

Inside: PLT Scheme C API

Version 4.2.5

Matthew Flatt

April 2, 2010

This manual describes PLT Scheme's C interface, which allows the interpreter to be extended by a dynamically-loaded library, or embedded within an arbitrary C/C++ program. The manual assumes familiarity with PLT Scheme as described in *Reference: PLT Scheme*.

For an alternative way of dealing with foreign code, see *FFI: PLT Scheme Foreign Interface*, which describes the `scheme/foreign` module for manipulating low-level libraries and structures purely through Scheme code.

Contents

1 Overview	5
1.1 CGC versus 3m	5
1.2 Writing MzScheme Extensions	5
1.2.1 CGC Extensions	5
1.2.2 3m Extensions	7
1.2.3 Declaring a Module in an Extension	8
1.3 Embedding MzScheme into a Program	9
1.3.1 CGC Embedding	9
1.3.2 3m Embedding	12
1.4 MzScheme and Threads	15
1.5 MzScheme, Unicode, Characters, and Strings	15
1.6 Integers	15
2 Values and Types	16
2.1 Standard Types	17
2.2 Global Constants	20
2.3 Strings	20
2.4 Value Functions	20
3 Memory Allocation	30
3.1 Cooperating with 3m	31
3.1.1 Tagged Objects	32
3.1.2 Local Pointers	33
3.1.3 Local Pointers and <code>mzc -xform</code>	37
3.1.4 Guiding <code>mzc -xform</code>	38

3.2	Memory Functions	40
4	Namespaces and Modules	48
5	Procedures	51
6	Evaluation	54
6.1	Top-level Evaluation Functions	54
6.2	Tail Evaluation	54
6.3	Multiple Values	55
6.4	Evaluation Functions	56
7	Exceptions and Escape Continuations	60
7.1	Temporarily Catching Error Escapes	61
7.2	Enabling and Disabling Breaks	64
7.3	Exception Functions	64
8	Threads	69
8.1	Integration with Threads	69
8.2	Allowing Thread Switches	70
8.3	Blocking the Current Thread	70
8.4	Threads in Embedded Scheme with Event Loops	71
8.4.1	Callbacks for Blocked Threads	72
8.5	Sleeping by Embedded Scheme	74
8.6	Thread Functions	75
9	Parameterizations	83
10	Continuation Marks	87

11 String Encodings	88
12 Bignums, Rationals, and Complex Numbers	93
13 Ports and the Filesystem	97
14 Structures	115
15 Security Guards	118
16 Custodians	119
17 Miscellaneous Utilities	121
18 Flags and Hooks	128
Index	130

1 Overview

1.1 CGC versus 3m

Before mixing any C code with MzScheme, first decide whether to use the **3m** variant of PLT Scheme, the **CGC** variant of PLT Scheme, or both:

- **3m** : the main variant of PLT Scheme, which uses *precise* garbage collection instead of conservative garbage collection, and it may move objects in memory during a collection.
- **CGC** : the original variant of PLT Scheme, where memory management depends on a *conservative* garbage collector. The conservative garbage collector can automatically find references to managed values from C local variables and (on some platforms) static variables.

At the C level, working with CGC can be much easier than working with 3m, but overall system performance is typically better with 3m.

1.2 Writing MzScheme Extensions

The process of creating an extension for 3m or CGC is essentially the same, but the process for 3m is most easily understood as a variant of the process for CGC.

1.2.1 CGC Extensions

To write a C/C++-based extension for PLT Scheme CGC, follow these steps:

- For each C/C++ file that uses PLT Scheme library functions, `#include` the file `"escheme.h"`.
This file is distributed with the PLT software in an `"include"` directory, but if `mzc` is used to compile, this path is found automatically.
- Define the C function `scheme_initialize`, which takes a `Scheme_Env*` namespace (see §4 “Namespaces and Modules”) and returns a `Scheme_Object*` Scheme value.
This initialization function can install new global primitive procedures or other values into the namespace, or it can simply return a Scheme value. The initialization function is called when the extension is loaded with `load-extension` (the first time); the return value from `scheme_initialize` is used as the return value for `load-extension`. The namespace provided to `scheme_initialize` is the current namespace when `load-extension` is called.

- Define the C function `scheme_reload`, which has the same arguments and return type as `scheme_initialize`.

This function is called if `load-extension` is called a second time (or more times) for an extension. Like `scheme_initialize`, the return value from this function is the return value for `load-extension`.

- Define the C function `scheme_module_name`, which takes no arguments and returns a `Scheme_Object*` value, either a symbol or `scheme_false`.

