds and `418` is in microseconds.

```
(seconds->date secs-n [local-time?]) → date*?
  secs-n : real?
  local-time? : any/c = #t
```

Takes *secs-n*, a platform-specific time in seconds returned by `current-seconds`, `file-or-directory-modify-seconds`, or 1/1000th of `current-inexact-milliseconds`, and returns an instance of the `date*` structure type. Note that *secs-n* can include fractions of a second. If *secs-n* is too small or large, the `exn:fail` exception is raised.

The resulting `date*` reflects the time according to the local time zone if *local-time?* is `#t`, otherwise it reflects a date in UTC.

```
(struct date (second
              minute
              hour
              day
              month
              year
              week-day
              year-day
              dst?
              time-zone-offset)
         #:extra-constructor-name make-date
         #:transparent)
```

```
  second : (integer-in 0 60)
  minute : (integer-in 0 59)
  hour : (integer-in 0 23)
  day : (integer-in 1 31)
  month : (integer-in 1 12)
  year : exact-integer?
  week-day : (integer-in 0 6)
  year-day : (integer-in 0 365)
  dst? : boolean?
  time-zone-offset : exact-integer?
```

Represents a date. The `second` field reaches `60` only for leap seconds. The `year-day` field reaches `365` only in leap years.

The `dst?` field is `#t` if the date reflects a daylight-saving adjustment. The `time-zone-offset` field reports the number of seconds east of UTC (GMT) for the current time zone (e.g., Pacific Standard Time is `-28800`), including any daylight-saving adjustment (e.g., Pacific Daylight Time is `-25200`). When a `date` record is generated by `seconds->date` with `#f` as the second argument, then the `dst?` and `time-zone-offset` fields are `#f` and `0`, respectively.

The `date` constructor accepts any value for `dst?` and converts any non-`#f` value to `#t`.

The value produced for the `time-zone-offset` field tends to be sensitive to the value of the `TZ` environment variable, especially on Unix platforms; consult the system documentation (usually under `tzset`) for details.

See also the `racket/date` library.

```
(struct date* date (nanosecond time-zone-name)
        #:extra-constructor-name make-date*)
  nanosecond : (integer-in 0 999999999)
  time-zone-name : (and/c string? immutable?)
```

Extends `date` with nanoseconds and a time zone name, such as `"MDT"`, `"Mountain Daylight Time"`, or `"UTC"`.

When a `date*` record is generated by `seconds->date` with `#f` as the second argument, then the `time-zone-name` field is `"UTC"`.

The `date*` constructor accepts a mutable string for `time-zone-name` and converts it to an immutable one.

```
(current-milliseconds) → exact-integer?
```

Like `current-inexact-milliseconds`, but coerced to a fixnum (possibly negative).

Since the result is a fixnum, the value increases only over a limited (though reasonably long) time on a 32-bit platform.

```
(current-process-milliseconds thread) → exact-integer?
  thread : (or/c thread? #f)
```

Returns an amount of processor time in fixnum milliseconds that has been consumed by the Racket process on the underlying operating system. (On Unix and Mac OS X, this includes both user and system time.) If *thread* is `#f`, the reported time is for all Racket threads, otherwise the result is specific to the time while *thread* ran. The precision of the result is platform-specific, and since the result is a fixnum, the value increases only over a limited (though reasonably long) time on a 32-bit platform.

```
(current-gc-milliseconds) → exact-integer?
```

Returns the amount of processor time in fixnum milliseconds that has been consumed by Racket's garbage collection so far. This time is a portion of the time reported by `(current-process-milliseconds)`, and is similarly limited.

```
                          list?
(time-apply proc lst) →   exact-integer?
                          exact-integer?
                          exact-integer?
  proc : procedure?
  lst : list?
```

Collects timing information for a procedure application.

Four values are returned: a list containing the result(s) of applying *proc* to the arguments in *lst*, the number of milliseconds of CPU time required to obtain this result, the number of "real" milliseconds required for the result, and the number of milliseconds of CPU time (included in the first result) spent on garbage collection.

The reliability of the timing numbers depends on the platform. If multiple Racket threads are running, then the reported time may include work performed by other threads.

```
(time expr)
```

Reports `time-apply`-style timing information for the evaluation of *expr* directly to the current output port. The result is the result of *expr*.

14.6.1 Date Utilities

```
(require racket/date)
```

The bindings documented in this section are provided by the racket/date library, not racket/base or racket.

```
(current-date) → date*?
```

An abbreviation for (seconds->date (current-seconds)).

```
(date->string date [time?]) → string?
  date : date?
  time? : any/c = #f
```

Converts a date to a string. The returned string contains the time of day only if *time?*. See also date-display-format.

```
                          (or/c 'american
                                'chinese
                                'german
(date-display-format) →         'indian
                                'irish
                                'iso-8601
                                'rfc2822
                                'julian)
(date-display-format format) → void?
           (or/c 'american
                 'chinese
                 'german
  format :       'indian
                 'irish
                 'iso-8601
                 'rfc2822
                 'julian)
```

Parameter that determines the date string format. The initial format is 'american.

```
(date->seconds date [local-time?]) → exact-integer?
  date : date?
  local-time? : any/c = #t
```

Finds the representation of a date in platform-specific seconds. The time-zone-offset field of *date* is ignored; the date is assumed to be in local time by default or in UTC if *local-time?* is #f. If the platform cannot represent the specified date, an error is signaled, otherwise an integer is returned.

```
(find-seconds  second
               minute
               hour
               day
               month
               year
              [local-time?]) → exact-integer?
  second : (integer-in 0 61)
  minute : (integer-in 0 59)
  hour : (integer-in 0 23)
  day : (integer-in 1 31)
  month : (integer-in 1 12)
  year : exact-nonnegative-integer?
  local-time? : any/c = #t
```

Finds the representation of a date in platform-specific seconds. The arguments correspond to the fields of the `date` structure—in local time by default or UTC if *local-time?* is `#f`. If the platform cannot represent the specified date, an error is signaled, otherwise an integer is returned.

```
(date->julian/scalinger date) → exact-integer?
  date : date?
```

Converts a date structure (up to 2099 BCE Gregorian) into a Julian date number. The returned value is not a strict Julian number, but rather Scalinger's version, which is off by one for easier calculations.

```
(julian/scalinger->string date-number) → string?
  date-number : exact-integer?
```

Converts a Julian number (Scalinger's off-by-one version) into a string.

14.7 Environment and Runtime Information

```
(getenv name) → (or/c string? #f)
  name : string?
```

Gets the value of an operating system environment variable. The *name* argument cannot contain a null character; if an environment variable named by *name* exists, its value is returned (as a string); otherwise, `#f` is returned.

```
(putenv name value) → boolean?
  name : string?
  value : string?
```

Sets the value of an operating system environment variable. The *name* and *value* arguments are strings that cannot contain a null character; the environment variable named by *name* is set to *value*. The return value is `#t` if the assignment succeeds, `#f` otherwise.

```
(system-type [mode]) → (or/c symbol? string? bytes?)
  mode : (or/c 'os 'gc 'link 'so-suffix 'machine) = 'os
```

Returns information about the operating system, build mode, or machine for a running Racket.

In `'os` mode, the possible symbol results are:

- `'unix`
- `'windows`
- `'macosx`

In `'gc` mode, the possible symbol results are:

- `'cgc`
- `'3m`

In `'link` mode, the possible symbol results are:

- `'static` (Unix)
- `'shared` (Unix)
- `'dll` (Windows)
- `'framework` (Mac OS X)

Future ports of Racket may expand the list of `'os`, `'gc`, and `'link` results.

In `'so-suffix` mode, then the result is a byte string that represents the file extension used for shared objects on the current platform. The byte string starts with a period, so it is suitable as a second argument to `path-replace-suffix`.

In `'machine` mode, then the result is a string, which contains further details about the current machine in a platform-specific format.

```
(system-language+country) → string?
```

Returns a string to identify the current user's language and country.

On Unix and Mac OS X, the string is five characters: two lowercase ASCII letters for the language, an underscore, and two uppercase ASCII letters for the country. On Windows, the string can be arbitrarily long, but the language and country are in English (all ASCII letters or spaces) separated by an underscore.

On Unix, the result is determined by checking the `LC_ALL`, `LC_TYPE`, and `LANG` environment variables, in that order (and the result is used if the environment variable's value starts with two lowercase ASCII letters, an underscore, and two uppercase ASCII letters, followed by either nothing or a period). On Windows and Mac OS X, the result is determined by system calls.

```
(system-library-subpath [mode]) → path?
  mode : (or/c 'cgc '3m #f) = (system-type 'gc)
```

Returns a relative directory path. This string can be used to build paths to system-specific files. For example, when Racket is running on Solaris on a Sparc architecture, the subpath starts `"sparc-solaris"`, while the subpath for Windows on an i386 architecture starts `"win32\\i386"`.

The optional *mode* argument specifies the relevant garbage-collection variant, which one of the possible results of `(system-type 'gc)`: `'cgc` or `'3m`. It can also be `#f`, in which case the result is independent of the garbage-collection variant.

```
(version) → (and/c string? immutable?)
```

Returns an string indicating the currently executing version of Racket.

```
(banner) → (and/c string? immutable?)
```

Returns an immutable string for Racket's start-up banner text (or the banner text for an embedding program, such as GRacket). The banner string ends with a newline.

```
(current-command-line-arguments)
 → (vectorof (and/c string? immutable?))
(current-command-line-arguments argv) → void?
  argv : (vectorof (and/c string? immutable?))
```

A parameter that is initialized with command-line arguments when Racket starts (not including any command-line arguments that were treated as flags for the system).

```
(current-thread-initial-stack-size) → exact-positive-integer?
(current-thread-initial-stack-size size) → void?
  size : exact-positive-integer?
```

A parameter that provides a hint about how much space to reserve for a newly created thread's local variables. The actual space used by a computation is affected by JIT compilation, but it is otherwise platform-independent.

```
(vector-set-performance-stats! results [thd]) → void?
  results : (and/c vector?
                   (not/c immutable?))
  thd : (or/c thread? #f) = #f
```

Sets elements in *results* to report current performance statistics. If *thd* is not `#f`, a particular set of thread-specific statistics are reported, otherwise a different set of global statics are reported.

For global statistics, up to 11 elements are set in the vector, starting from the beginning. If *results* has n elements where $n < 11$, then the n elements are set to the first n performance-statistics values. The reported statistics values are as follows, in the order that they are set within *results*:

- `0`: The same value as returned by `current-process-milliseconds`.
- `1`: The same value as returned by `current-milliseconds`.
- `2`: The same value as returned by `current-gc-milliseconds`.
- `3`: The number of garbage collections performed since start-up.
- `4`: The number of thread context switches performed since start-up.
- `5`: The number of internal stack overflows handled since start-up.
- `6`: The number of threads currently scheduled for execution (i.e., threads that are running, not suspended, and not unscheduled due to a synchronization).
- `7`: The number of syntax objects read from compiled code since start-up.
- `8`: The number of hash-table searches performed. When this counter reaches the maximum value of a fixnum, it overflows to the most negative fixnum.
- `9`: The number of additional hash slots searched to complete hash searches (using double hashing). When this counter reaches the maximum value of a fixnum, it overflows to the most negative fixnum.
- `10`: The number of bytes allocated for machine code that is not reported by `current-memory-use`.

For thread-specific statistics, up to 4 elements are set in the vector:

- `0`: `#t` if the thread is running, `#f` otherwise (same result as `thread-running?`).

- 1: #t if the thread has terminated, #f otherwise (same result as thread-dead?).
- 2: #t if the thread is currently blocked on a synchronizable event (or sleeping for some number of milliseconds), #f otherwise.
- 3: The number of bytes currently in use for the thread's continuation.

14.8 Command-Line Parsing

```
(require racket/cmdline)
```

The bindings documented in this section are provided by the racket/cmdline and racket libraries, but not racket/base.

```
(command-line optional-name-expr optional-argv-expr
              flag-clause ...
              finish-clause)
```

```
optional-name-expr =
                   | #:program name-expr

optional-argv-expr =
                   | #:argv argv-expr

      flag-clause = #:multi flag-spec ...
                  | #:once-each flag-spec ...
                  | #:once-any flag-spec ...
                  | #:final flag-spec ...
                  | #:help-labels string ...
                  | #:ps string ...

        flag-spec = (flags id ... help-spec body ...+)
                  | (flags => handler-expr help-expr)

            flags = flag-string
                  | (flag-string ...+)

        help-spec = string
                  | (string-expr ...+)

    finish-clause =
                  | #:args arg-formals body ...+
                  | #:handlers handlers-exprs

      arg-formals = rest-id
                  | (arg ...)
                  | (arg ...+ . rest-id)

              arg = id
                  | [id default-expr]

   handlers-exprs = finish-expr arg-strings-expr
                  | finish-expr arg-strings-expr help-expr
                  | finish-expr arg-strings-expr help-expr
                    unknown-expr
```

Parses a command line according to the specification in the *flag-clause*s.

The *name-expr*, if provided, should produce a path or string to be used as the program name for reporting errors when the command-line is ill-formed. It defaults to `(find-system-path 'run-file)`. When a path is provided, only the last element of the path is used to report an error.

The *argv-expr*, if provided, must evaluate to a list or a vector of strings. It defaults to

`(current-command-line-arguments)`.

The command-line is disassembled into flags, each possibly with flag-specific arguments, followed by (non-flag) arguments. Command-line strings starting with `-` or `+` are parsed as flags, but arguments to flags are never parsed as flags, and integers and decimal numbers that start with `-` or `+` are not treated as flags. Non-flag arguments in the command-line must appear after all flags and the flags' arguments. No command-line string past the first non-flag argument is parsed as a flag. The built-in `--` flag signals the end of command-line flags; any command-line string past the `--` flag is parsed as a non-flag argument.

A `#:multi`, `#:once-each`, `#:once-any`, or `#:final` clause introduces a set of command-line flag specifications. The clause tag indicates how many times the flag can appear on the command line:

- `#:multi` — Each flag specified in the set can be represented any number of times on the command line; i.e., the flags in the set are independent and each flag can be used multiple times.
- `#:once-each` — Each flag specified in the set can be represented once on the command line; i.e., the flags in the set are independent, but each flag should be specified at most once. If a flag specification is represented in the command line more than once, the `exn:fail` exception is raised.
- `#:once-any` — Only one flag specified in the set can be represented on the command line; i.e., the flags in the set are mutually exclusive. If the set is represented in the command line more than once, the `exn:fail` exception is raised.
- `#:final` — Like `#:multi`, except that no argument after the flag is treated as a flag. Note that multiple `#:final` flags can be specified if they have short names; for example, if `-a` is a `#:final` flag, then `-aa` combines two instances of `-a` in a single command-line argument.

A normal flag specification has four parts:

- *`flags`* — a flag string, or a set of flag strings. If a set of flags is provided, all of the flags are equivalent. Each flag string must be of the form `"-`*`x`*`"` or `"+`*`x`*`"` for some character *`x`*, or `"-`*`x`*`"` or `"++`*`x`*`"` for some sequence of characters *`x`*. An *`x`* cannot contain only digits or digits plus a single decimal point, since simple (signed) numbers are not treated as flags. In addition, the flags `"--"`, `"-h"`, and `"--help"` are predefined and cannot be changed.
- *`id`*s — identifier that are bound to the flag's arguments. The number of identifiers determines how many arguments can be provided on the command line with the flag, and the names of these identifiers will appear in the help message describing the flag. The *`id`*s are bound to string values in the *`body`*s for handling the flag.

- *help-spec* — a string or sequence of strings that describes the flag. This string is used in the help message generated by the handler for the built-in `-h` (or `--help`) flag. A single literal string can be provided, or any number of expressions that produce strings; in the latter case, strings after the first one are displayed on subsequent lines.

- *body*s — expressions that are evaluated when one of the *flags* appears on the command line. The flags are parsed left-to-right, and each sequence of *body*s is evaluated as the corresponding flag is encountered. When the *body*s are evaluated, the preceding *id*s are bound to the arguments provided for the flag on the command line.

A flag specification using `=>` escapes to a more general method of specifying the handler and help strings. In this case, the handler procedure and help string list returned by *handler-expr* and *help-expr* are used as in the *table* argument of `parse-command-line`.

A `#:help-labels` clause inserts text lines into the help table of command-line flags. Each string in the clause provides a separate line of text.

A `#:ps` clause inserts text lines at the end of the help output. Each string in the clause provides a separate line of text.

After the flag clauses, a final clause handles command-line arguments that are not parsed as flags:

- Supplying no finish clause is the same as supplying `#:args () (void)`.

- For an `#:args` finish clause, identifiers in *arg-formals* are bound to the leftover command-line strings in the same way that identifiers are bound for a `lambda` expression. Thus, specifying a single *id* (without parentheses) collects all of the leftover arguments into a list. The effective arity of the *arg-formals* specification determines the number of extra command-line arguments that the user can provide, and the names of the identifiers in *arg-formals* are used in the help string. When the command-line is parsed, if the number of provided arguments cannot be matched to identifiers in *arg-formals*, the `exn:fail` exception is raised. Otherwise, `args` clause's *body*s are evaluated to handle the leftover arguments, and the result of the last *body* is the result of the `command-line` expression.

- A `#:handlers` finish clause escapes to a more general method of handling the leftover arguments. In this case, the values of the expressions are used like the last two to four arguments `parse-command-line`.

Example:

```
(define verbose-mode (make-parameter #f))
(define profiling-on (make-parameter #f))
(define optimize-level (make-parameter 0))
(define link-flags (make-parameter null))
```

```
(define file-to-compile
  (command-line
   #:program "compiler"
   #:once-each
   [("-v" "--verbose") "Compile with verbose messages"
                       (verbose-mode #t)]
   [("-p" "--profile") "Compile with profiling"
                       (profiling-on #t)]
   #:once-any
   [("-o" "--optimize-1") "Compile with optimization level 1"
                          (optimize-level 1)]
   ["--optimize-2"        (; show help on separate lines
                           "Compile with optimization level 2,"
                           "which includes all of level 1")
                          (optimize-level 2)]
   #:multi
   [("-l" "--link-flags") lf ; flag takes one argument
                          "Add a flag <lf> for the linker"
                          (link-flags (cons lf (link-flags)))]
   #:args (filename) ; expect one command-line argument: <file-
name>
   ; return the argument as a filename to compile
   filename))
```

```
(parse-command-line  name
                     argv
                     table
                     finish-proc
                     arg-help-strs
                    [help-proc
                     unknown-proc]) → any
  name : (or/c string? path?)
  argv : (or/c (listof string?) (vectorof string?))
  table : (listof (cons/c symbol? list?))
  finish-proc : ((list?) () #:rest list? . ->* . any)
  arg-help-strs : (listof string?)
  help-proc : (string? . -> . any) = (lambda (str) ....)
  unknown-proc : (string? . -> . any) = (lambda (str) ...)
```

Parses a command-line using the specification in *table*. For an overview of command-line parsing, see the `command-line` form, which provides a more convenient notation for most purposes.

The *table* argument to this procedural form encodes the information in `command-line`'s clauses, except for the `args` clause. Instead, arguments are handled by the *finish-proc*

procedure, and help information about non-flag arguments is provided in *arg-help-strs*. In addition, the *finish-proc* procedure receives information accumulated while parsing flags. The *help-proc* and *unknown-proc* arguments allow customization that is not possible with `command-line`.

When there are no more flags, *finish-proc* is called with a list of information accumulated for command-line flags (see below) and the remaining non-flag arguments from the command-line. The arity of *finish-proc* determines the number of non-flag arguments accepted and required from the command-line. For example, if *finish-proc* accepts either two or three arguments, then either one or two non-flag arguments must be provided on the command-line. The *finish-proc* procedure can have any arity (see `procedure-arity`) except `0` or a list of `0`s (i.e., the procedure must at least accept one or more arguments).

The *arg-help-strs* argument is a list of strings identifying the expected (non-flag) command-line arguments, one for each argument. If an arbitrary number of arguments are allowed, the last string in *arg-help-strs* represents all of them.

The *help-proc* procedure is called with a help string if the `-h` or `--help` flag is included on the command line. If an unknown flag is encountered, the *unknown-proc* procedure is called just like a flag-handling procedure (as described below); it must at least accept one argument (the unknown flag), but it may also accept more arguments. The default *help-proc* displays the string and exits and the default *unknown-proc* raises the `exn:fail` exception.

A *table* is a list of flag specification sets. Each set is represented as a pair of two items: a mode symbol and a list of either help strings or flag specifications. A mode symbol is one of `'once-each`, `'once-any`, `'multi`, `'final`, `'help-labels`, or `'ps` with the same meanings as the corresponding clause tags in `command-line`. For the `'help-labels` or `'ps` mode, a list of help string is provided. For the other modes, a list of flag specifications is provided, where each specification maps a number of flags to a single handler procedure. A specification is a list of three items:

- A list of strings for the flags defined by the spec. See `command-line` for information about the format of flag strings.
- A procedure to handle the flag and its arguments when one of the flags is found on the command line. The arity of this handler procedure determines the number of arguments consumed by the flag: the handler procedure is called with a flag string plus the next few arguments from the command line to match the arity of the handler procedure. The handler procedure must accept at least one argument to receive the flag. If the handler accepts arbitrarily many arguments, all of the remaining arguments are passed to the handler. A handler procedure's arity must either be a number or an `arity-at-least` value.

