

The Racket Foreign Interface

Version 5.3.5

Eli Barzilay

June 18, 2013

`(require ffi/unsafe)`

The `ffi/unsafe` library enables the direct use of C-based APIs within Racket programs—without writing any new C code. From the Racket perspective, functions and data with a C-based API are *foreign*, hence the term *foreign interface*. Furthermore, since most APIs consist mostly of functions, the foreign interface is sometimes called a *foreign function interface*, abbreviated *FFI*.

Contents

1 Overview	4
1.1 Libraries, C Types, and Objects	4
1.2 Function-Type Bells and Whistles	6
1.3 By-Reference Arguments	6
1.4 C Structs	7
1.5 Pointers and Manual Allocation	8
1.6 Pointers and GC-Managed Allocation	9
1.7 Reliable Release of Resources	11
1.8 More Examples	12
2 Loading Foreign Libraries	13
3 C Types	17
3.1 Type Constructors	17
3.2 Numeric Types	18
3.3 Other Atomic Types	20
3.4 String Types	20
3.4.1 Primitive String Types	20
3.4.2 Fixed Auto-Converting String Types	21
3.4.3 Variable Auto-Converting String Type	21
3.4.4 Other String Types	22
3.5 Pointer Types	22
3.6 Function Types	24
3.6.1 Custom Function Types	29
3.7 C Struct Types	32
3.8 C Array Types	36
3.9 C Union Types	38
3.10 Enumerations and Masks	39
4 Pointer Functions	40
4.1 Pointer Dereferencing	41
4.2 Memory Management	44
4.3 Pointer Structure Property	47
5 Derived Utilities	49
5.1 Safe Homogenous Vectors	49
5.2 Safe C Vectors	56
5.3 Tagged C Pointer Types	58
5.4 Defining Bindings	60
5.5 Allocation and Finalization	62
5.6 Custodian Shutdown Registration	63
5.7 Atomic Execution	64
5.8 Speculatively Atomic Execution	65

5.9	Objective-C FFI	66
5.9.1	FFI Types and Constants	66
5.9.2	Syntactic Forms and Procedures	67
5.9.3	Raw Runtime Functions	71
5.9.4	Legacy Library	75
5.10	Cocoa Foundation	75
5.10.1	Strings	75
5.10.2	Allocation Pools	75
5.11	COM (Common Object Model)	76
5.11.1	COM Automation	77
5.11.2	COM Classes and Interfaces	86
5.11.3	ActiveX Controls	92
5.12	File Security-Guard Checks	94
5.13	Windows API Helpers	94
6	Miscellaneous Support	96
7	Unexported Primitive Functions	98
	Bibliography	100
	Index	101
	Index	101

1 Overview

Although using the FFI requires writing no new C code, it provides very little insulation against the issues that C programmers face related to safety and memory management. An FFI programmer must be particularly aware of memory management issues for data that spans the Racket–C divide. Thus, this manual relies in many ways on the information in *Inside: Racket C API*, which defines how Racket interacts with C APIs in general.

Since using the FFI entails many safety concerns that Racket programmers can normally ignore, the library name includes `unsafe`. Importing the library macro should be considered as a declaration that your code is itself unsafe, therefore can lead to serious problems in case of bugs: it is your responsibility to provide a safe interface. If your library provides an unsafe interface, then it should have `unsafe` in its name, too.

For more information on the motivation and design of the Racket FFI, see [Barzilay04].

1.1 Libraries, C Types, and Objects

To use the FFI, you must have in mind

- a particular library from which you want to access a function or value,
- a particular symbol exported by the file, and
- the C-level type (typically a function type) of the exported symbol.

The library corresponds to a file with a suffix such as `.dll`, `.so`, or `.dylib` (depending on the platform), or it might be a library within a `.framework` directory on Mac OS X.

Knowing the library's name and/or path is often the trickiest part of using the FFI. Sometimes, when using a library name without a path prefix or file suffix, the library file can be located automatically, especially on Unix. See `ffi-lib` for advice.

The `ffi-lib` function gets a handle to a library. To extract exports of the library, it's simplest to use `define-ffi-definer` from the `ffi/unsafe/define` library:

```
#lang racket/base
(require ffi/unsafe
         ffi/unsafe/define)

(define-ffi-definer define-curses (ffi-lib "libcurses"))
```

This `define-ffi-definer` declaration introduces a `define-curses` form for binding a Racket name to a value extracted from "libcurses"—which might be located at "/usr/lib/libcurses.so", depending on the platform.

To use `define-curses`, we need the names and C types of functions from "libcurses". We'll start by using the following functions:

```
WINDOW* initscr(void);
int waddstr(WINDOW *win, char *str);
int wrefresh(WINDOW *win);
int endwin(void);
```

We make these functions callable from Racket as follows:

```
(define _WINDOW-pointer (_cpointer 'WINDOW))

(define-curses initscr (_fun -> _WINDOW-pointer))
(define-curses waddstr (_fun _WINDOW-pointer _string -> _int))
(define-curses wrefresh (_fun _WINDOW-pointer -> _int))
(define-curses endwin (_fun -> _int))
```

By convention, an underscore prefix indicates a representation of a C type (such as `_int`) or a constructor of such representations (such as `_cpointer`).

The definition of `_WINDOW-pointer` creates a Racket value that reflects a C type via `_cpointer`, which creates a type representation for a pointer type—usually one that is opaque. The `'WINDOW` argument could have been any value, but by convention, we use a symbol matching the C base type.

Each `define-curses` form uses the given identifier as both the name of the library export and the Racket identifier to bind. The `(_fun ... -> ...)` part of each definition describes the C type of the exported function, since the library file does not encode that information for its exports. The types listed to the left of `->` are the argument types, while the type to the right of `->` is the result type. The pre-defined `_int` type naturally corresponds to the `int` C type, while `_string` corresponds to the `char*` type when it is intended as a string to read.

An optional `#:c-id` clause for `define-curses` can specify a name for the library export that is different from the Racket identifier to bind.

At this point, `initscr`, `waddstr`, `wrefresh`, and `endwin` are normal Racket bindings to Racket functions (that happen to call C functions), and so they can be exported from the defining module or called directly:

```
(define win (initscr))
(void (waddstr win "Hello"))
(void (wrefresh win))
(sleep 1)
(void (endwin))
```

1.2 Function-Type Bells and Whistles

Our initial use of functions like `waddstr` is sloppy, because we ignore return codes. C functions often return error codes, and checking them is a pain. A better approach is to build the check into the `waddstr` binding and raise an exception when the code is non-zero.

The `_fun` function-type constructor includes many options to help convert C functions to nicer Racket functions. We can use some of those features to convert return codes into either `#<void>` or an exception:

```
(define (check v who)
  (unless (zero? v)
    (error who "failed: ~a" v)))

(define-curses initscr (_fun -> _WINDOW-pointer))
(define-curses waddstr (_fun _WINDOW-pointer _string -> (r : _int)
  -> (check r 'waddstr)))
(define-curses wrefresh (_fun _WINDOW-pointer -> (r : _int)
  -> (check r 'wrefresh)))
(define-curses endwin (_fun -> (r : _int)
  -> (check r 'endwin)))
```

Using `(r : _int)` as a result type gives the local name `r` to the C function's result. This name is then used in the result post-processing expression that is specified after a second `->` in the `_fun` form.

1.3 By-Reference Arguments

To get mouse events from "libcurses", we must explicitly enable them through the `mousemask` function:

```
typedef unsigned long mmask_t;
#define BUTTON1_CLICKED 004L

mmask_t mousemask(mmask_t newmask, mmask_t *oldmask);
```

Setting `BUTTON1_CLICKED` in the mask enables button-click events. At the same time, `mousemask` returns the current mask by installing it into the pointer provided as its second argument.

Since these kinds of call-by-reference interfaces are common in C, `_fun` cooperates with a `_ptr` form to automatically allocate space for a by-reference argument and extract the value put there by the C function. Give the extracted value name to use in the post-processing expression. The post-processing expression can combine the by-reference result with the

function's direct result (which, in this case, reports a subset of the given mask that is actually supported).

```
(define _mmask_t _ulong)
(define-curses mousemask (_fun _mmask_t (o : (_ptr o _mmask_t))
                             -> (r : _mmask_t)
                             -> (values o r)))

(define BUTTON1_CLICKED 4)

(define-values (old supported) (mousemask BUTTON1_CLICKED))
```

1.4 C Structs

Assuming that mouse events are supported, the "libcurses" library reports them via `getmouse`, which accepts a pointer to a `MEVENT` struct to fill with mouse-event information:

```
typedef struct {
    short id;
    int x, y, z;
    mmask_t bstate;
} MEVENT;

int getmouse(MEVENT *event);
```

To work with `MEVENT` values, we use `define-cstruct`:

```
(define-cstruct _MEVENT ([id _short]
                        [x _int]
                        [y _int]
                        [z _int]
                        [bstate _mmask_t]))
```

This definition binds many names in the same way that `define-struct` binds many names: `_MEVENT` is a C type representing the struct type, `_MEVENT-pointer` is a C type representing a pointer to a `_MEVENT`, `make-MEVENT` constructs a `_MEVENT` value, `MEVENT-x` extracts the `x` fields from an `_MEVENT` value, and so on.

With this C struct declaration, we can define the function type for `getmouse`. The simplest approach is to define `getmouse` to accept an `_MEVENT-pointer`, and then explicitly allocate the `_MEVENT` value before calling `getmouse`:

```
(define-curses getmouse (_fun _MEVENT-pointer -> _int))

(define m (make-MEVENT 0 0 0 0 0))
(when (zero? (getmouse m))
```

```

; use m...
....)

```

For a more Racket-like function, use `(_ptr o _MEVENT)` and a post-processing expression:

```

(define-curses getmouse (_fun (m : (_ptr o _MEVENT))
  -> (r : _int)
  -> (and (zero? r) m)))

(waddstr win (format "click me fast..."))
(wrefresh win)
(sleep 1)

(define m (getmouse))
(when m
  (waddstr win (format "at ~a,~a"
    (MEVENT-x m)
    (MEVENT-y m)))

  (wrefresh win)
  (sleep 1))

(endwin)

```

The difference between `_MEVENT-pointer` and `_MEVENT` is crucial. Using `(_ptr o _MEVENT-pointer)` would allocate only enough space for a pointer to an `MEVENT` struct, which is not enough space for an `MEVENT` struct.

1.5 Pointers and Manual Allocation

To get text from the user instead of a mouse click, "libcurses" provides `wgetnstr`:

```
int wgetnstr(WINDOW *win, char *str, int n);
```

While the `char*` argument to `waddstr` is treated as a nul-terminated string, the `char*` argument to `wgetnstr` is treated as a buffer whose size is indicated by the final `int` argument. The C type `_string` does not work for such buffers.

One way to approach this function from Racket is to describe the arguments in their rawest form, using plain `_pointer` for the second argument to `wgetnstr`:

```
(define-curses wgetnstr (_fun _WINDOW-pointer _pointer _int -> _int))
```

To call this raw version of `wgetnstr`, allocate memory, zero it, and pass the size minus one (to leave room a nul terminator) to `wgetnstr`:


```
(define SIZE 256)
(define buffer (malloc 'raw SIZE))
(memset buffer 0 SIZE)

(void (wgetnstr win buffer (sub1 SIZE)))
```

When `wgetnstr` returns, it has written bytes to `buffer`. At that point, we can use `cast` to convert the value from a raw pointer to a string:

```
(cast buffer _pointer _string)
```

Conversion via the `_string` type causes the data referenced by the original pointer to be copied (and UTF-8 decoded), so the memory referenced by `buffer` is no longer needed. Memory allocated with `(malloc 'raw ...)` must be released with `free`:

```
(free buffer)
```

1.6 Pointers and GC-Managed Allocation

Instead of allocating `buffer` with `(malloc 'raw ...)`, we could have allocated it with `(malloc 'atomic ...)`:

```
(define buffer (malloc 'atomic SIZE))
```

Memory allocated with `'atomic` is managed by the garbage collector, so `free` is neither necessary nor allowed when the memory referenced by `buffer` is no longer needed. Instead, when `buffer` becomes inaccessible, the allocated memory will be reclaimed automatically.

Allowing the garbage collector (GC) to manage memory is usually preferable. It's easy to forget to call `free`, and exceptions or thread termination can easily skip a `free`.

At the same time, using GC-managed memory adds a different burden on the programmer: data managed by the GC may be moved to a new address as the GC compacts allocated objects to avoid fragmentation. C functions, meanwhile, expect to receive pointers to objects that will stay put.

Fortunately, unless a C function calls back into the Racket run-time system (perhaps through a function that is provided as an argument), no garbage collection will happen between the time that a C function is called and the time that the function returns.

Let's look a few possibilities related to allocation and pointers:

- Ok:

```
(define p (malloc 'atomic SIZE))
(wgetnstr win p (sub1 SIZE))
```

Although the data allocated by `malloc` can move around, `p` will always point to it, and no garbage collection will happen between the time that the address is extracted from `p` to pass to `wgetnstr` and the time that `wgetnstr` returns.

- Bad:

```
(define p (malloc 'atomic SIZE))
(define i (cast p _pointer _intptr))
(wgetnstr win (cast i _intptr _pointer) (sub1 SIZE))
```

The data referenced by `p` can move after the address is converted to an integer, in which case `i` cast back to a pointer will be the wrong address.

Obviously, casting a pointer to an integer is generally a bad idea, but the cast simulates another possibility, which is passing the pointer to a C function that retains the pointer in its own private store for later use. Such private storage is invisible to the Racket GC, so it has the same effect as casting the pointer to an integer.

- Ok:

```
(define p (malloc 'atomic SIZE))
(define p2 (ptr-add p 4))
(wgetnstr win p2 (- SIZE 5))
```

The pointer `p2` retains the original reference and only adds the `4` at the last minute before calling `wgetnstr` (i.e., after the point that garbage collection is allowed).

- Ok:

```
(define p (malloc 'atomic-interior SIZE))
(define i (cast p _pointer _intptr))
(wgetnstr win (cast i _intptr _pointer) (sub1 SIZE))
```

This is ok assuming that `p` itself stays accessible, so that the data it references isn't reclaimed. Allocating with `'atomic-interior` puts data at a particular address and keeps it there. A garbage collection will not change the address in `p`, and so `i` (cast back to a pointer) will always refer to the data.

Keep in mind that C struct constructors like `make-MEVENT` are effectively the same as `(malloc 'atomic ...)`; the result values can move in memory during a garbage collection. The same is true of byte strings allocated with `make-bytes`, which (as a convenience) can be used directly as a pointer value (unlike character strings, which are always copied for UTF-8 encoding or decoding).

For more information about memory management and garbage collection, see §3 “Memory Allocation” in *Inside: Racket C API*.

1.7 Reliable Release of Resources

Using GC-managed memory saves you from manual `free`s for plain memory blocks, but C libraries often allocate resources and require a matching call to a function that releases the resources. For example, "libcurses" supports windows on the screen that are created with `newwin` and released with `delwin`:

```
WINDOW *newwin(int lines, int ncols, int y, int x);
int delwin(WINDOW *win);
```

In a sufficiently complex program, ensuring that every `newwin` is paired with `delwin` can be challenging, especially if the functions are wrapped by otherwise safe functions that are provided from a library. A library that is intended to be safe for use in a sandbox, say, must protect against resource leaks within the Racket process as a whole when a sandboxed program misbehaves or is terminated.

The `ffi/unsafe/alloc` library provides functions to connect resource-allocating functions and resource-releasing functions. The library then arranges for finalization to release a resource if it becomes inaccessible (according to the GC) before it is explicitly released. At the same time, the library handles tricky atomicity requirements to ensure that the finalization is properly registered and never run multiple times.

Using `ffi/unsafe/alloc`, the `newwin` and `delwin` functions can be imported with `allocator` and `deallocator` wrappers, respectively:

```
(require ffi/unsafe/alloc)

(define-curses delwin (_fun _WINDOW-pointer -> _int)
  #:wrap (deallocator))

(define-curses newwin (_fun _int _int _int _int
                          -> _WINDOW-pointer)
  #:wrap (allocator delwin))
```

A `deallocator` wrapper makes a function cancel any existing finalizer for the function's argument. An `allocator` wrapper refers to the deallocator, so that the deallocator can be run if necessary by a finalizer.

If a resource is scarce or visible to end users, then custodian management is more appropriate than mere finalization as implemented by `allocator`. See the `ffi/unsafe/custodian` library.

1.8 More Examples

For more examples of common FFI patterns, see the defined interfaces in the "ffi/examples" collection. See also [Barzilay04].

2 Loading Foreign Libraries

The FFI is normally used by extracting functions and other objects from shared objects (a.k.a. *shared libraries* or *dynamically loaded libraries*). The `ffi-lib` function loads a shared object.

```
(ffi-lib? v) → boolean?  
v : any/c
```

Returns `#t` if `v` is a *foreign-library value*, `#f` otherwise.

```
(ffi-lib path  
  [version  
   #:get-lib-dirs get-lib-dirs  
   #:fail fail  
   #:global? global?]) → any  
path : (or/c path-string? #f)  
version : (or/c string? (listof (or/c string? #f)) #f) = #f  
get-lib-dirs : (-> (listof path?)) = get-lib-search-dirs  
fail : (or/c #f (-> any)) = #f  
global? : any/c = (eq? 'global (system-type 'so-mode))
```

Returns a foreign-library value or the result of `fail`. Normally,

- `path` is a path without a version or suffix (i.e., without ".dll", ".so", or ".dylib"); and
- `version` is a list of versions to try in order with `#f` (i.e., no version) as the last element of the list; for example, `'("2" #f)` indicates version 2 with a fallback to a versionless library.

A string or `#f` `version` is equivalent to a list containing just the string or `#f`, and an empty string (by itself or in a list) is equivalent to `#f`.

Beware of relying on versionless library names. On some platforms, versionless library names are provided only by development packages. At the same time, other platforms may require a versionless fallback. A list of version strings followed by `#f` is typically best for `version`.

Assuming that `path` is not `#f`, the result from `ffi-lib` represents the library found by the following search process:

- If `path` is not an absolute path, look in each directory reported by `get-lib-dirs`. In each directory, try `path` with the first version in `version`, adding a suitable suffix if `path` does not already end in the suffix, then try the second version in `version`, etc. (If `version` is an empty list, no paths are tried in this step.)

- Try the same filenames again, but without converting the path to an absolute path, which allows the operating system to use its own search paths. (If `version` is an empty list, no paths are tried in this step.)
- Try `path` without adding any version or suffix, and without converting to an absolute path.
- Try the version-adjusted filenames again, but relative to the current directory. (If `version` is an empty list, no paths are tried in this step.)
- Try `path` without adding any version or suffix, but converted to an absolute path relative to the current directory.

If none of the paths succeed and `fail` is a function, then `fail` is called in tail position. If `fail` is `#f`, an error is reported from trying the first path from the second bullet above or (if `version` is an empty list) from the third bullet above. A library file may exist but fail to load for some reason; the eventual error message will unfortunately name the fallback from the second or third bullet, since some operating systems offer no way to determine why a given library path failed.

If `path` is `#f`, then the resulting foreign-library value represents all libraries loaded in the current process, including libraries previously opened with `ffi-lib`. In particular, use `#f` to access C-level functionality exported by the run-time system (as described in *Inside: Racket C API*). The `version` argument is ignored when `path` is `#f`.

If `path` is not `#f`, `global?` is true, and the operating system supports opening a library in “global” mode so that the library’s symbols are used for resolving references from libraries that are loaded later, then global mode is used to open the library. Otherwise, the library is opened in “local” mode, where the library’s symbols are not made available for future resolution. This local-versus-global choice does not affect whether the library’s symbols are available via `(ffi-lib #f)`.

Due to the way the operating system performs dynamic binding, loaded libraries are associated with Racket (or DrRacket) for the duration of the process. Re-evaluating `ffi-lib` (or hitting the Run button in DrRacket) will not force a re-load of the corresponding library.

```
(get-ffi-obj objname lib type [failure-thunk]) → any
  objname : (or/c string? bytes? symbol?)
  lib : (or/c ffi-lib? path-string? #f)
  type : ctype?
  failure-thunk : (or/c (-> any) #f) = #f
```

Looks for `objname` in `lib` library. If `lib` is not a foreign-library value it is converted to one by calling `ffi-lib`. If `objname` is found in `lib`, it is converted to Racket using the given `type`. Types are described in §3 “C Types”; in particular the `get-ffi-obj` procedure is most often used with function types created with `_fun`.

Keep in mind that `get-ffi-obj` is an unsafe procedure; see §1 “Overview” for details.