The function should return a symbol when the effect of calling `scheme_initialize` and `scheme_reload` is only to declare a module with the returned name. This function is called when the extension is loaded to satisfy a `require` declaration.

The `scheme_module_name` function may be called before `scheme_initialize` and `scheme_reload`, after those functions, or both before and after, depending on how the extension is loaded and re-loaded.

- Compile the extension C/C++ files to create platform-specific object files.

The `mzc` compiler, which is distributed with PLT Scheme, compiles plain C files when the `-cc` flag is specified. More precisely, `mzc` does not compile the files itself, but it locates a C compiler on the system and launches it with the appropriate compilation flags. If the platform is a relatively standard Unix system, a Windows system with either Microsoft's C compiler or `gcc` in the path, or a Mac OS X system with Apple's developer tools installed, then using `mzc` is typically easier than working with the C compiler directly. Use the `-cgc` flag to indicate that the build is for use with PLT Scheme CGC.

- Link the extension C/C++ files with `"mzdyn.o"` (Unix, Mac OS X) or `"mzdyn.obj"` (Windows) to create a shared object. The resulting shared object should use the extension `".so"` (Unix), `".dll"` (Windows), or `".dylib"` (Mac OS X).

The `"mzdyn"` object file is distributed in the installation's `"lib"` directory. For Windows, the object file is in a compiler-specific sub-directory of `"plt\lib"`.

The `mzc` compiler links object files into an extension when the `-ld` flag is specified, automatically locating `"mzdyn"`. Again, use the `-cgc` flag with `mzc`.

- Load the shared object within Scheme using `(load-extension path)`, where `path` is the name of the extension file generated in the previous step.

Alternately, if the extension defines a module (i.e., `scheme_module_name` returns a symbol), then place the shared object in a special directory with a special name, so that it is detected by the module loader when `require` is used. The special directory is a platform-specific path that can be obtained by evaluating `(build-path "compiled" "native" (system-library-subpath))`; see `load/use-compiled` for more information. For example, if the shared object's name is `"example_ss.dll"`, then `(require "example.ss")` will be redirected to `"example_ss.dll"` if the latter is placed in the sub-directory `(build-path "compiled" "native" (system-library-subpath))` and if `"example.ss"` does not exist or has an earlier timestamp.

Note that `(load-extension path)` within a module does *not* introduce the extension's definitions into the module, because `load-extension` is a run-time operation. To introduce an extension's bindings into a module, make sure that the extension defines a module, put the extension in the platform-specific location as described above, and use `require`.

IMPORTANT: With PLT Scheme CGC, Scheme values are garbage collected using a conservative garbage collector, so pointers to Scheme objects can be kept in registers, stack variables, or structures allocated with `scheme_malloc`. However, static variables that contain pointers to collectable memory must be registered using `scheme_register_extension_global` (see §3 “Memory Allocation”).

As an example, the following C code defines an extension that returns `"hello world"` when it is loaded:

```
#include "escheme.h"
Scheme_Object *scheme_initialize(Scheme_Env *env) {
    return scheme_make_utf8_string("hello world");
}
Scheme_Object *scheme_reload(Scheme_Env *env) {
    return scheme_initialize(env); /* Nothing special for reload */
}
Scheme_Object *scheme_module_name() {
    return scheme_false;
}
```

Assuming that this code is in the file `"hw.c"`, the extension is compiled under Unix with the following two commands:

```
mzc -cgc -cc hw.c
mzc -cgc -ld hw.so hw.o
```

(Note that the `-cgc`, `-cc`, and `-ld` flags are each prefixed by two dashes, not one.)

The `"collects/mzscheme/examples"` directory in the PLT distribution contains additional examples.

1.2.2 3m Extensions

To build an extension to work with PLT Scheme 3m, the CGC instructions must be extended as follows:

- Adjust code to cooperate with the garbage collector as described in §3.1 “Cooperating with 3m”. Using `mzc` with the `-xform` might convert your code to implement part of

the conversion, as described in §3.1.3 “Local Pointers and `mzc -xform`”.

- In either your source in the in compiler command line, `#define MZ_PRECISE_GC` before including `"escheme.h"`. When using `mzc` with the `-cc` and `-3m` flags, `MZ_PRECISE_GC` is automatically defined.
- Link with `"mzdyn3m.o"` (Unix, Mac OS X) or `"mzdyn3m.obj"` (Windows) to create a shared object. When using `mzc`, use the `-ld` and `-3m` flags to link to these libraries.