  The return value from the handler is added to a list that is eventually passed to *finish-proc*. If the handler returns `#<void>`, no value is added onto this list. For all non-`#<void>` values returned by handlers, the order of the values in the list is the same as the order of the arguments on the command-line.

- A non-empty list for constructing help information for the spec. The first element of the list describes the flag; it can be a string or a non-empty list of strings, and in the latter case, each string is shown on its own line. Additional elements of the main list must be strings to name the expected arguments for the flag. The number of extra help strings provided for a spec must match the number of arguments accepted by the spec's handler procedure.

The following example is the same as the core example for `command-line`, translated to the procedural form:

```
(parse-command-line "compile" (current-command-line-arguments)
  `((once-each
     [("-v" "--verbose")
      ,(lambda (flag) (verbose-mode #t))
      ("Compile with verbose messages")]
     [("-p" "--profile")
      ,(lambda (flag) (profiling-on #t))
      ("Compile with profiling")])
    (once-any
     [("-o" "--optimize-1")
      ,(lambda (flag) (optimize-level 1))
      ("Compile with optimization level 1")]
     [("--optimize-2")
      ,(lambda (flag) (optimize-level 2))
      (("Compile with optimization level 2,"
        "which implies all optimizations of level 1"))])
    (multi
     [("-l" "--link-flags")
      ,(lambda (flag lf) (link-flags (cons lf (link-flags))))
      ("Add a flag <lf> for the linker" "lf")]))
   (lambda (flag-accum file) file)
   '("filename"))
```

15 Memory Management

15.1 Weak Boxes

A *weak box* is similar to a normal box (see §3.12 “Boxes”), but when the garbage collector (see §1.1.7 “Garbage Collection”) can prove that the content value of a weak box is only reachable via weak references, the content of the weak box is replaced with `#f`. A *weak reference* is a reference through a weak box, through a key reference in a weak hash table (see §3.13 “Hash Tables”), through a value in an ephemeron where the value can be replaced by `#f` (see §15.2 “Ephemerons”), or through a custodian (see §13.7 “Custodians”).

```
(make-weak-box v) → weak-box?
  v : any/c
```

Returns a new weak box that initially contains `v`.

```
(weak-box-value weak-box [gced-v]) → any/c
  weak-box : weak-box?
  gced-v : any/c = #f
```

Returns the value contained in *weak-box*. If the garbage collector has proven that the previous content value of *weak-box* was reachable only through a weak reference, then *gced-v* (which defaults to `#f`) is returned.

```
(weak-box? v) → boolean?
  v : any/c
```

Returns `#t` if `v` is a weak box, `#f` otherwise.

15.2 Ephemerons

An *ephemeron* [Hayes97] is a generalization of a weak box (see §15.1 “Weak Boxes”). Instead of just containing one value, an ephemeron holds two values: one that is considered the value of the ephemeron and another that is the ephemeron’s key. Like the value in a weak box, the value in an ephemeron may be replaced by `#f`, but when the *key* is no longer reachable (except possibly via weak references) instead of when the value is no longer reachable.

As long as an ephemeron’s value is retained, the reference is considered a non-weak reference. References to the key via the value are treated specially, however, in that the reference does not necessarily count toward the key’s reachability. A weak box can be seen as a specialization of an ephemeron where the key and value are the same.

One particularly common use of ephemerons is to combine them with a weak hash table (see §3.13 “Hash Tables”) to produce a mapping where the memory manager can reclaim key–value pairs even when the value refers to the key.

More precisely,

- the value in an ephemeron is replaced by `#f` when the automatic memory manager can prove that either the ephemeron or the key is reachable only through weak references (see §15.1 “Weak Boxes”); and
- nothing reachable from the value in an ephemeron counts toward the reachability of an ephemeron key (whether for the same ephemeron or another), unless the same value is reachable through a non-weak reference, or unless the value’s ephemeron key is reachable through a non-weak reference (see §15.1 “Weak Boxes” for information on weak references).

```
(make-ephemeron key v) → ephemeron?
  key : any/c
  v : any/c
```

Returns a new ephemeron whose key is *key* and whose value is initially *v*.

```
(ephemeron-value ephemeron [gced-v]) → any/c
  ephemeron : ephemeron?
  gced-v : any/c = #f
```

Returns the value contained in *ephemeron*. If the garbage collector has proven that the key for *ephemeron* is only weakly reachable, then the result is *gced-v* (which defaults to `#f`).

```
(ephemeron? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is an ephemeron, `#f` otherwise.

15.3 Wills and Executors

A *will executor* manages a collection of values and associated *will* procedures (a.k.a. *finalizers*). The will procedure for each value is ready to be executed when the value has been proven (by the garbage collector) to be unreachable, except through weak references (see §15.1 “Weak Boxes”) or as the registrant for other will executors. A will is useful for triggering clean-up actions on data associated with an unreachable value, such as closing a port embedded in an object when the object is no longer used.

Calling the `will-execute` or `will-try-execute` procedure executes a will that is ready in the specified will executor. A will executor is also a synchronizable event, so `sync` or `sync/timeout` can be used to detect when a will executor has ready wills. Wills are not executed automatically, because certain programs need control to avoid race conditions. However, a program can create a thread whose sole job is to execute wills for a particular executor.

If a value is registered with multiple wills (in one or multiple executors), the wills are readied in the reverse order of registration. Since readying a will procedure makes the value reachable again, the will must be executed and the value must be proven again unreachable through only weak references before another of the wills is readied or executed. However, wills for distinct unreachable values are readied at the same time, regardless of whether the values are reachable from each other.

A will executor's register is held non-weakly until after the corresponding will procedure is executed. Thus, if the content value of a weak box (see §15.1 "Weak Boxes") is registered with a will executor, the weak box's content is not changed to `#f` until all wills have been executed for the value and the value has been proven again reachable through only weak references.

```
(make-will-executor) → will-executor?
```

Returns a new will executor with no managed values.

```
(will-executor? v) → boolean?
  v : any/c
```

Returns `#t` if *v* is a will executor, `#f` otherwise.

```
(will-register executor v proc) → void?
  executor : will-executor?
  v : any/c
  proc : (any/c . -> . any)
```

Registers the value *v* with the will procedure *proc* in the will executor *executor*. When *v* is proven unreachable, then the procedure *proc* is ready to be called with *v* as its argument via `will-execute` or `will-try-execute`. The *proc* argument is strongly referenced until the will procedure is executed.

```
(will-execute executor) → any
  executor : will-executor?
```

Invokes the will procedure for a single "unreachable" value registered with the executor *executor*. The values returned by the will procedure are the result of the `will-execute` call. If no will is ready for immediate execution, `will-execute` blocks until one is ready.

```
(will-try-execute executor) → any
  executor : any/c
```

Like `will-execute` if a will is ready for immediate execution. Otherwise, `#f` is returned.

15.4 Garbage Collection

Set the `PLTDISABLEGC` environment variable (to any value) before Racket starts to disable garbage collection.

In Racket 3m (the main variant of Racket), each garbage collection logs a message (see §14.5 "Logging") at the `'debug` level. The data portion of the message is an instance of a `gc-info` prefab structure type with 10 fields as follows, but future versions of Racket may use a `gc-info` prefab structure with additional fields:

```
(struct gc-info (major? pre-amount pre-admin-amount code-amount
                        post-amount post-admin-amount
                        start-process-time end-process-time
                        start-time end-time)
  #:prefab)
```

The `major?` field indicates whether the collection was a "major" collection that inspects all memory or a "minor" collection that mostly inspects just recent allocations. The `pre-amount` field reports place-local memory use (i.e., not counting the memory use of child places) in bytes at the time that the garbage collection started. The `pre-admin-amount` is a larger number that includes memory use for the garbage collector's overhead (such as space on memory pages that is not yet used). The `code-amount` field reports additional memory use for generated native code (which is the same just before and after a garbage collection, since it is released via finalization). The `post-amount` and `post-admin-amount` fields correspond to `pre-amount` and `pre-admin-amount`, but after garbage collection. The `start-process-time` and `end-process-time` fields report processor time (in the sense of `current-process-milliseconds`) at the start and end of garbage collection; the difference is the processor time consumed by collection. The `start-time` and `end-time` fields report real time (in the sense of `current-inexact-milliseconds`) at the start and end of garbage collection; the difference is the real time consumed by garbage collection.

```
(collect-garbage) → void?
```

Forces an immediate garbage collection (unless garbage collection is disabled by setting `PLTDISABLEGC`). Some effectively unreachable data may remain uncollected, because the collector cannot prove that it is unreachable.

The collect-garbage procedure provides some control over the timing of collections, but garbage will obviously be collected even if this procedure is never called (unless garbage collection is disabled).

```
(current-memory-use [cust]) → exact-nonnegative-integer?
  cust : custodian? = #f
```

Returns an estimate of the number of bytes of memory occupied by reachable data from *cust*. This estimate is calculated by the last garbage collection, and can be 0 if none occurred (or if none occurred since the given custodian was created). The current-memory-use function does *not* perform a collection by itself; doing one before the call will generally decrease the result (or increase it from 0 if no collections happened yet).

If *cust* is not provided, the estimate is a total reachable from any custodians.

When Racket is compiled without support for memory accounting, the estimate is the same (i.e., all memory) for any individual custodian; see also custodian-memory-accounting-available?.

```
(dump-memory-stats v ...) → any
  v : any/c
```

Dumps information about memory usage to the low-level error port or console.

Various combinations of *v* arguments can control the information in a dump. The information that is available depends on your Racket build; check the end of a dump from a particular build to see if it offers additional information; otherwise, all *v*s are ignored.

16 Unsafe Operations

```
(require racket/unsafe/ops)
```

All functions and forms provided by `racket/base` and `racket` check their arguments to ensure that the arguments conform to contracts and other constraints. For example, `vector-ref` checks its arguments to ensure that the first argument is a vector, that the second argument is an exact integer, and that the second argument is between `0` and one less than the vector's length, inclusive.

Functions provided by `racket/unsafe/ops` are *unsafe*. They have certain constraints, but the constraints are not checked, which allows the system to generate and execute faster code. If arguments violate an unsafe function's constraints, the function's behavior and result is unpredictable, and the entire system can crash or become corrupted.

All of the exported bindings of `racket/unsafe/ops` are protected in the sense of `protect-out`, so access to unsafe operations can be prevented by adjusting the code inspector (see §13.10 "Code Inspectors").

16.1 Unsafe Numeric Operations

```
(unsafe-fx+ a b) → fixnum?
  a : fixnum?
  b : fixnum?
(unsafe-fx- a b) → fixnum?
  a : fixnum?
  b : fixnum?
(unsafe-fx* a b) → fixnum?
  a : fixnum?
  b : fixnum?
(unsafe-fxquotient a b) → fixnum?
  a : fixnum?
  b : fixnum?
(unsafe-fxremainder a b) → fixnum?
  a : fixnum?
  b : fixnum?
(unsafe-fxmodulo a b) → fixnum?
  a : fixnum?
  b : fixnum?
(unsafe-fxabs a) → fixnum?
  a : fixnum?
```

For fixnums: Like `+`, `-`, `*`, `quotient`, `remainder`, `modulo`, and `abs`, but constrained to consume fixnums and produce a fixnum result. The mathematical operation on `a` and `b` must

be representable as a fixnum. In the case of unsafe-fxquotient, unsafe-fxremainder, and unsafe-fxmodulo, b must not be 0.

```
(unsafe-fxand a b) → fixnum?
  a : fixnum?
  b : fixnum?
(unsafe-fxior a b) → fixnum?
  a : fixnum?
  b : fixnum?
(unsafe-fxxor a b) → fixnum?
  a : fixnum?
  b : fixnum?
(unsafe-fxnot a) → fixnum?
  a : fixnum?
(unsafe-fxlshift a b) → fixnum?
  a : fixnum?
  b : fixnum?
(unsafe-fxrshift a b) → fixnum?
  a : fixnum?
  b : fixnum?
```

For fixnums: Like bitwise-and, bitwise-ior, bitwise-xor, bitwise-not, and arithmetic-shift, but constrained to consume fixnums; the result is always a fixnum. The unsafe-fxlshift and unsafe-fxrshift operations correspond to arithmetic-shift, but require non-negative arguments; unsafe-fxlshift is a positive (i.e., left) shift, and unsafe-fxrshift is a negative (i.e., right) shift, where the number of bits to shift must be less than the number of bits used to represent a fixnum. In the case of unsafe-fxlshift, bits in the result beyond the number of bits used to represent a fixnum are effectively replaced with a copy of the high bit.

```
(unsafe-fx= a b) → boolean?
  a : fixnum?
  b : fixnum?
(unsafe-fx< a b) → boolean?
  a : fixnum?
  b : fixnum?
(unsafe-fx> a b) → boolean?
  a : fixnum?
  b : fixnum?
(unsafe-fx<= a b) → boolean?
  a : fixnum?
  b : fixnum?
(unsafe-fx>= a b) → boolean?
  a : fixnum?
  b : fixnum?
```

```
(unsafe-fxmin a b) → fixnum?
  a : fixnum?
  b : fixnum?
(unsafe-fxmax a b) → fixnum?
  a : fixnum?
  b : fixnum?
```

For fixnums: Like `=`, `<`, `>`, `<=`, `>=`, `min`, and `max`, but constrained to consume fixnums.

```
(unsafe-fl+ a b) → flonum?
  a : flonum?
  b : flonum?
(unsafe-fl- a b) → flonum?
  a : flonum?
  b : flonum?
(unsafe-fl* a b) → flonum?
  a : flonum?
  b : flonum?
(unsafe-fl/ a b) → flonum?
  a : flonum?
  b : flonum?
(unsafe-flabs a) → flonum?
  a : flonum?
```

For flonums: Unchecked versions of `fl+`, `fl-`, `fl*`, `fl/`, and `flabs`.

```
(unsafe-fl= a b) → boolean?
  a : flonum?
  b : flonum?
(unsafe-fl< a b) → boolean?
  a : flonum?
  b : flonum?
(unsafe-fl> a b) → boolean?
  a : flonum?
  b : flonum?
(unsafe-fl<= a b) → boolean?
  a : flonum?
  b : flonum?
(unsafe-fl>= a b) → boolean?
  a : flonum?
  b : flonum?
(unsafe-flmin a b) → flonum?
  a : flonum?
  b : flonum?
(unsafe-flmax a b) → flonum?
  a : flonum?
  b : flonum?
```

For flonums: Unchecked versions of `fl=`, `fl<`, `fl>`, `fl<=`, `fl>=`, `flmin`, and `flmax`.

```
(unsafe-flround a) → flonum?
  a : flonum?
(unsafe-flfloor a) → flonum?
  a : flonum?
(unsafe-flceiling a) → flonum?
  a : flonum?
(unsafe-fltruncate a) → flonum?
  a : flonum?
```

For flonums: Unchecked (potentially) versions of `flround`, `flfloor`, `flceiling`, and `fltruncate`. Currently, these bindings are simply aliases for the corresponding safe bindings.

```
(unsafe-flsin a) → flonum?
  a : flonum?
(unsafe-flcos a) → flonum?
  a : flonum?
(unsafe-fltan a) → flonum?
  a : flonum?
(unsafe-flasin a) → flonum?
  a : flonum?
(unsafe-flacos a) → flonum?
  a : flonum?
(unsafe-flatan a) → flonum?
  a : flonum?
(unsafe-fllog a) → flonum?
  a : flonum?
(unsafe-flexp a) → flonum?
  a : flonum?
(unsafe-flsqrt a) → flonum?
  a : flonum?
(unsafe-flexpt a b) → flonum?
  a : flonum?
  b : flonum?
```

For flonums: Unchecked (potentially) versions of `flsin`, `flcos`, `fltan`, `flasin`, `flacos`, `flatan`, `fllog`, `flexp`, `flsqrt`, and `flexpt`. Currently, some of these bindings are simply aliases for the corresponding safe bindings.

```
(unsafe-make-flrectangular a b)
 → (and/c complex? inexact? (not/c real?))
  a : flonum?
  b : flonum?
(unsafe-flreal-part a) → flonum?
  a : (and/c complex? inexact? (not/c real?))
```

```
(unsafe-flimag-part a) → flonum?
  a : (and/c complex? inexact? (not/c real?))
```

For flonums: Unchecked versions of `make-flrectangular`, `flreal-part`, and `flimag-part`.

```
(unsafe-fx->fl a) → flonum?
  a : fixnum?
(unsafe-fl->fx a) → fixnum?
  a : flonum?
```

Unchecked conversion of a fixnum to an integer flonum and vice versa. These are similar to the safe bindings `->fl` and `fl->exact-integer`, but further constrained to consume or produce a fixnum.

16.2 Unsafe Data Extraction

```
(unsafe-car p) → any/c
  p : pair?
(unsafe-cdr p) → any/c
  p : pair?
(unsafe-mcar p) → any/c
  p : mpair?
(unsafe-mcdr p) → any/c
  p : mpair?
(unsafe-set-mcar! p v) → void?
  p : mpair?
  v : any/c
(unsafe-set-mcdr! p v) → void?
  p : mpair?
  v : any/c
```

Unsafe variants of `car`, `cdr`, `mcar`, `mcdr`, `set-mcar!`, and `set-mcdr!`.

```
(unsafe-list-ref lst pos) → any/c
  lst : pair?
  pos : (and/c exact-nonnegative-integer? fixnum?)
(unsafe-list-tail lst pos) → any/c
  lst : any/c
  pos : (and/c exact-nonnegative-integer? fixnum?)
```

Unsafe variants of `list-ref` and `list-tail`, where *pos* must be a fixnum, and *lst* must start with at least `(add1 pos)` (for `unsafe-list-ref`) or *pos* (for `unsafe-list-tail`) pairs.

```
(unsafe-unbox b) → fixnum?
  b : box?
(unsafe-set-box! b k) → void?
  b : box?
  k : fixnum?
(unsafe-unbox* v) → any/c
  v : (and/c box? (not/c impersonator?))
(unsafe-set-box*! v val) → void?
  v : (and/c box? (not/c impersonator?))
  val : any/c
```

Unsafe versions of `unbox` and `set-box!`, where the `box*` variants can be faster but do not work on impersonators.

```
(unsafe-box*-cas! loc old new) → boolean?
  loc : box?
  old : any/c
  new : any/c
```

Unsafe version of `box-cas!`. Like `unsafe-set-box*!`, it does not work on impersonators.

```
(unsafe-vector-length v) → fixnum?
  v : vector?
(unsafe-vector-ref v k) → any/c
  v : vector?
  k : fixnum?
(unsafe-vector-set! v k val) → void?
  v : vector?
  k : fixnum?
  val : any/c
(unsafe-vector*-length v) → fixnum?
  v : (and/c vector? (not/c impersonator?))
(unsafe-vector*-ref v k) → any/c
  v : (and/c vector? (not/c impersonator?))
  k : fixnum?
(unsafe-vector*-set! v k val) → void?
  v : (and/c vector? (not/c impersonator?))
  k : fixnum?
  val : any/c
```

Unsafe versions of `vector-length`, `vector-ref`, and `vector-set!`, where the `vector*` variants can be faster but do not work on impersonators.

A vector's size can never be larger than a fixnum, so even `vector-length` always returns a fixnum.

```
(unsafe-string-length str) → fixnum?
  str : string?
(unsafe-string-ref str k)
 → (and/c char? (lambda (ch) (<= 0 (char->integer ch) 255)))
  str : string?
  k : fixnum?
(unsafe-string-set! str k ch) → void?
  str : (and/c string? (not/c immutable?))
  k : fixnum?
  ch : char?
```

Unsafe versions of `string-length`, `string-ref`, and `string-set!`. The `unsafe-string-ref` procedure can be used only when the result will be a Latin-1 character. A string's size can never be larger than a fixnum (so even `string-length` always returns a fixnum).

```
(unsafe-bytes-length bstr) → fixnum?
  bstr : bytes?
(unsafe-bytes-ref bstr k) → byte?
  bstr : bytes?
  k : fixnum?
(unsafe-bytes-set! bstr k b) → void?
  bstr : (and/c bytes? (not/c immutable?))
  k : fixnum?
  b : byte?
```

Unsafe versions of `bytes-length`, `bytes-ref`, and `bytes-set!`. A bytes's size can never be larger than a fixnum (so even `bytes-length` always returns a fixnum).

```
(unsafe-flvector-length v) → fixnum?
  v : flvector?
(unsafe-flvector-ref v k) → any/c
  v : flvector?
  k : fixnum?
(unsafe-flvector-set! v k x) → void?
  v : flvector?
  k : fixnum?
  x : flonum?
```

Unsafe versions of `flvector-length`, `flvector-ref`, and `flvector-set!`. A flvector's size can never be larger than a fixnum (so even `flvector-length` always returns a fixnum).

```
(unsafe-f64vector-ref vec k) → flonum?
  vec : f64vector?
  k : fixnum?
```

```
(unsafe-f64vector-set! vec k n) → void?
  vec : f64vector?
  k : fixnum?
  n : flonum?
```

Unsafe versions of `f64vector-ref` and `f64vector-set!`.

```
(unsafe-s16vector-ref vec k) → (integer-in -32768 32767)
  vec : s16vector?
  k : fixnum?
(unsafe-s16vector-set! vec k n) → void?
  vec : s16vector?
  k : fixnum?
  n : (integer-in -32768 32767)
```

Unsafe versions of `s16vector-ref` and `s16vector-set!`.