If the name is not found, and `failure-thunk` is provided, it is used to produce a return value. For example, a failure thunk can be provided to report a specific error if a name is not found:

```
(define foo
  (get-ffi-obj "foo" foolib (_fun _int -> _int)
    (lambda ()
      (error 'foolib
        "installed foolib does not provide \"foo\""))))
```

The default (also when `failure-thunk` is provided as `#f`) is to raise an exception.

```
(set-ffi-obj! objname lib type new) → void?
  objname : (or/c string? bytes? symbol?)
  lib : (or/c ffi-lib? path-string? #f)
  type : ctype?
  new : any/c
```

Looks for `objname` in `lib` similarly to `get-ffi-obj`, but then it stores the given `new` value into the library, converting it to a C value. This can be used for setting library customization variables that are part of its interface, including Racket callbacks.

```
(make-c-parameter objname lib type) → (and/c (-> any)
                                             (any/c -> void?))
  objname : (or/c string? bytes? symbol?)
  lib : (or/c ffi-lib? path-string? #f)
  type : ctype?
```

Returns a parameter-like procedure that can either reference the specified foreign value, or set it. The arguments are handled as in `get-ffi-obj`.

A parameter-like function is useful in case Racket code and library code interact through a library value. Although `make-c-parameter` can be used with any time, it is not recommended to use this for foreign functions, since each reference through the parameter will construct the low-level interface before the actual call.

```
(define-c id lib-expr type-expr)
```

Defines `id` behave like a Racket binding, but `id` is actually redirected through a parameter-like procedure created by `make-c-parameter`. The `id` is used both for the Racket binding and for the foreign name.

```
(ffi-obj-ref objname lib [failure-thunk]) → any
  objname : (or/c string? bytes? symbol?)
  lib : (or/c ffi-lib? path-string? #f)
  failure-thunk : (or/c (-> any) #f) = #f
```

Returns a pointer for the specified foreign name, calls *failure-thunk* if the name is not found, or raises an exception if *failure-thunk* is *#f*.

Normally, *get-ffi-obj* should be used, instead.

3 C Types

C types are the main concept of the FFI, either primitive types or user-defined types. The FFI deals with primitive types internally, converting them to and from C types. A user type is defined in terms of existing primitive and user types, along with conversion functions to and from the existing types.

3.1 Type Constructors

```
(make-ctype type scheme-to-c c-to-scheme) → ctype?  
type : ctype?  
scheme-to-c : (or/c #f (any/c . -> . any))  
c-to-scheme : (or/c #f (any/c . -> . any))
```

Creates a new C type value whose representation for foreign code is the same as *type*'s. The given conversions functions convert to and from the Racket representation of *type*. Either conversion function can be *#f*, meaning that the conversion for the corresponding direction is the identity function. If both functions are *#f*, *type* is returned.

```
(ctype? v) → boolean?  
v : any/c
```

Returns *#t* if *v* is a C type, *#f* otherwise.

```
(ctype-sizeof type) → exact-nonnegative-integer?  
type : ctype?  
(ctype-alignof type) → exact-nonnegative-integer?  
type : ctype?
```

Returns the size or alignment of a given *type* for the current platform.

```
(ctype->layout type)  
→ (flat-rec-contract rep symbol? (listof rep))  
type : ctype?
```

Returns a value to describe the eventual C representation of the type. It can be any of the following symbols:

```
'int8 'uint8 'int16 'uint16 'int32 'uint32 'int64 'uint64  
'float 'double 'bool 'void 'pointer 'fpointer  
'bytes 'string/ucs-4 'string/utf-16
```

The result can also be a list, which describes a C struct whose element representations are provided in order within the list. Finally, the result can be a vector of size 2 containing an element representation followed by an exact-integer count.

```
(compiler-sizeof sym) → exact-nonnegative-integer?  
  sym : (or/c symbol? (listof symbol?))
```

Possible values for *sym* are `'int`, `'char`, `'short`, `'long`, `'*`, `'void`, `'float`, `'double`, or lists of symbols, such as `'(long long)`. The result is the size of the corresponding type according to the C `sizeof` operator for the current platform. The `compiler-sizeof` operation should be used to gather information about the current platform, such as defining alias type like `_int` to a known type like `_int32`.

3.2 Numeric Types

```
_int8 : ctype?  
_sint8 : ctype?  
_uint8 : ctype?  
_int16 : ctype?  
_sint16 : ctype?  
_uint16 : ctype?  
_int32 : ctype?  
_sint32 : ctype?  
_uint32 : ctype?  
_int64 : ctype?  
_sint64 : ctype?  
_uint64 : ctype?
```

The basic integer types at various sizes. The `s` or `u` prefix specifies a signed or an unsigned integer, respectively; the ones with no prefix are signed.

```
_byte : ctype?  
_sbyte : ctype?  
_ubyte : ctype?
```

The `_sbyte` and `_ubyte` types are aliases for `_sint8` and `_uint8`, respectively. The `_byte` type is like `_ubyte`, but adds 256 to a negative Racket value that would work as a `_sbyte` (i.e., it casts signed bytes to unsigned bytes).

```
_word : ctype?  
_sword : ctype?  
_uword : ctype?
```

The `_sword` and `_uword` types are aliases for `_sint16` and `_uint16`, respectively. The `_word` type is like `_uword`, but coerces negative values in the same way as `_byte`.

```
_short : ctype?  
_short : ctype?  
_ushort : ctype?  
_int : ctype?  
_sint : ctype?  
_uint : ctype?  
_long : ctype?  
_slong : ctype?  
_ulong : ctype?  
_llong : ctype?  
_sllong : ctype?  
_ullong : ctype?  
_intptr : ctype?  
_sintptr : ctype?  
_uintptr : ctype?
```

Aliases for basic integer types. The `_short` aliases correspond to `_int16`. The `_int` aliases correspond to `_int32`. The `_long` aliases correspond to either `_int32` or `_int64`, depending on the platform. Similarly, the `_intptr` aliases correspond to either `_int32` or `_int64`, depending on the platform.

```
_size : ctype?  
_ssize : ctype?  
_ptrdiff : ctype?  
_intmax : ctype?  
_uintmax : ctype?
```

More aliases for basic integer types. The `_size` and `_uintmax` types are aliases for `_uintptr`, and the rest are aliases for `_intptr`.

```
_fixnum : ctype?  
_ufixnum : ctype?
```

For cases where speed matters and where you know that the integer is small enough, the types `_fixnum` and `_ufixnum` are similar to `_intptr` and `_uintptr` but assume that the quantities fit in Racket's immediate integers (i.e., not bignums).

```
_fixint : ctype?  
_ufixint : ctype?
```

Similar to `_fixnum/_ufixnum`, but based on `_int/_uint` instead of `_intptr/_uintptr`, and coercions from C are checked to be in range.

```
_float : ctype?  
_double : ctype?  
_double* : ctype?
```

The `_float` and `_double` types represent the corresponding C types. Both single- and double-precision Racket numbers are accepted for conversion via both `_float` and `_double`, while both `_float` and `_double` coerce C values to double-precision Racket numbers. The type `_double*` coerces any Racket real number to a C `double`.

`| _longdouble : ctype?`

Represents the `long double` type on platforms where it is supported, in which case Racket `extflonums` convert to and from `long double` values.

3.3 Other Atomic Types

`| _bool : ctype?`

Translates `#f` to a 0 `_int`, and any other value to 1.

`| _void : ctype?`

Indicates a Racket `#<void>` return value, and it cannot be used to translate values to C. This type cannot be used for function inputs.

3.4 String Types

3.4.1 Primitive String Types

The five primitive string types correspond to cases where a C representation matches Racket's representation without encodings.

The form `_bytes` form can be used type for Racket byte strings, which corresponds to C's `char*` type. In addition to translating byte strings, `#f` corresponds to the NULL pointer.

`| _string/ucs-4 : ctype?`

A type for Racket's native Unicode strings, which are in UCS-4 format. These correspond to the C `mzchar*` type used by Racket. As usual, the types treat `#f` as NULL and vice versa.

`| _string/utf-16 : ctype?`

Unicode strings in UTF-16 format. As usual, the types treat `#f` as NULL and vice versa.

`| _path : ctype?`

Simple `char*` strings, corresponding to Racket's paths. As usual, the types treat `#f` as NULL and vice versa.

Beware that changing the current directory via `current-directory` does not change the OS-level current directory as seen by foreign library functions. Paths normally should be converted to absolute form using `path->complete-path` (which uses the `current-directory` parameter) before passing them to a foreign function.

```
| _symbol : ctype?
```

Simple `char*` strings as Racket symbols (encoded in UTF-8). Return values using this type are interned as symbols.

3.4.2 Fixed Auto-Converting String Types

```
| _string/utf-8 : ctype?  
| _string/latin-1 : ctype?  
| _string/locale : ctype?
```

Types that correspond to (character) strings on the Racket side and `char*` strings on the C side. The bridge between the two requires a transformation on the content of the string. As usual, the types treat `#f` as NULL and vice versa.

```
| _string*/utf-8 : ctype?  
| _string*/latin-1 : ctype?  
| _string*/locale : ctype?
```

Similar to `_string/utf-8`, etc., but accepting a wider range of values: Racket byte strings are allowed and passed as is, and Racket paths are converted using `path->bytes`.

3.4.3 Variable Auto-Converting String Type

The `_string/ucs-4` type is rarely useful when interacting with foreign code, while using `_bytes` is somewhat unnatural, since it forces Racket programmers to use byte strings. Using `_string/utf-8`, etc., meanwhile, may prematurely commit to a particular encoding of strings as bytes. The `_string` type supports conversion between Racket strings and `char*` strings using a parameter-determined conversion.

```
| _string : ctype?
```

Expands to a use of the `default-_string-type` parameter. The parameter's value is consulted when `_string` is evaluated, so the parameter should be set before any interface definition that uses `_string`.

```
(default-_string-type) → ctype?  
(default-_string-type type) → void?  
  type : ctype?
```

A parameter that determines the current meaning of `_string`. It is initially set to `_string*/utf-8`. If you change it, do so *before* interfaces are defined.

3.4.4 Other String Types

```
| _file : ctype?
```

Like `_path`, but when values go from Racket to C, `cleanse-path` is used on the given value. As an output value, it is identical to `_path`.

```
| _bytes/eof : ctype?
```

Similar to the `_bytes` type, except that a foreign return value of NULL is translated to a Racket `eof` value.

```
| _string/eof : ctype?
```

Similar to the `_string` type, except that a foreign return value of NULL is translated to a Racket `eof` value.

3.5 Pointer Types

```
| _pointer : ctype?
```

Corresponds to Racket *C pointer* values. These pointers can have an arbitrary Racket object attached as a type tag. The tag is ignored by built-in functionality; it is intended to be used by interfaces. See §5.3 “Tagged C Pointer Types” for creating pointer types that use these tags for safety. A `#f` value is converted to NULL and vice versa.

The address referenced by a `_pointer` value must not refer to memory managed by the garbage collector (unless the address corresponds to a value that supports interior pointers and that is otherwise referenced to preserve the value from garbage collection). The reference is not traced or updated by the garbage collector.

The `equal?` predicate equates C pointers (including pointers for `_gcpointer` and possibly containing an offset) when they refer to the same address—except for C pointers that are instances of structure types with the `prop:cpointer` property, in which case the equality rules of the relevant structure types apply.

`| _gcpointer : ctype?`

Like `_pointer`, but for a C pointer value that can refer to memory managed by the garbage collector.

Although a `_gcpointer` can reference to memory that is not managed by the garbage collector, beware of using an address that might eventually become managed by the garbage collector. For example, if a reference is created by `malloc` with `'raw` and released by `free`, then the `free` may allow the memory formerly occupied by the reference to be used later by the garbage collector.

The `cpointer-gcable?` function returns `#t` for a `cpointer` generated via the `_gcpointer` type, while it generates `#f` for a `cpointer` generated via the `_cpointer` type.

`| _racket : ctype?`
`| _scheme : ctype?`

A type that can be used with any Racket object; it corresponds to the `Scheme_Object*` type of Racket's C API (see *Inside: Racket C API*). The `_racket` or `_scheme` type is useful only for libraries that are aware of Racket's C API.

As a result type with a function type, `_racket` or `_scheme` permits multiple values, but multiple values are not allowed in combination with a true value for `#:in-original-place?` or `#:async-apply` in `_cprocedure` or `_fun`.

`| _fpointer : ctype?`

Similar to `_pointer`, except that when `_fpointer` is used as the type for `get-ffi-obj` or `ffi-obj-ref`, then a level of indirection is skipped. Furthermore, for a C pointer value from `get-ffi-obj` or `ffi-obj-ref` using `_fpointer`, `ptr-ref` on the pointer as a `_fpointer` simply returns the pointer instead of dereferencing it. Like `_pointer`, `_fpointer` treats `#f` as NULL and vice versa.

A type generated by `_cprocedure` or `_fun` builds on `_fpointer`, and normally `_cprocedure` or `_fun` should be used instead of `_fpointer`.

`| (_or-null ctype) → ctype?`
`ctype : ctype?`

Creates a type that is like `ctype`, but `#f` is converted to NULL and vice versa. The given `ctype` must have the same C representation as `_pointer`, `_gcpointer`, or `_fpointer`.

`| (_gcable ctype) → ctype?`
`ctype : ctype?`

Creates a type that is like `ctype`, but whose base representation is like `_gcpointer` instead of `_pointer`. The given `ctype` must have a base representation like `_pointer` or `_gcpointer` (and in the later case, the result is the `ctype`).

3.6 Function Types

```
(_cprocedure input-types
             output-type
             [#:abi abi
             #:atomic? atomic?
             #:async-apply async-apply
             #:in-original-place? in-original-place?
             #:save-errno save-errno
             #:wrapper wrapper
             #:keep keep]) → any
input-types : (list ctype?)
output-type : ctype?
abi : (or/c #f 'default 'stdcall 'sysv) = #f
atomic? : any/c = #f
async-apply : (or/c #f ((-> any) . -> . any)) = #f
in-original-place? : any/c = #f
save-errno : (or/c #f 'posix 'windows) = #f
wrapper : (or/c #f (procedure? . -> . procedure?)) = #f
keep : (or/c boolean? box? (any/c . -> . any/c)) = #t
```

A type constructor that creates a new function type, which is specified by the given `input-types` list and `output-type`. Usually, the `_fun` syntax (described below) should be used instead, since it manages a wide range of complicated cases.

The resulting type can be used to reference foreign functions (usually `ffi-objs`, but any pointer object can be referenced with this type), generating a matching foreign `callout` object. Such objects are new primitive procedure objects that can be used like any other Racket procedure. As with other pointer types, `#f` is treated as a NULL function pointer and vice versa.

A type created with `_cprocedure` can also be used for passing Racket procedures to foreign functions, which will generate a foreign function pointer that calls to the given Racket `callback` procedure. There are no restrictions on the Racket procedure; in particular, its lexical context is properly preserved.

The optional `abi` keyword argument determines the foreign ABI that is used. Supplying `#f` or `'default` indicates the platform-dependent default. The other possible values—`'stdcall` and `'sysv` (i.e., “cdecl”)—are currently supported only for 32-bit Windows; using them on other platforms raises an exception. See also `ffi/winapi`.

For callouts to foreign functions with the generated type:

- If `save-errno` is `'posix`, then the value of `errno` is saved (specific to the current thread) immediately after a foreign function callout returns. The saved value is accessible through `saved-errno`. If `save-errno` is `'windows`, then the value of `GetLastError()` is saved for later use via `saved-errno`; the `'windows` option is available only on Windows (on other platforms `saved-errno` will return 0). If `save-errno` is `#f`, no error value is saved automatically.

The error-recording support provided by `save-errno` is needed because the Racket runtime system may otherwise preempt the current Racket thread and itself call functions that set error values.

- If `wrapper` is not `#f`, it takes the callout that would otherwise be generated and returns a replacement procedure. Thus, `wrapper` acts a hook to perform various argument manipulations before the true callout is invoked, and it can return different results (for example, grabbing a value stored in an “output” pointer and returning multiple values).
- If `in-original-place?` is true, then when a foreign callout procedure with the generated type is called in any Racket place, the procedure is called from the original Racket place. Use this mode for a foreign function that is not thread-safe at the C level, which means that it is not place-safe at the Racket level. Callbacks from place-unsafe code back into Racket at a non-original place typically will not work, since the place of the Racket code may have a different allocator than the original place.
- Values that are provided to a callout (i.e., the underlying callout, and not the replacement produced by a `wrapper`, if any) are always considered reachable by the garbage collector until the called foreign function returns. If the foreign function invokes Racket callbacks, however, beware that values managed by the Racket garbage collector might be moved in memory by the garbage collector.
- A callout object is finalized internally. Beware of trying to use a callout object that is reachable only from a finalized object, since the two objects might be finalized in either order.

For callbacks to Racket functions with the generated type:

- The `keep` argument provides control over reachability by the garbage collector of the underlying value that foreign code see as a plain C function. Additional care must be taken in case the foreign code might retain the callback function, in which case the callback value must remain reachable or else the held callback will become invalid. The possible values of `keep` are as follows:
 - `#t` — the callback stays in memory as long as the converted Racket function is reachable. This mode is the default, as it is fine in most cases. Note that each Racket function can hold onto only one callback value through this mode, so it is not suitable for a function used multiple times as a retained callback.

- `#f` — the callback value is not held. This mode may be useful for a callback that is only used for the duration of the foreign call; for example, the comparison function argument to the standard C library `qsort` function is only used while `qsort` is working, and no additional references to the comparison function are kept. Use this option only in such cases, when no holding is necessary and you want to avoid the extra cost.
 - A box holding `#f` or any other non-list value — the callback value is stored in the box, overriding any non-list value that was in the box (making it useful for holding a single callback value). When you know that the callback is no longer needed, you can “release” the callback value by changing the box contents or by allowing the box itself to become unreachable. This mode can be useful if the box is held for a dynamic extent that corresponds to when the callback is needed; for example, you might encapsulate some foreign functionality in a Racket class or a unit, and keep the callback box as a field in new instances or instantiations of the unit.
 - A box holding `null` (or any list) — similar to a box holding a non-list value, except that new callback values are `consed` onto the contents of the box. This mode is therefore useful in cases when a Racket function is used in multiple callbacks (that is, sent to foreign code to hold onto multiple times) and all callbacks should be retained together.
 - A one-argument function — the function is invoked with the callback value when it is generated. This mode allows you to explicitly manage reachability of the generated callback closure.
- If `wrapper` is not `#f`, it takes the procedure to be converted into a callback and returns a replacement procedure to be invoked as the callback. Thus, `wrapper` acts a hook to perform various argument manipulations before a Racket callback function is called, and it can return different results to the foreign caller.

The callback value’s reachability (and its interaction with `keep` is based on the original function for the callback, not the result of `wrapper`).

- If `atomic?` is true, then when a Racket procedure is given this procedure type and called as a callback from foreign code, then the Racket process is put into atomic mode while evaluating the Racket procedure body.

In atomic mode, other Racket threads do not run, so the Racket code must not call any function that potentially blocks on synchronization with other threads, or else it may lead to deadlock. In addition, the Racket code must not perform any potentially blocking operation (such as *I/O*), it must not raise an uncaught exception, it must not perform any escaping continuation jumps, and its non-tail recursion must be minimal to avoid C-level stack overflow; otherwise, the process may crash or misbehave.

- If an `async-apply` procedure is provided, then a Racket callback procedure with the generated procedure type can be applied in a foreign thread (i.e., an OS-level thread other than the one used to run Racket). The call in the foreign thread is transferred to the OS-level thread that runs Racket, but the Racket-level thread (in the sense of

`thread`) is unspecified; the job of the provided `async-apply` procedure is to arrange for the callback procedure to be run in a suitable Racket thread.

The given `async-apply` procedure is applied to a thunk that encapsulates the specific callback invocation, and the foreign OS-level thread blocks until the thunk is called and completes; the thunk must be called exactly once, and the callback invocation must return normally. The given `async-apply` procedure itself is called in atomic mode (see `atomic?` above).

If the callback is known to complete quickly, requires no synchronization, and works independent of the Racket thread in which it runs, then it is safe for the given `async-apply` procedure to apply the thunk directly. Otherwise, the given `async-apply` procedure must arrange for the thunk to be applied in a suitable Racket thread sometime after the given `async-apply` procedure itself returns; if the thunk raises an exception or synchronizes within an unsuitable Racket-level thread, it can deadlock or otherwise damage the Racket process.