For a relatively simple extension `"hw.c"`, the extension is compiled under Unix for 3m with the following three commands:

```
mzc -xform hw.c

mzc -3m -cc hw.3m.c

mzc -3m -ld hw.so hw.o
```

Some examples in `"collects/mzscheme/examples"` work with `MzScheme3m` in this way. A few examples are manually instrumented, in which case the `-xform` step should be skipped.

1.2.3 Declaring a Module in an Extension

To create an extension that behaves as a module, return a symbol from `scheme_module_name`, and have `scheme_initialize` and `scheme_rename` declare a module using `scheme_primitive_module`.

For example, the following extension implements a module named `hello` that exports a binding `greeting`:

```
#include "escheme.h"

Scheme_Object *scheme_initialize(Scheme_Env *env) {
    Scheme_Env *mod_env;
    mod_env = scheme_primitive_module(scheme_intern_symbol("hi"),
                                     env);
    scheme_add_global("greeting",
                    scheme_make_utf8_string("hello"),
                    mod_env);
    scheme_finish_primitive_module(mod_env);
    return scheme_void;
}

Scheme_Object *scheme_reload(Scheme_Env *env) {
    return scheme_initialize(env); /* Nothing special for reload */
}
```

```

}

Scheme_Object *scheme_module_name() {
    return scheme_intern_symbol("hi");
}

```

This extension could be compiled for 3m on Mac OS X for i386, for example, using the following sequence of `mzc` commands:

```

mzc -xform hi.c

mzc -3m -cc hi.3m.c

mkdir -p compiled/native/i386-macosx/3m

mzc -3m -ld compiled/native/i386-macosx/3m/hi_ss.dylib hi_3m.o

```

The resulting module can be loaded with

```
(require "hi.ss")
```

1.3 Embedding MzScheme into a Program

Like creating extensions, the embedding process for PLT Scheme CGC or PLT Scheme 3m is essentially the same, but the process for PLT Scheme 3m is most easily understood as a variant of the process for PLT Scheme CGC.

1.3.1 CGC Embedding

To embed PLT Scheme CGC in a program, follow these steps:

- Locate or build the PLT Scheme CGC libraries. Since the standard distribution provides 3m libraries, only, you will most likely have to build from source.

Under Unix, the libraries are `"libmzscheme.a"` and `"libmzgc.a"` (or `"libmzscheme.so"` and `"libmzgc.so"` for a dynamic-library build, with `"libmzscheme.la"` and `"libmzgc.la"` files for use with `libtool`). Building from source and installing places the libraries into the installation's `"lib"` directory. Be sure to build the CGC variant, since the default is 3m.

Under Windows, stub libraries for use with Microsoft tools are `"libmzschx.lib"` and `"libmzgcx.lib"` (where `x` represents the version number) are in a compiler-specific directory in `"plt\lib"`. These libraries identify the bindings that are provided by `"libmzschx.dll"` and `"libmzgcx.dll"` — which are typically installed

in "plt\lib". When linking with Cygwin, link to "libmzschx.dll" and "libmzgcx.dll" directly. At run time, either "libmzschx.dll" and "libmzgcx.dll" must be moved to a location in the standard DLL search path, or your embedding application must "delayload" link the DLLs and explicitly load them before use. ("MzScheme.exe" and "MrEd.exe" use the latter strategy.)

Under Mac OS X, dynamic libraries are provided by the "PLT_MzScheme" framework, which is typically installed in "lib" sub-directory of the installation. Supply `-framework PLT_MzScheme` to `gcc` when linking, along with `-F` and a path to the "lib" directory. Beware that CGC and 3m libraries are installed as different versions within a single framework, and installation marks one version or the other as the default (by setting symbolic links); install only CGC to simplify accessing the CGC version within the framework. At run time, either "PLT_MzScheme.framework" must be moved to a location in the standard framework search path, or your embedding executable must provide a specific path to the framework (possibly an executable-relative path using the Mach-O `@executable_path` prefix).

- For each C/C++ file that uses MzScheme library functions, `#include` the file "scheme.h".

The C preprocessor symbol `SCHEME_DIRECT_EMBEDDED` is defined as 1 when "scheme.h" is `#included`, or as 0 when "escheme.h" is `#included`.

The "scheme.h" file is distributed with the PLT software in the installation's "include" directory. Building and installing from source also places this file in the installation's "include" directory.