```
(unsafe-u16vector-ref vec k) → (integer-in 0 65535)
  vec : u16vector?
  k : fixnum?
(unsafe-u16vector-set! vec k n) → void?
  vec : u16vector?
  k : fixnum?
  n : (integer-in 0 65535)
```

Unsafe versions of `u16vector-ref` and `u16vector-set!`.

```
(unsafe-struct-ref v k) → any/c
  v : any/c
  k : fixnum?
(unsafe-struct-set! v k val) → void?
  v : any/c
  k : fixnum?
  val : any/c
(unsafe-struct*-ref v k) → any/c
  v : (not/c impersonator?)
  k : fixnum?
(unsafe-struct*-set! v k val) → void?
  v : (not/c impersonator?)
  k : fixnum?
  val : any/c
```

Unsafe field access and update for an instance of a structure type, where the `struct*` variants can be faster but do not work on impersonators. The index *k* must be between `0` (inclusive) and the number of fields in the structure (exclusive). In the case of `unsafe-struct-set!`, the field must be mutable.

17 Running Racket

17.1 Running Racket or GRacket

The core Racket run-time system is available in two main variants:

- Racket, which provides the primitives libraries on which `racket/base` is implemented. On Unix and Mac OS X, the executable is called `racket`. On Windows, the executable is called `Racket.exe`.
- GRacket, which is a GUI variant of `racket` to the degree that the system distinguishes them. On Unix, the executable is called `gracket`, and single-instance flags and X11-related flags are handled and communicated specially to the `racket/gui/base` library. On Windows, the executable is called `GRacket.exe`, and it is a GUI application (as opposed to a console application) that implements singe-instance support. On Mac OS X, the `gracket` script launches `GRacket.app`.

17.1.1 Initialization

On start-up, the top-level environment contains no bindings—not even `#%app` for function application. Primitive modules with names that start with `#%` are defined, but they are not meant for direct use, and the set of such modules can change. For example, the `'#%kernel` module is eventually used to bootstrap the implementation of `racket/base`.

The first action of Racket or GRacket is to initialize `current-library-collection-paths` to the result of `(find-library-collection-paths` *`pre-extras extras`*`)`, where *`pre-extras`* is normally `null` and *`extras`* are extra directory paths provided in order in the command line with `-S`/`--search`. An executable created from the Racket or GRacket executable can embed paths used as *`pre-extras`*.

Racket and GRacket next `require` `racket/init` and `racket/gui/init`, respectively, but only if the command line does not specify a `require` flag (`-t`/`--require`, `-l`/`--lib`, or `-u`/`--require-script`) before any `eval`, `load`, or read-eval-print-loop flag (`-e`/`--eval`, `-f`/`--load`, `-r`/`--script`, `-m`/`--main`, or `-i`/`--repl`). The initialization library can be changed with the `-I` configuration option. The `'configure-runtime` property of the initialization library's language is used before the library is instantiated; see §17.1.5 "Language Run-Time Configuration".

After potentially loading the initialization module, expression `eval`s, files `load`s, and module `require`s are executed in the order that they are provided on the command line. If any raises an uncaught exception, then the remaining `eval`s, `load`s, and `require`s are skipped. If the first `require` precedes any `eval` or `load` so that the initialization library is skipped, then the `'configure-runtime` property of the required module's library language is used before the module is instantiated; see §17.1.5 "Language Run-Time Configuration".

After running all command-line expressions, files, and modules, Racket or GRacket then starts a read-eval-print loop for interactive evaluation if no command line flags are provided other than configuration options. If any command-line argument is provided that is not a configuration option, then the read-eval-print-loop is not started, unless the `-i`/`--repl` flag is provided on the command line to specifically re-enable it. In addition, just before the command line is started, Racket loads the file `(find-system-path 'init-file)` and GRacket loads the file `(find-graphical-system-path 'init-file)` is loaded, unless the `-q`/`--no-init-file` flag is specified on the command line.

Finally, before Racket or GRacket exits, it calls the procedure that is the current value of `executable-yield-handler` in the main thread, unless the `-V`/`--no-yield` command-line flag is specified. Requiring `racket/gui/base` sets this parameter call `(racket 'yield)`.

17.1.2 Exit Status

The default exit status for a Racket or GRacket process is non-zero if an error occurs during a command-line `eval` (via `-e`, etc.), `load` (via `-f`, `-r`, etc.), or `require` (via `--l`, `-t`, etc.), but only when no read-eval-print loop is started. Otherwise, the default exit status is `0`.

In all cases, a call to `exit` (when the default exit handler is in place) can end the process with a specific status value.

17.1.3 Init Libraries

```
(require racket/init)
```

The `racket/init` library is the default start-up library for Racket. It re-exports the `racket`, `racket/enter` and `racket/help` libraries, and it sets `current-print` to use `pretty-print`.

```
(require racket/gui/init)
```

The `racket/gui/init` library is the default start-up library for GRacket. It re-exports the `racket/init` and `racket/gui/base` libraries, and it sets `current-load` to use `text-editor-load-handler`.

```
(require racket/language-info)
```

The `racket/language-info` library provides a `get-info` function that takes any value and returns another function; the returned function takes a key value and a default value, and it returns `'(#(racket/runtime-config configure #f))` if the key is `'configure-runtime` or the default value otherwise.

See also §17.3.6 “Module-Handling Configuration” in ***The Racket Guide***.

The vector `'#(racket/language-info get-info #f)` is suitable for attaching to a module as its language info to get the same language information as the `racket/base` language.

```
(require racket/runtime-config)
```

The `racket/runtime-config` library provides a `configure` function that returns another function; the returned function takes an value ans set `print-as-expression` to `#t`.

The vector `#(racket/runtime-config configure #f)` is suitable as a member of a list of runtime-configuration specification (as returned by a module's language-information function for the key `'configure-runtime`) to obtain the same runtime configuration as for the `racket/base` language.

```
(require racket/kernel)
```

The `racket/kernel` library is a minimal initial environment intended for scripts that require low startup time. It re-exports the bindings of the `'#%kernel`, along with a binding for `#%top-interaction`.

17.1.4 Command Line

The Racket and GRacket executables recognize the following command-line flags:

- File and expression options:
  - `-e` ⟨*expr*⟩ or `--eval` ⟨*expr*⟩ : `eval`s ⟨*expr*⟩. The results of the evaluation are printed via `current-print`.
  - `-f` ⟨*file*⟩ or `--load` ⟨*file*⟩ : `load`s ⟨*file*⟩; if ⟨*file*⟩ is `"-"`, then expressions are read and evaluated from standard input.
  - `-t` ⟨*file*⟩ or `--require` ⟨*file*⟩ : `requires` ⟨*file*⟩, and then `requires (submod (file "`⟨*file*⟩`") main)` if available.
  - `-l` ⟨*path*⟩ or `--lib` ⟨*path*⟩ : `requires (lib "`⟨*path*⟩`")`, and then `requires (submod (lib "`⟨*path*⟩`") main)` if available.
  - `-p` ⟨*package*⟩ : `requires (planet "`⟨*package*⟩`")`, and then `requires (submod (planet "`⟨*package*⟩`") main)` if available.
  - `-r` ⟨*file*⟩ or `--script` ⟨*file*⟩ : `load`s ⟨*file*⟩ as a script. This flag is like `-t` ⟨*file*⟩ plus `-N` ⟨*file*⟩ to set the program name and `--` to cause all further command-line elements to be treated as non-flag arguments.

    Despite its name, `--script` is not usually used for Unix scripts. See §19.2 "Scripts" for more information on scripts.
  - `-u` ⟨*file*⟩ or `--require-script` ⟨*file*⟩ : `requires` ⟨*file*⟩ as a script; This flag is like `-t` ⟨*file*⟩ plus `-N` ⟨*file*⟩ to set the program name and `--` to cause all further command-line elements to be treated as non-flag arguments.
  - `-k` ⟨*n*⟩ ⟨*m*⟩ ⟨*p*⟩ : Loads code embedded in the executable from file position ⟨*n*⟩ to ⟨*m*⟩ and from ⟨*m*⟩ to ⟨*p*⟩. (On Mac OS X, ⟨*n*⟩, ⟨*m*⟩, and ⟨*p*⟩ are relative to

a `__PLTSCHEME` segment in the executable.) The first range is loaded in every new place, and any modules declared in that range are considered predefined in the sense of `module-predefined?`. This option is normally embedded in a stand-alone binary that also embeds Racket code.

– `-m` or `--main` : Evaluates a call to `main` as bound in the top-level environment. All of the command-line arguments that are not processed as options (i.e., the arguments put into `current-command-line-arguments`) are passed as arguments to `main`. The results of the call are printed via `current-print`.

The call to `main` is constructed as an expression `(main arg-str ...)` where the lexical context of the expression gives `#%app` and `#%datum` bindings as `#%plain-app` and `#%datum`, but the lexical context of `main` is the top-level environment.

- Interaction options:

  – `-i` or `--repl` : Runs an interactive read-eval-print loop, using either `read-eval-print-loop` (Racket) or `graphical-read-eval-print-loop` (GRacket) after showing `(banner)` and loading `(find-system-path 'init-file)`. For GRacket, supply the `-z`/`--text-repl` configuration option to use `read-eval-print-loop` instead of `graphical-read-eval-print-loop`.

  – `-n` or `--no-lib` : Skips requiring the initialization library (i.e., `racket/init` or `racket/gui/init`, unless it is changed with the `-I` flag) when not otherwise disabled.

  – `-v` or `--version` : Shows `(banner)`.

  – `-K` or `--back` : GRacket, Mac OS X only; leave application in the background.

  – `-V` `--no-yield` : Skips final `executable-yield-handler` action, which normally waits until all frames are closed, etc. in the main eventspace before exiting for programs that use `racket/gui/base`.

- *Configuration options*:

  – `-c` or `--no-compiled` : Disables loading of compiled byte-code `".zo"` files, by initializing `current-compiled-file-paths` to `null`.

  – `-q` or `--no-init-file` : Skips loading `(find-system-path 'init-file)` for `-i`/`--repl`.

  – `-z` or `--text-repl` : GRacket only; changes `-i`/`--repl` to use `textual-read-eval-print-loop` instead of `graphical-read-eval-print-loop`.

  – `-I` ⟨*path*⟩ : Sets `(lib "`⟨*path*⟩`")` as the path to `require` to initialize the namespace, unless namespace initialization is disabled. Using this flag can effectively set the language for the read-eval-print loop and other top-level evaluation.

  – `-X` ⟨*dir*⟩ or `--collects` ⟨*dir*⟩ : Sets ⟨*dir*⟩ as the path to the main collection of libraries by making `(find-system-path 'collects-dir)` produce ⟨*dir*⟩. If ⟨*dir*⟩ is an empty string, then `(find-system-path 'collects-dir)` returns `"."`, but `current-library-collection-paths` is initialized to the empty list and `use-collection-link-paths` is initialized to `#f`.

  - `-S` ⟨*dir*⟩ or `--search` ⟨*dir*⟩ : Adds ⟨*dir*⟩ to the default library collection search path after the main collection directory. If the `-S`/`--dir` flag is supplied multiple times, the search order is as supplied.
  - `-A` ⟨*dir*⟩ or `--addon` ⟨*dir*⟩ : Sets the directory that is returned by `(find-system-path 'addon-dir)`.
  - `-C` ⟨*file*⟩ or `--links` ⟨*file*⟩ : Sets the user-specific collection links file path that is returned by `(find-system-path 'links-file)`; see also §17.2.1 “Collection Links”.
  - `-U` or `--no-user-path` : Omits user-specific paths in the search for collections, C libraries, etc. by initializing the `use-user-specific-search-paths` parameter to `#f`.
  - `-N` ⟨*file*⟩ or `--name` ⟨*file*⟩ : sets the name of the executable as reported by `(find-system-path 'run-file)` to ⟨*file*⟩.
  - `-J` ⟨*name*⟩ or `--wm-class` ⟨*name*⟩ : GRacket, Unix only; sets the `WM_CLASS` program class to ⟨*name*⟩ (while the `WM_CLASS` program name is derived from the executable name or a `-N`/`--name` argument).
  - `-j` or `--no-jit` : Disables the native-code just-in-time compiler by setting the `eval-jit-enabled` parameter to `#f`.
  - `-d` or `--no-delay` : Disables on-demand parsing of compiled code and syntax objects by setting the `read-on-demand-source` parameter to `#f`.
  - `-b` or `--binary` : Requests binary mode, instead of text mode, for the process’s input, out, and error ports. This flag currently has no effect, because binary mode is always used.
  - `-W` ⟨*level*⟩ or `--warn` ⟨*level*⟩ : Sets the logging level for writing events to the original error port. The possible ⟨*level*⟩ values are the same as for the `PLTSTDERR` environment variable. See §14.5 “Logging” for more information.
  - `-L` ⟨*level*⟩ or `--syslog` ⟨*level*⟩ : Sets the logging level for writing events to the system log. The possible ⟨*level*⟩ values are the same as for the `PLTSYSLOG` environment variable. See §14.5 “Logging” for more information.

- Meta options:
  - `--` : No argument following this flag is itself used as a flag.
  - `-h` or `--help` : Shows information about the command-line flags and start-up process and exits, ignoring all other flags.

If at least one command-line argument is provided, and if the first one after any configuration option is not a flag, then a `-u`/`--require-script` flag is implicitly added before the first non-flag argument.

If no command-line arguments are supplied other than configuration options, then the `-i`/`--repl` flag is effectively added.

For GRacket on Unix, the follow flags are recognized when they appear at the beginning of the command line, and they count as configuration options (i.e., they do not disable the read-eval-print loop or prevent the insertion of `-u`/`--require-script`):

- `-display` ⟨*display*⟩ : Sets the X11 display to use.
- `-geometry` ⟨*arg*⟩, `-bg` ⟨*arg*⟩, `-background` ⟨*arg*⟩, `-fg` ⟨*arg*⟩, `-foreground` ⟨*arg*⟩, `-fn` ⟨*arg*⟩, `-font` ⟨*arg*⟩, `-iconic`, `-name` ⟨*arg*⟩, `-rv`, `-reverse`, `+rv`, `-selectionTimeout` ⟨*arg*⟩, `-synchronous`, `-title` ⟨*arg*⟩, `-xnllanguage` ⟨*arg*⟩, or `-xrm` ⟨*arg*⟩ : Standard X11 arguments that are mostly ignored but accepted for compatibility with other X11 programs. The `-synchronous` flag behaves in the usual way.
- `-singleInstance` : If an existing GRacket is already running on the same X11 display, if it was started on a machine with the same hostname, and if it was started with the same name as reported by `(find-system-path 'run-file)`—possibly set with the `-N`/`--name` command-line argument—then all non-option command-line arguments are treated as filenames and sent to the existing GRacket instance via the application file handler (see `application-file-handler`).

Similarly, on Mac OS X, a leading switch starting with `-psn_` is treated as a special configuration option. It indicates that Finder started the application, so the current input, output, and error output are redirected to a GUI window.

Multiple single-letter switches (the ones preceded by a single `-`) can be collapsed into a single switch by concatenating the letters, as long as the first switch is not `--`. The arguments for each switch are placed after the collapsed switches (in the order of the switches). For example,

`-ifve` ⟨*file*⟩ ⟨*expr*⟩

and

`-i -f` ⟨*file*⟩ `-v -e` ⟨*expr*⟩

are equivalent. If a collapsed `--` appears before other collapsed switches in the same collapsed set, it is implicitly moved to the end of the collapsed set.

Extra arguments following the last option are available from the `current-command-line-arguments` parameter.

17.1.5 Language Run-Time Configuration

See also §17.3.6 “Module-Handling Configuration” in ***The Racket Guide***.

When a module is implemented using `#lang`, the language after `#lang` can specify configuration actions to perform when a module using the language is the main module of a program. The language specifies run-time configuration by

- attaching a `'module-language` syntax property to the module as read from its source (see `module` and `module-compiled-language-info`);
- having the function indicated by the `'module-language` syntax property recognize the `'configure-runtime` key, for which it returns a list of vectors; each vector must have the form `(vector mp name val)` where *mp* is a module path, *name* is a symbol, and *val* is an arbitrary value; and
- having each function called as `((dynamic-require mp name) val)` configure the run-time environment, typically by setting parameters such as `current-print`.

A `'configure-runtime` query returns a list of vectors, instead of directly configuring the environment, so that the indicated modules to be bundled with a program when creating a stand-alone executable; see §3 "`raco exe`: Creating Stand-Alone Executables" in *raco: Racket Command-Line Tools*.

For information on defining a new `#lang` language, see `syntax/module-reader`.

17.2 Libraries and Collections

A *library* is `module` declaration for use by multiple programs. Racket further groups libraries into *collections* that can be easily distributed and added to a local Racket installation.

Some libraries are distributed via PLaneT packages. Such libraries are referenced through a `planet` module path (see `require`) and are downloaded by Racket on demand.

Other collections are distributed with Racket, in which case each collection is a directory that is located in a `"collects"` directory relative to the Racket executable. A collection can also be installed in a user-specific directory. More generally, the search path for installed collections can be configured through the `current-library-collection-paths` parameter. Finally, the location of collections can be specified through the collection links files; see §17.2.1 "Collection Links" for more information. In all of these cases, the collections are referenced through `lib` paths (see `require`) or symbolic shorthands.

For example, the following module uses the `"getinfo.rkt"` library module from the `"setup"` collection, and the `"cards.rkt"` library module from the `"games"` collection's `"cards"` subcollection:

```
#lang racket
(require (lib "setup/getinfo.rkt")
         (lib "games/cards/cards.rkt"))
....
```

This example is more compactly and more commonly written using symbolic shorthands:

```
#lang racket
(require setup/getinfo
         games/cards/cards)
....
```

When an identifier *id* is used in a `require` form, it is converted to `(lib` *rel-string*`)` where *rel-string* is the string form of *id*.

A *rel-string* in `(lib` *rel-string*`)` consists of one or more path elements that name collections, and then a final path element that names a library file; the path elements are separated by `/`. If *rel-string* contains no `/`s, then `/main.rkt` is implicitly appended to the path. If *rel-string* contains `/` but does not end with a file suffix, then `.rkt` is implicitly appended to the path.

The translation of a `planet` or `lib` path to a `module` declaration is determined by the module name resolver, as specified by the `current-module-name-resolver` parameter.

For the default module name resolver, the search path for collections is determined by the content of `(find-system-path 'links-file)` (if it exists) and the `current-library-collection-paths` parameter. The collection links and then list of paths in `current-library-collection-paths` is searched from first to last to locate the first that contains *rel-string*. In other words, the filesystem tree for each element in the link table and search path is spliced together with the filesystem trees of other path elements. Some Racket tools rely on unique resolution of module path names, so an installation and `current-library-collection-paths` configuration should not allow multiple files to match the same collection and file name.

The value of the `current-library-collection-paths` parameter is initialized in the Racket executable to the result of `(find-library-collection-paths)`.

```
(find-library-collection-paths [pre-extras
                                post-extras]) → (listof path?)
  pre-extras : (listof path-string?) = null
  post-extras : (listof path-string?) = null
```

Produces a list of paths as follows:

- The path produced by `(build-path (find-system-path 'addon-dir) (version) "collects")` is the first element of the default collection path list, unless the value of the `use-user-specific-search-paths` parameter is `#f`.
- Extra directories provided in *pre-extras* are included next to the default collection path list, converted to complete paths relative to the executable.
- If the directory specified by `(find-system-path 'collects-dir)` is absolute, or if it is relative (to the executable) and it exists, then it is added to the end of the default collection path list.

- Extra directories provided in *post-extras* are included last in the default collection path list, converted to complete paths relative to the executable.
- If the PLTCOLLECTS environment variable is defined, it is combined with the default list using path-list-string->path-list. If it is not defined, the default collection path list (as constructed by the first three bullets above) is used directly.

  Note that on Unix and Mac OS X, paths are separated by :, and on Windows by ;. Also, path-list-string->path-list splices the default paths at an empty path, for example, with many Unix shells you can set PLTCOLLECTS to ":`pwd`", "`pwd`:", or "`pwd`" to specify search the current directory after, before, or instead of the default paths, respectively.

```
(collection-file-path file collection ...+) → path?
  file : path-string?
  collection : path-string?
(collection-file-path file
                      collection ...+
                      #:fail fail-proc) → any
  file : path-string?
  collection : path-string?
  fail-proc : (string? . -> . any)
```

Returns the path to the file indicated by *file* in the collection specified by the *collections*, where the second *collection* (if any) names a sub-collection, and so on.

If *file* is not found, but *file* ends in ".rkt" and a file with the suffix ".ss" exists, then the directory of the ".ss" file is used. If *file* is not found and the ".rkt"/".ss" conversion does not apply, but a directory corresponding to the *collection*s is found, then a path using the first such directory is returned.

Finally, if the collection is not found, and if *fail-proc* is provided, then *fail-proc* is applied to an error message (that does not start "collection-file-path:" or otherwise claim a source), and its result is the result of collection-file-path. If *fail-proc* is not provided and the collection is not found, then the exn:fail:filesystem exception is raised.

```
(collection-path collection ...+) → path?
  collection : path-string?
(collection-path collection
                 ...+
                 #:fail fail-proc) → any
  collection : path-string?
  fail-proc : (string? . -> . any)
```

Like collection-file-path, but without a specified file name, so that the first directory indicated by *collection*s is returned. The collection-file-path function normally should be used, instead, to support splicing of library-collection trees at the file level.