Foreign-thread detection to trigger `async-apply` works only when Racket is compiled with OS-level thread support, which is the default for many platforms. If a callback with an `async-apply` is called from foreign code in the same OS-level thread that runs Racket, then the `async-apply` wrapper is not used.

```
(_fun fun-option ... maybe-args type-spec ... -> type-spec
  maybe-wrapper)

fun-option = #:abi abi-expr
             | #:save-errno save-errno-expr
             | #:keep keep-expr
             | #:atomic? atomic?-expr
             | #:async-apply async-apply-expr
             | #:in-original-place? in-original-place?-expr

maybe-args =
  | (id ...) ::
  | id ::
  | (id ... . id) ::

type-spec = type-expr
           | (id : type-expr)
           | (type-expr = value-expr)
           | (id : type-expr = value-expr)

maybe-wrapper =
  | -> output-expr
```

Creates a new function type. The `_fun` form is a convenient syntax for the `_cprocedure` type constructor. In its simplest form, only the input `type-exprs` and the output `type-`

expr are specified, and each types is a simple expression, which creates a straightforward function type.

For example,

```
(_fun _int _string -> _int)
```

specifies a function that receives an integer and a string, and returns an integer.

In its full form, the `_fun` syntax provides an IDL-like language that can be used to create a wrapper function around the primitive foreign function. These wrappers can implement complex foreign interfaces given simple specifications. The full form of each of the type specifications can include an optional label and an expression. If a `= value-expr` is provided, then the resulting function will be a wrapper that calculates the argument for that position itself, meaning that it does not expect an argument for that position. The expression can use previous arguments if they were labeled with `id :`. In addition, the result of a function call need not be the value returned from the foreign call: if the optional `output-expr` is specified, or if an expression is provided for the output type, then this specifies an expression that will be used as a return value. This expression can use any of the previous labels, including a label given for the output which can be used to access the actual foreign return value.

In rare cases where complete control over the input arguments is needed, the wrapper's argument list can be specified as `maybe-args`, in any form (including a "rest" argument). Identifiers in this place are related to type labels, so if an argument is there is no need to use an expression.

For example,

```
(_fun (n s) :: (s : _string) (n : _int) -> _int)
```

specifies a function that receives an integer and a string, but the foreign function receives the string first.

```
(function-ptr ptr-or-proc fun-type) → cpointer?  
ptr-or-proc : (or cpointer? procedure?)  
fun-type : ctype?
```

Casts `ptr-or-proc` to a function pointer of type `fun-type`.

```
->
```

A literal used in `_fun` forms. (It's unfortunate that this literal has the same name as `->` from `racket/contract`, but it's a different binding.)

3.6.1 Custom Function Types

The behavior of the `_fun` type can be customized via *custom function types*, which are pieces of syntax that can behave as C types and C type constructors, but they can interact with function calls in several ways that are not possible otherwise. When the `_fun` form is expanded, it tries to expand each of the given type expressions, and ones that expand to certain keyword-value lists interact with the generation of the foreign function wrapper. This expansion makes it possible to construct a single wrapper function, avoiding the costs involved in compositions of higher-order functions.

Custom function types are macros that expand to a sequence (`key: val ...`), where each `key:` is from a short list of known keys. Each key interacts with generated wrapper functions in a different way, which affects how its corresponding argument is treated:

- `type:` specifies the foreign type that should be used, if it is `#f` then this argument does not participate in the foreign call.
- `expr:` specifies an expression to be used for arguments of this type, removing it from wrapper arguments.
- `bind:` specifies a name that is bound to the original argument if it is required later (e.g., `_box` converts its associated value to a C pointer, and later needs to refer back to the original box).
- `1st-arg:` specifies a name that can be used to refer to the first argument of the foreign call (good for common cases where the first argument has a special meaning, e.g., for method calls).
- `prev-arg:` similar to `1st-arg:`, but refers to the previous argument.
- `pre:` a pre-foreign code chunk that is used to change the argument's value.
- `post:` a similar post-foreign code chunk.
- `keywords:` specifies keyword/value expressions that will be used with the surrounding `_fun` form. (Note: the keyword/value sequence follows `keywords:`, not parenthesized.)

The `pre:` and `post:` bindings can be of the form (`id => expr`) to use the existing value. Note that if the `pre:` expression is not (`id => expr`), then it means that there is no input for this argument to the `_fun`-generated procedure. Also note that if a custom type is used as an output type of a function, then only the `post:` code is used.

Most custom types are meaningful only in a `_fun` context, and will raise a syntax error if used elsewhere. A few such types can be used in non-`_fun` contexts: types which use only `type:`, `pre:`, `post:`, and no others. Such custom types can be used outside a `_fun` by expanding them into a usage of `make-ctype`, using other keywords makes this impossible, because it means that the type has specific interaction with a function call.

```
(define-fun-syntax id transformer-expr)
```

Binds *id* as a custom function type as well as a syntax transformer (i.e, macro). The type is expanded by applying the procedure produced by *transformer-expr* to a use of the custom function type.

For instance, the following defines a new type that automatically coerces the input number to an inexact form which is compatible with the `_float` type.

```
(define-fun-syntax _float*  
  (syntax-id-rules (_float*)  
    [(_float*) (type: _float pre: (x => (+ 0.0 x)))]))  
  
(_fun _float* -> _bool)
```

```
_|_?
```

A custom function type that is a marker for expressions that should not be sent to the foreign function. Use this to bind local values in a computation that is part of an ffi wrapper interface, or to specify wrapper arguments that are not sent to the foreign function (e.g., an argument that is used for processing the foreign output).

```
(_ptr mode type-expr)  
  
mode = i  
      | o  
      | io
```

Creates a C pointer type, where *mode* indicates input or output pointers (or both). The *mode* can be one of the following:

- *i* — indicates an *input* pointer argument: the wrapper arranges for the function call to receive a value that can be used with the `type` and to send a pointer to this value to the foreign function. After the call, the value is discarded.
- *o* — indicates an *output* pointer argument: the foreign function expects a pointer to a place where it will save some value, and this value is accessible after the call, to be used by an extra return expression. If `_ptr` is used in this mode, then the generated wrapper does not expect an argument since one will be freshly allocated before the call.
- *io* — combines the above into an *input/output* pointer argument: the wrapper gets the Racket value, allocates and set a pointer using this value, and then references the value after the call. The “`_ptr`” name can be confusing here: it means that the foreign

function expects a pointer, but the generated wrapper uses an actual value. (Note that if this is used with structs, a struct is created when calling the function, and a copy of the return value is made too—which is inefficient, but ensures that structs are not modified by C code.)

For example, the `_ptr` type can be used in output mode to create a foreign function wrapper that returns more than a single argument. The following type:

```
(_fun (i : (_ptr o _int))
      -> (d : _double)
      -> (values d i))
```

creates a function that calls the foreign function with a fresh integer pointer, and use the value that is placed there as a second return value.

`|` `_box`

A custom function type similar to a `(_ptr io type)` argument, where the input is expected to be a box holding an appropriate value, which is unboxed on entry and modified accordingly on exit.

```
(_list mode type maybe-len)
      mode = i
           | o
           | io
maybe-len =
           | len-expr
```

A custom function type that is similar to `_ptr`, except that it is used for converting lists to/from C vectors. The optional `maybe-len` argument is needed for output values where it is used in the post code, and in the pre code of an output mode to allocate the block. (If the length is 0, then NULL is passed in and an empty list is returned.) In either case, it can refer to a previous binding for the length of the list which the C function will most likely require.

`|` `(_vector mode type maybe-len)`

A custom function type like `_list`, except that it uses Racket vectors instead of lists.

`|` `_bytes`
`(_bytes o len-expr)`

A custom function type that can be used by itself as a simple type for a byte string as a C pointer. Alternatively, the second form is for a pointer return value, where the size should be explicitly specified.

There is no need for other modes analogous to those of `_ptr`: input or input/output would be just like `_bytes`, since the string carries its size information.

3.7 C Struct Types

```
(make-cstruct-type types [abi alignment]) → ctype?  
types : (listof ctype?)  
abi : (or/c #f 'default 'stdcall 'sysv) = #f  
alignment : (or/c #f 1 2 4 8 16) = #f
```

The primitive type constructor for creating new C struct types. These types are actually new primitive types; they have no conversion functions associated. The corresponding Racket objects that are used for structs are pointers, but when these types are used, the value that the pointer *refers to* is used, rather than the pointer itself. This value is basically made of a number of bytes that is known according to the given list of `types` list.

If `alignment` is `#f`, then the natural alignment of each type in `types` is used for its alignment within the struct type. Otherwise, `alignment` is used for all struct type members.

```
(_list-struct [#:alignment alignment]  
             type ...+) → ctype?  
alignment : (or/c #f 1 2 4 8 16) = #f  
type : ctype?
```

A type constructor that builds a struct type using `make-cstruct-type` function and wraps it in a type that marshals a struct as a list of its components. Note that space for structs must to be allocated; the converter for a `_list-struct` type immediately allocates and uses a list from the allocated space, so it is inefficient. Use `define-cstruct` below for a more efficient approach.

```
(define-cstruct id/sup ([field-id type-expr] ...) property ...)  
  
id/sup = _id  
       | (_id super-id)  
  
property = #:alignment alignment-expr  
          | #:property prop-expr val-expr  
          | #:no-equal
```

Defines a new C struct type, but unlike `_list-struct`, the resulting type deals with C structs in binary form, rather than marshaling them to Racket values. The syntax is similar to `define-struct`, providing accessor functions for raw struct values (which are pointer objects); the `_id` must start with `_`, and at most one `#:alignment` can be supplied.

The resulting bindings are as follows:

- `_id` : the new C type for this struct.
- `_id-pointer`: a pointer type that should be used when a pointer to values of this struct are used.
- `_id-pointer/null`: like `_id-pointer`, but allowing NULL pointers (as represented on the Racket side by `#f`).
- `id?`: a predicate for the new type.
- `id-tag`: the tag object that is used with instances. The tag object may be the symbol form of `id` or a list of symbols containing the `id` symbol and other symbols, such as the `super-id` symbol.
- `make-id` : a constructor, which expects an argument for each field.
- `id-field-id` : an accessor function for each `field-id`; if the field has a C struct type, then the result of the accessor is a pointer to the field within the enclosing structure, rather than a copy of the field.
- `set-id-field-id!` : a mutator function for each `field-id`.
- `id`: structure-type information compatible with `struct-out` or `match` (but not `struct` or `define-struct`); currently, this information is correct only when no `super-id` is specified.
- `id->list`, `list->id` : a function that converts a struct into a list of field values and vice versa.
- `id->list*`, `list*->id` : like `id->list`, `list->id`, but fields that are structs are recursively unpacked to lists or packed from lists.
- `struct:cpointer:id`: only when a `#:property` is specified — a structure type that corresponds to a wrapper to reflect properties (see below).

Objects of the new type are actually C pointers, with a type tag that is the symbol form of `id` or a list that contains the symbol form of `id`. Since structs are implemented as pointers, they can be used for a `_pointer` input to a foreign function: their address will be used. To make this a little safer, the corresponding cpointer type is defined as `_id-pointer`. The `_id` type should not be used when a pointer is expected, since it will cause the struct to be copied rather than use the pointer value, leading to memory corruption.

Instances of the new type are not normally Racket structure instances. However, if at least one `#:property` modifier is specified, then struct creation and coercions from `_id` variants wrap a non-NULL C pointer representation in a Racket structure that has the specified properties. The wrapper Racket structure also has a `prop:cpointer` property, so that wrapped C pointers can be treated the same as unwrapped C pointers. If a `super-id` is provided and it corresponds to a C struct type with a wrapper structure type, then the wrapper structure type is a subtype of `super-id`'s wrapper structure type. If a `#:property` modifier is

specified, `#:no-equal` is not specified, and if `prop:equal+hash` is not specified as any `#:property`, then the `prop:equal+hash` property is automatically implemented for the wrapper structure type to use `ptr-equal?`.

If the first field is itself a C struct type, its tag will be used in addition to the new tag. This feature supports common cases of object inheritance, where a sub-struct is made by having a first field that is its super-struct. Instances of the sub-struct can be considered as instances of the super-struct, since they share the same initial layout. Using the tag of an initial C struct field means that the same behavior is implemented in Racket; for example, accessors and mutators of the super-struct can be used with the new sub-struct. See the example below.

Providing a *super-id* is shorthand for using an initial field named *super-id* and using `_super-id` as its type. Thus, the new struct will use `_super-id`'s tag in addition to its own tag, meaning that instances of *_id* can be used as instances of `_super-id`. Aside from the syntactic sugar, the constructor function is different when this syntax is used: instead of expecting a first argument that is an instance of `_super-id`, the constructor will expect arguments for each of `_super-id`'s fields, in addition for the new fields. This adjustment of the constructor is, again, in analogy to using a supertype with `define-struct`.

Structs are allocated as atomic blocks, which means that the garbage collector ignores their content. Thus, struct fields can hold only non-pointer values, pointers to memory outside the GC's control, and otherwise-reachable pointers to immobile GC-managed values (such as those allocated with `malloc` and `'internal` or `'internal-atomic`).

As an example, consider the following C code:

```
typedef struct { int x; char y; } A;
typedef struct { A a; int z; } B;

A* makeA() {
  A *p = malloc(sizeof(A));
  p->x = 1;
  p->y = 2;
  return p;
}

B* makeB() {
  B *p = malloc(sizeof(B));
  p->a.x = 1;
  p->a.y = 2;
  p->z = 3;
  return p;
}

char gety(A* a) {
  return a->y;
}
```

```
}
```

Using the simple `_list-struct`, you might expect this code to work:

```
(define makeB
  (get-ffi-obj 'makeB "foo.so"
    (_fun -> (_list-struct (_list-struct _int _byte) _int))))
(makeB) ; should return '((1 2) 3)
```

The problem here is that `makeB` returns a pointer to the struct rather than the struct itself. The following works as expected:

```
(define makeB
  (get-ffi-obj 'makeB "foo.so" (_fun -> _pointer)))
(ptr-ref (makeB) (_list-struct (_list-struct _int _byte) _int))
```

As described above, `_list-structs` should be used in cases where efficiency is not an issue. We continue using `define-cstruct`, first define a type for `A` which makes it possible to use `makeA`:

```
(define-cstruct _A ([x _int] [y _byte]))
(define makeA
  (get-ffi-obj 'makeA "foo.so"
    (_fun -> _A-pointer))) ; using _A is a memory-corrupting bug!
(define a (makeA))
(list a (A-x a) (A-y a))
; produces an A containing 1 and 2
```

Using `gety` is also simple:

```
(define gety
  (get-ffi-obj 'gety "foo.so"
    (_fun _A-pointer -> _byte)))
(gety a) ; produces 2
```

We now define another `C` struct for `B`, and expose `makeB` using it:

```
(define-cstruct _B ([a _A] [z _int]))
(define makeB
  (get-ffi-obj 'makeB "foo.so"
    (_fun -> _B-pointer)))
(define b (makeB))
```

We can access all values of `b` using a naive approach:

```
(list (A-x (B-a b)) (A-y (B-a b)) (B-z b))
```

but this is inefficient as it allocates and copies an instance of A on every access. Inspecting the tags (`cpointer-tag b`) we can see that A's tag is included, so we can simply use its accessors and mutators, as well as any function that is defined to take an A pointer:

```
(list (A-x b) (A-y b) (B-z b))  
(gety b)
```

Constructing a B instance in Racket requires allocating a temporary A struct:

```
(define b (make-B (make-A 1 2) 3))
```

To make this more efficient, we switch to the alternative `define-cstruct` syntax, which creates a constructor that expects arguments for both the super fields and the new ones:

```
(define-cstruct (_B _A) ([z _int]))  
(define b (make-B 1 2 3))
```

3.8 C Array Types

```
(make-array-type type count) → ctype?  
type : ctype?  
count : exact-nonnegative-integer?
```

The primitive type constructor for creating new C array types. Like C struct types, array types are new primitive types with no conversion functions associated. When used as a function argument or return type, array types behave like pointer types; otherwise, array types behave like struct types (i.e., a struct with as many fields as the array has elements), particularly when used for a field within a struct type.

Since an array is treated like a struct, `casting` a pointer type to an array type does not work. Instead, use `ptr-ref` with a pointer, an array type constructed with `_array`, and index 0 to convert a pointer to a Racket representation that works with `array-ref` and `array-set!`.

```
(_array type count ...+) → ctype?  
type : ctype?  
count : exact-nonnegative-integer?
```

Creates an array type whose Racket representation is an array that works with `array-ref` and `array-set!`. The array is not copied; the Racket representation is backed by the underlying C representation.

Supply multiple *counts* for a multidimensional array. Since C uses row-major order for arrays, `(_array t n m)` is equivalent to `(_array (_array t m) n)`, which is different from an array of pointers to arrays.

```
(array? v) → boolean?  
v : any/c
```

Returns `#t` if *v* is a Racket representation of a C value via `_array`, `#f` otherwise.

```
(array-ref a i ...+) → any/c  
a : array?  
i : exact-nonnegative-integer?
```

Extracts an element from an array. Use multiple *i* indices for a multidimensional array access; using fewer indices than the array dimension produces a sub-array.

```
(array-set! a i ...+ v) → void?  
a : array?  
i : exact-nonnegative-integer?  
v : any/c
```

Sets an element in an array. Use multiple *i* indices for a multidimensional array update; using fewer indices than the array dimension sets a sub-array (i.e., *v* must be an array of the same size as the sub-array and *v* is copied into the sub-array).

```
(array-ptr a) → cpointer?  
a : array?
```

Extracts the pointer for an array's storage.

```
(array-length a) → exact-nonnegative-integer?  
a : array?
```

Extracts the length of an array. For a multidimensional array, the result is still a single number; extract an element to get a sub-array to get the length of the next dimension, and so on.

```
(_array/list type count ...+) → ctype?  
type : ctype?  
count : exact-nonnegative-integer?
```

Like `_array`, but the Racket representation is a list (or list of lists for a multidimensional array) of elements copied to and from an underlying C array.

```
(_array/vector type count ...+) → ctype?  
  type : ctype?  
  count : exact-nonnegative-integer?
```

Like `_array`, but the Racket representation is a vector (or vector of vectors for a multidimensional array) of elements copied to and from an underlying C array.

3.9 C Union Types

```
(make-union-type type ...+) → ctype?  
  type : ctype?
```

The primitive type constructor for creating new C union types. Like C struct types, union types are new primitive types with no conversion functions associated. Unions are always treated like structs.

```
(_union type ...+) → ctype?  
  type : ctype?
```

Creates a union type whose Racket representation is a union that works with `union-ref` and `union-set!`. The union is not copied; the Racket representation is backed by the underlying C representation.

```
(union? v) → boolean?  
  v : any/c
```

Returns `#t` if `v` is a Racket representation of a C value via `_union`, `#f` otherwise.

```
(union-ref u i) → any/c  
  u : union?  
  i : exact-nonnegative-integer?
```

Extracts a variant from a union.

```
(union-set! u i v) → void?  
  u : union?  
  i : exact-nonnegative-integer?  
  v : any/c
```

Sets a variant in a union..

```
(union-ptr u) → cpointer?  
  u : array?
```

Extracts the pointer for a union's storage.

3.10 Enumerations and Masks

Although the constructors below are describes as procedures, they are implemented as syntax, so that error messages can report a type name where the syntactic context implies one.

```
(_enum symbols [basetype #:unknown unknown]) → ctype?  
  symbols : list?  
  basetype : ctype? = _ufixint  
  unknown : any/c = (lambda (x) (error ....))
```

Takes a list of symbols and generates an enumeration type. The enumeration maps between a symbol in the given *symbols* list and corresponding integers, counting from 0.

The list *symbols* can also set the values of symbols by putting '=' and an exact integer after the symbol. For example, the list '(x y = 10 z)' maps 'x' to 0, 'y' to 10, and 'z' to 11.

The *basetype* argument specifies the base type to use.

The *unknown* argument specifies the result of converting an unknown integer from the foreign side: it can be a one-argument function to be applied on the integer, or a value to return instead. The default is to throw an exception.

```
(_bitmask symbols [basetype]) → ctype?  
  symbols : (or symbol? list?)  
  basetype : ctype? = _uint
```

Similar to `_enum`, but the resulting mapping translates a list of symbols to a number and back, using `bitwise-ior`. A single symbol is equivalent to a list containing just the symbol. The default *basetype* is `_uint`, since high bits are often used for flags.

4 Pointer Functions

```
(cpointer? v) → boolean?  
v : any/c
```

Returns `#t` if `v` is a C pointer or a value that can be used as a pointer: `#f` (used as a NULL pointer), byte strings (used as memory blocks), or a structure instance with the `prop:cpointer` structure type property. Returns `#f` for other values.

```
(ptr-equal? cptr1 cptr2) → boolean?  
cptr1 : cpointer?  
cptr2 : cpointer?
```

Compares the values of the two pointers. Two different Racket pointer objects can contain the same pointer.

If the values are both pointers that are not represented by `#f`, a byte string, a callback, a pointer based on `_fpointer`, or a structure with the `prop:cpointer` property, then the `ptr-equal?` comparison is the same as using `equal?`.

```
(ptr-add cptr offset [type]) → cpointer?  
cptr : cpointer?  
offset : exact-integer?  
type : ctype? = _byte
```

Returns a cpointer that is like `cptr` offset by `offset` instances of `ctype`.