```
(current-library-collection-paths)
 → (listof (and/c path? complete-path?))
(current-library-collection-paths paths) → void?
  paths : (listof (and/c path? complete-path?))
```

Parameter that determines a list of complete directory paths for library collections used by `require`. See §17.2 “Libraries and Collections” for more information.

```
(use-user-specific-search-paths) → boolean?
(use-user-specific-search-paths on?) → void?
  on? : any/c
```

Parameter that determines whether user-specific paths, which are in the directory produced by `(find-system-path 'addon-dir)`, are included in search paths for collections and other files. For example, the initial value of `find-library-collection-paths` omits the user-specific collection directory when this parameter's value is `#f`.

```
(use-collection-link-paths) → boolean?
(use-collection-link-paths on?) → void?
  on? : any/c
```

Parameter that determines whether collection links files are used to locate collections.

17.2.1 Collection Links

The *collection links files* are used by `collection-file-path`, `collection-path`, and the default module name resolver to locate collections before trying the `(current-library-collection-paths)` search path, but only if the `use-collection-link-paths` parameter is set to `#t`. Furthermore, a user-specific collection links file takes precedence over an installation-wide collection links file, but the user-specific collection links file is used only the `use-user-specific-search-paths` parameter is set to `#t`.

The path of the user-specific collection links file is by `(find-system-path 'links-file)`, while an installation-wide collection links file is `"links.rktd"` in the `"config"` collection within the installation's main collection directory. Each collection links file is cached by Racket, but the file is re-read if its timestamp changes.

Each collection links file is `read` with default reader parameter settings to obtain a list. Every element of the list must be a link specification with one of the forms `(list string path)`, `(list string path regexp)`, `(list 'root path)`, or `(list 'root regexp)`. A *string* names a top-level collection, in which case *path* is a path that can be used as the collection's path (directly, as opposed to a subdirectory of *path* named by *string*). A `'root` entry, in contrast, acts like an path in `(current-library-collection-paths)`.

If *path* is a relative path, it is relative to the directory containing the collection links file. If *regexp* is specified in a link, then the link is used only if `(regexp-match? `*regexp*` (version))` produces a true result.

A single top-level collection can have multiple links in a collection links file, and any number of `'root` entries can appear. The corresponding paths are effectively spliced together, since the paths are tried in order to locate a file or sub-collection.

The `raco link` command-link tool can display, install, and remove links in the collection links file. See §2 "`raco link`: Library Collection Links" in *raco: Racket Command-Line Tools* for more information.

17.3 Interactive Help

```
(require racket/help)
```

The bindings documented in this section are provided by the `racket/help` and `racket/init` libraries, which means that they are available when the Racket executable is started with no command-line arguments. They are not provided by `racket/base` or `racket`.

```
help
(help string ...)
(help id)
(help id #:from module-path)
(help #:search datum ...)
```

For general help, see the main documentation page.

The `help` form searches the documentation and opens a web browser (using the user's selected browser) to display the results.

See `net/sendurl` for information on how the user's browser is launched to display help information.

A simple `help` or `(help)` form opens the main documentation page.

The `(help `*string* `...)` form—using literal strings, as opposed to expressions that produce strings—performs a string-matching search. For example,

```
(help "web browser" "firefox")
```

searches the documentation index for references that include the phrase "web browser" or "firefox."

A `(help `*id*`)` form looks for documentation specific to the current binding of *id*. For example,

```
(require net/url)
(help url->string)
```

opens a web browser to show the documentation for `url->string` from the `net/url` library.

For the purposes of `help`, a `for-label` require introduces a binding without actually executing the `net/url` library—for cases when you want to check documentation, but cannot or do not want to run the providing module.

```
(require racket/gui) ; does not work in racket
(require (for-label racket/gui)) ; ok in racket
(help frame%)
```

If *id* has no for-label and normal binding, then `help` lists all libraries that are known to export a binding for *id*.

The `(help` *id* `#:from` *module-path*`)` variant is similar to `(help` *id*`)`, but using only the exports of *module-path*. (The *module-path* module is required `for-label` in a temporary namespace.)

```
(help frame% #:from racket/gui) ; equivalent to the above
```

The `(help #:search` *datum* `...)` form is similar to `(help` *string* `...)`, where any non-string form of *datum* is converted to a string using `display`. No *datum* is evaluated as an expression.

For example,

```
(help #:search "web browser" firefox)
```

also searches the documentation index for references that include the phrase “web browser” or “firefox.”

17.4 Interactive Module Loading

```
(require racket/enter)
```

The bindings documented in this section are provided by the `racket/enter` and `racket/init` libraries, which means that they are available when the Racket executable is started with no command-line arguments. They are not provided by `racket/base` or `racket`.

```
(enter! module-path)
(enter! #f)
(enter! module-path flag ...+)

flag = #:quiet
     | #:verbose-reload
     | #:verbose
     | #:dont-re-require-enter
```

Intended for use in a REPL, such as when `racket` is started in interactive mode. When a *module-path* is provided (in the same sense as for `require`), the corresponding module is loaded or invoked, and the current namespace is changed to the body of the module via `module->namespace`. When `#f` is provided, then the current namespace is restored to the original one.

If invoking *module-path* requires loading any files, then modification dates of the files are recorded. If the file is modified, then a later `enter!` re-loads the module from source; see also §1.1.10.2 "Module Re-declarations". Similarly if a later `enter!` transitively `requires` a modified module, then the required module is re-loaded. Re-loading support works only for modules that are first loaded (either directly or indirectly through transitive `requires`) via `enter!`.

Additional *flag*s can customize aspects of `enter!`:

- When `enter!` loads or re-loads a module from a file, it can print a message to `(current-error-port)`. Use the `#:verbose` flag to print a message about such loads and re-loads, `#:verbose-reload` to print a message only for re-loaded modules, and `#:quiet` for no printouts. The default reporting corresponds to `#:verbose-reload`.
- After switching namespaces to the designated module, `enter!` automatically requires `racket/enter` into the namespace, so that `enter!` can be used to switch namespaces again. In some cases, requiring `racket/enter` might not be desirable (e.g., in a tool that uses `racket/enter`); use the `#:dont-re-require-enter` flag to disable the require.

17.5 Debugging

Racket's built-in debugging support is limited to context (i.e., "stack trace") information that is printed with an exception. In some cases, disabling the JIT compiler can affect context information. The `errortrace` library supports more consistent (independent of the JIT compiler) and precise context information. The `racket/trace` library provides simple tracing support. Finally, the DrRacket programming environment provides much more debugging support.

17.5.1 Tracing

```
(require racket/trace)
```

The bindings documented in this section are provided by the `racket/trace` library, not `racket/base` or `racket`.

The `racket/trace` library mimics the tracing facility available in Chez Scheme.

```
(trace id ...)
```

Each *id* must be bound to a procedure in the environment of the `trace` expression. Each *id* is `set!`ed to a new procedure that traces procedure calls and returns by printing the arguments and results of the call via `current-trace-notify`. If multiple values are returned, each value is displayed starting on a separate line.

When traced procedures invoke each other, nested invocations are shown by printing a nesting prefix. If the nesting depth grows to ten and beyond, a number is printed to show the actual nesting depth.

The `trace` form can be used on an identifier that is already traced. In this case, assuming that the variable's value has not been changed, `trace` has no effect. If the variable has been changed to a different procedure, then a new trace is installed.

Tracing respects tail calls to preserve loops, but its effect may be visible through continuation marks. When a call to a traced procedure occurs in tail position with respect to a previous traced call, then the tailness of the call is preserved (and the result of the call is not printed for the tail call, because the same result will be printed for an enclosing call). Otherwise, however, the body of a traced procedure is not evaluated in tail position with respect to a call to the procedure.

The result of a `trace` expression is `#<void>`.

Examples:

```
> (define (f x) (if (zero? x) 0 (add1 (f (sub1 x)))))

> (trace f)

> (f 10)
>(f 10)
> (f 9)
> >(f 8)
> > (f 7)
> > >(f 6)
> > > (f 5)
> > > >(f 4)
```

```
> > > > (f 3)
> > > > >(f 2)
> > > > > (f 1)
> > > >[10] (f 0)
< < < <[10] 0
< < < < < 1
< < < < <2
< < < < 3
< < < <4
< < < 5
< < <6
< < 7
< <8
< 9
<10
10
```

```
(untrace id ...)
```

Undoes the effects of the `trace` form for each *id*, `set!`ing each *id* back to the untraced procedure, but only if the current value of *id* is a traced procedure. If the current value of a *id* is not a procedure installed by `trace`, then the variable is not changed.

The result of an `untrace` expression is `#<void>`.

```
(current-trace-notify) → (string? . -> . any)
(current-trace-notify proc) → void?
  proc : (string? . -> . any)
```

A parameter that determines the way that trace output is displayed. The string given to *proc* is a trace; it does not end with a newline, but it may contain internal newlines. Each call or result is converted into a string using `pretty-print`. The parameter's default value prints the given string followed by a newline to `(current-output-port)`.

```
(trace-call id proc #:<kw> kw-arg ...) → any/c
  id : symbol?
  proc : procedure?
  kw-arg : any/c
```

Calls *proc* with the arguments supplied in `args`, and possibly using keyword arguments. Also prints out the trace information during the call, as described above in the docs for `trace`, using *id* as the name of *proc*.

```
                                (-> symbol?
                                    list?
(current-trace-print-args) →        (listof keyword?)
                                    list?
                                    number?
                                    void?)
(current-trace-print-args trace-print-args) → void?
                      (-> symbol?
                          list?
  trace-print-args :      (listof keyword?)
                          list?
                          number?
                          void?)
```

The value of this parameter is invoked to print out the arguments of a traced call. It receives the name of the function, the function's ordinary arguments, its keywords, the values of the keywords, and a number indicating the depth of the call.

```
                                  (-> symbol?
                                      list?
(current-trace-print-results) →       number?
                                      any)
(current-trace-print-results trace-print-results) → void?
                        (-> symbol?
                            list?
  trace-print-results :     number?
                            any)
```

The value of this parameter is invoked to print out the results of a traced call. It receives the name of the function, the function's results, and a number indicating the depth of the call.

```
(current-prefix-in) → string?
(current-prefix-in prefix) → void?
  prefix : string?
```

This string is used by the default value of `current-trace-print-args` indicating that the current line is showing the a call to a traced function.

It defaults to `">"`.

```
(current-prefix-out) → string?
(current-prefix-out prefix) → void?
  prefix : string?
```

This string is used by the default value of `current-trace-print-results` indicating that the current line is showing the result of a traced call.

It defaults to `"<"`.

Bibliography

[Danvy90] Olivier Danvy and Andre Filinski, "Abstracting Control," LISP and Functional Programming, 1990.

[Felleisen88a] Matthias Felleisen, "The theory and practice of first-class prompts," Principles of Programming Languages, 1988.

[Felleisen88] Matthias Felleisen, Mitch Wand, Dan Friedman, and Bruce Duba, "Abstract Continuations: A Mathematical Semantics for Handling Full Functional Jumps," LISP and Functional Programming, 1988.

[Friedman95] Daniel P. Friedman, C. T. Haynes, and R. Kent Dybvig, "Exception system proposal," web page, 1995. `http://www.cs.indiana.edu/scheme-repository/doc.proposals.exceptions.html`

[Gasbichler02] Martin Gasbichler and Michael Sperber, "Processes vs. User-Level Threads in Scsh," Workshop on Scheme and Functional Programming, 2002.

[Gunter95] Carl Gunter, Didier Remy, and Jon Rieke, "A Generalization of Exceptions and Control in ML-like Languages," Functional Programming Languages and Computer Architecture, 1995.

[Hayes97] Barry Hayes, "Ephemerons: a New Finalization Mechanism," Object-Oriented Languages, Programming, Systems, and Applications, 1997.

[Hieb90] Robert Hieb and R. Kent Dybvig, "Continuations and Concurrency," Principles and Practice of Parallel Programming, 1990.

[L'Ecuyer02] Pierre L'Ecuyer, Richard Simard, E. Jack Chen, and W. David Kelton, "An Object-Oriented Random-Number Package With Many Long Streams and Substreams," Operations Research, 50(6), 2002.

[Queinnec91] Queinnec and Serpette, "A Dynamic Extent Control Operator for Partial Continuations," Principles of Programming Languages, 1991.

[Shan04] Ken Shan, "Shift to Control," Workshop on Scheme and Functional Programming, 2004.

[Sperber07] Michael Sperber, R. Kent Dybvig, Matthew Flatt, and Anton van Straaten (editors), "The Revised^6 Report on the Algorithmic Language Scheme." 2007. `http://www.r6rs.org/`

[Sitaram90] Dorai Sitaram and Matthias Felleisen, "Control Delimiters and Their Hierarchies," *Lisp and Symbolic Computation*, 1990.

[Sitaram93] Dorai Sitaram, "Handling Control," Programming Language Design and Implementation, 1993.

[SRFI-42] Sebastian Egner, "SRFI-42: Eager Comprehensions," SRFI, 2003. `http://srfi.schemers.org/srfi-42/`

[Waddell99] Oscar Waddell and R. Kent Dybvig, "Extending the Scope of Syntactic Abstraction," Principles of Programming Languages, 1999.

Index

", 54
#! , 56
#!, 61
#!/, 56
#", 55
#%, 50
#%app, 103
#%datum, 101
#%expression, 101
'#%kernel, 959
'#%kernel, 961
#%module-begin, 75
#%plain-app, 104
#%plain-lambda, 108
#%plain-module-begin, 75
#%provide, 95
#%require, 93
#%stratified-body, 144
#%top, 102
#%top-interaction, 140
#%variable-reference, 102
#&, 58
#', 56
#,, 56
#,@, 56
#0#, 60
#0=, 60
#:, 59
#;, 56
#<<, 55
#\, 58
#`, 56
#b, 51
#ci, 50
#cs, 50
#d, 51
#e, 51
#F, 52
#f, 52
#false, 52
#hash, 58
#hasheq, 58
#hasheqv, 58
#i, 51
#lang, 61
#o, 51
#px, 59
#reader, 60
#rx, 59
#T, 52
#t, 52
#true, 52
#x, 51
#|, 56
%, 557
', 55
(, 52
), 52
*, 160
*channel-get, 625
*channel-put, 625
*channel?, 625
+, 159
+inf.0, 152
+inf.f, 152
+nan.0, 152
+nan.f, 152
+rv, 964
,, 56
,@, 56
-, 159
--, 963
--addon, 963
--back, 962
--binary, 963
--collects, 962
--eval, 961
--help, 963
--lib, 961
--links, 963
--load, 961
--main, 962

--name, 963
--no-compiled, 962
--no-delay, 963
--no-init-file, 962
--no-jit, 963
--no-lib, 962
--no-user-path, 963
--no-yield, 962
--repl, 962
--require, 961
--require-script, 961
--script, 961
--search, 963
--syslog, 963
--text-repl, 962
--version, 962
--warn, 963
--wm-class, 963
->, 476
->*, 477
->*m, 435
->d, 480
->dm, 435
->fl, 187
->i, 478
->m, 435
-A, 963
-b, 963
-background, 964
-bg, 964
-C, 963
-c, 962
-d, 963
-display, 964
-e, 961
-f, 961
-fg, 964
-fn, 964
-font, 964
-foreground, 964
-geometry, 964
-h, 963
-I, 962
-i, 962
-iconic, 964
-inf.0, 152
-inf.f, 152
-J, 963
-j, 963
-K, 962
-k, 961
-L, 963
-l, 961
-m, 962
-N, 963
-n, 962
-name, 964
-p, 961
-psn_, 964
-q, 962
-r, 961
-reverse, 964
-rv, 964
-S, 963
-selectionTimeout, 964
-singleInstance, 964
-synchronous, 964
-t, 961
-title, 964
-U, 963
-u, 961
-V, 962
-v, 962
-W, 963
-X, 962
-xnllanguage, 964
-xrm, 964
-z, 962
., 52
..., 644
".racketrc", 887
/, 160
'3m, 936
:do-in, 136

;, 56
<, 167
</c, 467
<=, 168
<=/c, 467
=, 167
=/c, 467
==, 527
=>, 116
>, 168
>/c, 467
>=, 168
>=/c, 467
[, 52
\, 49
\", 54
\', 54
\⟨*digit*$_8$⟩$^{\{1,3\}}$, 54
\⟨*newline*⟩, 55
\\, 54
\a, 54
\b, 54
\e, 54
\f, 54
\n, 54
\r, 54
\t, 54
\u⟨*digit*$_{16}$⟩$^{\{1,4\}}$, 54
\U⟨*digit*$_{16}$⟩$^{\{1,8\}}$, 55
\u⟨*digit*$_{16}$⟩$^{\{4,4\}}$\u⟨*digit*$_{16}$⟩$^{\{4,4\}}$, 54
\v, 54
\x⟨*digit*$_{16}$⟩$^{\{1,2\}}$, 54
], 52
_, 645
`, 55
abort, 557
abort-current-continuation, 551
abort/cc, 557
'aborts, 222
abs, 163
absent, 435
absolute-path?, 875
abstract, 410
accessor, 376
acos, 172
add-between, 287
add1, 163
Additional Byte String Functions, 224
Additional Control Operators, 556
Additional Custom-Output Support, 779
Additional Higher-Order Functions, 370
Additional List Functions and Synonyms, 282
Additional Matching Forms, 523
Additional Promise Kinds, 548
Additional provide Forms, 100
Additional require Forms, 96
Additional String Functions, 206
Additional Syntactic Constraints, 241
Additional Vector Functions, 300
'addon-dir, 887
after-seconds, 615
after-seconds%, 623
alarm-evt, 579
all-defined-out, 87
all-from-out, 87
always-evt, 579
and, 117
and/c, 467
andmap, 267
angle, 174
any, 475
'any, 748
'any-one, 748
any/c, 466
append, 266
'append, 703
append*, 288
append-map, 290
applicable structure, 376
apply, 358
argmax, 291
argmin, 291
Arithmetic, 159

arithmetic-shift, 176
arity-at-least, 365
arity-at-least-value, 365
arity-at-least?, 365
armed, 682
asin, 171
assf, 275
assignment transformers, 39
Assignment: set! and set!-values, 126
assoc, 274
association list, 337
assq, 275
assv, 274
Async Bidirectional Channels, 630
async-bi-channel-get, 630
async-bi-channel-put, 631
async-bi-channel?, 630
async-channel-get, 585
async-channel-put, 585
async-channel-put-evt, 585
async-channel-try-get, 585
async-channel?, 584
atan, 172
Attaching Contracts to Values, 489
augment, 408
augment*, 412
augment-final, 409
augment-final*, 412
augmenting, 397
augride, 409
augride*, 412
automatic fields, 374
"AUX", 883
available, 27
banner, 937
base environment, 32
base phase, 45
Basic Pretty-Print Options, 774
begin, 124
begin-encourage-inline, 144
begin-for-syntax, 125
begin0, 125
between/c, 468
'binary, 701
binding, 32
binds, 32
Bitwise Operations, 174
bitwise-and, 175
bitwise-bit-field, 176
bitwise-bit-set?, 175
bitwise-ior, 174
bitwise-not, 175
bitwise-xor, 175
blame object, 493
Blame Objects, 500
blame-add-context, 500
blame-contract, 502
blame-negative, 502
blame-original?, 502
blame-positive, 502
blame-replace-negative, 502
blame-source, 502
blame-swap, 502
blame-swapped?, 502
blame-value, 502
blame?, 500
block, 143
'block, 698
Blocks: block, 143
Boolean Aliases, 150
boolean=?, 150
boolean?, 145
booleans, 145
Booleans and Equality, 145
bound, 32
bound-identifier=?, 651
box, 306
box, 306
box-cas!, 307
box-immutable, 306
box-immutable/c, 470
box/c, 470
box?, 306
Boxes, 306

break, 565
break-enabled, 567
break-evaluator, 866
break-thread, 571
Breaks, 565
Buffered Asynchronous Channels, 584
build-chaperone-contract-property, 506
build-compound-type-name, 499
build-contract-property, 507
build-distributed-launch-path, 624
build-flat-contract-property, 506
build-list, 264
build-path, 874
build-path/convention-type, 875
build-string, 199
build-vector, 300
Building New Contract Combinators, 493
Built-in Exception Types, 542
Byte and String Input, 747
Byte and String Output, 756
byte converter, 220
byte string, 209
Byte String Comparisons, 214
Byte String Constructors, Selectors, and Mutators, 209
byte strings, parsing, 54
byte strings, immutable, 209
byte strings, concatenate, 213
Byte Strings, 209
byte-pregexp, 244
byte-pregexp?, 243
byte-ready?, 755
byte-regexp, 244
byte-regexp?, 243
byte?, 210
bytes, 210
bytes, 209
Bytes to Bytes Encoding Conversion, 220
Bytes to/from Characters, Decoding and Encoding, 215
bytes->immutable-bytes, 210
bytes->list, 213
bytes->path, 872
bytes->path-element, 873
bytes->string/latin-1, 216
bytes->string/locale, 216
bytes->string/utf-8, 215
bytes-append, 213
bytes-append*, 225
bytes-close-converter, 221
bytes-convert, 222
bytes-convert-end, 223
bytes-converter?, 224
bytes-copy, 212
bytes-copy!, 212
bytes-fill!, 213
bytes-join, 225
bytes-length, 211
bytes-no-nuls?, 927
bytes-open-converter, 220
bytes-ref, 211
bytes-set!, 211
bytes-utf-8-index, 219
bytes-utf-8-length, 218
bytes-utf-8-ref, 219
bytes<?, 215
bytes=?, 214
bytes>?, 215
bytes?, 209
caaaar, 278
caaadr, 278
caaar, 276
caadar, 278
caaddr, 279
caadr, 276
caar, 275
cadaar, 279
cadadr, 279
cadar, 277
caddar, 279
cadddr, 280
caddr, 277
cadr, 275

call-by-value, 25
call-in-nested-thread, 570
call-in-sandbox-context, 868
call-with-break-parameterization, 567
call-with-composable-continuation, 552
call-with-continuation-barrier, 553
call-with-continuation-prompt, 550
call-with-current-continuation, 552
call-with-custodian-shutdown, 867
call-with-escape-continuation, 553
call-with-exception-handler, 538
call-with-file-lock/timeout, 905
call-with-immediate-continuation-mark, 563
call-with-input-bytes, 734
call-with-input-file, 704
call-with-input-file*, 705
call-with-input-string, 733
call-with-killing-threads, 867
call-with-limits, 869
call-with-output-bytes, 733
call-with-output-file, 705
call-with-output-file*, 705
call-with-output-string, 733
call-with-parameterization, 592
call-with-semaphore, 584
call-with-semaphore/enable-break, 584
call-with-trusted-sandbox-configuration, 856
call-with-values, 529
call/cc, 552
call/comp, 557
call/ec, 553
call/prompt, 556
'can-update, 703
car, 263
case, 118
case->, 481
case->m, 435
case-insensitive, 50
case-lambda, 108
case-sensitivity, 50
'cc, 231
cdaaar, 280
cdaadr, 280
cdaar, 277
cdadar, 280
cdaddr, 280
cdadr, 277
cdar, 276
cddaar, 281
cddadr, 281
cddar, 277
cdddar, 281
cddddr, 281
cdddr, 278
cddr, 276
cdr, 263
ceiling, 165
'certify-mode, 683
'cf, 231
'cgc, 936
Chaining Reader Language, 62
channel, 581
channel-get, 582
channel-put, 582
channel-put-evt, 582
channel-put-evt?, 582
channel-try-get, 582
channel?, 581
Channels, 581
chaperone, 826
Chaperone Constructors, 835
chaperone contract property, 507
chaperone-box, 836
chaperone-continuation-mark-key, 840
chaperone-contract-property?, 508
chaperone-contract?, 511
chaperone-evt, 838
chaperone-hash, 837

chaperone-of?, 828
chaperone-procedure, 835
chaperone-prompt-tag, 839
chaperone-struct, 835
chaperone-struct-type, 837
chaperone-vector, 836
chaperone?, 828
char->integer, 226
char-alphabetic?, 229
char-blank?, 231
char-ci<=?, 228
char-ci<?, 228
char-ci=?, 228
char-ci>=?, 229
char-ci>?, 229
char-downcase, 232
char-foldcase, 232
char-general-category, 231
char-graphic?, 230
char-iso-control?, 231
char-lower-case?, 230
char-numeric?, 230
char-punctuation?, 230
char-ready?, 755
char-symbolic?, 230
char-title-case?, 230
char-titlecase, 232
char-upcase, 231
char-upper-case?, 230
char-utf-8-length, 226
char-whitespace?, 230
char<=?, 227
char<?, 227
char=?, 226
char>=?, 228
char>?, 227
char?, 225
Character Comparisons, 226
Character Conversions, 231
Characters, 225
Characters, 225
Characters and Scalar Values, 225
check-duplicate-identifier, 652
checked-procedure-check-and-extract, 369
checked-struct-info?, 395
choice-evt, 578
class, 401
class, 397
class*, 399
class->interface, 443
class-field-accessor, 427
class-field-mutator, 427
class-info, 447
class/c, 432
class/derived, 414
class?, 441
Classes and Objects, 397
Classifications, 229
cleanse, 871
cleanse-path, 877
'client, 843
close-input-port, 696
close-output-port, 696
'cn, 231
'co, 231
Code Inspectors, 850
code point, 225
coerce-chaperone-contract, 500
coerce-chaperone-contracts, 500
coerce-contract, 499
coerce-contract/f, 500
coerce-contracts, 499
coerce-flat-contract, 500
coerce-flat-contracts, 500
collect-garbage, 949
Collection Links, 968
collection links files, 968
collection-file-path, 967
collection-path, 967
collections, 965
'collects-dir, 888
column locations, 699
column numbers, 699

"COM1", 883
"COM2", 883
"COM3", 883
"COM4", 883
"COM5", 883
"COM6", 883
"COM7", 883
"COM8", 883
"COM9", 883
combine-in, 79
combine-out, 90
Command Line, 961
command-line, 939
Command-Line Parsing, 939
Compilation, 44
compilation handler, 814
compile, 814
compile-allow-set!-undefined, 815
compile-context-preservation-enabled, 815
compile-enforce-module-constants, 815
compile-syntax, 814
compiled, 44
Compiled Modules and References, 820
compiled-expression?, 815
compiled-load handler, 811
compiled-module-expression?, 822
'compiler-hint:cross-module-inline, 120
'complete, 222
'complete, 224
complete-path?, 876
completion value, 604
Complex Numbers, 172
complex numbers, 152
complex?, 153
composable continuation, 28
compose, 359
compose1, 359
compound-unit, 455
compound-unit/infer, 457
"CON", 883
Concurrency and Parallelism, 569
cond, 115
Conditionals: if, cond, and, and or, 115
Configuration options, 962
'configure-runtime, 965
Configuring Default Handling, 540
conjugate, 184
connect-to-named-place, 615
connection%, 622
cons, 262
cons/c, 471
cons?, 282
const, 370
Constructing Graphs: shared, 112
constructor, 375
context, 37
continuation, 18
continuation barrier, 29
continuation frames, 28
Continuation Frames and Marks, 28
Continuation Marks, 561
continuation marks, 28
Continuation Marks: with-continuation-mark, 137
continuation-mark-key/c, 474
continuation-mark-key?, 564
continuation-mark-set->context, 564
continuation-mark-set->list, 562
continuation-mark-set->list*, 563
continuation-mark-set-first, 563
continuation-mark-set?, 564
continuation-marks, 562
continuation-prompt-available?, 554
continuation-prompt-tag?, 554
continuation?, 554
Continuations, 550
'continues, 222
'continues, 224
contract, 492
contract property, 507
Contract Utilities, 510

contract-first-order, 510
contract-first-order-passes?, 510
contract-name, 511
contract-out, 489
contract-proc, 514
contract-projection, 512
contract-property?, 508
contract-random-generate, 515
contract-stronger?, 510
contract-struct, 485
contract?, 510
contracted, 453
Contracts, 465
Contracts, 465
Contracts as structs, 505
control, 558
Control Flow, 529
control-at, 558
control0, 559
control0-at, 560
convert-relative-module-path, 672
convert-stream, 747
copy-directory/files, 899
copy-file, 892
copy-port, 747
Copying and Updating Structures, 390
Copying Streams, 747
cos, 171
cosh, 184
count, 289
Counting Positions, Lines, and Columns, 699
create-place-node, 613
Creating Classes, 399
Creating formatted identifiers, 688
Creating Interfaces, 398
Creating Loggers, 928
Creating Objects, 422
Creating Ports, 734
Creating Structure Types, 379
Creating Threads, 569
Creating Units, 450
creation tree, 597
creation-tree-pict, 598
'cs, 231
cupto, 561
current custodian, 30
current logger, 928
current namespace, 45
current-blame-format, 503
current-break-parameterization, 567
current-code-inspector, 851
current-command-line-arguments, 937
current-command-line-arguments, 964
current-compile, 814
current-continuation-marks, 562
current-contract-region, 492
current-custodian, 845
current-date, 934
current-directory, 893
current-drive, 893
current-error-port, 697
current-eval, 807
current-evt-pseudo-random-generator, 581
current-future, 593
current-gc-milliseconds, 933
current-get-interaction-input-port, 813
current-inexact-milliseconds, 931
current-input-port, 697
current-inspector, 848
current-library-collection-paths, 968
current-load, 808
current-load-extension, 810
current-load-relative-directory, 812
current-load/use-compiled, 811
current-locale, 695
current-logger, 929
current-memory-use, 950
current-milliseconds, 932
current-module-declare-name, 820

current-module-declare-source, 820
current-module-name-resolver, 818
current-namespace, 800
current-output-port, 697
current-parameterization, 592
current-prefix-in, 974
current-prefix-out, 974
current-preserved-thread-cell-values, 588
current-print, 814
current-process-milliseconds, 933
current-prompt-read, 813
current-pseudo-random-generator, 178
current-read-interaction, 813
current-reader-guard, 764
current-readtable, 764
current-recorded-disappeared-uses, 691
current-require-module-path, 672
current-seconds, 931
current-security-guard, 844
current-subprocess-custodian-mode, 922
current-syntax-context, 690
current-thread, 570
current-thread-group, 847
current-thread-initial-stack-size, 937
current-trace-notify, 973
current-trace-print-args, 974
current-trace-print-results, 974
current-write-relative-directory, 771
curry, 371
curryr, 372
custodian, 30
custodian box, 31
custodian-box-value, 846
custodian-box?, 846
custodian-limit-memory, 845
custodian-managed-list, 845
custodian-memory-accounting-available?, 845
custodian-require-memory, 845
custodian-shutdown-all, 844
custodian?, 844
Custodians, 30
Custodians, 844
Custom Ports, 711
custom ports, 711
custom-print-quotable-accessor, 789
custom-print-quotable?, 789
custom-write-accessor, 789
custom-write?, 789
Customizing Evaluators, 856
Data-structure Contracts, 465
Datatypes, 145
date, 931
Date Utilities, 933
date*, 932
date*-nanosecond, 932
date*-time-zone-name, 932
date*?, 932
date->julian/scalinger, 935
date->seconds, 934
date->string, 934
date-day, 931
date-display-format, 934
date-dst?, 931
date-hour, 931
date-minute, 931
date-month, 931
date-second, 931
date-time-zone-offset, 931
date-week-day, 931
date-year, 931
date-year-day, 931
date?, 931
datum, 48
datum->syntax, 648
datum-intern-literal, 649
Debugging, 971
Declaring Paths Needed at Run Time, 894
default-continuation-prompt-tag,

551
DEFAULT-ROUTER-PORT, 626
define, 119
Define Remote Server, 628
define*, 142
define*-syntax, 142
define*-syntaxes, 142
define*-values, 142
define-cast, 628
define-compound-unit, 458
define-compound-unit/infer, 458
define-contract-struct, 486
define-for-syntax, 122
define-generics, 386
define-local-member-name, 420
define-match-expander, 525
define-member-name, 421
define-named-remote-server, 628
define-namespace-anchor, 799
define-opt/c, 512
define-package, 140
define-provide-syntax, 124
define-remote-server, 628
define-require-syntax, 123
define-rpc, 628
define-runtime-module-path, 897
define-runtime-module-path-index, 896
define-runtime-path, 894
define-runtime-path-list, 896
define-runtime-paths, 896
define-sequence-syntax, 135
define-serializable-class, 440
define-serializable-class*, 439
define-serializable-struct, 794
define-serializable-struct/versions, 796
define-signature, 452
define-signature-form, 461
define-state, 628
define-struct, 378
define-struct/contract, 491
define-struct/derived, 379
define-syntax, 120
define-syntax-parameter, 676
define-syntax-rule, 644
define-syntaxes, 121
define-unit, 457
define-unit-binding, 458
define-unit-from-context, 459
define-unit/contract, 462
define-unit/new-import-export, 460
define-unit/s, 460
define-values, 120
define-values-for-export, 453
define-values-for-syntax, 123
define-values/invoke-unit, 455
define-values/invoke-unit/infer, 459
define/augment, 413
define/augment-final, 414
define/augride, 413
define/contract, 491
define/final-prop, 509
define/generic, 387
define/overment, 413
define/override, 413
define/override-final, 413
define/private, 414
define/public, 412
define/public-final, 413
define/pubment, 413
define/subexpression-pos-prop, 509
define/with-syntax, 689
Defining Structure Types: struct, 375
Definitions: define, define-syntax, ..., 119
degrees->radians, 183
delay, 547
delay/idle, 549
delay/name, 548
delay/strict, 549
delay/sync, 549
delay/thread, 549

Delayed Evaluation, 547
'delete, 843
delete-directory, 893
delete-directory/files, 900
delete-file, 890
delimited continuation, 28
Delimiters and Dispatch, 48
denominator, 166
depth marker, 636
derived class, 397
Deriving New Iteration Forms, 133
deserialize, 790
deserialize-module-guard, 793
'desk-dir, 887
dict->list, 350
dict-can-functional-set?, 339
dict-can-remove-keys?, 338