The resulting cpointer keeps the base pointer and offset separate. The two pieces are combined at the last minute before any operation on the pointer, such as supplying the pointer to a foreign function. In particular, the pointer and offset are not combined until after all allocation leading up to a foreign-function call; if the called function does not itself call anything that can trigger a garbage collection, it can safely use pointers that are offset into the middle of a GCable object.

```
(offset-ptr? cptr) → boolean?  
cptr : cpointer?
```

A predicate for cpointers that have an offset, such as pointers that were created using `ptr-add`. Returns `#t` even if such an offset happens to be 0. Returns `#f` for other cpointers and non-cpointers.

```
(ptr-offset cptr) → exact-integer?  
cptr : cpointer?
```

Returns the offset of a pointer that has an offset. The resulting offset is always in bytes.


```
(cpointer-gcable? cptr) → boolean?  
  cptr : cpointer?
```

Returns `#t` if `cptr` is treated as a reference to memory that is managed by the garbage collector, `#f` otherwise.

4.1 Pointer Dereferencing

```
(set-ptr-offset! cptr offset [ctype]) → void?  
  cptr : cpointer?  
  offset : exact-integer?  
  ctype : ctype? = _byte
```

Sets the offset component of an offset pointer. The arguments are used in the same way as `ptr-add`. If `cptr` has no offset, the `exn:fail:contract` exception is raised.

```
(ptr-add! cptr offset [ctype]) → void?  
  cptr : cpointer?  
  offset : exact-integer?  
  ctype : ctype? = _byte
```

Like `ptr-add`, but destructively modifies the offset contained in a pointer. The same operation could be performed using `ptr-offset` and `set-ptr-offset!`.

```
(ptr-ref cptr type [offset]) → any  
  cptr : cpointer?  
  type : ctype?  
  offset : exact-nonnegative-integer? = 0  
(ptr-ref cptr type abs-tag offset) → any  
  cptr : cpointer?  
  type : ctype?  
  abs-tag : (one-of/c 'abs)  
  offset : exact-nonnegative-integer?  
(ptr-set! cptr type val) → void?  
  cptr : cpointer?  
  type : ctype?  
  val : any/c  
(ptr-set! cptr type offset val) → void?  
  cptr : cpointer?  
  type : ctype?  
  offset : exact-nonnegative-integer?  
  val : any/c
```

```
(ptr-set! cptr type abs-tag offset val) → void?
  cptr : cpointer?
  type : ctype?
  abs-tag : (one-of/c 'abs)
  offset : exact-nonnegative-integer?
  val : any/c
```

The `ptr-ref` procedure returns the object referenced by `cptr`, using the given `type`. The `ptr-set!` procedure stores the `val` in the memory `cptr` points to, using the given `type` for the conversion.

In each case, `offset` defaults to 0 (which is the only value that should be used with `ffi-obj` objects, see §7 “Unexported Primitive Functions”). If an `offset` index is non-0, the value is read or stored at that location, considering the pointer as a vector of `types` — so the actual address is the pointer plus the size of `type` multiplied by `offset`. In addition, a `'abs` flag can be used to use the `offset` as counting bytes rather than increments of the specified `type`.

Beware that the `ptr-ref` and `ptr-set!` procedure do not keep any meta-information on how pointers are used. It is the programmer’s responsibility to use this facility only when appropriate. For example, on a little-endian machine:

```
> (define block (malloc _int 5))
> (ptr-set! block _int 0 196353)
> (map (lambda (i) (ptr-ref block _byte i)) '(0 1 2 3))
(1 255 2 0)
```

In addition, `ptr-ref` and `ptr-set!` cannot detect when offsets are beyond an object’s memory bounds; out-of-bounds access can easily lead to a segmentation fault or memory corruption.

```
(memmove cptr src-cptr count [type]) → void?
  cptr : cpointer?
  src-cptr : cpointer?
  count : exact-nonnegative-integer?
  type : ctype? = _byte
(memmove cptr offset src-cptr count [type]) → void?
  cptr : cpointer?
  offset : exact-integer?
  src-cptr : cpointer?
  count : exact-nonnegative-integer?
  type : ctype? = _byte
```

```

(memmove cptr
         offset
         src-cptr
         src-offset
         count
         [type]) → void?
cptr : cpointer?
offset : exact-integer?
src-cptr : cpointer?
src-offset : exact-integer?
count : exact-nonnegative-integer?
type : ctype? = _byte

```

Copies to *cptr* from *src-cptr*. The destination pointer can be offset by an optional *offset*, which is in *type* instances. The source pointer can be similarly offset by *src-offset*. The number of bytes copied from source to destination is determined by *count*, which is in *type* instances when supplied.

```

(memcpy cptr src-cptr count [type]) → void?
cptr : cpointer?
src-cptr : cpointer?
count : exact-nonnegative-integer?
type : ctype? = _byte
(memcpy cptr offset src-cptr count [type]) → void?
cptr : cpointer?
offset : exact-integer?
src-cptr : cpointer?
count : exact-nonnegative-integer?
type : ctype? = _byte
(memcpy cptr
         offset
         src-cptr
         src-offset
         count
         [type]) → void?
cptr : cpointer?
offset : exact-integer?
src-cptr : cpointer?
src-offset : exact-integer?
count : exact-nonnegative-integer?
type : ctype? = _byte

```

Like `memmove`, but the result is undefined if the destination and source overlap.

```

(memset cptr byte count [type]) → void?
cptr : cpointer?

```

```

byte : byte?
count : exact-nonnegative-integer?
type : ctype? = _byte
(memset cptr offset byte count [type]) → void?
cptr : cpointer?
offset : exact-integer?
byte : byte?
count : exact-nonnegative-integer?
type : ctype? = _byte

```

Similar to `memmove`, but the destination is uniformly filled with `byte` (i.e., an exact integer between 0 and 255 inclusive). When a `type` argument is present, the result is that of a call to `memset` with no `type` argument and the `count` multiplied by the size associated with the `type`.

```

(cpointer-tag cptr) → any
cptr : cpointer?

```

Returns the Racket object that is the tag of the given `cptr` pointer.

```

(set-cpointer-tag! cptr tag) → void?
cptr : cpointer?
tag : any/c

```

Sets the tag of the given `cptr`. The `tag` argument can be any arbitrary value; other pointer operations ignore it. When a `cpointer` value is printed, its tag is shown if it is a symbol, a byte string, a string. In addition, if the tag is a pair holding one of these in its `car`, the `car` is shown (so that the tag can contain other information).

4.2 Memory Management

For general information on C-level memory management with Racket, see *Inside: Racket C API*.

```

(malloc bytes-or-type
 [type-or-bytes
  cptr
  mode
  fail-mode]) → cpointer?
bytes-or-type : (or/c (and/c exact-nonnegative-integer? fixnum?)
                      ctype?)
type-or-bytes : (or/c (and/c exact-nonnegative-integer? fixnum?)
                      ctype?)
               = absent

```

```

cptr : cpointer? = absent
      (one-of/c 'nonatomic 'stubborn 'uncollectable
mode :      'eternal 'interior 'atomic-interior = absent
          'raw)
fail-mode : (one-of/c 'failok) = absent

```

Allocates a memory block of a specified size using a specified allocation. The result is a `cpointer` to the allocated memory, or `#f` if the requested size is zero. Although not reflected above, the four arguments can appear in any order, since they are all different types of Racket objects; a size specification is required at minimum:

- If a C type `bytes-or-type` is given, its size is used to the block allocation size.
- If an integer `bytes-or-type` is given, it specifies the required size in bytes.
- If both `bytes-or-type` and `type-or-bytes` are given, then the allocated size is for a vector of values (the multiplication of the size of the C type and the integer).
- If a `cptr` pointer is given, its content is copied to the new block.
- A symbol `mode` argument can be given, which specifies what allocation function to use. It should be one of `'nonatomic` (uses `scheme_malloc` from Racket's C API), `'atomic` (`scheme_malloc_atomic`), `'stubborn` (`scheme_malloc_stubborn`), `'uncollectable` (`scheme_malloc_uncollectable`), `'eternal` (`scheme_malloc_eternal`), `'interior` (`scheme_malloc_allow_interior`), `'atomic-interior` (`scheme_malloc_atomic_allow_interior`), or `'raw` (uses the operating system's `malloc`, creating a GC-invisible block).
- If an additional `'failok` flag is given, then `scheme_malloc_fail_ok` is used to wrap the call.

If no mode is specified, then `'nonatomic` allocation is used when the type is a `_gcpointer-` or `_scheme-`based type, and `'atomic` allocation is used otherwise.

```

(free cptr) → void
cptr : cpointer?

```

Uses the operating system's `free` function for `'raw`-allocated pointers, and for pointers that a foreign library allocated and we should free. Note that this is useful as part of a finalizer (see below) procedure hook (e.g., on the Racket pointer object, freeing the memory when the pointer object is collected, but beware of aliasing).

```

(end-stubborn-change cptr) → void?
cptr : cpointer?

```

Uses `scheme_end_stubborn_change` on the given stubborn-allocated pointer.

```
(malloc-immobile-cell v) → cpointer?  
v : any/c
```

Allocates memory large enough to hold one arbitrary (collectable) Racket value, but that is not itself collectable or moved by the memory manager. The cell is initialized with *v*; use the type `_scheme` with `ptr-ref` and `ptr-set!` to get or set the cell's value. The cell must be explicitly freed with `free-immobile-cell`.

```
(free-immobile-cell cptr) → void?  
cptr : cpointer?
```

Frees an immobile cell created by `malloc-immobile-cell`.

```
(register-finalizer obj finalizer) → void?  
obj : any/c  
finalizer : (any/c . -> . any)
```

Registers a finalizer procedure `finalizer-proc` with the given *obj*, which can be any Racket (GC-able) object. The finalizer is registered with a will executor; see `make-will-executor`. The finalizer is invoked when *obj* is about to be collected. See also `register-custodian-shutdown`.

The finalizer is invoked in a thread that is in charge of triggering will executors for `register-finalizer`. The given *finalizer* procedure should generally not rely on the environment of the triggering thread, such as its output ports or custodians, except that relying on a default logger is reasonable.

Finalizers are mostly intended to be used with `cpointer` objects (for freeing unused memory that is not under GC control), but it can be used with any Racket object—even ones that have nothing to do with foreign code. Note, however, that the finalizer is registered for the *Racket* object that represents the pointer. If you intend to free a pointer object, then you must be careful to not register finalizers for two `cpointers` that point to the same address. Also, be careful to not make the finalizer a closure that holds on to the object.

For example, suppose that you're dealing with a foreign function that returns a C string that you should free. Here is an attempt at creating a suitable type:

```
(define _sixteen-bytes/free  
  (make-ctype _pointer  
    #f ; a Racket bytes can be used as a pointer  
    (lambda (x)  
      (let ([b (make-sized-byte-string x 16)])  
        (register-finalizer x free)  
        b))))
```

The above code is wrong: the finalizer is registered for `x`, which is no longer needed after the byte string is created. Changing the example to register the finalizer for `b` corrects the problem, but then `free` is invoked on `b` instead of on `x`. In the process of fixing this problem, we might be careful and log a message for debugging:

```
(define _sixteen-bytes/free
  (make-ctype _pointer
    #f ; a Racket bytes can be used as a pointer
    (lambda (x)
      (let ([b (make-sized-byte-string x 16)])
        (register-finalizer b
          (lambda (ignored)
            (log-debug (format "Releasing ~s\n" b))
            (free x)))
          b))))
```

Now, we never see any logged event. The problem is that the finalizer is a closure that keeps a reference to `b`. Instead of referencing the value that is finalized, use the input argument to the finalizer; simply changing `ignored` to `b` above solves the problem. (Removing the debugging message also avoids the problem, since the finalization procedure would then not close over `b`.)

```
(make-sized-byte-string cptr length) → bytes?
  cptr : cpointer?
  length : exact-nonnegative-integer?
```

Returns a byte string made of the given pointer and the given length. No copying is done. This can be used as an alternative to make pointer values accessible in Racket when the size is known.

If `cptr` is an offset pointer created by `ptr-add`, the offset is immediately added to the pointer. Thus, this function cannot be used with `ptr-add` to create a substring of a Racket byte string, because the offset pointer would be to the middle of a collectable object (which is not allowed).

4.3 Pointer Structure Property

```
prop:cpointer : struct-type-property?
```

A structure type property that causes instances of a structure type to work as C pointer values. The property value must be either an exact non-negative integer indicating an immutable field in the structure (which must, in turn, be initialized to a C pointer value), a procedure that takes the structure instance and returns a C pointer value, or a C pointer value.

The `prop:cpointer` property allows a structure instance to be used transparently as a C pointer value, or it allows a C pointer value to be transparently wrapped by a structure that may have additional values or properties.

5 Derived Utilities

5.1 Safe Homogenous Vectors

```
(require ffi/vector)
```

Homogenous vectors are similar to C vectors (see §5.2 “Safe C Vectors”), except that they define different types of vectors, each with a fixed element type. An exception is the `u8` family of bindings, which are just aliases for byte-string bindings; for example, `make-u8vector` is an alias for `make-bytes`.

```
(make-u8vector len) → u8vector?  
  len : exact-nonnegative-integer?  
(u8vector val ...) → u8vector?  
  val : byte?  
(u8vector? v) → boolean?  
  v : any/c  
(u8vector-length vec) → exact-nonnegative-integer?  
  vec : u8vector?  
(u8vector-ref vec k) → byte?  
  vec : u8vector?  
  k : exact-nonnegative-integer?  
(u8vector-set! vec k val) → void?  
  vec : u8vector?  
  k : exact-nonnegative-integer?  
  val : byte?  
(list->u8vector lst) → u8vector?  
  lst : (listof byte?)  
(u8vector->list vec) → (listof byte?)  
  vec : u8vector?  
(u8vector->cpointer vec) → cpointer?  
  vec : u8vector?
```

Like `_cvector`, but for vectors of `_uint8` elements. These are aliases for `byte` operations, where `u8vector->cpointer` is the identity function.

```
(_u8vector mode maybe-len)  
_u8vector
```

Like `_cvector`, but for vectors of `_uint8` elements.

```
(make-s8vector len) → s8vector?  
  len : exact-nonnegative-integer?  
(s8vector val ...) → s8vector?
```

```

    val : (integer-in -128 127)
(s8vector? v) → boolean?
  v : any/c
(s8vector-length vec) → exact-nonnegative-integer?
  vec : s8vector?
(s8vector-ref vec k) → (integer-in -128 127)
  vec : s8vector?
  k : exact-nonnegative-integer?
(s8vector-set! vec k val) → void?
  vec : s8vector?
  k : exact-nonnegative-integer?
  val : (integer-in -128 127)
(list->s8vector lst) → s8vector?
  lst : (listof (integer-in -128 127))
(s8vector->list vec) → (listof (integer-in -128 127))
  vec : s8vector?
(s8vector->cpointer vec) → cpointer?
  vec : s8vector?

```

Like `make-vector`, etc., but for `_int8` elements. The `s8vector->cpointer` function extracts a plain pointer to the underlying array.

```

(_s8vector mode maybe-len)
_s8vector

```

Like `_cvector`, but for vectors of `_int8` elements.

```

(make-s16vector len) → s16vector?
  len : exact-nonnegative-integer?
(s16vector val ...) → s16vector?
  val : (integer-in -32768 32767)
(s16vector? v) → boolean?
  v : any/c
(s16vector-length vec) → exact-nonnegative-integer?
  vec : s16vector?
(s16vector-ref vec k) → (integer-in -32768 32767)
  vec : s16vector?
  k : exact-nonnegative-integer?
(s16vector-set! vec k val) → void?
  vec : s16vector?
  k : exact-nonnegative-integer?
  val : (integer-in -32768 32767)
(list->s16vector lst) → s16vector?
  lst : (listof (integer-in -32768 32767))

```

```
(s16vector->list vec) → (listof (integer-in -32768 32767))
  vec : s16vector?
(s16vector->cpointer vec) → cpointer?
  vec : s16vector?
```

Like `make-vector`, etc., but for `_int16` elements. The `s16vector->cpointer` function extracts a plain pointer to the underlying array.

```
(_s16vector mode maybe-len)
_s16vector
```

Like `_cvector`, but for vectors of `_int16` elements.

```
(make-u16vector len) → u16vector?
  len : exact-nonnegative-integer?
(u16vector val ...) → u16vector?
  val : (integer-in 0 65535)
(u16vector? v) → boolean?
  v : any/c
(u16vector-length vec) → exact-nonnegative-integer?
  vec : u16vector?
(u16vector-ref vec k) → (integer-in 0 65535)
  vec : u16vector?
  k : exact-nonnegative-integer?
(u16vector-set! vec k val) → void?
  vec : u16vector?
  k : exact-nonnegative-integer?
  val : (integer-in 0 65535)
(list->u16vector lst) → u16vector?
  lst : (listof (integer-in 0 65535))
(u16vector->list vec) → (listof (integer-in 0 65535))
  vec : u16vector?
(u16vector->cpointer vec) → cpointer?
  vec : u16vector?
```

Like `make-vector`, etc., but for `_uint16` elements. The `u16vector->cpointer` function extracts a plain pointer to the underlying array.

```
(_u16vector mode maybe-len)
_u16vector
```

Like `_cvector`, but for vectors of `_uint16` elements.

```

(make-s32vector len) → s32vector?
  len : exact-nonnegative-integer?
(s32vector val ...) → s32vector?
  val : (integer-in -2147483648 2147483647)
(s32vector? v) → boolean?
  v : any/c
(s32vector-length vec) → exact-nonnegative-integer?
  vec : s32vector?
(s32vector-ref vec k) → (integer-in -2147483648 2147483647)
  vec : s32vector?
  k : exact-nonnegative-integer?
(s32vector-set! vec k val) → void?
  vec : s32vector?
  k : exact-nonnegative-integer?
  val : (integer-in -2147483648 2147483647)
(list->s32vector lst) → s32vector?
  lst : (listof (integer-in -2147483648 2147483647))
(s32vector->list vec)
→ (listof (integer-in -2147483648 2147483647))
  vec : s32vector?
(s32vector->cpointer vec) → cpointer?
  vec : s32vector?

```

Like `make-vector`, etc., but for `_int32` elements. The `s32vector->cpointer` function extracts a plain pointer to the underlying array.

```

(_s32vector mode maybe-len)
_s32vector

```

Like `_cvector`, but for vectors of `_int32` elements.

```

(make-u32vector len) → u32vector?
  len : exact-nonnegative-integer?
(u32vector val ...) → u32vector?
  val : (integer-in 0 4294967295)
(u32vector? v) → boolean?
  v : any/c
(u32vector-length vec) → exact-nonnegative-integer?
  vec : u32vector?
(u32vector-ref vec k) → (integer-in 0 4294967295)
  vec : u32vector?
  k : exact-nonnegative-integer?
(u32vector-set! vec k val) → void?

```

```

vec : u32vector?
k : exact-nonnegative-integer?
val : (integer-in 0 4294967295)
(list->u32vector lst) → u32vector?
lst : (listof (integer-in 0 4294967295))
(u32vector->list vec) → (listof (integer-in 0 4294967295))
vec : u32vector?
(u32vector->cpointer vec) → cpointer?
vec : u32vector?

```

Like `make-vector`, etc., but for `_uint32` elements. The `u32vector->cpointer` function extracts a plain pointer to the underlying array.

```

(_u32vector mode maybe-len)
_u32vector

```

Like `_cvector`, but for vectors of `_uint32` elements.

```

(make-s64vector len) → s64vector?
len : exact-nonnegative-integer?
(s64vector val ...) → s64vector?
val : (integer-in -9223372036854775808 9223372036854775807)
(s64vector? v) → boolean?
v : any/c
(s64vector-length vec) → exact-nonnegative-integer?
vec : s64vector?
(s64vector-ref vec k)
→ (integer-in -9223372036854775808 9223372036854775807)
vec : s64vector?
k : exact-nonnegative-integer?
(s64vector-set! vec k val) → void?
vec : s64vector?
k : exact-nonnegative-integer?
val : (integer-in -9223372036854775808 9223372036854775807)
(list->s64vector lst) → s64vector?
lst : (listof (integer-in -9223372036854775808 9223372036854775807))
(s64vector->list vec)
→ (listof (integer-in -9223372036854775808 9223372036854775807))
vec : s64vector?
(s64vector->cpointer vec) → cpointer?
vec : s64vector?

```

Like `make-vector`, etc., but for `_int64` elements. The `s64vector->cpointer` function extracts a plain pointer to the underlying array.

```
(_s64vector mode maybe-len)
_s64vector
```

Like `_cvector`, but for vectors of `_int64` elements.

```
(make-u64vector len) → u64vector?
  len : exact-nonnegative-integer?
(u64vector val ...) → u64vector?
  val : (integer-in 0 18446744073709551615)
(u64vector? v) → boolean?
  v : any/c
(u64vector-length vec) → exact-nonnegative-integer?
  vec : u64vector?
(u64vector-ref vec k) → (integer-in 0 18446744073709551615)
  vec : u64vector?
  k : exact-nonnegative-integer?
(u64vector-set! vec k val) → void?
  vec : u64vector?
  k : exact-nonnegative-integer?
  val : (integer-in 0 18446744073709551615)
(list->u64vector lst) → u64vector?
  lst : (listof (integer-in 0 18446744073709551615))
(u64vector->list vec)
→ (listof (integer-in 0 18446744073709551615))
  vec : u64vector?
(u64vector->cpointer vec) → cpointer?
  vec : u64vector?
```

Like `make-vector`, etc., but for `_uint64` elements. The `u64vector->cpointer` function extracts a plain pointer to the underlying array.

```
(_u64vector mode maybe-len)
_u64vector
```

Like `_cvector`, but for vectors of `_uint64` elements.

```
(make-f32vector len) → f32vector?
  len : exact-nonnegative-integer?
(f32vector val ...) → f32vector?
  val : real?
(f32vector? v) → boolean?
  v : any/c
```

```

(f32vector-length vec) → exact-nonnegative-integer?
  vec : f32vector?
(f32vector-ref vec k) → real?
  vec : f32vector?
  k : exact-nonnegative-integer?
(f32vector-set! vec k val) → void?
  vec : f32vector?
  k : exact-nonnegative-integer?
  val : real?
(list->f32vector lst) → f32vector?
  lst : (listof real?)
(f32vector->list vec) → (listof real?)
  vec : f32vector?
(f32vector->cpointer vec) → cpointer?
  vec : f32vector?

```

Like `make-vector`, etc., but for `_float` elements. The `f32vector->cpointer` function extracts a plain pointer to the underlying array.

```

(_f32vector mode maybe-len)
_f32vector

```

Like `_cvector`, but for vectors of `_float` elements.

```

(make-f64vector len) → f64vector?
  len : exact-nonnegative-integer?
(f64vector val ...) → f64vector?
  val : real?
(f64vector? v) → boolean?
  v : any/c
(f64vector-length vec) → exact-nonnegative-integer?
  vec : f64vector?
(f64vector-ref vec k) → real?
  vec : f64vector?
  k : exact-nonnegative-integer?
(f64vector-set! vec k val) → void?
  vec : f64vector?
  k : exact-nonnegative-integer?
  val : real?
(list->f64vector lst) → f64vector?
  lst : (listof real?)
(f64vector->list vec) → (listof real?)
  vec : f64vector?
(f64vector->cpointer vec) → cpointer?

```

```
vec : f64vector?
```

Like `make-vector`, etc., but for `_double*` elements. The `f64vector->cpointer` function extracts a plain pointer to the underlying array.

```
(_f64vector mode maybe-len)  
_f64vector
```

Like `_cvector`, but for vectors of `_double*` elements.

```
(make-f80vector len) → f80vector?  
  len : exact-nonnegative-integer?  
(f80vector val ...) → f80vector?  
  val : extflonum?  
(f80vector? v) → boolean?  
  v : any/c  
(f80vector-length vec) → exact-nonnegative-integer?  
  vec : f80vector?  
(f80vector-ref vec k) → extflonum?  
  vec : f80vector?  
  k : exact-nonnegative-integer?  
(f80vector-set! vec k val) → void?  
  vec : f80vector?  
  k : exact-nonnegative-integer?  
  val : extflonum?  
(list->f80vector lst) → f80vector?  
  lst : (listof extflonum?)  
(f80vector->list vec) → (listof extflonum?)  
  vec : f80vector?  
(f80vector->cpointer vec) → cpointer?  
  vec : f80vector?
```

Like `make-vector`, etc., but for `_longdouble` elements. The `f80vector->cpointer` function extracts a plain pointer to the underlying array.

```
(_f80vector mode maybe-len)  
_f80vector
```

Like `_cvector`, but for vectors of `_longdouble` elements.

5.2 Safe C Vectors

```
(require ffi/cvector)
```



```
(require ffi/unsafe/cvector)
```

The `ffi/unsafe/cvector` library exports the bindings of this section. The `ffi/cvector` library exports the same bindings, except for the unsafe `make-cvector*` operation.

The `cvector` form can be used as a type C vectors (i.e., a pointer to a memory block).

```
(_cvector mode type maybe-len)  
_cvector
```

Like `_bytes`, `_cvector` can be used as a simple type that corresponds to a pointer that is managed as a safe C vector on the Racket side. The longer form behaves similarly to the `_list` and `_vector` custom types, except that `_cvector` is more efficient; no Racket list or vector is needed.

```
(make-cvector type length) → cvector?  
  type : ctype?  
  length : exact-nonnegative-integer?
```

Allocates a C vector using the given `type` and `length`. The resulting vector is not guaranteed to contain any particular values.

```
(cvector type val ...) → cvector?  
  type : ctype?  
  val : any/c
```

Creates a C vector of the given `type`, initialized to the given list of `vals`.

```
(cvector? v) → boolean?  
  v : any/c
```

Returns `#t` if `v` is a C vector, `#f` otherwise.

```
(cvector-length cvec) → exact-nonnegative-integer?  
  cvec : cvector?
```

Returns the length of a C vector.

```
(cvector-type cvec) → ctype?  
  cvec : cvector?
```

Returns the C type object of a C vector.

```
(cvector-ptr cvec) → cpointer?  
  cvec : cvector?
```

Returns the pointer that points at the beginning block of the given C vector.

```
(cvector-ref cvec k) → any
  cvec : cvector?
  k : exact-nonnegative-integer?
```

References the *k*th element of the *cvec* C vector. The result has the type that the C vector uses.

```
(cvector-set! cvec k val) → void?
  cvec : cvector?
  k : exact-nonnegative-integer?
  val : any
```

Sets the *k*th element of the *cvec* C vector to *val*. The *val* argument should be a value that can be used with the type that the C vector uses.

```
(cvector->list cvec) → list?
  cvec : cvector?
```

Converts the *cvec* C vector object to a list of values.

```
(list->cvector lst type) → cvector?
  lst : list?
  type : ctype?
```

Converts the list *lst* to a C vector of the given *type*.

```
(make-cvector* cptr type length) → cvector?
  cptr : any/c
  type : ctype?
  length : exact-nonnegative-integer?
```

Constructs a C vector using an existing pointer object. This operation is not safe, so it is intended to be used in specific situations where the *type* and *length* are known.

5.3 Tagged C Pointer Types

The unsafe `cpointer-has-tag?` and `cpointer-push-tag!` operations manage tags to distinguish pointer types.

```
(_cpointer tag
  [ptr-type
  scheme-to-c
  c-to-scheme]) → ctype
```

```

tag : any/c
ptr-type : ctype? = _pointer
scheme-to-c : (any/c . -> . any/c) = values
c-to-scheme : (any/c . -> . any/c) = values
(_cpointer/null tag
  [ptr-type
   scheme-to-c
   c-to-scheme]) → ctype
tag : any/c
ptr-type : ctype? = _pointer
scheme-to-c : (any/c . -> . any/c) = values
c-to-scheme : (any/c . -> . any/c) = values

```

Construct a kind of a pointer that gets a specific tag when converted to Racket, and accept only such tagged pointers when going to C. An optional `ptr-type` can be given to be used as the base pointer type, instead of `_pointer`.

Although any value can be used as a tag, by convention the symbol form of a type name—without a leading underscore—is used as the tag. For example, a pointer type `_animal` would normally use `'animal` as the key.

Pointer tags are checked with `cpointer-has-tag?` and changed with `cpointer-push-tag!`, which means that other tags are preserved on an existing pointer value. Specifically, if a base `ptr-type` is given and is itself a `_cpointer`, then the new type will handle pointers that have the new tag in addition to `ptr-type`'s tag(s). When the tag is a pair, its first value is used for printing, so the most recently pushed tag which corresponds to the inheriting type is displayed.

The `_cpointer/null` function is similar to `_cpointer`, except that it tolerates NULL pointers both going to C and back. Note that NULL pointers are represented as `#f` in Racket, so they are not tagged.

```

(define-cpointer-type _id)
(define-cpointer-type _id ptr-type-expr)
(define-cpointer-type _id ptr-type-expr
  scheme-to-c-expr c-to-scheme-expr)

```

A macro version of `_cpointer` and `_cpointer/null`, using the defined name for a tag symbol, and defining a predicate too. The `_id` must start with `_`.

The optional expressions produce optional arguments to `_cpointer`.

In addition to defining `_id` to a type generated by `_cpointer`, `_id/null` is bound to a type produced by `_cpointer/null` type. Finally, `id?` is defined as a predicate, and `id-tag` is defined as an accessor to obtain a tag. The tag is the symbol form of `id`.

```

(cpointer-has-tag? cptr tag) → boolean?
  cptr : any/c
  tag : any/c
(cpointer-push-tag! cptr tag) → void
  cptr : any/c
  tag : any/c

```

These two functions treat pointer tags as lists of tags. As described in §4 “Pointer Functions”, a pointer tag does not have any role, except for Racket code that uses it to distinguish pointers; these functions treat the tag value as a list of tags, which makes it possible to construct pointer types that can be treated as other pointer types, mainly for implementing inheritance via upcasts (when a struct contains a super struct as its first element).

The `cpointer-has-tag?` function checks whether if the given `cptr` has the `tag`. A pointer has a tag `tag` when its tag is either `eq?` to `tag` or a list that contains (in the sense of `memq`) `tag`.

The `cpointer-push-tag!` function pushes the given `tag` value on `cptr`'s tags. The main properties of this operation are: (a) pushing any tag will make later calls to `cpointer-has-tag?` succeed with this tag, and (b) the pushed tag will be used when printing the pointer (until a new value is pushed). Technically, pushing a tag will simply set it if there is no tag set, otherwise push it on an existing list or an existing value (treated as a single-element list).

5.4 Defining Bindings

```

(require ffi/unsafe/define)

(define-ffi-definer define-id ffi-lib-expr
  option ...)

option = #:provide provide-id
         | #:define core-define-id
         | #:default-make-fail default-make-fail-expr

```

Binds `define-id` as a definition form to extract bindings from the library produced by `ffi-lib-expr`. The syntax of `define-id` is

```

(define-id id type-expr
  bind-option ...)

bind-option = #:c-id c-id
             | #:wrap wrap-expr
             | #:make-fail make-fail-expr
             | #:fail fail-expr

```

A `define-id` form binds `id` by extracting a binding with the name `c-id` from the library produced by `ffi-lib-expr`, where `c-id` defaults to `id`. The other options support further wrapping and configuration:

- Before the extracted result is bound as `id`, it is passed to the result of `wrap-expr`, which defaults to `values`. Expressions such as `(allocator delete)` or `(deallocator)` are useful as `wrap-exprs`.
- The `#:make-fail` and `#:fail` options are mutually exclusive; if `make-fail-expr` is provided, it is applied to `id` to obtain the last argument to `get-ffi-obj`; if `fail-expr` is provided, it is supplied directly as the last argument to `get-ffi-obj`. The `make-not-available` function is useful as `make-fail-expr` to cause a use of `id` to report an error when it is applied if `c-id` was not found in the foreign library.

If `provide-id` is provided to `define-ffi-definer`, then `define-id` also provides its binding using `provide-id`. The `provide-protected` form is usually a good choice for `provide-id`.

If `core-define-id` is provided to `define-ffi-definer`, then `code-define-id` is used in place of `define` in the expansion of `define-id` for each binding.

If `default-make-fail-expr` is provided to `define-ffi-definer`, it serves as the default `#:make-fail` value for `define-id`.

For example,

```
(define-ffi-definer define-gtk gtk-lib)
```

binds `define-gtk` to extract FFI bindings from `gtk-lib`, so that `gtk_rc_parse` could be bound as

```
(define-gtk gtk_rc_parse (_fun _path -> _void))
```

If `gtk_rc_parse` is not found, then `define-gtk` reports an error immediately. If `define-gtk` is instead defined with

```
(define-ffi-definer define-gtk gtk-lib
  #:default-make-fail make-not-available)
```

then if `gtk_rc_parse` is not found in `gtk-lib`, an error is reported only when `gtk_rc_parse` is called.

```
(make-not-available name) → (:#rest list? -> any/c)
  name : symbol?
```

Returns a procedure that takes any number of arguments and reports an error message from *name*. This function is intended for using with `#:make-fail` or `#:default-make-fail` in `define-ffi-definer`

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

Equivalent to `(provide (protect-out provide-spec ...))`. The `provide-protected` identifier is useful with `#:provide` in `define-ffi-definer`.

5.5 Allocation and Finalization

```
(require ffi/unsafe/alloc)
```

The `ffi/unsafe/alloc` library provides utilities for ensuring that values allocated through foreign functions are reliably deallocated.

```
((allocator dealloc) alloc) → procedure?  
dealloc : (any/c . -> . any)  
alloc : procedure?
```

Produces a procedure that behaves like `alloc`, but the result of `alloc` is given a finalizer that calls `dealloc` on the result if it is not otherwise freed through a deallocator (as designated with `deallocator`). In addition, `alloc` is called in atomic mode (see `start-atomic`); its result is received and registered in atomic mode, so that the result is reliably freed.

The `dealloc` procedure itself need not be specifically designated a deallocator (via `deallocator`). If a deallocator is called explicitly, it need not be the same as `dealloc`.

```
((deallocator [get-arg]) dealloc) → procedure?  
get-arg : (list? . -> . any/c) = car  
dealloc : procedure?  
(releaser [get-arg]) dealloc) → procedure?  
get-arg : (list? . -> . any/c) = car  
dealloc : procedure?
```

Produces a procedure that behaves like `dealloc`. The `dealloc` procedure is called in atomic mode (see `start-atomic`), and the reference count on one of its arguments is decremented; if the reference count reaches zero, no finalizer associated by an `allocator`- or `referencer`-wrapped procedure is invoked when the value becomes inaccessible.

The optional `get-arg` procedure determines which of `dealloc`'s arguments correspond to the released object; `get-arg` receives a list of arguments passed to `dealloc`, so the default `car` selects the first one.

The `releaser` procedure is a synonym for `deallocator`.

```

((retainer release [get-arg]) retain) → procedure?
  release : (any/c . -> . any)
  get-arg : (list? . -> . any/c) = car
  retain : procedure?

```

Produces a procedure that behaves like `retain`. The procedure is called in atomic mode (see `start-atomic`), and the reference count on one of its arguments is incremented, with `release` recorded as the corresponding release procedure to be called by the finalizer on the retained object (unless some deallocator, as wrapped by `deallocate`, is explicitly called first).

The optional `get-arg` procedure determines which of `retain`'s arguments correspond to the retained object; `get-arg` receives a list of arguments passed to `retain`, so the default `car` selects the first one.

The `release` procedure itself need not be specifically designated a deallocator (via `deallocator`). If a deallocator is called explicitly, it need not be the same as `release`.

5.6 Custodian Shutdown Registration

```
(require ffi/unsafe/custodian)
```

The `ffi/unsafe/custodian` library provides utilities for registering shutdown callbacks with custodians.

```

(register-custodian-shutdown v
  callback
  [custodian
   #:at-exit? at-exit?
   #:weak? weak?]) → cpointer?

v : any/c
callback : (any/c . -> . any)
custodian : custodian? = (current-custodian)
at-exit? : any/c = #f
weak? : any/c = #f

```

Registers `callback` to be applied (in atomic mode and an unspecified Racket thread) to `v` when `custodian` is shutdown. The result is a pointer that can be supplied to `unregister-custodian-shutdown` to remove the registration.

If `at-exit?` is true, then `callback` is applied when Racket exits, even if the custodian is not explicitly shut down.

If `weak?` is true, then `callback` may not be called if `v` is determined to be unreachable during garbage collection. The value `v` is always weakly held by the custodian, even if `weak?` is `#f`; see `scheme_add_managed` for more information.

Normally, `weak?` should be false. To trigger actions based on finalization or custodian shutdown—whichever happens first—leave `weak?` as `#f` and have a finalizer cancel the shutdown action via `unregister-custodian-shutdown`. Otherwise, a not-yet-run finalizer may remain pending after the custodian is shutdown.

```
(unregister-custodian-shutdown v
                               registration) → void?
v : any/c
registration : _cpointer
```

Cancels a custodian-shutdown registration, where `registration` is a previous result from `register-custodian-shutdown` applied to `v`.

5.7 Atomic Execution

```
(require ffi/unsafe/atomic)
```

```
(start-atomic) → void?
(end-atomic) → void?
```

Disables and enables context switches and delivery of break exceptions at the level of Racket threads. Calls to `start-atomic` and `end-atomic` can be nested.

Using `call-as-atomic` is somewhat safer, in that `call-as-atomic` correctly catches exceptions and re-raises them after exiting atomic mode. For simple uses where exceptions need not be handled, however, `start-atomic` and `end-atomic` are faster.

```
(start-breakable-atomic) → void?
(end-breakable-atomic) → void?
```

Like `start-atomic` and `end-atomic`, but the delivery of break exceptions is not suspended. To ensure that a call to `start-atomic` is reliably paired with a call to `end-atomic`, use `dynamic-wind` pre- and post thunks or some other context where breaks are disabled. These variants are not faster than plain `start-atomic` and `end-atomic`.

```
(call-as-atomic thunk) → any
thunk : (-> any)
```

Calls `thunk` in atomic mode. If `thunk` raises an exception, the exception is caught and re-raised after exiting atomic mode.

When `call-as-atomic` is used in the dynamic extent of `call-as-atomic`, then `thunk` is simply called directly (as a non-tail call).

```
(call-as-nonatomic thunk) → any
  thunk : (-> any)
```

Within the dynamic extent of a `call-as-atomic`, calls `thunk` in non-atomic mode. Beware that the current thread may be suspended or terminated by other threads during `thunk`, in which case the call never returns.

When used not in the dynamic extent of `call-as-atomic`, `call-as-nonatomic` raises `exn:fail:contract`.

```
(in-atomic-mode?) → boolean?
```

Returns `#t` if Racket context switches are disabled, `#f` otherwise.

5.8 Speculatively Atomic Execution

```
(require ffi/unsafe/try-atomic)
```

The `ffi/unsafe/try-atomic` supports atomic execution that can be suspended and resumed in non-atomic mode if it takes too long or if some external event causes the attempt to be abandoned.

```
(call-as-nonatomic-retry-point thunk) → any
  thunk : (-> any)
```

Calls `thunk` in atomic mode (see `start-atomic` and `end-atomic`) while allowing `thunk` to use `try-atomic`. Any incomplete computations started with `try-atomic` are run non-atomically after `thunk` returns. The result of `thunk` is used as the result of `call-as-nonatomic-retry-point`.

```
(try-atomic thunk
  default-val
  [#:should-give-up? give-up-proc
   #:keep-in-order? keep-in-order?]) → any
  thunk : (-> any)
  default-val : any/c
  give-up-proc : (-> any/c) = run-200-milliseconds
  keep-in-order? : any/c = #t
```

Within the dynamic extent of a `call-as-nonatomic-retry-point` call, attempts to run `thunk` in the existing atomic mode. The `give-up-proc` procedure is called periodically to

determine whether atomic mode should be abandoned; the default *give-up-proc* returns true after 200 milliseconds. If atomic mode is abandoned, the computation is suspended, and *default-val* is returned, instead. The computation is resumed later by the enclosing `call-as-nonatomic-retry-point` call.

If *keep-in-order?* is true, then if `try-atomic` is called after an earlier computation was suspended for the same `call-as-nonatomic-retry-point` call, then *thunk* is immediately enqueued for completion by `call-as-nonatomic-retry-point` and *default-val* is returned.

The *give-up-proc* callback is polled only at points where the level of atomic-mode nesting (see `start-atomic`, `start-breakable-atomic`, and `call-as-atomic`) is the same as at the point of calling `try-atomic`.

If *thunk* aborts the current continuation using (`default-continuation-prompt-tag`), the abort is suspended the resumed by the enclosing `call-as-nonatomic-retry-point`. Escapes to the context of the call to *thunk* using any other prompt tag or continuation are blocked (using `dynamic-wind`) and simply return (`void`) from *thunk*.

5.9 Objective-C FFI

```
(require ffi/unsafe/objc)
```

The `ffi/unsafe/objc` library builds on `ffi/unsafe` to support interaction with Objective-C.