dict-count, 346
dict-for-each, 346
dict-has-key?, 341
dict-iter-contract, 352
dict-iterate-first, 346
dict-iterate-key, 347
dict-iterate-next, 347
dict-iterate-value, 347
dict-key-contract, 352
dict-keys, 349
dict-map, 345
dict-mutable?, 338
dict-ref, 342
dict-ref!, 343
dict-remove, 345
dict-remove!, 344
dict-set, 340
dict-set!, 339
dict-set*, 341
dict-set*!, 340
dict-update, 344
dict-update!, 343
dict-value-contract, 352
dict-values, 349
dict?, 338
Dictionaries, 337
dictionary, 337
Directories, 893
directory-exists?, 893
directory-list, 893
'disappeared-binding, 680
'disappeared-use, 680
disarm, 682
Opening a null output port, 737
'dispatch-macro, 782
Dispatch: case, 118
display, 766
display-lines, 732
display-lines-to-file, 899
display-to-file, 898
displayln, 767
displayln/f, 627
Distributed Places, 607
Distributed Places MPI, 631
distributed-launch-path, 624
distributed-place-wait, 620
division by inexact zero, 152
'dll, 936
do, 137
Do Loops, 137
'doc-dir, 887
double-flonum?, 156
drop, 286
drop-right, 286
dump-memory-stats, 950
dup-input-port, 740
dup-output-port, 740
dye packs, 682
dynamic extension, 810
dynamic extent, 18
Dynamic Module Access, 824
dynamic-place, 603
dynamic-place*, 604
dynamic-require, 824
dynamic-require-for-syntax, 825
dynamic-send, 425
dynamic-wind, 554

eighth, 284
else, 116
empty, 282
empty-sequence, 326
empty-stream, 330
empty?, 282
'enclosing-module-name, 72
Encodings and Locales, 694
enter!, 971
environment, 32
Environment and Runtime Information, 935
eof, 697
eof-evt, 744
eof-object?, 697
ephemeron, 946
ephemeron-value, 947
ephemeron?, 947
Ephemerons, 946
eprintf, 768
eq-hash-code, 315
eq?, 147
equal-hash-code, 315
equal-secondary-hash-code, 315
equal<%>, 439
equal?, 146
equal?/recur, 147
eqv-hash-code, 315
eqv?, 146
error, 531
'error, 222
'error, 703
error display handler, 540
error escape handler, 540
error message convention, 530
Error Message Conventions, 530
Error reporting, 690
error value conversion handler, 541
error-display-handler, 540
error-escape-handler, 540
error-print-context-length, 541
error-print-source-location, 542
error-print-width, 541
error-value->string-handler, 541
escape continuation, 29
eval, 807
eval-jit-enabled, 815
eval-syntax, 808
Evaluation and Compilation, 807
evaluation handler, 807
Evaluation Model, 18
evaluation order, 103
evaluator-alive?, 865
even?, 157
event-container<%>, 616
Events, 574
every-seconds, 615
evt?, 577
'exact, 920
exact number, 152
exact->inexact, 158
exact-ceiling, 184
exact-floor, 184
exact-integer?, 155
exact-nonnegative-integer?, 155
exact-positive-integer?, 155
exact-round, 184
exact-truncate, 184
exact?, 158
except, 453
except-in, 78
except-out, 88
exception handler, 30
Exceptions, 30
Exceptions, 530
Exceptions, 30
'exec-file, 888
executable-yield-handler, 568
'execute, 843
'execute, 891
Execution Timeline, 596
execution timeline, 596
'exists, 843
exit, 567
exit handler, 568

Exit Status, 960
exit-handler, 568
Exiting, 567
exn, 542
exn-continuation-marks, 542
exn-message, 542
exn:break, 545
exn:break-continuation, 545
exn:break?, 545
exn:fail, 542
exn:fail:contract, 542
exn:fail:contract:arity, 542
exn:fail:contract:arity?, 542
exn:fail:contract:blame, 503
exn:fail:contract:blame-object, 503
exn:fail:contract:blame?, 503
exn:fail:contract:continuation, 543
exn:fail:contract:continuation?, 543
exn:fail:contract:divide-by-zero, 542
exn:fail:contract:divide-by-zero?, 542
exn:fail:contract:non-fixnum-result, 543
exn:fail:contract:non-fixnum-result?, 543
exn:fail:contract:variable, 543
exn:fail:contract:variable-id, 543
exn:fail:contract:variable?, 543
exn:fail:contract?, 542
exn:fail:filesystem, 544
exn:fail:filesystem:exists, 544
exn:fail:filesystem:exists?, 544
exn:fail:filesystem:version, 544
exn:fail:filesystem:version?, 544
exn:fail:filesystem?, 544
exn:fail:network, 544
exn:fail:network?, 544
exn:fail:object, 448
exn:fail:object?, 448
exn:fail:out-of-memory, 544
exn:fail:out-of-memory?, 544
exn:fail:read, 543
exn:fail:read-srclocs, 543
exn:fail:read:eof, 544
exn:fail:read:eof?, 544
exn:fail:read:non-char, 544
exn:fail:read:non-char?, 544
exn:fail:read?, 543
exn:fail:resource-resource, 869
exn:fail:resource?, 869
exn:fail:sandbox-terminated-reason, 856
exn:fail:sandbox-terminated?, 856
exn:fail:syntax, 543
exn:fail:syntax-exprs, 543
exn:fail:syntax:unbound, 543
exn:fail:syntax:unbound?, 543
exn:fail:syntax?, 543
exn:fail:unsupported, 545
exn:fail:unsupported?, 545
exn:fail:user, 545
exn:fail:user?, 545
exn:fail?, 542
exn:misc:match?, 525
exn:srclocs-accessor, 546
exn:srclocs?, 546
exn?, 542
exp, 170
expand, 684
expand, 31
expand-export, 673
expand-import, 670
expand-once, 685
expand-syntax, 685
expand-syntax-once, 685
expand-syntax-to-top-form, 686
expand-to-top-form, 685
expand-user-path, 877
Expanding Top-Level Forms, 684
Expansion, 34
Expansion (Parsing), 34
Expansion Context, 37

Expansion Steps, 35
explode-path, 880
export, 675
export, 454
export-local-id, 675
export-mode, 675
export-orig-stx, 675
export-out-sym, 675
export-protect?, 675
export?, 675
expression context, 37
Expression Wrapper: #%expression, 101
expt, 169
extend, 398
Extending match, 525
Extending the Syntax of Signatures, 461
extends, 454
extension-load handler, 810
externalizable<%>, 440
Extra Constants and Functions, 182
extract-struct-info, 396
false, 150
false/c, 468
false?, 150
fasl->s-exp, 797
Fast-Load Serialization, 797
fcontrol, 557
field, 404
Field and Method Access, 424
field-bound?, 427
field-names, 447
Fields, 427
Fields, 416
fifth, 283
file, 84
File Inclusion, 687
File Ports, 701
file->bytes, 897
file->bytes-lines, 898
file->lines, 898
file->list, 898
file->string, 897
file->value, 898
file-exists?, 889
file-name-from-path, 880
file-or-directory-identity, 892
file-or-directory-modify-seconds, 890
file-or-directory-permissions, 891
file-position, 699
file-size, 892
file-stream port, 701
file-stream-buffer-mode, 698
file-stream-port?, 697
filename-extension, 880
Files, 889
Filesystem, 886
filesystem-root-list, 894
filter, 269
filter-map, 289
filter-not, 290
filter-read-input-port, 742
filtered-in, 97
filtered-out, 100
finalizers, 947
find-executable-path, 888
find-files, 900
find-library-collection-paths, 966
find-relative-path, 880
find-seconds, 935
find-system-path, 886
findf, 274
first, 282
fixnum, 153
Fixnum Arithmetic, 191
Fixnum Vectors, 193
fixnum?, 156
Fixnums, 190
fl*, 186
fl+, 186
fl-, 186
fl->exact-integer, 188
fl->fx, 192
fl/, 186

fl<, 186
fl<=, 186
fl=, 186
fl>, 186
fl>=, 186
flabs, 186
flacos, 187
flasin, 187
flat contract, 465
flat contract property, 507
flat-contract, 465
flat-contract-predicate, 511
flat-contract-property?, 508
flat-contract?, 511
flat-murec-contract, 475
flat-named-contract, 465
flat-rec-contract, 475
flatan, 187
flatten, 288
flceiling, 187
flcos, 187
flexp, 187
flexpt, 187
flfloor, 187
flimag-part, 188
fllog, 187
flmax, 186
flmin, 186
floating-point-bytes->real, 181
Flonum Arithmetic, 186
Flonum Vectors, 188
flonum?, 156
Flonums, 185
flonums, 152
floor, 165
flreal-part, 188
flround, 186
flsin, 187
flsqrt, 187
fltan, 187
fltruncate, 187
flush-output, 698
flvector, 188
flvector, 188
flvector-copy, 189
flvector-length, 189
flvector-ref, 189
flvector-set!, 189
flvector?, 188
fold-files, 901
foldl, 268
foldr, 269
for, 128
for*, 132
for*/and, 133
for*/first, 133
for*/flvector, 190
for*/fold, 133
for*/fold/derived, 134
for*/fxvector, 194
for*/hash, 133
for*/hasheq, 133
for*/hasheqv, 133
for*/last, 133
for*/list, 133
for*/lists, 133
for*/or, 133
for*/product, 133
for*/set, 358
for*/seteq, 358
for*/seteqv, 358
for*/sum, 133
for*/vector, 133
for-clause-syntax-protect, 137
for-each, 268
for-label, 93
for-meta, 93
for-syntax, 93
for-template, 93
for/and, 130
for/first, 131
for/flvector, 190
for/fold, 132
for/fold/derived, 133

for/fxvector, 194
for/hash, 130
for/hasheq, 130
for/hasheqv, 130
for/last, 132
for/list, 129
for/lists, 131
for/or, 130
for/product, 131
for/set, 358
for/seteq, 358
for/seteqv, 358
for/sum, 131
for/vector, 129
force, 548
format, 768
format-id, 688
format-symbol, 689
fourth, 283
fprintf, 767
'framework, 936
free-identifier=?, 651
free-label-identifier=?, 652
free-template-identifier=?, 652
free-transformer-identifier=?, 652
fsemaphore-count, 595
fsemaphore-post, 595
fsemaphore-try-wait?, 595
fsemaphore-wait, 595
fsemaphore?, 595
Fully Expanded Programs, 34
fully-expanded, 35
function contract, 476
Function Contracts, 476
future, 593
future, 593
Future Creation Tree, 597
Future Performance Logging, 600
future semaphore, 594
future-event, 600
future-event-action, 600
future-event-future-id, 600
future-event-prim-name, 600
future-event-proc-id, 600
future-event-time-id, 600
future-event-user-data, 600
future-event?, 600
future-visualizer, 595
future-visualizer/trace, 598
future?, 594
Futures, 592
futures trace, 598
Futures Tracing, 598
Futures Visualizer, 595
futures visualizer, 595
futures-enabled?, 593
fx*, 191
fx+, 191
fx-, 191
fx->fl, 192
fx<, 192
fx<=, 192
fx=, 192
fx>, 192
fx>=, 192
fxabs, 191
fxand, 191
fxior, 191
fxlshift, 192
fxmax, 192
fxmin, 192
fxmodulo, 191
fxnot, 192
fxquotient, 191
fxremainder, 191
fxrshift, 192
fxvector, 193
fxvector, 193
fxvector-copy, 194
fxvector-length, 193
fxvector-ref, 193
fxvector-set!, 194
fxvector?, 193
fxxor, 191

Garbage Collection, 949
Garbage Collection, 23
garbage collection, 23
'gc, 936
gc-info, 949
gcd, 164
gen:custom-write, 787
gen:dict, 350
gen:equal+hash, 148
gen:stream, 332
generate-member-key, 421
generate-temporaries, 650
generate-temporary, 692
Generating A Unit from Context, 459
generator, 333
generator, 333
generator-state, 336
generator?, 333
Generators, 333
generic, 428
generic, 428
generic interface, 386
Generic Interfaces, 386
Generic Numerics, 159
generic?, 442
Generics, 428
gensym, 235
get, 627
get-error-output, 867
get-field, 427
get-first-place, 619
get-first-place-channel, 619
get-log-prefix, 619
get-output, 867
get-output-bytes, 709
get-output-string, 709
get-pid, 616
get-preference, 902
get-sc-id, 617
get-uncovered-expressions, 867
get-user-custodian, 866
getenv, 935
global port print handler, 772
global-port-print-handler, 772
gracket, 959
GRacket.app, 959
GRacket.exe, 959
greatest common divisor, 164
group-execute-bit, 907
group-read-bit, 907
group-write-bit, 907
guard-evt, 579
Guarded Evaluation: when and unless, 125
gui?, 869
handle-evt, 578
handle-evt?, 579
Handling Exceptions, 538
has-contract?, 511
hash, 309
hash, 307
hash table, 307
Hash Tables, 307
hash->list, 313
hash-copy, 315
hash-count, 314
hash-eq?, 308
hash-equal?, 308
hash-eqv?, 308
hash-for-each, 314
hash-has-key?, 312
hash-iterate-first, 314
hash-iterate-key, 314
hash-iterate-next, 314
hash-iterate-value, 315
hash-keys, 313
hash-map, 313
hash-placeholder?, 293
hash-ref, 311
hash-ref!, 311
hash-remove, 313
hash-remove!, 312
hash-set, 310
hash-set!, 310
hash-set*, 311

hash-set*!, 310
hash-update, 312
hash-update!, 312
hash-values, 313
hash-weak?, 309
hash/c, 473
hash?, 308
hasheq, 309
hasheqv, 309
help, 969
here strings, 55
heredoc, 54
HOME, 886
'home-dir, 886
HOMEDRIVE, 886
HOMEPATH, 886
identifier, 32
identifier-binding, 653
identifier-label-binding, 654
identifier-prune-lexical-context, 650
identifier-prune-to-source-module, 650
identifier-remove-from-definition-context, 662
identifier-template-binding, 654
identifier-transformer-binding, 654
identifier?, 645
Identifiers and Binding, 32
identity, 370
IEEE floating-point numbers, 152
if, 115
imag-part, 173
Immutable Cyclic Data, 292
immutable?, 147
impersonate-box, 831
impersonate-continuation-mark-key, 834
impersonate-hash, 832
impersonate-procedure, 829
impersonate-prompt-tag, 833
impersonate-struct, 830
impersonate-vector, 831
impersonator, 826
Impersonator Constructors, 829
Impersonator Properties, 841
impersonator properties, 827
impersonator property accessor, 841
impersonator property descriptor, 841
impersonator property predicate, 841
impersonator-contract?, 511
impersonator-of?, 828
impersonator-prop:application-mark, 841
impersonator-prop:contracted, 505
impersonator-property-accessor-procedure?, 841
impersonator-property?, 841
impersonator?, 828
Impersonators and Chaperones, 826
implementation?, 445
implementation?/c, 438
implements, 397
implies, 151
import, 454
import, 671
import-local-id, 671
import-mode, 671
import-orig-mode, 671
import-orig-stx, 671
import-req-mode, 671
import-source, 671
import-source-mod-path-stx, 671
import-source-mode, 671
import-source?, 671
import-src-mod-path, 671
import-src-sym, 671
import?, 671
Importing and Exporting: require and provide, 75
in-bytes, 319
in-bytes-lines, 320
in-cycle, 323
in-dict, 348

in-dict-keys, 348
in-dict-pairs, 349
in-dict-values, 348
in-directory, 322
in-flvector, 189
in-fxvector, 194
in-generator, 335
in-hash, 320
in-hash-keys, 321
in-hash-pairs, 321
in-hash-values, 321
in-indexed, 322
in-input-port-bytes, 320
in-input-port-chars, 320
in-lines, 320
in-list, 318
in-mlist, 318
in-naturals, 318
in-parallel, 323
in-port, 320
in-producer, 322
in-range, 318
in-sequences, 322
in-set, 357
in-stream, 330
in-string, 319
in-value, 322
in-values*-sequence, 323
in-values-sequence, 323
in-vector, 319
include, 687
include-at/relative-to, 687
include-at/relative-to/reader, 688
include/reader, 687
indexed-future-event, 599
indexed-future-event-event, 599
indexed-future-event-index, 599
indexed-future-event?, 599
inexact number, 152
inexact->exact, 158
inexact-real?, 155
inexact?, 158
Inferred Linking, 456
Inferred Value Names, 47
'inferred-name, 47
infinite-generator, 335
infinite?, 185
infinity, 152
infix, 53
Information on Expanded Modules, 686
inherit, 410
inherit-field, 404
inherit/inner, 411
inherit/super, 411
inheritance, 397
Inherited and Superclass Methods, 418
init, 403
Init Libraries, 960
init-depend, 454
'init-dir, 887
init-field, 403
'init-file, 887
init-rest, 404
Initialization, 959
Initialization Variables, 414
initiate, 317
inner, 418
inode, 708
inode, 892
Input and Output, 694
input ports, pattern matching, 235
input-port-append, 734
input-port?, 696
inspect, 403
inspector, 847
inspector?, 847
instanceof/c, 436
instantiate, 423
instantiates, 26
instantiation, 26
integer->char, 226
integer->integer-bytes, 180
integer-bytes->integer, 180
integer-in, 468

integer-length, 177
integer-sqrt, 169
integer-sqrt/remainder, 169
integer?, 154
integers, 152
Interacting with Evaluators, 865
interaction port handler, 813
Interaction Wrapper: #%top-interaction, 140
Interactive Help, 969
Interactive Module Loading, 970
interface, 398
interface, 397
interface*, 399
interface->method-names, 446
interface-extension?, 445
interface?, 442
Internal and External Names, 419
Internal Definitions, 40
internal-definition context, 37
Internal-Definition Limiting: #%stratified-body, 144
internal-definition-context-apply, 692
internal-definition-context-seal, 662
internal-definition-context?, 661
interned, 48
Introducing Bindings, 37
invoke-unit, 454
invoke-unit/infer, 458
invoked, 450
Invoking Units, 454
is-a?, 444
is-a?/c, 438
Iteration and Comprehension Forms, 128
Iterations and Comprehensions: for, for/list, ..., 127
iterator, 316
JIT, 816
julian/scalinger->string, 935
keyword, 260
keyword->string, 261
keyword-apply, 360
Keyword-Argument Conversion Introspection, 675
keyword<?, 261
keyword?, 261
Keywords, 260
Keywords and Arity, 360
kill-evaluator, 866
kill-thread, 571
label phase level, 32
lambda, 104
LANG, 937
Language Model, 18
Language Run-Time Configuration, 964
last, 285
last-pair, 285
launch-place, 619
lazy, 547
Lazy Data-structure Contracts, 485
LC_ALL, 937
LC_TYPE, 937
lcm, 164
least common multiple, 164
Legacy Contracts, 513
length, 265
let, 108
let*, 109
let*-values, 110
let-syntax, 111
let-syntaxes, 111
let-values, 110
let/cc, 553
let/ec, 553
letrec, 109
letrec-syntax, 111
letrec-syntaxes, 111
letrec-syntaxes+values, 111
letrec-values, 110
'lexical, 653
lexical information, 33
lexical scoping, 26

lib, 82
liberal expansion, 670
liberal-define-context?, 669
Libraries and Collections, 965
library, 965
Library Extensions, 527
Limiting Scope: define-package, open-package, ..., 140
'line, 698
line locations, 699
line numbers, 699
Line-Output Hook, 776
'linefeed, 748
link, 454
'link, 936
link-exists?, 889
linked, 450
Linking Units and Creating Compound Units, 455
'links-file, 887
list, 264
list, 261
List Filtering, 269
List Iteration, 266
List Operations, 265
List Searching, 272
list*, 264
list->bytes, 214
list->mlist, 295
list->set, 358
list->seteq, 358
list->seteqv, 358
list->string, 199
list->vector, 299
list-ref, 265
list-tail, 265
list/c, 471
list?, 263
listof, 471
Literals: quote and #%datum, 100
'll, 231
'lm, 231
'lo, 231
load, 810
load handler, 808
load-extension, 810
load-on-demand-enabled, 816
load-relative, 810
load-relative-extension, 810
load/cd, 810
load/use-compiled, 812
local, 112
Local Binding with Splicing Body, 677
Local Binding: let, let*, letrec, ..., 108
local bindings, 32
Local Definitions: local, 112
local variable, 26
local-expand, 658
local-expand/capture-lifts, 660
local-require, 85
local-transformer-expand, 660
local-transformer-expand/capture-lifts, 660
locale, 695
Locale-Specific String Operations, 205
locale-string-encoding, 224
Locating Paths, 886
location, 24
Locations: #%variable-reference, 102
log, 170
log receiver, 928
log-debug, 930
log-error, 930
log-fatal, 930
log-info, 930
log-level?, 929
log-max-level, 930
log-message, 616
log-message, 929
log-receiver?, 930
log-to-parent, 630
log-warning, 930
logger, 928
logger-name, 929

logger?, 928
Logging, 928
Logging Events, 929
logical operators, 174
LOGNAME, 886
Low-level Contract Boundaries, 492
"LPT1", 883
"LPT2", 883
"LPT3", 883
"LPT4", 883
"LPT5", 883
"LPT6", 883
"LPT7", 883
"LPT8", 883
"LPT9", 883
'lt, 231
'lu, 231
'machine, 936
'macosx, 936
macro, 39
Macro-Introduced Bindings, 42
Macros, 635
magnitude, 173
make-arity-at-least, 365
make-async-bi-channel, 630
make-async-channel, 585
make-base-empty-namespace, 799
make-base-namespace, 799
make-bytes, 210
make-channel, 581
make-chaperone-contract, 496
make-continuation-mark-key, 562
make-continuation-prompt-tag, 551
make-contract, 496
make-custodian, 844
make-custodian-box, 846
make-custom-hash, 352
make-date, 931
make-date*, 932
make-derived-parameter, 591
make-deserialize-info, 796
make-directory, 893
make-directory*, 902
make-do-sequence, 324
make-empty-namespace, 799
make-ephemeron, 947
make-evaluator, 851
make-exn, 542
make-exn:break, 545
make-exn:fail, 542
make-exn:fail:contract, 542
make-exn:fail:contract:arity, 542
make-exn:fail:contract:blame, 503
make-exn:fail:contract:continuation, 543
make-exn:fail:contract:divide-by-zero, 542
make-exn:fail:contract:non-fixnum-result, 543
make-exn:fail:contract:variable, 543
make-exn:fail:filesystem, 544
make-exn:fail:filesystem:exists, 544
make-exn:fail:filesystem:version, 544
make-exn:fail:network, 544
make-exn:fail:object, 448
make-exn:fail:out-of-memory, 544
make-exn:fail:read, 543
make-exn:fail:read:eof, 544
make-exn:fail:read:non-char, 544
make-exn:fail:syntax, 543
make-exn:fail:syntax:unbound, 543
make-exn:fail:unsupported, 545
make-exn:fail:user, 545
make-export, 675
make-file-or-directory-link, 892
make-flat-contract, 497
make-flrectangular, 188
make-flvector, 189
make-fsemaphore, 594
make-future-event, 600
make-fxvector, 193
make-generic, 428

make-handle-get-preference-locked, 905
make-hash, 309
make-hash-placeholder, 293
make-hasheq, 309
make-hasheq-placeholder, 293
make-hasheqv, 309
make-hasheqv-placeholder, 293
make-immutable-custom-hash, 352
make-immutable-hash, 310
make-immutable-hasheq, 310
make-immutable-hasheqv, 310
make-impersonator-property, 841
make-import, 671
make-import-source, 671
make-indexed-future-event, 599
make-input-port, 711
make-input-port/read-to-peek, 735
make-inspector, 847
make-keyword-procedure, 363
make-known-char-range-list, 231
make-limited-input-port, 736
make-list, 285
make-lock-file-name, 906
make-log-receiver, 930
make-logger, 929
make-mixin-contract, 438
make-module-evaluator, 852
make-none/c, 512
make-object, 423
make-output-port, 724
make-parameter, 589
make-parameter-rename-transformer, 677
make-pipe, 710
make-pipe-with-specials, 737
make-placeholder, 293
make-polar, 173
make-prefab-struct, 392
make-proj-contract, 513
make-provide-pre-transformer, 674
make-provide-transformer, 674