The library supports Objective-C interaction in two layers. The upper layer provides syntactic forms for sending messages and deriving subclasses. The lower layer is a thin wrapper on the Objective-C runtime library functions. Even the upper layer is unsafe and relatively low-level compared to normal Racket libraries, because argument and result types must be declared in terms of FFI C types (see §3.1 “Type Constructors”).

5.9.1 FFI Types and Constants

```
|_id : ctype?
```

The type of an Objective-C object, an opaque pointer.

```
|_Class : ctype?
```

The type of an Objective-C class, which is also an `_id`.

```
|_Protocol : ctype?
```

The type of an Objective-C protocol, which is also an `_id`.

```
| _SEL : ctype?
```

The type of an Objective-C selector, an opaque pointer.

```
| _BOOL : ctype?
```

The Objective-C boolean type. Racket values are converted for C in the usual way: `#f` is false and any other value is true. C values are converted to Racket booleans.

```
| YES : boolean?
```

Synonym for `#t`

```
| NO : boolean?
```

Synonym for `#f`

5.9.2 Syntactic Forms and Procedures

```
(tell result-type obj-expr method-id)
(tell result-type obj-expr arg ...)

result-type =
  | #:type ctype-expr

arg = method-id arg-expr
      | method-id #:type ctype-expr arg-expr
```

Sends a message to the Objective-C object produced by `obj-expr`. When a type is omitted for either the result or an argument, the type is assumed to be `_id`, otherwise it must be specified as an FFI C type (see §3.1 “Type Constructors”).

If a single `method-id` is provided with no arguments, then `method-id` must not end with `:`; otherwise, each `method-id` must end with `:`.

Examples:

```
> (tell NSString alloc)
#<cpointer:id>
> (tell (tell NSString alloc)
      initWithUTF8String: #:type _string "Hello")
```

```
#<pointer:id>
```

```
(tellv obj-expr method-id)
(tellv obj-expr arg ...)
```

Like `tell`, but with a result type `_void`.

```
(import-class class-id ...)
```

Defines each `class-id` to the class (a value with FFI type `_Class`) that is registered with the string form of `class-id`. The registered class is obtained via `objc_lookupClass`.

Example:

```
> (import-class NSString)
```

A class accessed by `import-class` is normally declared as a side effect of loading a foreign library. For example, if you want to import the class `NSString` on Mac OS X, the "Foundation" framework must be loaded, first. Beware that if you use `import-class` in DrRacket or a module that requires `racket/gui/base`, then "Foundation" will have been loaded into the Racket process already. To avoid relying on other libraries to load "Foundation", explicitly load it with `ffi-lib`:

```
> (ffi-lib
   "/System/Library/Frameworks/Foundation.framework/Foundation")
> (import-class NSString)
```

```
(import-protocol protocol-id ...)
```

Defines each `protocol-id` to the protocol (a value with FFI type `_Protocol`) that is registered with the string form of `protocol-id`. The registered class is obtained via `objc_getProtocol`.

Example:

```
> (import-protocol NSCoder)
```

```
(define-objc-class class-id superclass-expr
  maybe-mixins
  maybe-protocols
  [field-id ...]
  method ...)
```

```

maybe-mixins =
  | #:mixins (mixin-expr ...)

maybe-protocols =
  | #:protocols (protocol-expr ...)

  method = (mode result-ctype-expr (method-id) body ...+)
            | (mode result-ctype-expr (arg ...+) body ...+)

  mode = +
          | -
          | +a
          | -a

  arg = method-id [ctype-expr arg-id]

```

Defines *class-id* as a new, registered Objective-C class (of FFI type `_Class`). The *superclass-expr* should produce an Objective-C class or `#f` for the superclass. An optional *#:mixins* clause can specify mixins defined with `define-objc-mixin`. An optional *#:protocols* clause can specify Objective-C protocols to be implemented by the class.

Each *field-id* is an instance field that holds a Racket value and that is initialized to `#f` when the object is allocated. The *field-ids* can be referenced and `set!` directly when the method *bodys*. Outside the object, they can be referenced and set with `get-ivar` and `set-ivar!`.

Each *method* adds or overrides a method to the class (when *mode* is `-` or `-a`) to be called on instances, or it adds a method to the meta-class (when *mode* is `+` or `+a`) to be called on the class itself. All result and argument types must be declared using FFI C types (see §3.1 “Type Constructors”). When *mode* is `+a` or `-a`, the method is called in atomic mode (see `_cprocedure`).

If a *method* is declared with a single *method-id* and no arguments, then *method-id* must not end with `.`. Otherwise, each *method-id* must end with `.`.

If the special method `dealloc` is declared for mode `-`, it must not call the superclass method, because a `(super-tell dealloc)` is added to the end of the method automatically. In addition, before `(super-tell dealloc)`, space for each *field-id* within the instance is deallocated.

Example:

```

> (define-objc-class MyView NSView
  [bm] ; <- one field
  (- _racket (swapBitwmap: [_racket new-bm])
    (begin0 bm (set! bm new-bm)))

```

```
(- _void (drawRect: [_NSRect exposed-rect])
  (super-tell drawRect: exposed-rect)
  (draw-bitmap-region bm exposed-rect))
(- _void (dealloc)
  (when bm (done-with-bm bm))))
```

```
(define-objc-mixin (class-id superclass-id)
  maybe-mixins
  maybe-protocols
  [field-id ...]
  method ...)
```

Like `define-objc-class`, but defines a mixin to be combined with other method definitions through either `define-objc-class` or `define-objc-mixin`. The specified *field-ids* are not added by the mixin, but must be a subset of the *field-ids* declared for the class to which the methods are added.

```
| self
```

When used within the body of a `define-objc-class` or `define-objc-mixin` method, refers to the object whose method was called. This form cannot be used outside of a `define-objc-class` or `define-objc-mixin` method.

```
| (super-tell result-type method-id)
| (super-tell result-type arg ...)
```

When used within the body of a `define-objc-class` or `define-objc-mixin` method, calls a superclass method. The *result-type* and *arg* sub-forms have the same syntax as in `tell`. This form cannot be used outside of a `define-objc-class` or `define-objc-mixin` method.

```
| (get-ivar obj-expr field-id)
```

Extracts the Racket value of a field in a class created with `define-objc-class`.

```
| (set-ivar! obj-expr field-id value-expr)
```

Sets the Racket value of a field in a class created with `define-objc-class`.

```
| (selector method-id)
```

Returns a selector (of FFI type `_SEL`) for the string form of *method-id*.

Example:

```
> (tellv button setAction: #:type _SEL (selector terminate:))
```

```
(objc-is-a? obj cls) → boolean?  
  obj : _id  
  cls : _Class
```

Check whether *obj* is an instance of the Objective-C class *cls*.

5.9.3 Raw Runtime Functions

```
(objc_lookUpClass s) → (or/c _Class #f)  
  s : string?
```

Finds a registered class by name.

```
(objc_getProtocol s) → (or/c _Protocol #f)  
  s : string?
```

Finds a registered protocol by name.

```
(sel_registerName s) → _SEL  
  s : string?
```

Interns a selector given its name in string form.

```
(objc_allocateClassPair cls s extra) → _Class  
  cls : _Class  
  s : string?  
  extra : integer?
```

Allocates a new Objective-C class.

```
(objc_registerClassPair cls) → void?  
  cls : _Class
```

Registers an Objective-C class.

```
(object_getClass obj) → _Class  
  obj : _id
```

Returns the class of an object (or the meta-class of a class).

```
(class_addMethod cls
                 sel
                 imp
                 type
                 type-encoding) → boolean?
cls : _Class
sel : _SEL
imp : procedure?
type : ctype?
type-encoding : string?
```

Adds a method to a class. The *type* argument must be a FFI C type (see §3.1 “Type Constructors”) that matches both *imp* and the not Objective-C type string *type-encoding*.

```
(class_addIvar cls
               name
               size
               log-alignment
               type-encoding) → boolean?
cls : _Class
name : string?
size : exact-nonnegative-integer?
log-alignment : exact-nonnegative-integer?
type-encoding : string?
```

Adds an instance variable to an Objective-C class.

```
(object_getInstanceVariable obj name) → _Ivar any/c
obj : _id
name : string?
```

Gets the value of an instance variable whose type is *_pointer*.

```
(object_setInstanceVariable obj name val) → _Ivar
obj : _id
name : string?
val : any/c
```

Sets the value of an instance variable whose type is *_pointer*.

```
_Ivar : ctype?
```

The type of an Objective-C instance variable, an opaque pointer.


```

((objc_msgSend/typed types) obj sel arg) → any/c
  types : (vector/c result-ctype arg-ctype ...)
  obj : _id
  sel : _SEL
  arg : any/c

```

Calls the Objective-C method on `_id` named by `sel`. The `types` vector must contain one more than the number of supplied `args`; the first FFI C type in `type` is used as the result type.

```

((objc_msgSendSuper/typed types)
  super
  sel
  arg) → any/c
  types : (vector/c result-ctype arg-ctype ...)
  super : _objc_super
  sel : _SEL
  arg : any/c

```

Like `objc_msgSend/typed`, but for a super call.

```

(make-objc_super id super) → _objc_super
  id : _id
  super : _Class
  _objc_super : ctype?

```

Constructor and FFI C type use for super calls.

Contents

5.9.4 Legacy Library

```
(require ffi/objc)
```

The `ffi/objc` library is a deprecated entry point to `ffi/unsafe/objc`. It exports only safe operations directly, and unsafe operations are imported using `objc-unsafe!`.

```
| (objc-unsafe!)
```

Analogous to `(unsafe!)`, makes unsafe bindings of `ffi/unsafe/objc` available in the importing module.

5.10 Cocoa Foundation

The `ffi/unsafe/nsalloc` and `ffi/unsafe/nsstring` libraries provide basic facilities for working with Cocoa and/or Mac OS X Foundation libraries (usually along with `ffi/objc`).

5.10.1 Strings

```
(require ffi/unsafe/nsstring)
```

```
| _NSString : ctype?
```

A type that converts between Racket strings and `NSString` (a.k.a. `CFString`) values. That is, use `_NSString` as a type for a foreign-function `NSString` argument or result.

The `_NSString` conversion keeps a weak mapping from Racket strings to converted strings, so that converting the same string (in the `equal?` sense) multiple times may avoid allocating multiple `NSString` objects.

5.10.2 Allocation Pools

```
(require ffi/unsafe/nsalloc)
```

Calling any Foundation API that allocates requires an `NSAutoreleasePool` installed. The `ffi/unsafe/nsalloc` library provides a function and shorthand syntactic form for setting up such a context. (The `_NSString` type creates an autorelease pool implicitly while converting from/to a Racket string, however.)

```
| (call-with-autorelease thunk) → any  
|   thunk : (-> any)
```

Calls `thunk` in atomic mode and with a fresh `NSAutoreleasePool` that is released after `thunk` returns.

`(with-autorelease expr)`

A shorthand for `(call-with-autorelease (lambda () expr))`.

5.11 COM (Common Object Model)

The `ffi/com` and `ffi/unsafe/com` libraries support COM interaction in two layers. The safe upper layer provides functions for creating COM objects and dynamically constructing method calls based on COM automation (i.e., reflective information provided by the object). The unsafe lower layer provides a syntactic form and functions for working more directly with COM objects and interfaces.

A *COM object* instantiates a particular *COM class*. A COM class can be specified in either of two ways:

- A *CLSID* (class id), which is represented as a GUID. A *GUID* (globally unique identifier) is a 16-byte structure. GUIDs are typically written in string forms such as `"{A3B0AF9E-2AB0-11D4-B6D2-0060089002FE}"`. The `string->guid` and `guid->string` convert between string and GUID forms. The `string->clsid` function is the same as `string->guid`, but its use suggests that the resulting GUID is to be used as a CLSID.
- A *ProgID* is a human-readable name, such as `"MzCom.MzObj.5.2.0.7"`, which includes a version number. The version number can be omitted in a ProgID, in which case the most recent available version is used. The operating system provides a mapping between ProgIDs and CLSIDs that is available via `progid->clsid` and `clsid->progid`.

A COM object can be instantiated either on the local machine or on a remote machine. The latter relies on the operating system's *DCOM* (distributed COM) support.

Each COM object supports some number of *COM interfaces*. A COM interface has a programmatic name, such as `IDispatch`, that corresponds to a C-layer protocol. Each interface also has an *IID* (interface id) that is represented as a GUID such as `"{00020400-0000-0000-C000-00000000046}"`. Direct calls to COM methods require extracting a suitable interface pointer from an object using `QueryInterface` and the desired IID; the result is effectively cast it to a pointer to a dispatch-table pointer, where the dispatch table has a statically known size and foreign-function content. The `define-com-interface` form simplifies description and use of interface pointers. The COM automation layer uses a fixed number of reflection interfaces internally, notably `IDispatch`, to call methods by name and with safe argument marshaling.

5.11.1 COM Automation

```
(require ffi/com)
```

The `ffi/com` library builds on COM automation to provide a safe use of COM objects that support the `IDispatch` interface.

GUIDs, CLSIDs, IIDs, and ProgIDs

```
(guid? v) → boolean?  
  v : any/c  
(clsid? v) → boolean?  
  v : any/c  
(iid? v) → boolean?  
  v : any/c
```

Returns `#t` if `v` is a structure representing a GUID, `#f` otherwise. The `clsid?` and `iid?` functions are the same as `guid?`.

A GUID corresponds to a `_GUID` structure at the unsafe layer.

```
(string->guid str) → guid?  
  str : string?  
(string->clsid str) → clsid?  
  str : string?  
(string->iid str) → iid?  
  str : string?
```

Converts a string of the form `"{00000000-0000-0000-0000-0000000000}"`, where each 0 can be a hexadecimal digit, to a GUID. If `str` does not have the expected form, the `exn:fail` exception is raised.

The `string->clsid` and `string->iid` functions are the same as `string->guid`.

```
(guid->string g) → string?  
  g : guid?
```

Converts a GUID to its string form.

```
(guid=? g1 g2) → boolean?  
  g1 : guid?  
  g2 : guid?
```

Determines whether `g1` and `g2` represent the same GUID.

The `ffi/com` library is based on the `MysterX` library by Paul Steckler. `MysterX` is included with Racket but deprecated, and it will be replaced in the next version with a partial compatibility library that redirects to this one.

```
(progid->clsid progid) → clsid?
  progid : string?
(clsid->progid clsid) → (or/c string? #f)
  clsid : clsid?
```

Converts a ProgID to a CLSID or vice versa. Not every COM class has a ProgID, so the result of `clsid->progid` can be `#f`.

The `progid->clsid` function accepts a versionless ProgID, in which case it produces the CLSID of the most recent available version. The `clsid->progid` function always produces a ProgID with its version.

COM Objects

```
(com-object? obj) → boolean?
  obj : com-object?
```

Returns `#t` if the argument represents a COM object, `#f` otherwise.

```
(com-create-instance clsid-or-progid [where]) → com-object?
  clsid-or-progid : (or/c clsid? string?)
  where : (or/c (one-of/c 'local 'remote) string?) = 'local
```

Returns an instance of the COM class specified by `clsid-or-progid`, which is either a CLSID or a ProgID.

The optional `where` argument indicates a location for running the instance, and may be `'local`, `'remote`, or a string indicating a machine name. See §5.11.1.7 “Remote COM servers (DCOM)” for more information.

An object can be created this way for any COM class, but functions such as `com-invoke` work only if the object supports the IDispatch COM automation interface.

The resulting object is registered with the current custodian, which retains a reference to the object until it is released with `com-release` or the custodian is shut down.

```
(com-release obj) → void?
  obj : com-object?
```

Releases the given COM object. The given `obj` is subsequently unusable, and the underlying COM object is destroyed unless its reference count has been incremented (via COM methods or unsafe operations).

If `obj` has already been released, `com-release` has no effect.

```
(com-get-active-object clsid-or-progid) → com-object?  
  clsid-or-progid : (or/c clsid? string?)
```

Like `com-create-instance`, but gets an existing active object (always local) instead of creating a new one.

```
(com-object-clsid obj) → clsid?  
  obj : com-object?
```

Returns the "CLSID" of the COM class instantiated by `obj`, or raises an error if the COM class is not known.

```
(com-object-set-clsid! obj clsid) → void?  
  obj : com-object?  
  clsid : clsid?
```

Sets the COM CLSID for `obj` to `clsid`. This is useful when COM event-handling procedures can obtain only ambiguous information about the object's COM class.

```
(com-object-eq? obj1 obj2) → boolean?  
  obj1 : com-object?  
  obj2 : com-object?
```

Returns `#t` if `obj1` and `obj2` refer to the same COM object, `#f` otherwise.

If two references to a COM object are the same according to `com-object-eq?`, then they are also the same according to `equal?`. Two `com-object-eq?` references are not necessarily `eq?`, however.

```
(com-type? v) → boolean?  
  v : any/c
```

Returns `#t` if `v` represents reflective information about a COM object's type, `#f` otherwise.

```
(com-object-type obj) → com-type?  
  obj : com-object?
```

Returns a representation of a COM object's type that is independent of the object itself.

```
(com-type=? t1 t2) → boolean?  
  t1 : com-type?  
  t2 : com-type?
```

Returns `#t` if `t1` and `t2` represent the same type information, `#f` otherwise.

COM Methods

```
(com-methods obj/type) → (listof string?)
  obj/type : (or/c com-object? com-type?)
```

Returns a list of strings indicating the names of methods on *obj/type*.

```
(com-method-type obj/type method-name)
→ (list/c '-> (listof type-description?)
      type-description?)
  obj/type : (or/c com-object? com-type?)
  method-name : string?
```

Returns a list indicating the type of the specified method in *obj/type*. The list after the '->' represents the argument types, and the final value represents the result type. See §5.11.1.8 “COM Types” for more information.

```
(com-invoke obj method-name v ...) → any/c
  obj : com-object?
  method-name : string?
  v : any/c
```

Invokes *method-name* on *obj* with *vs* as the arguments. The special value `com-omit` may be used for optional arguments, which useful when values are supplied for arguments after the omitted argument(s).

The types of arguments are determined via `com-method-type`, if possible, and `type-describe` wrappers in the *vs* are simply replaced with the values that they wrap. If the types are not available from `com-method-type`, then types are inferred for each *v* with attention to descriptions in any `type-describe` wrappers in *v*.

```
com-omit : any/c
```

A constant for use with `com-invoke` in place of an optional argument.

COM Properties

```
(com-get-properties obj/type) → (listof string?)
  obj/type : (or/c com-object? com-type?)
```

Returns a list of strings indicating the names of readable properties in *obj/type*.

```
(com-get-property-type obj/type
  property-name)
→ (list/c '-> '() type-description?)
  obj/type : (or/c com-object? com-type?)
  property-name : string?
```


Returns a type for *property-name* like a result of *com-method*, where the result type corresponds to the property value type. See §5.11.1.8 “COM Types” for information on the symbols.

```
(com-get-property obj property ...+) → any/c
  obj : com-object?
  property : (or/c string?
             (cons/c string? list?))
```

Returns the value of the final property by following the indicated path of *properties*, where each intermediate property must be a COM object.

Each *property* is either a property-name string or a list that starts with a property-name string and continues with arguments for a parameterized property.

```
(com-get-property* obj property v ...) → any/c
  obj : com-object?
  property : string?
  v : any/c
```

Returns the value of a parameterized property, which behaves like a method and accepts the *vs* as arguments (like *com-invoke*). When no *vs* are provided, *com-get-property** is the same as *com-get-property*.

```
(com-set-properties obj/type) → (listof string?)
  obj/type : (or/c com-object? com-type?)
```

Returns a list of strings indicating the names of writeable properties in *obj/type*.

```
(com-set-property-type obj/type
  property-name)
→ (list/c '-> (list/c type-description?) 'void)
  obj/type : (or/c com-object? com-type?)
  property-name : string?
```

Returns a type for *property-name* like a result of *com-method*, where the sole argument type corresponds to the property value type. See §5.11.1.8 “COM Types” for information on the symbols.

```
(com-set-property! obj property ...+ v) → void?
  obj : com-object?
  property : (or/c string?
             (cons/c string? list?))
  v : any/c
```

Sets the value of the final property in *obj* to *v* by following the *property*s, where the value of each intermediate property must be a COM object. A *property* can be a list instead of a string to represent a parameterized property and its arguments.

The type of the property is determined via *com-property-type*, if possible, and *type-describe* wrappers in *v* are then replaced with the values that they wrap. If the type is not available from *com-property-type*, then a type is inferred for *v* with attention to the descriptions in any *type-describe* wrappers in *v*.