make-pseudo-random-generator, 178
make-reader-graph, 292
make-readtable, 781
make-rectangular, 172
make-rename-transformer, 656
make-require-transformer, 670
make-resolved-module-path, 817
make-security-guard, 842
make-semaphore, 583
make-serialize-info, 796
make-set!-transformer, 655
make-shared-bytes, 214
make-shared-flvector, 190
make-shared-fxvector, 195
make-sibling-inspector, 848
make-special-comment, 787
make-srcloc, 547
make-string, 196
make-struct-field-accessor, 383
make-struct-field-mutator, 383
make-struct-info, 395
make-struct-type, 380
make-struct-type-property, 384
make-syntax-delta-introducer, 668
make-syntax-introducer, 667
make-temporary-file, 902
make-tentative-pretty-print-output-port, 779
make-thread-cell, 587
make-thread-group, 847
make-vector, 298
make-weak-box, 946
make-weak-custom-hash, 352
make-weak-hash, 309
make-weak-hasheq, 310
make-weak-hasheqv, 310
make-will-executor, 948
Managing Ports, 696
Manipulating Paths, 871
map, 266
mappend, 296
mappend!, 296

massoc, 297
massq, 297
massv, 297
match, 516
match expander, 525
match*, 523
match*/derived, 527
match-define, 524
match-define-values, 525
match-equality-test, 527
match-lambda, 523
match-lambda*, 523
match-lambda**, 524
match-let, 524
match-let*, 524
match-let*-values, 524
match-let-values, 524
match-letrec, 524
match/derived, 527
match/values, 523
matching-identifiers-in, 96
matching-identifiers-out, 100
max, 163
'mc, 231
mcar, 294
mcdr, 294
mcons, 294
'me, 231
member, 273
member-name-key, 421
member-name-key-hash-code, 421
member-name-key=?, 421
member-name-key?, 421
memf, 273
Memory Management, 946
memq, 273
memv, 273
merge-input, 737
message router, 607
message-router, 609
Method Definitions, 417
'method-arity-error, 360
'method-arity-error, 107
method-in-interface?, 445
Methods, 417
Methods, 424
mfor-each, 297
min, 163
Miscellaneous, 869
Miscellaneous utilities, 692
mixin, 428
mixin, 429
mixin-contract, 438
Mixins, 428
mlength, 295
mlist, 295
mlist->list, 295
mlist-ref, 296
mlist-tail, 296
mlist?, 295
mlistof, 297
mmap, 296
mmember, 297
mmemq, 297
mmemv, 297
'mn, 231
module, 71
module binding, 32
module context, 37
module name resolver, 818
Module Names and Loading, 817
module path, 817
module path index, 820
module path resolver, 818
Module Phases and Visits, 41
Module Re-declarations, 27
module registry, 45
module*, 74
module+, 75
module->exports, 826
module->imports, 826
module->language-info, 825
module->namespace, 804
module-begin context, 37

module-compiled-exports, 823
module-compiled-imports, 823
module-compiled-language-info, 824
module-compiled-name, 822
module-compiled-submodules, 822
module-declared?, 825
'module-direct-for-syntax-requires, 686
'module-direct-for-template-requires, 686
'module-direct-requires, 686
'module-indirect-provides, 687
'module-language, 824
'module-language, 74
module-level variable, 26
module-path-index-join, 822
module-path-index-resolve, 821
module-path-index-split, 821
module-path-index-submodule, 822
module-path-index?, 821
module-path?, 818
module-predefined?, 826
module-provide-protected?, 805
'module-syntax-provides, 686
'module-variable-provides, 686
modules, re-define, 27
modules, imports, 75
modules, exports, 75
Modules and Module-Level Variables, 26
Modules: module, module*, ..., 71
modulo, 162
More File and Directory Utilities, 897
More Path Utilities, 879
More Port Constructors, Procedures, and Events, 731
mpair?, 294
mr-connect-to, 626
mr-spawn-remote-node, 625
mr-supervise-named-dynamic-place-at, 626
mreverse, 296
mreverse!, 296
multi-in, 99
Multiple Return Values, 19
Multiple Values, 529
multiple values, 19
'must-truncate, 703
mutable list, 294
Mutable List Functions, 295
mutable pair, 294
Mutable Pair Constructors and Selectors, 294
Mutable Pairs and Lists, 294
mutator, 376
nack-guard-evt, 579
named-place-typed-channel%, 627
namespace, 45
namespace-anchor->empty-namespace, 800
namespace-anchor->namespace, 800
namespace-anchor?, 799
namespace-attach-module, 803
namespace-attach-module-declaration, 803
namespace-base-phase, 800
namespace-mapped-symbols, 802
namespace-module-identifier, 800
namespace-module-registry, 804
namespace-require, 802
namespace-require/constant, 802
namespace-require/copy, 802
namespace-require/expansion-time, 802
namespace-set-variable-value!, 801
namespace-symbol->identifier, 800
namespace-syntax-introduce, 804
namespace-undefine-variable!, 802
namespace-unprotect-module, 804
namespace-variable-value, 801
namespace?, 799
Namespaces, 799
Namespaces, 45
nan?, 185
nand, 151

natural-number/c, 468
'nd, 231
negate, 371
negative?, 157
Nested Contract Boundaries, 490
Networking, 907
never-evt, 579
new, 423
new-prompt, 561
new-∀/c, 484
new-∃/c, 484
newline, 756
ninth, 284
'nl, 231
'no, 231
node%, 618
node-get-first-place, 620
node-send-exit, 620
nodes, 607
'nominal-id, 657
non-empty-listof, 471
'non-terminating-macro, 781
'non-terminating-macro, 782
'none, 682
'none, 698
none/c, 466
nor, 151
normal-case-path, 878
normalize-path, 880
not, 145
not-a-number, 152
'not-free-identifier=?, 657
'not-free-identifier=?, 657
'not-provide-all-defined, 657
not/c, 467
"NUL", 883
null, 263
null?, 262
Number Comparison, 167
Number Types, 153
Number–String Conversions, 179
number->string, 179
number?, 153
numbers, parsing, 51
numbers, machine representations, 179
numbers, little-endian, 179
numbers, floating-point, 179
numbers, converting, 179
numbers, big-endian, 179
Numbers, 152
numbers, 152
numerator, 166
Object and Class Contracts, 432
Object Equality and Hashing, 438
Object Identity and Comparisons, 23
Object Printing, 441
Object Serialization, 439
object%, 399
Object, Class, and Interface Utilities, 441
object->vector, 443
object-contract, 436
object-info, 447
object-interface, 444
object-method-arity-includes?, 446
object-name, 849
object/c, 435
object=?, 443
object?, 441
objects, 21
Objects and Imperative Update, 21
Obligation Information in Check Syntax, 508
odd?, 157
one-of/c, 468
only, 453
only-in, 77
only-meta-in, 79
'opaque, 682
open, 453
open*-package, 142
open-input-bytes, 708
open-input-file, 701
open-input-output-file, 704
open-input-string, 708

open-output-bytes, 709
open-output-file, 702
open-output-nowhere, 737
open-output-string, 709
open-package, 140
Operating System, 871
opt/c, 512
or, 117
or/c, 466
order-of-magnitude, 185
'orig-dir, 888
'origin, 679
ormap, 267
'os, 936
other-execute-bit, 907
other-read-bit, 907
other-write-bit, 907
output-port?, 696
overment, 407
overment*, 412
override, 406
override*, 412
override-final, 408
override-final*, 412
overriding, 397
package-begin, 141
package-exported-identifiers, 142
package-original-identifiers, 143
package?, 142
pair, 261
Pair Accessor Shorthands, 275
Pair Constructors and Selectors, 262
pair?, 262
Pairs and Lists, 261
parameter procedure, 29
parameter-procedure=?, 592
parameter/c, 473
parameter?, 592
parameterization, 30
parameterization?, 592
parameterize, 590
parameterize*, 591
parameterize-break, 567
Parameters, 29
Parameters, 589
Parameters, 29
Parametric Contracts, 483
parametric->/c, 483
'paren-shape, 53
parse, 35
parse-command-line, 943
parsed, 35
Partial Expansion, 40
partial expansion, 40
partition, 290
PATH, 889
path, 871
path element, 881
path->bytes, 873
path->complete-path, 876
path->directory-path, 876
path->string, 872
path-add-suffix, 879
path-convention-type, 874
path-element->bytes, 874
path-element->string, 873
path-element?, 881
path-for-some-system?, 872
path-list-string->path-list, 888
path-only, 881
path-replace-suffix, 879
path-string?, 871
path-up, 98
path?, 871
pathlist-closure, 900
Paths, 871
Pattern Matching, 516
pattern matching, 235
pattern variable, 636
Pattern variables, 689
Pattern-Based Syntax Matching, 635
'pc, 231
'pd, 231
'pe, 231

peek-byte, 754
peek-byte-or-special, 754
peek-bytes, 751
peek-bytes!, 752
peek-bytes!-evt, 745
peek-bytes-avail!, 752
peek-bytes-avail!*, 753
peek-bytes-avail!-evt, 746
peek-bytes-avail!/enable-break, 753
peek-bytes-evt, 745
peek-char, 754
peek-char-or-special, 754
peek-string, 751
peek-string!, 751
peek-string!-evt, 746
peek-string-evt, 746
peeking-input-port, 738
Per-Symbol Special Printing, 775
Performance Hints: begin-encourage-inline, 144
'pf, 231
phase level, 32
Phases, 26
phases, 26
pi, 182
'pi, 231
pi.f, 182
pipe, 710
pipe-content-length, 710
Pipes, 710
place, 605
place, 602
place channels, 602
place descriptor, 603
place%, 621
place*, 605
place-break, 606
place-channel, 606
place-channel-get, 606
place-channel-put, 606
place-channel-put/get, 606
place-channel?, 603
place-dead-evt, 605
place-enabled?, 603
place-kill, 606
place-message-allowed?, 607
place-socket-bridge%, 617
place-wait, 605
place?, 603
placeholder-get, 293
placeholder-set!, 293
placeholder?, 293
Places, 602
planet, 84
PLT_DELAY_FROM_ZO, 764
PLT_VALIDATE_COMPILE, 814
PLTADDONDIR, 887
PLTCOLLECTS, 967
PLTDISABLEGC, 949
PLTLINKSFILE, 887
PLTNOMZJIT, 816
PLTSTDERR, 928
PLTSYSLOG, 928
'po, 231
poll, 574
poll-guard-evt, 579
Port Buffers and Positions, 697
port display handler, 771
Port Events, 744
port positions, 699
port print handler, 771
port read handler, 765
Port String and List Conversions, 731
port write handler, 771
port->bytes, 732
port->bytes-lines, 732
port->lines, 732
port->list, 731
port->string, 731
port-closed-evt, 696
port-closed?, 696
port-commit-peeked, 755
port-count-lines!, 700
port-count-lines-enabled, 701

port-display-handler, 771
port-file-identity, 707
port-file-unlock, 707
port-next-location, 700
port-no?, 626
port-print-handler, 771
port-progress-evt, 754
port-provides-progress-evts?, 754
port-read-handler, 765
port-try-file-lock?, 707
port-write-handler, 771
port-writes-atomic?, 759
port-writes-special?, 759
port?, 696
ports, flushing, 698
Ports, 694
Ports, 694
position, 699
positive?, 156
Powers and Roots, 169
pre-expand-export, 673
predicate, 376
predicate/c, 482
'pref-dir, 886
'pref-file, 887
prefab, 374
prefab-key->struct-type, 393
prefab-key?, 394
prefab-struct-key, 392
preferences-lock-file-mode, 904
prefix, 454
prefix-in, 78
prefix-out, 89
pregexp, 244
pregexp?, 243
preserved, 29
Pretty Printing, 772
pretty-display, 773
pretty-format, 773
pretty-print, 772
pretty-print-.-symbol-without-bars, 774
pretty-print-abbreviate-read-macros, 775
pretty-print-columns, 774
pretty-print-current-style-table, 775
pretty-print-depth, 774
pretty-print-exact-as-decimal, 774
pretty-print-extend-style-table, 775
pretty-print-handler, 773
pretty-print-newline, 776
pretty-print-post-print-hook, 778
pretty-print-pre-print-hook, 778
pretty-print-print-hook, 778
pretty-print-print-line, 776
pretty-print-remap-stylable, 776
pretty-print-show-inexactness, 774
pretty-print-size-hook, 777
pretty-print-style-table?, 775
pretty-printing, 779
pretty-write, 773
primitive-closure?, 369
primitive-result-arity, 369
primitive?, 369
print, 766
print handler, 814
print-as-expression, 770
print-boolean-long-form, 770
print-box, 769
print-graph, 769
print-hash-table, 770
print-mpair-curly-braces, 769
print-pair-curly-braces, 769
print-reader-abbreviations, 770
print-struct, 769
print-syntax-width, 770
print-unreadable, 769
print-vector-length, 770
printable/c, 468
printable<%>, 441
Printer Extension, 787
printf, 768

printf/f, 627
Printing Booleans, 64
Printing Boxes, 68
Printing Characters, 68
Printing Compiled Code, 69
Printing Hash Tables, 67
Printing Keywords, 68
Printing Numbers, 63
Printing Pairs and Lists, 64
Printing Paths, 69
Printing Regular Expressions, 69
Printing Strings, 65
Printing Structures, 66
Printing Symbols, 63
Printing Unreadable Values, 69
Printing Vectors, 66
println, 767
private, 409
private*, 412
"PRN", 883
Procedure Applications and #%app, 103
Procedure Applications and Local Variables, 24
Procedure Expressions: lambda and case-lambda, 104
procedure->method, 360
procedure-arity, 361
procedure-arity-includes/c, 473
procedure-arity-includes?, 362
procedure-arity?, 361
procedure-extract-target, 367
procedure-keywords, 363
procedure-reduce-arity, 362
procedure-reduce-keyword-arity, 364
procedure-rename, 359
procedure-struct-type?, 367
procedure?, 358
Procedures, 358
process, 925
process*, 926
process*/ports, 927
process/ports, 926
Processes, 919
processes, 596
processor-count, 594
progress-evt?, 755
promise, 547
promise-forced?, 548
promise-running?, 548
promise/c, 476
promise?, 547
prompt, 558
prompt, 28
prompt read handler, 813
prompt tag, 28
prompt-at, 558
prompt-tag/c, 474
prompt0, 559
prompt0-at, 560
Prompts, Delimited Continuations, and Barriers, 28
prop:arity-string, 368
prop:chaperone-contract, 505
prop:checked-procedure, 368
prop:contract, 505
prop:contracted, 505
prop:custom-print-quotable, 789
prop:custom-write, 788
prop:dict, 351
prop:dict/contract, 351
prop:equal+hash, 150
prop:evt, 580
prop:exn:srclocs, 545
prop:flat-contract, 505
prop:impersonator-of, 834
prop:input-port, 710
prop:legacy-match-expander, 527
prop:liberal-define-context, 669
prop:match-expander, 526
prop:output-port, 711
prop:procedure, 365
prop:provide-pre-transformer, 674
prop:provide-transformer, 674
prop:rename-transformer, 657

prop:require-transformer, 670
prop:sequence, 325
prop:serializable, 796
prop:set!-transformer, 656
prop:stream, 333
prop:struct-auto-info, 396
prop:struct-info, 396
proper-subset?, 356
property accessor, 384
property predicate, 384
protect-out, 90
'protected, 680
provide, 85
provide Macros, 123
provide pre-transformer, 673
provide transformer, 673
provide Transformers, 673
provide-pre-transformer?, 674
provide-signature-elements, 461
provide-transformer?, 674
provide/contract, 490
proxy design pattern, 448
'ps, 231
pseudo-random-generator->vector, 178
pseudo-random-generator-vector?, 179
pseudo-random-generator?, 178
public, 405
public*, 411
public-final, 406
public-final*, 412
pubment, 405
pubment*, 411
put-input, 867
put-preferences, 904
putenv, 935
quasiquote, 138
Quasiquoting: quasiquote, unquote, and unquote-splicing, 138
quasisyntax, 643
quasisyntax/loc, 643
quotable, 62
quote, 100
quote-syntax, 139
quote-syntax/prune, 644
quotient, 161
quotient/remainder, 162
quoting depth, 62
racket, 1
racket, 959
racket-path, 624
Racket.exe, 959
racket/async-channel, 584
racket/base, 1
racket/block, 143
racket/bool, 150
racket/bytes, 224
racket/class, 397
racket/cmdline, 939
racket/contract, 465
racket/contract/base, 513
racket/contract/base, 513
racket/contract/combinator, 493
racket/contract/parametric, 483
racket/contract/region, 490
racket/control, 556
racket/date, 933
racket/dict, 337
racket/enter, 970
racket/fasl, 797
racket/file, 897
racket/fixnum, 190
racket/flonum, 185
racket/function, 370
racket/future, 592
racket/generator, 333
racket/generic, 386
racket/gui/init, 960
racket/help, 969
racket/include, 687
racket/init, 960
racket/kernel, 960
racket/keyword-transform, 675

racket/language-info, 960
racket/list, 282
racket/load, 816
racket/local, 112
racket/match, 516
racket/math, 182
racket/mpair, 295
racket/package, 140
racket/path, 879
racket/performance-hint, 144
racket/place, 602
racket/place/define-remote-server, 628
racket/place/distributed, 607
racket/place/distributed/rmpi, 631
racket/place/private/async-bi-channel, 630
racket/port, 731
racket/pretty, 772
racket/promise, 547
racket/provide, 100
racket/provide-syntax, 123
racket/provide-transform, 673
racket/require, 96
racket/require-syntax, 123
racket/require-transform, 670
racket/runtime-config, 960
racket/runtime-path, 894
racket/sandbox, 851
racket/sequence, 326
racket/serialize, 789
racket/set, 353
racket/shared, 112
racket/signature, 463
racket/splicing, 677
racket/stream, 329
racket/string, 206
racket/struct-info, 394
racket/stxparam, 676
racket/stxparam-exptime, 677
racket/surrogate, 448
racket/syntax, 688
racket/system, 924
racket/tcp, 907
racket/trace, 972
racket/trait, 429
racket/udp, 912
racket/unit, 450
racket/unit-exptime, 463
racket/unsafe/ops, 951
racket/vector, 300
"racketrc.rktl", 887
radians->degrees, 183
raise, 531
raise-argument-error, 533
raise-arguments-error, 534
raise-arity-error, 536
raise-blame-error, 503
raise-contract-error, 514
raise-mismatch-error, 536
raise-range-error, 535
raise-result-error, 534
raise-syntax-error, 537
raise-type-error, 536
raise-user-error, 532
Raising Exceptions, 531
random, 177
Random generation, 515
Random Numbers, 177
random-seed, 177
range, 291
rational numbers, 152
rational?, 154
rationalize, 166
read, 759
'read, 843
'read, 891
read, 31
read interaction handler, 813
read-accept-bar-quote, 763
read-accept-box, 762
read-accept-compiled, 763
read-accept-dot, 763
read-accept-graph, 763

read-accept-infix-dot, 763
read-accept-lang, 764
read-accept-quasiquote, 763
read-accept-reader, 764
read-byte, 748
read-byte-or-special, 753
read-bytes, 749
read-bytes!, 749
read-bytes!-evt, 744
read-bytes-avail!, 750
read-bytes-avail!*, 750
read-bytes-avail!-evt, 745
read-bytes-avail!/enable-break, 750
read-bytes-evt, 744
read-bytes-line, 748
read-bytes-line-evt, 745
read-case-sensitive, 762
read-char, 747
read-char-or-special, 753
read-curly-brace-as-paren, 762
read-decimal-as-inexact, 763
read-eval-print-loop, 812
read-language, 761
read-line, 748
read-line-evt, 745
read-on-demand-source, 764
read-square-bracket-as-paren, 762
read-string, 749
read-string!, 749
read-string!-evt, 745
read-string-evt, 745
read-syntax, 759
read-syntax/recursive, 760
read/recursive, 759
reader, 62
Reader Extension, 780
reader extension procedures, 780
reader language, 761
reader macro, 781
Reader-Extension Procedures, 786
Reading, 759
Reading Booleans, 52
Reading Boxes, 58
Reading Characters, 58
Reading Comments, 56
Reading Graph Structure, 60
Reading Hash Tables, 58
Reading Keywords, 59
Reading Numbers, 51
Reading Pairs and Lists, 52
Reading Quotes, 55
Reading Regular Expressions, 59
Reading Strings, 54
Reading Structures, 57
Reading Symbols, 50
Reading Vectors, 56
Reading via an Extension, 60
readtable, 780
readtable-mapping, 782
readtable?, 781
Readtables, 780
real numbers, 152
real->decimal-string, 180
real->double-flonum, 159
real->floating-point-bytes, 181
real->single-flonum, 159
real-in, 468
real-part, 173
real?, 154
rearm, 682
Receiving Logged Events, 930
record-disappeared-uses, 691
Recording disappeared uses, 691
recursive-contract, 512
redex, 18
reencode-input-port, 739
reencode-output-port, 739
Reflecting on Primitives, 369
Reflection and Security, 799
regexp, 243
Regexp Constructors, 243
Regexp Matching, 245
Regexp Splitting, 257
Regexp Substitution, 257

Regexp Syntax, 236
regexp value, 236
regexp-match, 245
regexp-match*, 247
regexp-match-evt, 746
regexp-match-exact?, 251
regexp-match-peek, 252
regexp-match-peek-immediate, 253
regexp-match-peek-positions, 252
regexp-match-peek-positions*, 254
regexp-match-peek-positions-immediate, 253
regexp-match-peek-positions-immediate/end, 256
regexp-match-peek-positions/end, 256
regexp-match-positions, 249
regexp-match-positions*, 250
regexp-match-positions/end, 255
regexp-match/end, 254
regexp-match?, 251
regexp-max-lookbehind, 245
regexp-quote, 244
regexp-replace, 258
regexp-replace*, 259
regexp-replace-quote, 260
regexp-replaces, 260
regexp-split, 257
regexp-try-match, 249
regexp?, 243
regexps, 235
register, 616
Regular Expressions, 235
Regular expressions, 235
'relative, 878
relative-in, 79
relative-path?, 876
relocate-input-port, 741
relocate-output-port, 741
remainder, 161
remote-connect, 620
remote-connection%, 620
remote-node%, 618
remove, 270
remove*, 271
remove-duplicates, 289
remq, 270
remq*, 271
remv, 270
remv*, 271
rename, 454
rename transformer, 39
rename-file-or-directory, 890
rename-in, 78
rename-inner, 411
rename-out, 88
rename-super, 411
rename-transformer-target, 657
rename-transformer?, 656
REPL, 812
'replace, 703
require, 75
require Macros, 123
require transformer, 670
require Transformers, 670
require-transformer?, 671
reset, 559
reset-at, 559
reset0, 559
reset0-at, 560
resolve-path, 877
resolved, 821
resolved module path, 817
resolved-module-path-name, 818
resolved-module-path?, 817
Resolving Module Names, 817
respawn-and-fire%, 622
rest, 283
restart-every, 615
restarter%, 623
'return, 748
'return-linefeed, 748
reverse, 266
rmpi-allreduce, 632

rmpi-barrier, 632
rmpi-broadcast, 632
rmpi-build-default-config, 633
rmpi-cnt, 631
rmpi-comm, 631
rmpi-comm-channels, 631
rmpi-comm-cnt, 631
rmpi-comm-id, 631
rmpi-comm?, 631
rmpi-finish, 633