COM Events

```
(com-events obj/type) → (listof string?)  
  obj/type : (or/c com-object? com-type?)
```

Returns a list of strings indicating the names of events on *obj/type*.

```
(com-event-type obj/type event-name)  
→ (list/c '-> (listof type-description?) 'void)  
  obj/type : (or/c com-object? com-type?)  
  event-name : string?
```

Returns a list indicating the type of the specified events in *obj/type*. The list after the *'->* represents the argument types. See §5.11.1.8 “COM Types” for more information.

```
(com-event-executor? v) → boolean?  
  v : any/c
```

Returns *#t* if *v* is a *COM event executor*, which queues event callbacks. A COM event executor *com-ev-ex* is a synchronizable event in the sense of *sync*, and (*sync com-ev-ex*) returns a thunk for a ready callback.

```
(com-make-event-executor) → com-event-executor?
```

Creates a fresh COM event executor for use with *com-register-event-callback*.

```
(com-register-event-callback obj  
                             name  
                             proc  
                             com-ev-ex) → void?  
  
  obj : com-object?  
  name : string?  
  proc : procedure?  
  com-ev-ex : com-event-executor?
```

Registers a callback for the event named by *name* in *obj*. When the event fires, an invocation of *proc* to event arguments (which depends on *obj* and *name*) is queued in *com-ev-ex*.

Synchronizing on `com-ev-ex` produces a thunk that applies `proc` to the event arguments and returns the result.

Only one callback can be registered for each `obj` and `name` combination.

Registration of event callbacks relies on prior registration of the COM class implemented by "myssink.dll" as distributed with Racket. (The DLL is the same for all Racket versions.)

```
(com-unregister-event-callback obj name) → void?  
  obj : com-object?  
  name : string?
```

Removes any existing callback for `name` in `obj`.

Interface Pointers

```
(com-object-get-iunknown obj) → com-iunknown?  
  obj : com-object?  
(com-object-get-idispatch obj) → com-idispatch?  
  obj : com-object?
```

Extracts an IUnknown or IDispatch pointer from `obj`. The former succeeds for any COM object that has not been released via `com-release`. The latter succeeds only when the COM object supports IDispatch, otherwise `exn:fail` is raised.

```
(com-iunknown? v) → boolean?  
  v : any/c
```

Returns `#t` if `v` corresponds to an unsafe `_IUnknown-pointer`, `#f` otherwise. Every COM interface extends IUnknown, so `com-iunknown?` returns `#t` for every interface pointers.

```
(com-idispatch? v) → boolean?  
  v : any/c
```

Returns `#t` if `v` corresponds to an unsafe IDispatch, `#f` otherwise.

Remote COM servers (DCOM)

The optional `where` argument to `com-create-instance` can be `'remote`. In that case, the server instance is run at the location given by the Registry key

```
HKEY_CLASSES_ROOT\AppID\⟨CLSID⟩\RemoteServerName
```

where `⟨CLSID⟩` is the CLSID of the application. This key may be set using the `dcomcnfg` utility. From `dcomcnfg`, pick the application to be run on the Applications tab, then click

on the Properties button. On the Location tab, choose Run application on the following computer, and enter the machine name.

To run a COM remote server, the registry on the client machine must contain an entry at

```
HKEY_CLASSES_ROOT\CLSID\<CLSID>
```

where *<CLSID>* is the CLSID for the server. The server application itself need not be installed on the client machine.

There are a number of configuration issues relating to DCOM. See

<http://www.distribucon.com/dcom95.html>

for more information on how to setup client and server machines for DCOM.

COM Types

In the result of a function like `com-method-type`, symbols are used to represent various atomic types:

- `'int` — a 32-bit signed integer
- `'unsigned-int` — a 32-bit unsigned integer
- `'short-int` — a 16-bit signed integer
- `'unsigned-short` — a 16-bit unsigned integer
- `'signed-char` — an 8-bit signed integer
- `'char` — an 8-bit unsigned integer
- `'long-long` — a 64-bit signed integer
- `'unsigned-long-long` — a 64-bit unsigned integer
- `'float` — a 32-bit floating-point number
- `'double` — a 64-bit floating-point number
- `'currency` — an exact number that, when multiplied by 10,000, is a 64-bit signed integer
- `'boolean` — a boolean
- `'string` — a string
- `'date` — a `date` or `date*`

- `'com-object` — a COM object as in `com-object?`
- `'iunknown` — like `'com-object`, but also accepts an IUnknown pointer as in `com-iunknown?`
- `'com-enumeration` — a 32-bit signed integer
- `'any` — any of the above, or an array when not nested in an array type
- `'...` — treated like `'any`, but when it appears at the end of the sequence of types for arguments, allows the preceding type 0 or more times
- `'void` — no value

A type symbol wrapped in a list with `'box`, such as `'(box int)`, is a call-by-reference argument. A box supplied for the argument is updated with a new value when the method returns.

A type wrapped in a list with `'opt`, such as `'(opt (box int))`, is an optional argument. The argument can be omitted or replaced with `com-omit`.

A type wrapped in a list with `'array` and a positive exact integer, such as `'(array 7 int)`, represents a vector of values to be used as a COM array. A `'?` can be used in place of the length integer to support a vector of any length. Array types with non-`'?` lengths can be nested to specify a multidimensional array as represented by nested vectors.

A type wrapped in a list with `'variant`, such as `'(variant (array 7 int))`, is the same as the wrapped type, but a `'variant` wrapper within an `'array` type prevents construction of another array dimension. For example, `'(array 2 (array 3 int))` is a two-dimensional array of integers, but `'(array 2 (variant (array 3 int)))` is a one-dimensional array whose elements are one-dimensional arrays of integers.

When type information is not available, functions like `com-invoke` infer type descriptions from arguments. Inference chooses `'boolean` for booleans; the first of `'int`, `'unsigned-int`, `'long-long`, `'unsigned-long-long` that fits for an exact integer; `'double` for inexact real numbers; `'string` for a string; `'com-object` and `'iunknown` for corresponding COM object references; and an `'array` type for a vector, where the element type is inferred from vector values, resorting to `'any` if any two elements have different inferred types.

```
(type-description? v) → boolean?
  v : any/c
```

Return `#t` if `v` is a COM argument or result type description as above, `#f` otherwise.

```
(type-described? v) → boolean?
  v : any/c
(type-describe v desc) → type-described?
  v : any/c
  desc : type-description?
```

```
(type-described-value td) → any/c
  td : type-described?
(type-described-description td) → type-description?
  td : type-described?
```

The `type-described?` predicate recognizes wrappers produced with `type-describe`, and `type-described-value` and `type-described-description` extract the value and description parts of a `type-describe` value.

A `type-describe` wrapper combines a base value with a type description. The description is used instead of an automatically inferred COM argument type when no type is available for from COM automation a method for `com-invoke` or a property for `com-set-property!`. A wrapper can be placed on an immediate value, or it can be on a value within a box or vector.

Class Display Names

```
(require ffi/com-registry)
```

The `ffi/com-registry` library provides a mapping from coclass names to CLSIDs for compatibility with the older `MysterX` interface.

A *coclass* name corresponds to the display name of a COM class; the display name is not uniquely mapped to a COM class, and some COM classes have no display name.

```
(com-all-coclasses) → (listof string?)
```

Returns a list of coclass strings for all COM classes registered on a system.

```
(com-all-controls) → (listof string?)
```

Returns a list of coclass strings for all COM classes in the system registry that have the `"Control"` subkey.

```
(coclass->clsid coclass) → clsid?
  coclass : string?
(clsid->coclass clsid) → string?
  clsid : clsid?
```

Converts a coclass string to/from a CLSID. This conversion is implemented by an enumeration of COM classes from the system registry.

5.11.2 COM Classes and Interfaces

```
(require ffi/unsafe/com)
```

The `ffi/unsafe/com` library exports all of `ffi/com`, and it also supports direct, FFI-based calls to COM object methods.

Describing COM Interfaces

```
(define-com-interface (_id _super-id)
  ([method-id ctype-expr maybe-alloc-spec] ...))

maybe-alloc-spec =
  | #:release-with-function function-id
  | #:release-with-method method-id
  | #:releases
```

Defines `_id` as an interface that extends `_super-id`, where `_super-id` is often `_IUnknown`, and that includes methods named by `method-id`. The `_id` and `_super-id` identifiers must start with an underscore. A `_super-id _vt` must also be defined for deriving a virtual-method table type.

The order of the `method-ids` must match the specification of the COM interface, not including methods inherited from `_super-id`. Each method type produced by `ctype-expr` that is not `_fpointer` must be a function type whose first argument is the “self” pointer, usually constructed with `_mfun` or `_hmfun`.

The `define-com-interface` form binds `_id`, `id ?`, `_id -pointer`, `_id _vt` (for the virtual-method table), `_id _vt-pointer`, and `method-id` for each method whose `ctype-expr` is not `_fpointer`. (In other words, use `_fpointer` as a placeholder for methods of the interface that you do not need to call.) An instance of the interface will have type `_id -pointer`. Each defined `method-id` is bound to a function-like macro that expects a `_id -pointer` as its first argument and the method arguments as the remaining arguments.

A `maybe-alloc-spec` describes allocation and finalization information for a method along the lines of `ffi/unsafe/alloc`. If the `maybe-alloc-spec` is `#:release-with-function function-id`, then `function-id` is used to deallocate the result produced by the method, unless the result is explicitly deallocated before it becomes unreachable; for example, `#:release-with-function Release` is suitable for a method that returns a COM interface reference that must be eventually released. The `#:release-with-method method-id` form is similar, except that the deallocator is a method on the same object as the allocating method (i.e., one of the other `method-ids` or an inherited method). A `#:releases` annotation indicates that a method is a deallocator (so that a value should not be automatically deallocated if it is explicitly deallocated using the method).

See §5.11.2.4 “COM Interface Example” for an example using `define-com-interface`.

Obtaining COM Interface References

```
(QueryInterface iunknown
  iid
  intf-pointer-type) → (or/c cpointer? #f)
```

```

iunknown : com-iunknown?
iid : iid?
intf-pointer-type : ctype?

```

Attempts to extract a COM interface pointer for the given COM object. If the object does not support the requested interface, the result is `#f`, otherwise it is cast to the type `intf-pointer-type`.

Specific IIDs and `intf-pointer-types` go together. For example, `IID_IUnknown` goes with `_IUnknown-pointer`.

For a non-`#f` result, `Release` function is the automatic deallocator for the resulting pointer. The pointer is register with a deallocator after the cast to `intf-pointer-type`, which is why `QueryInterface` accepts the `intf-pointer-type` argument (since a cast generates a fresh reference).

```

(AddRef iunknown) → exact-positive-integer?
  iunknown : com-iunknown?
(Release iunknown) → exact-nonnegative-integer?
  iunknown : com-iunknown?

```

Increments or decrements the reference count on `iunknown`, returning the new reference count and releasing the interface reference if the count goes to zero.

```

(make-com-object iunknown
                 clsid
                 [#:manage? manage?]) → com-object?
  iunknown : com-iunknown?
  clsid : (or/c clsid? #f)
  manage? : any/c = #t

```

Converts a COM object into an object that can be used with the COM automation functions, such as `com-invoke`.

If `manage?` is true, the resulting object is registered with the current custodian and a finalizer to call `com-release` when the custodian is shut down or when the object becomes inaccessible.

COM FFI Helpers

```

(_wfun fun-option ... maybe-args type-spec ... -> type-spec
  maybe-wrapper)

```

Like `_fun`, but adds `#:abi winapi`.

```

(_mfun fun-option ... maybe-args type-spec ... -> type-spec
  maybe-wrapper)

```


Like `_wfun`, but adds a `_pointer` type (for the “self” argument of a method) as the first argument `type-spec`.

```
(_hfun fun-option ... type-spec ... -> id output-expr)
```

Like `_wfun`, but for a function that returns an `_HRESULT`. If the result is not zero, then an error is raised using `windows-error` and using `id` as the name of the failed function. Otherwise, `output-expr` (as in a `maybe-racket` for `_fun`) determines the result.

```
(_hmfun fun-option ... type-spec ... -> id output-expr)
```

Like `_hfun`, but like `_mfun` in that `_pointer` is added for the first argument.

```
_GUID : ctype?  
_GUID-pointer : ctype?  
_HRESULT : ctype?  
_LCID : ctype?
```

Some C types that commonly appear in COM interface specifications.

```
LOCALE_SYSTEM_DEFAULT : exact-integer?
```

The usual value for a `_LCID` argument.

```
(SysFreeString str) → void?  
  str : _pointer  
(SysAllocStringLen content len) → cpointer?  
  content : _pointer  
  len : integer?
```

COM interfaces often require or return strings that must be allocated or freed as system strings.

When receiving a string value, `cast` it to `_string/utf-16` to extract a copy of the string, and then free the original pointer with `SysFreeString`.

```
IID_NULL : iid?  
IID_IUnknown : iid?
```

Commonly used IIDs.

```
_IUnknown : ctype?  
_IUnknown-pointer : ctype?  
_IUnknown_vt : ctype?
```

Types for the IUnknown COM interface.

```
(windows-error msg hresult) → any
  msg : string?
  hresult : exact-integer?
```

Raises an exception. The *msg* string provides the base error message, but *hresult* and its human-readable interpretation (if available) are added to the message.

COM Interface Example

Here's an example using the Standard Component Categories Manager to enumerate installed COM classes that are in the different system-defined categories. The example illustrates instantiating a COM class by CLSID, describing COM interfaces with `define-com-interface`, and using allocation specifications to ensure that resources are reclaimed even if an error is encountered or the program is interrupted.

```
#lang racket/base
(require ffi/unsafe
         ffi/unsafe/com)

(provide show-all-classes)

; The function that uses COM interfaces defined further below:

(define (show-all-classes)
  (define ccm
    (com-create-instance CLSID_StdComponentCategoriesMgr))
  (define icat (QueryInterface (com-object-get-iunknown ccm)
                               IID_ICatInformation
                               _ICatInformation-pointer))
  (define eci (EnumCategories icat LOCALE_SYSTEM_DEFAULT))
  (for ([catinfo (in-producer (lambda () (Next/ci eci)) #f)])
    (printf "~a:\n"
            (cast (array-ptr (CATEGORYINFO-szDescription catinfo))
                  _pointer
                  _string/utf-16)))
  (define eg
    (EnumClassesOfCategories icat (CATEGORYINFO-catid catinfo)))
  (for ([guid (in-producer (lambda () (Next/g eg)) #f)])
    (printf " ~a\n" (or (clsid->progid guid)
                        (guid->string guid))))
  (Release eg))
(Release eci)
(Release icat))
```

```

; The class to instantiate:

(define CLSID_StdComponentCategoriesMgr
  (string->clsid "{0002E005-0000-0000-C000-000000000046}"))

; Some types and variants to match the specification:

(define _ULONG _ulong)
(define _CATID _GUID)
(define _REFCATID _GUID-pointer)
(define-cstruct _CATEGORYINFO ([catid _CATID]
                               [lcid _LCID]
                               [szDescription (_array _short 128)]))

; ---- IEnumGUID ----

(define IID_IEnumGUID
  (string->iid "{0002E000-0000-0000-C000-000000000046}"))

(define-com-interface (_IEnumGUID _IUnknown)
  ([Next/g (_mfun (_ULONG = 1) ; simplified to just one
                 (guid : (_ptr o _GUID))
                 (got : (_ptr o _ULONG))
                 -> (r : _HRESULT)
                 -> (cond
                    [(zero? r) guid]
                    [(= r 1) #f]
                    [else (windows-error "Next/g failed" r)]))]
   [Skip _fpointer]
   [Reset _fpointer]
   [Clone _fpointer]))

; ---- IEnumCATEGORYINFO ----

(define IID_IEnumCATEGORYINFO
  (string->iid "{0002E011-0000-0000-C000-000000000046}"))

(define-com-interface (_IEnumCATEGORYINFO _IUnknown)
  ([Next/ci (_mfun (_ULONG = 1) ; simplified to just one
                  (catinfo : (_ptr o _CATEGORYINFO))
                  (got : (_ptr o _ULONG))
                  -> (r : _HRESULT)
                  -> (cond
                     [(zero? r) catinfo]
                     [(= r 1) #f]
                     [else (windows-error "Next/ci

```

```

failed" r)))]
  [Skip _fpointer]
  [Reset _fpointer]
  [Clone _fpointer]))

; ---- ICatInformation ----

(define IID_ICatInformation
  (string->iid "{0002E013-0000-0000-C000-000000000046}"))

(define-com-interface (_ICatInformation _IUnknown)
  ([EnumCategories (_hmfun _LCID
    (p : (_ptr o _IEnumCATEGORYINFO-
pointer))
    -> EnumCategories p)]
  [GetCategoryDesc (_hmfun _REFCATID _LCID
    (p : (_ptr o _pointer))
    -> GetCategoryDesc
    (begin0
      (cast p _pointer _string/utf-16)
      (SysFreeString p)))]
  [EnumClassesOfCategories (_hmfun (_ULONG = 1) ; simplified
    _REFCATID
    (_ULONG = 0) ; simplified
    (_pointer = #f)
    (p : (_ptr o
      _IEnumGUID-
pointer))
    -> EnumClassesOfCategories p)
  #:release-with-function Release]
  [IsClassOfCategories _fpointer]
  [EnumImplCategoriesOfClass _fpointer]
  [EnumReqCategoriesOfClass _fpointer]))

```

5.11.3 ActiveX Controls

An ActiveX control is a COM object that needs a container to manage its graphical representation. Although `ffi/com` does not provide direct support for ActiveX controls, you can use `ffi/com` to drive Internet Explorer as an ActiveX container.

The following code demonstrates using Internet Explorer to instantiate the “Sysmon” ActiveX control that is included with Windows.

```

#lang racket
(require ffi/com

```

```

xml)

;; The control we want to run:
(define control-progid "Sysmon")

;; Start IE:
(define ie (com-create-instance "InternetExplorer.Application.1"))

;; Set up an event callback so that we know when the initial document
;; is ready:
(define ex (com-make-event-executor))
(void (thread (lambda () (let loop () ((sync ex)) (loop))))))
(define ready (make-semaphore))
(com-register-event-callback ie "DocumentComplete"
                             (lambda (doc url) (semaphore-
post ready))
                             ex)

;; Navigate to an empty URL to get an initial document:
(com-invoke ie "Navigate" "")
(semaphore-wait ready)
(define doc (com-get-property ie "Document"))

;; Install HTML to show the ActiveX control:
(com-invoke doc "write"
             (xexpr->string
              '(html
                (head (title "Demo"))
                (body
                 (object ((class "object")
                         (CLASSID ,(format
                                     "CLSID:~a"
                                     (let ([s (guid->string
                                               (progid->clsid
                                               control-progid))])
                                     (string-length s))
                                     (substring s 1 (sub1 len))))))))))

;; must remove curly braces:
(define len
  (string-length s))
(substring s 1 (sub1 len))))))

;; Configure the IE window and show it:
(com-set-property! ie "MenuBar" #f)
(com-set-property! ie "ToolBar" 0)
(com-set-property! ie "StatusBar" #f)
(com-set-property! ie "Visible" #t)

```

```

;; Extract the ActiveX control from the IE document:
(define ctl (com-get-property
             (com-invoke (com-invoke doc "getElementsByTagName" "object")
                          "item"
                          0)
             "object"))

;; At this point, 'ctl' is the ActiveX control;
;; demonstrate by getting a list of method names:
(com-methods ctl)

```

5.12 File Security-Guard Checks

```

(require ffi/file)

(security-guard-check-file who path perms) → void?
  who : symbol?
  path : path-string?
  perms : (listof (or/c 'read 'write 'execute 'delete 'exists))

```

Checks whether ([current-security-guard](#)) permits access to the file specified by *path* with the permissions *perms*. See [make-security-guard](#) for more information on *perms*.

The symbol *who* should be the name of the function on whose behalf the security check is performed; it is passed to the security guard to use in access-denied errors.

```

(_file/guard perms [who]) → ctype?
  perms : (listof (or/c 'read 'write 'execute 'delete 'exists))
  who : symbol? = '_file/guard

```

Like [_file](#) and [_path](#), but conversion from Racket to C first completes the path using [path->complete-path](#) then cleanses it using [cleanse-path](#), then checks that the current security guard grants access on the resulting complete path with *perms*. As an output value, identical to [_path](#).

```

_file/r : ctype?
_file/rw : ctype?

```

Equivalent to ([_file/guard 'read](#)) and ([_file/guard 'read write](#)), respectively.

5.13 Windows API Helpers

```

(require ffi/winapi)

```

`win64?` : `boolean?`

Indicates whether the current platform is 64-bit Windows: `#t` if so, `#f` otherwise.

`winapi` : `(or/c 'stdcall 'default)`

Suitable for use as an ABI specification for a Windows API function: `'stdcall` on 32-bit Windows, `'default` on 64-bit Windows or any other platform.