rmpi-id, 631
rmpi-init, 632
rmpi-launch, 633
rmpi-partition, 633
rmpi-recv, 631
rmpi-reduce, 632
rmpi-send, 631
round, 164
RTT, 598
'run-file, 888
'running, 921
Running Racket, 959
Running Racket or GRacket, 959
runtime-paths, 897
s-exp, 61
s-exp->fasl, 797
S-Expression Reader Language, 61
'same, 874
sandbox-coverage-enabled, 859
sandbox-error-output, 858
sandbox-eval-handlers, 864
sandbox-eval-limits, 863
sandbox-exit-handler, 862
sandbox-gui-available, 860
sandbox-init-hook, 857
sandbox-input, 857
sandbox-make-code-inspector, 865
sandbox-make-inspector, 865
sandbox-make-logger, 865
sandbox-make-namespace, 860
sandbox-memory-limit, 862
sandbox-namespace-specs, 860
sandbox-network-guard, 862
sandbox-output, 858
sandbox-override-collection-paths, 861
sandbox-path-permissions, 861
sandbox-propagate-breaks, 859
sandbox-propagate-exceptions, 859
sandbox-reader, 857
sandbox-security-guard, 861
Sandboxed Evaluation, 851
'sc, 231
scalar value, 225
scope, 32
second, 283
seconds->date, 931
security guard, 842
Security Guards, 842
security-guard?, 842
select, 574
self, 402
semaphore, 582
semaphore-peek-evt, 583
semaphore-peek-evt?, 584
semaphore-post, 583
semaphore-try-wait?, 583
semaphore-wait, 583
semaphore-wait/enable-break, 583
semaphore?, 583
Semaphores, 582
send, 424
send*, 425
send+, 426
send-exit, 620
send-generic, 428
send-new-place-channel-to-named-dest, 625
send/apply, 425
send/keyword-apply, 425
sequence, 316
Sequence Combinations, 326
Sequence Conversion, 326
Sequence Predicate and Constructors, 318

sequence->generator, 337
sequence->list, 327
sequence->repeated-generator, 337
sequence->stream, 326
sequence-add-between, 329
sequence-andmap, 328
sequence-append, 327
sequence-count, 328
sequence-filter, 328
sequence-fold, 328
sequence-for-each, 328
sequence-generate, 326
sequence-generate*, 326
sequence-length, 327
sequence-map, 327
sequence-ormap, 328
sequence-ref, 327
sequence-tail, 327
sequence?, 318
Sequences, 316
Sequences and Streams, 315
Sequencing: begin, begin0, and begin-for-syntax, 124
serializable-struct, 794
serializable-struct/versions, 794
serializable?, 789
Serialization, 789
serialize, 789
serialized=?, 793
'server, 843
set, 561
set, 353
set, 353
set!, 126
set!-transformer-procedure, 656
set!-transformer?, 655
set!-values, 127
set->list, 358
set-add, 354
set-box!, 307
set-count, 354
set-empty?, 353
set-eq?, 357
set-equal?, 357
set-eqv?, 357
set-eval-handler, 866
set-eval-limits, 866
set-field!, 427
set-for-each, 357
set-intersect, 355
set-map, 356
set-mcar!, 294
set-mcdr!, 295
set-member?, 354
set-on-channel!, 621
set-port-next-location!, 701
set-remove, 354
set-subtract, 355
set-symmetric-difference, 355
set-union, 354
set/c, 357
set=?, 355
set?, 357
seteq, 353
seteqv, 353
Sets, 353
seventh, 284
sgn, 183
shadowing, 32
shadows, 32
shared, 112
'shared, 936
shared memory space, 602
shared-bytes, 214
shared-flvector, 190
shared-fxvector, 194
shell-execute, 922
ShellExecute, 923
shift, 559
shift-at, 559
shift0, 559
shift0-at, 560
show-visualizer, 596
shuffle, 290

signature, 450
signature-members, 464
Simple Subprocesses, 924
simple-form-path, 881
simplify-path, 877
sin, 170
single-flonum?, 156
Single-Signature Modules, 463
Single-Unit Modules, 462
sinh, 184
sixth, 284
'sk, 231
sleep, 572
'sm, 231
'so, 231
'so-suffix, 936
socket-connection%, 617
some-system-path->string, 881
sort, 272
spawn, 560
spawn-node-at, 624
spawn-node-supervise-dynamic-place-at, 610
spawn-node-supervise-place-thunk-at, 612
spawn-node-with-dynamic-place-at, 610
spawn-node-with-place-thunk-at, 611
spawn-nodes/join, 625
spawn-remote-racket-node, 612
spawned-process%, 616
Special Comments, 787
special-comment-value, 787
special-comment?, 787
special-filter-input-port, 743
splicing-let, 677
splicing-let-syntax, 677
splicing-let-syntaxes, 678
splicing-let-values, 677
splicing-letrec, 677
splicing-letrec-syntax, 678
splicing-letrec-syntaxes, 678
splicing-letrec-syntaxes+values, 678
splicing-letrec-values, 677
splicing-local, 678
splicing-syntax-parameterize, 678
split-at, 286
split-at-right, 287
split-path, 878
splitter, 560
sqr, 183
sqrt, 169
square root, 169
srcloc, 547
srcloc-column, 547
srcloc-line, 547
srcloc-position, 547
srcloc-source, 547
srcloc-span, 547
srcloc?, 547
ssh-bin-path, 624
stack dump, 564
stack trace, 564
start-future-tracing!, 599
start-message-router/thread, 626
'static, 936
stop-after, 323
stop-before, 323
stop-future-tracing!, 599
stream, 330
stream, 329
stream->list, 330
stream-add-between, 332
stream-andmap, 331
stream-append, 331
stream-cons, 329
stream-count, 332
stream-empty?, 329
stream-filter, 332
stream-first, 329
stream-fold, 331
stream-for-each, 331
stream-length, 330

stream-map, 331
stream-ormap, 331
stream-ref, 330
stream-rest, 329
stream-tail, 330
stream?, 329
Streams, 329
string, 196
string, 195
String Comparisons, 200
String Constructors, Selectors, and Mutators, 195
String Conversions, 203
string port, 708
String Ports, 708
string->bytes/latin-1, 217
string->bytes/locale, 217
string->bytes/utf-8, 217
string->immutable-string, 196
string->keyword, 261
string->list, 199
string->number, 179
string->path, 872
string->path-element, 873
string->some-system-path, 881
string->symbol, 234
string->uninterned-symbol, 234
string->unreadable-symbol, 235
string-append, 199
string-append*, 206
string-ci<=?, 202
string-ci<?, 201
string-ci=?, 201
string-ci>=?, 202
string-ci>?, 202
string-copy, 197
string-copy!, 198
string-downcase, 203
string-fill!, 198
string-foldcase, 204
string-join, 206
string-len/c, 468
string-length, 196
string-locale-ci<?, 205
string-locale-ci=?, 205
string-locale-ci>?, 205
string-locale-downcase, 206
string-locale-upcase, 205
string-locale<?, 205
string-locale=?, 205
string-locale>?, 205
string-no-nuls?, 927
string-normalize-nfc, 204
string-normalize-nfd, 204
string-normalize-nfkc, 204
string-normalize-nfkd, 204
string-normalize-spaces, 207
string-ref, 196
string-replace, 207
string-set!, 197
string-split, 208
string-titlecase, 203
string-trim, 208
string-upcase, 203
string-utf-8-length, 218
string<=?, 200
string<?, 200
string=?, 200
string>=?, 201
string>?, 201
string?, 195
strings, pattern matching, 235
strings, parsing, 54
strings, immutable, 195
strings, concatenate, 199
Strings, 195
struct, 375
struct*, 527
struct->vector, 391
struct-accessor-procedure?, 392
struct-auto-info-lists, 396
struct-auto-info?, 396
struct-constructor-procedure?, 391
struct-copy, 390

struct-field-index, 378
struct-info, 848
struct-info?, 395
struct-mutator-procedure?, 392
struct-out, 89
struct-predicate-procedure?, 392
struct-type-info, 848
struct-type-make-constructor, 849
struct-type-make-predicate, 849
struct-type-property-accessor-procedure?, 386
struct-type-property/c, 486
struct-type-property?, 386
struct-type?, 391
struct/c, 471
struct/ctc, 461
struct/dc, 472
struct:arity-at-least, 365
struct:date, 931
struct:date*, 932
struct:exn, 542
struct:exn:break, 545
struct:exn:fail, 542
struct:exn:fail:contract, 542
struct:exn:fail:contract:arity, 542
struct:exn:fail:contract:blame, 503
struct:exn:fail:contract:continuation, 543
struct:exn:fail:contract:divide-by-zero, 542
struct:exn:fail:contract:non-fixnum-result, 543
struct:exn:fail:contract:variable, 543
struct:exn:fail:filesystem, 544
struct:exn:fail:filesystem:exists, 544
struct:exn:fail:filesystem:version, 544
struct:exn:fail:network, 544
struct:exn:fail:object, 448
struct:exn:fail:out-of-memory, 544
struct:exn:fail:read, 543
struct:exn:fail:read:eof, 544
struct:exn:fail:read:non-char, 544
struct:exn:fail:syntax, 543
struct:exn:fail:syntax:unbound, 543
struct:exn:fail:unsupported, 545
struct:exn:fail:user, 545
struct:export, 675
struct:future-event, 600
struct:import, 671
struct:import-source, 671
struct:indexed-future-event, 599
struct:rmpi-comm, 631
struct:srcloc, 547
struct:struct-info, 396
struct?, 391
Structural Matching, 459
structure, 374
Structure Inspectors, 847
structure subtypes, 374
structure type, 374
structure type descriptor, 375
Structure Type Properties, 384
structure type property, 384
Structure Type Property Contracts, 486
structure type property descriptor, 384
Structure Type Transformer Binding, 394
Structure Utilities, 391
structures, equality, 374
Structures, 374
Structures as Ports, 710
Sub-expression Evaluation and Continuations, 18
sub1, 163
subbytes, 212
subclass?, 444
subclass?/c, 438
submod, 85
'submodule, 74
'submodule, 74
submodule, 28
Submodules, 28
subprocess, 919

subprocess-group-enabled, 922
subprocess-kill, 921
subprocess-pid, 922
subprocess-status, 921
subprocess-wait, 921
subprocess?, 922
subset?, 356
substring, 197
subtract-in, 96
super, 418
super-instantiate, 424
super-make-object, 423
super-new, 424
superclass, 397
supervise-dynamic-place-at, 613
supervise-named-dynamic-place-at, 614
supervise-named-place-thunk-at, 614
supervise-place-thunk-at, 613
supervise-process-at, 614
supervise-thread-at, 615
surrogate, 448
Surrogates, 448
Suspending, Resuming, and Killing Threads, 570
'SW_HIDE, 923
'sw_hide, 923
'SW_MAXIMIZE, 923
'sw_maximize, 923
'SW_MINIMIZE, 923
'sw_minimize, 923
'SW_RESTORE, 923
'sw_restore, 923
'SW_SHOW, 923
'sw_show, 923
'SW_SHOWDEFAULT, 923
'sw_showdefault, 923
'SW_SHOWMAXIMIZED, 924
'sw_showmaximized, 924
'SW_SHOWMINIMIZED, 924
'sw_showminimized, 924
'SW_SHOWMINNOACTIVE, 924
'sw_showminnoactive, 924
'SW_SHOWNA, 924
'sw_showna, 924
'SW_SHOWNOACTIVATE, 924
'sw_shownoactivate, 924
'SW_SHOWNORMAL, 924
'sw_shownormal, 924
symbol, 233
symbol->string, 234
symbol-interned?, 233
symbol-unreadable?, 234
symbol=?, 150
symbol?, 233
symbols, unique, 233
symbols, generating, 233
Symbols, 233
symbols, 469
sync, 577
sync-events, 618
sync/enable-break, 578
sync/timeout, 578
sync/timeout/enable-break, 578
synchronizable event, 574
Synchronization, 574
synchronization result, 574
Synchronizing Thread State, 572
syntactic form, 36
Syntactic Forms, 71
syntax, 640
syntax mark, 39
Syntax Model, 31
syntax object, 33
Syntax Object Bindings, 651
Syntax Object Content, 645
Syntax Object Properties, 679
Syntax Objects, 33
syntax pair, 647
syntax parameter, 676
Syntax Parameter Inspection, 677
Syntax Parameters, 676
syntax property, 679
Syntax Quoting: quote-syntax, 139

Syntax Taints, 681
syntax transformer, 39
Syntax Transformers, 655
Syntax Utilities, 688
syntax->datum, 648
syntax->list, 647
syntax-arm, 683
syntax-case, 635
syntax-case*, 638
syntax-column, 646
syntax-disarm, 684
syntax-e, 647
syntax-id-rules, 644
syntax-line, 645
syntax-local-bind-syntaxes, 661
syntax-local-certifier, 667
syntax-local-context, 665
syntax-local-eval, 692
syntax-local-expand-expression, 659
syntax-local-get-shadower, 666
syntax-local-introduce, 667
syntax-local-lift-context, 664
syntax-local-lift-expression, 663
syntax-local-lift-module-end-declaration, 664
syntax-local-lift-provide, 665
syntax-local-lift-require, 664
syntax-local-lift-values-expression, 663
syntax-local-make-definition-context, 661
syntax-local-make-delta-introducer, 668
syntax-local-module-defined-identifiers, 669
syntax-local-module-exports, 666
syntax-local-module-required-identifiers, 669
syntax-local-name, 665
syntax-local-phase-level, 666
syntax-local-provide-certifier, 675
syntax-local-require-certifier, 672
syntax-local-submodules, 666
syntax-local-transforming-module-provides?, 669
syntax-local-value, 662
syntax-local-value/immediate, 663
syntax-local-value/record, 691
syntax-original?, 646
syntax-parameter-value, 677
syntax-parameterize, 676
syntax-pattern-variable?, 645
syntax-position, 646
syntax-procedure-alias-property, 675
syntax-procedure-converted-arguments-property, 676
syntax-property, 680
syntax-property-symbol-keys, 681
syntax-protect, 683
syntax-rearm, 684
syntax-recertify, 650
syntax-rules, 644
syntax-shift-phase-level, 649
syntax-source, 645
syntax-source-module, 646
syntax-span, 646
syntax-taint, 684
syntax-tainted?, 683
syntax-track-origin, 681
syntax-transforming-module-expression?, 667
syntax-transforming?, 667
syntax/c, 471
syntax/loc, 643
syntax?, 645
'sys-dir, 888
system, 924
system*, 925
system*/exit-code, 925
system-big-endian?, 182
system-idle-evt, 579
system-language+country, 936
system-library-subpath, 937
system-path-convention-type, 874

system-type, 936
system/exit-code, 925
tag, 454
Tail Position, 18
tail position, 18
'taint-mode, 682
'taint-mode, 682
tainted, 681
take, 285
take-right, 286
tamper status, 681
tan, 171
tanh, 184
tc-get, 627
TCP, 907
TCP listener, 908
TCP port, 911
tcp-abandon-port, 911
tcp-accept, 909
tcp-accept-evt, 910
tcp-accept-ready?, 910
tcp-accept/enable-break, 910
tcp-addresses, 911
tcp-close, 910
tcp-connect, 908
tcp-connect/enable-break, 909
tcp-listen, 907
tcp-listener?, 910
tcp-port?, 911
TEMP, 887
'temp-dir, 887
template environment, 32
tentative-pretty-print-port-cancel, 780
tentative-pretty-print-port-transfer, 779
tenth, 284
terminal-port?, 697
'terminating-macro, 781
'text, 702
The Printer, 62
The Racket Reference, 1
The racket/load Language, 816
The Reader, 48
the-unsupplied-arg, 482
third, 283
this, 402
this%, 402
thread, 569
Thread Cells, 587
thread cells, 29
thread descriptor, 569
thread group, 846
Thread Groups, 846
Thread Mailboxes, 573
thread-cell-ref, 587
thread-cell-set!, 587
thread-cell-values?, 589
thread-cell?, 587
thread-dead-evt, 572
thread-dead?, 572
thread-group?, 847
Thread-Local Storage, 587
thread-receive, 573
thread-receive-evt, 573
thread-resume, 571
thread-resume-evt, 572
thread-rewind-receive, 574
thread-running?, 572
thread-send, 573
thread-suspend, 570
thread-suspend-evt, 573
thread-try-receive, 573
thread-wait, 572
thread/suspend-to-kill, 570
thread?, 569
threads, run state, 572
threads, breaking, 565
threads, breaking, 572
Threads, 29
Threads, 569
threads, 29
thunk, 370
thunk*, 370

Time, 931
time, 933
time-apply, 933
timeline-events, 599
timeline-pict, 597
TMP, 887
TMPDIR, 887
top-level binding, 32
top-level context, 37
top-level variable, 26
Top-Level Variables, 19
touch, 593
trace, 972
trace-call, 973
trace-futures, 598
trace-futures-thunk, 598
Tracing, 972
trait, 429
trait, 429
trait->mixin, 430
trait-alias, 432
trait-exclude, 431
trait-exclude-field, 432
trait-rename, 432
trait-rename-field, 432
trait-sum, 430
trait?, 430
Traits, 429
transformer binding, 38
Transformer Bindings, 38
transformer environment, 32
Transformer Helpers, 463
'transparent, 682
'transparent-binding, 682
transplant-input-port, 741
transplant-output-port, 742
Trigonometric Functions, 170
true, 150
truncate, 165
'truncate, 703
'truncate/replace, 703
UDP, 912
UDP socket, 912
udp-addresses, 918
udp-bind!, 912
udp-bound?, 917
udp-close, 916
udp-connect!, 913
udp-connected?, 917
udp-open-socket, 912
udp-receive!, 915
udp-receive!*, 916
udp-receive!-evt, 918
udp-receive!/enable-break, 916
udp-receive-ready-evt, 917
udp-send, 914
udp-send*, 914
udp-send-evt, 918
udp-send-ready-evt, 917
udp-send-to, 913
udp-send-to*, 914
udp-send-to-evt, 917
udp-send-to/enable-break, 915
udp-send/enable-break, 915
udp?, 916
unbound, 32
unbox, 306
uncaught-exception handler, 539
uncaught-exception-handler, 539
unconstrained-domain->, 482
uninterned, 233
unit, 450
unit contract, 462
Unit Contracts, 462
Unit Utilities, 461
unit-from-context, 459
unit-static-signatures, 463
unit/c, 462
unit/new-import-export, 460
unit/s, 460
unit?, 461
Units, 450
Units, 450
'unix, 874

'unix, 936
Unix and Mac OS X Paths, 882
unless, 126
unquote, 139
unquote-splicing, 139
unreadable symbol, 233
unsafe, 951
Unsafe Data Extraction, 955
Unsafe Numeric Operations, 951
Unsafe Operations, 951
unsafe-box*-cas!, 956
unsafe-bytes-length, 957
unsafe-bytes-ref, 957
unsafe-bytes-set!, 957
unsafe-car, 955
unsafe-cdr, 955
unsafe-f64vector-ref, 957
unsafe-f64vector-set!, 958
unsafe-fl*, 953
unsafe-fl+, 953
unsafe-fl-, 953
unsafe-fl->fx, 955
unsafe-fl/, 953
unsafe-fl<, 953
unsafe-fl<=, 953
unsafe-fl=, 953
unsafe-fl>, 953
unsafe-fl>=, 953
unsafe-flabs, 953
unsafe-flacos, 954
unsafe-flasin, 954
unsafe-flatan, 954
unsafe-flceiling, 954
unsafe-flcos, 954
unsafe-flexp, 954
unsafe-flexpt, 954
unsafe-flfloor, 954
unsafe-flimag-part, 955
unsafe-fllog, 954
unsafe-flmax, 953
unsafe-flmin, 953
unsafe-flreal-part, 954
unsafe-flround, 954
unsafe-flsin, 954
unsafe-flsqrt, 954
unsafe-fltan, 954
unsafe-fltruncate, 954
unsafe-flvector-length, 957
unsafe-flvector-ref, 957
unsafe-flvector-set!, 957
unsafe-fx*, 951
unsafe-fx+, 951
unsafe-fx-, 951
unsafe-fx->fl, 955
unsafe-fx<, 952
unsafe-fx<=, 952
unsafe-fx=, 952
unsafe-fx>, 952
unsafe-fx>=, 952
unsafe-fxabs, 951
unsafe-fxand, 952
unsafe-fxior, 952
unsafe-fxlshift, 952
unsafe-fxmax, 953
unsafe-fxmin, 953
unsafe-fxmodulo, 951
unsafe-fxnot, 952
unsafe-fxquotient, 951
unsafe-fxremainder, 951
unsafe-fxrshift, 952
unsafe-fxxor, 952
unsafe-list-ref, 955
unsafe-list-tail, 955
unsafe-make-flrectangular, 954
unsafe-mcar, 955
unsafe-mcdr, 955
unsafe-s16vector-ref, 958
unsafe-s16vector-set!, 958
unsafe-set-box!, 956
unsafe-set-box*!, 956
unsafe-set-mcar!, 955
unsafe-set-mcdr!, 955
unsafe-string-length, 957
unsafe-string-ref, 957

unsafe-string-set!, 957
unsafe-struct*-ref, 958
unsafe-struct*-set!, 958
unsafe-struct-ref, 958
unsafe-struct-set!, 958
unsafe-u16vector-ref, 958
unsafe-u16vector-set!, 958
unsafe-unbox, 956
unsafe-unbox*, 956
unsafe-vector*-length, 956
unsafe-vector*-ref, 956
unsafe-vector*-set!, 956
unsafe-vector-length, 956
unsafe-vector-ref, 956
unsafe-vector-set!, 956
'unshadowable, 666
unsupplied-arg?, 483
unsyntax, 643
unsyntax-splicing, 643
untrace, 973
'up, 874
'update, 703
use-collection-link-paths, 968
use-compiled-file-paths, 812
use-user-specific-search-paths, 968
USER, 886
user-execute-bit, 906
user-read-bit, 906
user-write-bit, 906
USERPROFILE, 886
UTF-8-permissive, 220
Utilities for Building New Combinators, 510
value, 18
Value Output Hook, 777
value-contract, 511
values, 529
variable, 26
variable reference, 102
Variable References and #%top, 101
variable-reference->empty-namespace, 805
variable-reference->module-base-phase, 806
variable-reference->module-declaration-inspector, 806
variable-reference->module-path-index, 806
variable-reference->module-source, 806
variable-reference->namespace, 805
variable-reference->phase, 806
variable-reference->resolved-module-path, 806
variable-reference-constant?, 805
variable-reference?, 805
Variables and Locations, 26
vector, 298
vector, 298
vector->immutable-vector, 299
vector->list, 299
vector->pseudo-random-generator, 178
vector->pseudo-random-generator!, 178
vector->values, 300
vector-append, 301
vector-argmax, 305
vector-argmin, 305
vector-copy, 303
vector-copy!, 299
vector-count, 304
vector-drop, 302
vector-drop-right, 302
vector-fill!, 299
vector-filter, 304
vector-filter-not, 304
vector-immutable, 298
vector-immutable/c, 470
vector-immutableof, 469
vector-length, 298
vector-map, 301
vector-map!, 301
vector-member, 305
vector-memq, 306
vector-memv, 305

vector-ref, 299
vector-set!, 299
vector-set*!, 301
vector-set-performance-stats!, 938
vector-split-at, 303
vector-split-at-right, 303
vector-take, 302
vector-take-right, 302
vector/c, 469
vector?, 298
vectorof, 469
Vectors, 298
version, 937
visits, 41
visualize-futures, 595
visualize-futures-thunk, 595
void, 373
Void and Undefined, 373
void?, 373
wait-for-die, 622
weak box, 946
Weak Boxes, 946
weak references, 23
weak-box-value, 946
weak-box?, 946
when, 125
will, 947
will executor, 947
will-execute, 948
will-executor?, 948
will-register, 948
will-try-execute, 949
Wills and Executors, 947
'windows, 874
'windows, 936
Windows Path Conventions, 882
with-continuation-mark, 137
with-contract, 490
with-disappeared-uses, 691
with-handlers, 539
with-handlers*, 540
with-input-from-bytes, 734
with-input-from-file, 706
with-input-from-string, 734
with-limits, 869
with-method, 426
with-output-to-bytes, 733
with-output-to-file, 706
with-output-to-string, 733
with-syntax, 639
with-syntax*, 693
would-be-future, 594
wrap-evt, 578
writable<%>, 441
write, 765
'write, 843
'write, 891
write-byte, 756
write-bytes, 756
write-bytes-avail, 757
write-bytes-avail*, 757
write-bytes-avail-evt, 758
write-bytes-avail/enable-break, 757
write-char, 756
write-flush, 627
write-special, 758
write-special-avail*, 758
write-special-evt, 758
write-string, 756
write-to-file, 899
Writing, 765
wrong-syntax, 690
xor, 151
yield, 334
zero?, 156
'zl, 231
'zp, 231
'zs, 231
{, 52
|, 48
}, 52
λ, 105