6 Miscellaneous Support

```
(list->cblock lst type) → any
  lst : list?
  type : ctype?
```

Allocates a memory block of an appropriate size, and initializes it using values from *lst* and the given *type*. The *lst* must hold values that can all be converted to C values according to the given *type*.

```
(vector->cblock vec type) → any
  vec : vector?
  type : type?
```

Like `list->cblock`, but for Racket vectors.

```
(vector->cpointer vec) → cpointer?
  vec : vector?
```

Returns a pointer to an array of `_scheme` values, which is the internal representation of *vec*.

```
(flvector->cpointer flvec) → cpointer?
  flvec : flvector?
```

Returns a pointer to an array of `_double` values, which is the internal representation of *flvec*.

```
(saved-errno) → exact-integer?
```

Returns the value most recently saved (in the current thread) after a foreign call with a non-`#f` `#:save-errno` option (see `_fun` and `_cprocedure`).

```
(lookup-errno sym) → exact-integer?
  sym : (or/c 'EINTR 'EEXIST 'EAGAIN)
```

Returns a platform-specific value corresponding to a Posix `errno` symbol. The set of supported symbols is likely to expand in the future.

```
(cast v from-type to-type) → any/c
  v : any/c
  from-type : ctype?
  to-type : ctype?
```


Converts *v* from a value matching *from-type* to a value matching *to-type*, where `(ctype-sizeof from-type)` matches `(ctype-sizeof to-type)`.

The conversion is roughly equivalent to

```
(let ([p (malloc from-type)])
  (ptr-set! p from-type v)
  (ptr-ref p to-type))
```

If *v* is a cpointer, `(cpointer-gcable? v)` is true, and *from-type* and *to-type* are both based on `_pointer` or `_gpointer`, then *from-type* is implicitly converted with `_gcable` to ensure that the result cpointer is treated as referring to memory that is managed by the garbage collector.

If *v* is a pointer with an offset component (e.g., from `ptr-add`), `(cpointer-gcable? v)` is true, and the result is a cpointer, then the result pointer has the same offset component as *v*. If `(cpointer-gcable? v)` is false, then any offset is folded into the pointer base for the result.

```
(cblock->list cblock type length) → list?
  cblock : any/c
  type : ctype?
  length : exact-nonnegative-integer?
```

Converts C *cblock*, which is a vector of *types*, to a Racket list. The arguments are the same as in the `list->cblock`. The *length* must be specified because there is no way to know where the block ends.

```
(cblock->vector cblock type length) → vector?
  cblock : any/c
  type : ctype?
  length : exact-nonnegative-integer?
```

Like `cblock->vector`, but for Racket vectors.

7 Unexported Primitive Functions

Parts of the `ffi/unsafe` library are implemented by the Racket built-in `'%foreign` module. The `'%foreign` module is not intended for direct use, but it exports the following procedures (among others).

```
(ffi-obj objname lib) → ffi-obj?
  objname : (or/c string? bytes? symbol?)
  lib : (or/c ffi-lib? path-string? #f)
```

Pulls out a foreign object from a library, returning a value that can be used as a C pointer. If `lib` is a path or string, then `ffi-lib` is used to create a library object.

```
(ffi-obj? x) → boolean?
  x : any/c
(ffi-obj-lib obj) → ffi-lib?
  obj : ffi-obj?
(ffi-obj-name obj) → string?
  obj : ffi-obj?
```

A predicate for objects returned by `ffi-obj`, and accessor functions that return its corresponding library object and name. These values can also be used as C pointer objects.

```
(ctype-basetype type) → (or/c ctype? #f)
  type : ctype?
(ctype-scheme->c type) → procedure?
  type : ctype?
(ctype-c->scheme type) → procedure?
  type : ctype?
```

Accessors for the components of a C type object, made by `make-ctype`. The `ctype-basetype` selector returns a symbol for primitive types that names the type, a list of ctypes for cstructs, and another ctype for user-defined ctypes.

```
(ffi-call ptr
  in-types
  out-type
  [abi]
  save-errno?
  orig-place?) → procedure?
  ptr : cpointer?
  in-types : (listof ctype?)
  out-type : ctype?
  abi : (or/c #f 'default 'stdcall 'sysv) = #f
  save-errno? : any/c
  orig-place? : any/c
```

The primitive mechanism that creates Racket “callout” values for `_cprocedure`. The given `ptr` is wrapped in a Racket-callable primitive function that uses the types to specify how values are marshaled.

```
(ffi-callback proc
              in-types
              out-type
              [abi
              atomic?
              async-apply]) → ffi-callback?
proc : any/c
in-types : any/c
out-type : any/c
abi : (or/c #f 'default 'stdcall 'sysv) = #f
atomic? : any/c = #f
async-apply : (or/c #f ((-> any) . -> . any)) = #f
```

The symmetric counterpart of `ffi-call`. It receives a Racket procedure and creates a callback object, which can also be used as a C pointer.

```
(ffi-callback? x) → boolean?
x : any/c
```

A predicate for callback values that are created by `ffi-callback`.

Bibliography

- [Barzilay04] Eli Barzilay and Dmitry Orlovsky, “Foreign Interface for PLT Scheme,”
Workshop on Scheme and Functional Programming, 2004.

Index

->, 28
_?, 30
_array, 36
_array/list, 37
_array/vector, 38
_bitmask, 39
_BOOL, 67
_bool, 20
_box, 31
_byte, 18
_bytes, 31
_bytes/eof, 22
_Class, 66
_cpointer, 58
_cpointer/null, 59
_cprocedure, 24
_cvector, 57
_double, 19
_double*, 19
_enum, 39
_f32vector, 55
_f64vector, 56
_f80vector, 56
_file, 22
_file/guard, 94
_file/r, 94
_file/rw, 94
_fixint, 19
_fixnum, 19
_float, 19
_fpointer, 23
_fun, 27
_gcable, 23
_gcpointer, 23
_GUID, 89
_GUID-pointer, 89
_hfun, 89
_hmfun, 89
_HRESULT, 89
_id, 66
_int, 19
_int16, 18
_int32, 18
_int64, 18
_int8, 18
_intmax, 19
_intptr, 19
_IUnknown, 89
_IUnknown-pointer, 89
_IUnknown_vt, 89
_Ivar, 72
_LCID, 89
_list, 31
_list-struct, 32
_llong, 19
_long, 19
_longdouble, 20
_mfun, 88
_NSString, 75
_objc_super, 73
_or-null, 23
_path, 20
_pointer, 22
_Protocol, 66
_ptr, 30
_ptrdiff, 19
_racket, 23
_s16vector, 51
_s32vector, 52
_s64vector, 54
_s8vector, 50
_sbyte, 18
_scheme, 23
_SEL, 67
_short, 19
_sint, 19
_sint16, 18
_sint32, 18
_sint64, 18
_sint8, 18
_sintptr, 19
_size, 19

- [_sllong](#), 19
- [_slong](#), 19
- [_sshort](#), 19
- [_ssize](#), 19
- [_string](#), 21
- [_string*/latin-1](#), 21
- [_string*/locale](#), 21
- [_string*/utf-8](#), 21
- [_string/eof](#), 22
- [_string/latin-1](#), 21
- [_string/locale](#), 21
- [_string/ucs-4](#), 20
- [_string/utf-16](#), 20
- [_string/utf-8](#), 21
- [_sword](#), 18
- [_symbol](#), 21
- [_u16vector](#), 51
- [_u32vector](#), 53
- [_u64vector](#), 54
- [_u8vector](#), 49
- [_ubyte](#), 18
- [_ufixint](#), 19
- [_ufixnum](#), 19
- [_uint](#), 19
- [_uint16](#), 18
- [_uint32](#), 18
- [_uint64](#), 18
- [_uint8](#), 18
- [_uintmax](#), 19
- [_uintptr](#), 19
- [_ullong](#), 19
- [_ulong](#), 19
- [_union](#), 38
- [_ushort](#), 19
- [_uword](#), 18
- [_vector](#), 31
- [_void](#), 20
- [_wfun](#), 88
- [_word](#), 18
- ActiveX Controls, 92
- [AddRef](#), 88
- Allocation and Finalization, 62
- Allocation Pools, 75
- [allocator](#), 62
- [array-length](#), 37
- [array-ptr](#), 37
- [array-ref](#), 37
- [array-set!](#), 37
- [array?](#), 37
- ['atomic](#), 45
- Atomic Execution, 64
- ['atomic-interior](#), 45
- By-Reference Arguments, 6
- C Array Types, 36
- C pointer*, 22
- C Struct Types, 32
- C Structs, 7
- C Types, 17
- C types*, 17
- C Union Types, 38
- [call-as-atomic](#), 64
- [call-as-nonatomic](#), 65
- [call-as-nonatomic-retry-point](#), 65
- [call-with-autorelease](#), 75
- callback*, 24
- callout*, 24
- [cast](#), 96
- [cblock->list](#), 97
- [cblock->vector](#), 97
- CFString, 75
- Class Display Names, 86
- [class_addIvar](#), 72
- [class_addMethod](#), 72
- CLSID*, 76
- [clsid->coclass](#), 86
- [clsid->progid](#), 78
- [clsid?](#), 77
- coclass*, 86
- [coclass->clsid](#), 86
- Cocoa Foundation, 75
- COM (Common Object Model), 76
- COM Automation, 77
- COM class*, 76
- COM Classes and Interfaces, 86

COM event executor, 82
 COM Events, 82
 COM FFI Helpers, 88
 COM Interface Example, 90
COM interfaces, 76
 COM Methods, 79
COM object, 76
 COM Objects, 78
 COM Properties, 80
 COM Types, 84
com-all-coclasses, 86
com-all-controls, 86
com-create-instance, 78
com-event-executor?, 82
com-event-type, 82
com-events, 82
com-get-active-object, 79
com-get-properties, 80
com-get-property, 81
*com-get-property**, 81
com-get-property-type, 80
com-idispatch?, 83
com-invoke, 80
com-iunknown?, 83
com-make-event-executor, 82
com-method-type, 80
com-methods, 80
com-object-clsid, 79
com-object-eq?, 79
com-object-get-idispatch, 83
com-object-get-iunknown, 83
com-object-set-clsid!, 79
com-object-type, 79
com-object?, 78
com-omit, 80
com-register-event-callback, 82
com-release, 78
com-set-properties, 81
com-set-property!, 81
com-set-property-type, 81
com-type=?, 79
com-type?, 79
com-unregister-event-callback, 83
compiler-sizeof, 18
cpointer-gcable?, 41
cpointer-has-tag?, 60
cpointer-push-tag!, 60
cpointer-tag, 44
cpointer?, 40
ctype->layout, 17
ctype-alignof, 17
ctype-basetype, 98
ctype-c->scheme, 98
ctype-scheme->c, 98
ctype-sizeof, 17
ctype?, 17
 Custodian Shutdown Registration, 63
 Custom Function Types, 29
custom function types, 29
cvector, 57
cvector->list, 58
cvector-length, 57
cvector-ptr, 57
cvector-ref, 58
cvector-set!, 58
cvector-type, 57
cvector?, 57
DCOM, 76
deallocator, 62
default-_string-type, 22
define-c, 15
define-com-interface, 87
define-cpointer-type, 59
define-cstruct, 32
define-ffi-definer, 60
define-fun-syntax, 30
define-objc-class, 68
define-objc-mixin, 70
 Defining Bindings, 60
 Derived Utilities, 49
 Describing COM Interfaces, 87
 dynamically loaded libraries, 13
end-atomic, 64
end-breakable-atomic, 64

- [end-stubborn-change](#), 45
- Enumerations and Masks, 39
- [errno](#), 25
- ['eternal](#), 45
- [f32vector](#), 54
- [f32vector->cpointer](#), 55
- [f32vector->list](#), 55
- [f32vector-length](#), 55
- [f32vector-ref](#), 55
- [f32vector-set!](#), 55
- [f32vector?](#), 54
- [f64vector](#), 55
- [f64vector->cpointer](#), 55
- [f64vector->list](#), 55
- [f64vector-length](#), 55
- [f64vector-ref](#), 55
- [f64vector-set!](#), 55
- [f64vector?](#), 55
- [f80vector](#), 56
- [f80vector->cpointer](#), 56
- [f80vector->list](#), 56
- [f80vector-length](#), 56
- [f80vector-ref](#), 56
- [f80vector-set!](#), 56
- [f80vector?](#), 56
- ['failok](#), 45
- FFI*, 1
- FFI Types and Constants, 66
- [ffi-call](#), 98
- [ffi-callback](#), 99
- [ffi-callback?](#), 99
- [ffi-lib](#), 13
- [ffi-lib?](#), 13
- [ffi-obj](#), 98
- [ffi-obj-lib](#), 98
- [ffi-obj-name](#), 98
- [ffi-obj-ref](#), 15
- [ffi-obj?](#), 98
- [ffi/com](#), 77
- [ffi/com-registry](#), 86
- [ffi/cvector](#), 56
- [ffi/file](#), 94
- [ffi/objc](#), 75
- [ffi/unsafe](#), 1
- [ffi/unsafe/alloc](#), 62
- [ffi/unsafe/atomic](#), 64
- [ffi/unsafe/com](#), 86
- [ffi/unsafe/custodian](#), 63
- [ffi/unsafe/cvector](#), 56
- [ffi/unsafe/define](#), 60
- [ffi/unsafe/nsalloc](#), 75
- [ffi/unsafe/nsstring](#), 75
- [ffi/unsafe/objc](#), 66
- [ffi/unsafe/try-atomic](#), 65
- [ffi/vector](#), 49
- [ffi/winapi](#), 94
- File Security-Guard Checks, 94
- Fixed Auto-Converting String Types, 21
- [flvector->cpointer](#), 96
- foreign-library value*, 13
- [free](#), 45
- [free-immobile-cell](#), 46
- Function Types, 24
- [function-ptr](#), 28
- Function-Type Bells and Whistles, 6
- [get-ffi-obj](#), 14
- [get-ivar](#), 70
- GetLastError, 25
- GUID*, 76
- [guid->string](#), 77
- [guid=?](#), 77
- [guid?](#), 77
- GUIDs, CLSIDs, IIDs, and ProgIDs, 77
- IDispatch, 77
- IID*, 76
- [iid?](#), 77
- [IID_IUnknown](#), 89
- [IID_NULL](#), 89
- [import-class](#), 68
- [import-protocol](#), 68
- [in-atomic-mode?](#), 65
- Interface Pointers, 83
- ['interior](#), 45
- Legacy Library, 75

Libraries, C Types, and Objects, 4

- [list->cblock](#), 96
- [list->cvector](#), 58
- [list->f32vector](#), 55
- [list->f64vector](#), 55
- [list->f80vector](#), 56
- [list->s16vector](#), 50
- [list->s32vector](#), 52
- [list->s64vector](#), 53
- [list->s8vector](#), 50
- [list->u16vector](#), 51
- [list->u32vector](#), 53
- [list->u64vector](#), 54
- [list->u8vector](#), 49

Loading Foreign Libraries, 13

[LOCALE_SYSTEM_DEFAULT](#), 89

[lookup-errno](#), 96

[make-array-type](#), 36

[make-c-parameter](#), 15

[make-com-object](#), 88

[make-construct-type](#), 32

[make-ctype](#), 17

[make-cvector](#), 57

[make-cvector*](#), 58

[make-f32vector](#), 54

[make-f64vector](#), 55

[make-f80vector](#), 56

[make-not-available](#), 61

[make-objc_super](#), 73

[make-s16vector](#), 50

[make-s32vector](#), 52

[make-s64vector](#), 53

[make-s8vector](#), 49

[make-sized-byte-string](#), 47

[make-u16vector](#), 51

[make-u32vector](#), 52

[make-u64vector](#), 54

[make-u8vector](#), 49

[make-union-type](#), 38

[malloc](#), 44

[malloc-immobile-cell](#), 46

[memcpy](#), 43

[memmove](#), 42

Memory Management, 44

[memset](#), 43

Miscellaneous Support, 96

More Examples, 12

MysterX, 77

[NO](#), 67

['nonatomic](#), 45

[NSString](#), 75

Numeric Types, 18

[objc-is-a?](#), 71

[objc-unsafe!](#), 75

[objc_allocateClassPair](#), 71

[objc_getProtocol](#), 71

[objc_lookupClass](#), 71

[objc_msgSend/typed](#), 73

[objc_msgSendSuper/typed](#), 73

[objc_registerClassPair](#), 71

[object_getClass](#), 71

[object_getInstanceVariable](#), 72

[object_setInstanceVariable](#), 72

Objective-C FFI, 66

Obtaining COM Interface References, 87

[offset_ptr?](#), 40

Other Atomic Types, 20

Other String Types, 22

Overview, 4

Pointer Dereferencing, 41

Pointer Functions, 40

Pointer Structure Property, 47

Pointer Types, 22

Pointers and GC-Managed Allocation, 9

Pointers and Manual Allocation, 8

Primitive String Types, 20

ProgID, 76

[progid->clsid](#), 78

[prop:cpointer](#), 47

[provide-protected](#), 62

[ptr-add](#), 40

[ptr-add!](#), 41

[ptr-equal?](#), 40

[ptr-offset](#), 40

- `ptr-ref`, 41
- `ptr-set!`, 41
- `QueryInterface`, 87
- `'raw`, 45
- Raw Runtime Functions, 71
- `register-custodian-shutdown`, 63
- `register-finalizer`, 46
- Release, 88
- `releaser`, 62
- Reliable Release of Resources, 11
- Remote COM servers (DCOM), 83
- `retainer`, 63
- `s16vector`, 50
- `s16vector->cpointer`, 51
- `s16vector->list`, 51
- `s16vector-length`, 50
- `s16vector-ref`, 50
- `s16vector-set!`, 50
- `s16vector?`, 50
- `s32vector`, 52
- `s32vector->cpointer`, 52
- `s32vector->list`, 52
- `s32vector-length`, 52
- `s32vector-ref`, 52
- `s32vector-set!`, 52
- `s32vector?`, 52
- `s64vector`, 53
- `s64vector->cpointer`, 53
- `s64vector->list`, 53
- `s64vector-length`, 53
- `s64vector-ref`, 53
- `s64vector-set!`, 53
- `s64vector?`, 53
- `s8vector`, 49
- `s8vector->cpointer`, 50
- `s8vector->list`, 50
- `s8vector-length`, 50
- `s8vector-ref`, 50
- `s8vector-set!`, 50
- `s8vector?`, 50
- Safe C Vectors, 56
- Safe Homogenous Vectors, 49
- `saved-errno`, 96
- `security-guard-check-file`, 94
- `sel_registerName`, 71
- `selector`, 70
- `self`, 70
- `set-cpointer-tag!`, 44
- `set-ffi-obj!`, 15
- `set-ivar!`, 70
- `set-ptr-offset!`, 41
- shared libraries, 13
- shared objects, 13
- Speculatively Atomic Execution, 65
- `start-atomic`, 64
- `start-breakable-atomic`, 64
- String Types, 20
- `string->clsid`, 77
- `string->guid`, 77
- `string->iid`, 77
- Strings, 75
- `'stubborn`, 45
- `super-tell`, 70
- Syntactic Forms and Procedures, 67
- `SysAllocStringLen`, 89
- `SysFreeString`, 89
- Tagged C Pointer Types, 58
- `tell`, 67
- `tellv`, 68
- The Racket Foreign Interface, 1
- `try-atomic`, 65
- Type Constructors, 17
- `type-describe`, 85
- `type-described-description`, 86
- `type-described-value`, 86
- `type-described?`, 85
- `type-description?`, 85
- `u16vector`, 51
- `u16vector->cpointer`, 51
- `u16vector->list`, 51
- `u16vector-length`, 51
- `u16vector-ref`, 51
- `u16vector-set!`, 51
- `u16vector?`, 51

- [u32vector](#), 52
- [u32vector->cpointer](#), 53
- [u32vector->list](#), 53
- [u32vector-length](#), 52
- [u32vector-ref](#), 52
- [u32vector-set!](#), 52
- [u32vector?](#), 52
- [u64vector](#), 54
- [u64vector->cpointer](#), 54
- [u64vector->list](#), 54
- [u64vector-length](#), 54
- [u64vector-ref](#), 54
- [u64vector-set!](#), 54
- [u64vector?](#), 54
- [u8vector](#), 49
- [u8vector->cpointer](#), 49
- [u8vector->list](#), 49
- [u8vector-length](#), 49
- [u8vector-ref](#), 49
- [u8vector-set!](#), 49
- [u8vector?](#), 49
- ['uncollectable](#), 45
- Unexported Primitive Functions, 98
- [union-ptr](#), 38
- [union-ref](#), 38
- [union-set!](#), 38
- [union?](#), 38
- [unregister-custodian-shutdown](#), 64
- Variable Auto-Converting String Type, 21
- [vector->cblock](#), 96
- [vector->cpointer](#), 96
- [win64?](#), 95
- [winapi](#), 95
- Windows API Helpers, 94
- [windows-error](#), 90
- [with-autorelease](#), 76
- [YES](#), 67