- Start your main program through the `scheme_main_setup` (or `scheme_main_stack_setup`) trampoline, and put all uses of MzScheme functions inside the function passed to `scheme_main_setup`. The `scheme_main_setup` function registers the current C stack location with the memory manager. It also creates the initial namespace `Scheme_Env*` by calling `scheme_basic_env` and passing the result to the function provided to `scheme_main_setup`. (The `scheme_main_stack_setup` trampoline registers the C stack with the memory manager without creating a namespace.)
- Configure the namespace by adding module declarations. The initial namespace contains declarations only for a few primitive modules, such as `'%kernel`, and no bindings are imported into the top-level environment.

To embed a module like `scheme/base` (along with all its dependencies), use `mzc -c-mods`, which generates a C file that contains modules in bytecode form as encapsulated in a static array. The generated C file defines a `declare_modules` function that takes a `Scheme_Env*`, installs the modules into the environment, and adjusts the module name resolver to access the embedded declarations.

Alternately, use `scheme_set_collects_path` and `scheme_init_collection_paths` to configure and install a path for finding modules at run time.

- Access Scheme through `scheme_dynamic_require`, `scheme_load`, `scheme_eval`, and/or other functions described in this manual.

If the embedding program configures built-in parameters in a way that should be considered part of the default configuration, then call `scheme_seal_parameters` afterward. The snapshot of parameter values taken by `scheme_seal_parameters` is used for certain privileged operations, such as installing a PLaneT package.

- Compile the program and link it with the MzScheme libraries.

With PLT Scheme CGC, Scheme values are garbage collected using a conservative garbage collector, so pointers to Scheme objects can be kept in registers, stack variables, or structures allocated with `scheme_malloc`. In an embedding application on some platforms, static variables are also automatically registered as roots for garbage collection (but see notes below specific to Mac OS X and Windows).

For example, the following is a simple embedding program which evaluates all expressions provided on the command line and displays the results, then runs a [read-eval-print](#) loop. Run

```
mzc -c-mods base.c ++lib scheme/base
```

to generate "base.c", which encapsulates `scheme/base` and all of its transitive imports (so that they need not be found separately a run time).

```
#include "scheme.h"

#include "base.c"

static int run(Scheme_Env *e, int argc, char *argv[])
{
    Scheme_Object *curout;
    int i;
    mz_jmp_buf * volatile save, fresh;

    /* Declare embedded modules in "base.c": */
    declare_modules(e);

    scheme_namespace_require(scheme_intern_symbol("scheme/base"));

    curout = scheme_get_param(scheme_current_config(),
                              MZCONFIG_OUTPUT_PORT);

    for (i = 1; i < argc; i++) {
        save = scheme_current_thread->error_buf;
        scheme_current_thread->error_buf = &fresh;
        if (scheme_setjmp(scheme_error_buf)) {
```

```

        scheme_current_thread->error_buf = save;
        return -1; /* There was an error */
    } else {
        Scheme_Object *v, *a[2];
        v = scheme_eval_string(argv[i], e);
        scheme_display(v, curout);
        scheme_display(scheme_make_char('\n'), curout);
        /* read-eval-print loop, uses initial Scheme_Env: */
        a[0] = scheme_intern_symbol("scheme/base");
        a[1] = scheme_intern_symbol("read-eval-print-loop");
        scheme_apply(scheme_dynamic_require(2, a), 0, NULL);
        scheme_current_thread->error_buf = save;
    }
}
return 0;
}

int main(int argc, char *argv[])
{
    return scheme_main_setup(1, run, argc, argv);
}

```

Under Mac OS X, or under Windows when MzScheme is compiled to a DLL using Cygwin, the garbage collector cannot find static variables automatically. In that case, `scheme_main_setup` must be called with a non-zero first argument.

Under Windows (for any other build mode), the garbage collector finds static variables in an embedding program by examining all memory pages. This strategy fails if a program contains multiple Windows threads; a page may get unmapped by a thread while the collector is examining the page, causing the collector to crash. To avoid this problem, call `scheme_main_setup` with a non-zero first argument.

When an embedding application calls `scheme_main_setup` with a non-zero first argument, it must register each of its static variables with `MZ_REGISTER_STATIC` if the variable can contain a GCable pointer. For example, if `curout` above is made `static`, then `MZ_REGISTER_STATIC(curout)` should be inserted before the call to `scheme_get_param`.

When building an embedded MzSchemeCGC to use SenoraGC (SGC) instead of the default collector, `scheme_main_setup` must be called with a non-zero first argument. See §3 “Memory Allocation” for more information.

1.3.2 3m Embedding

MzScheme3m can be embedded mostly the same as MzScheme, as long as the embedding program cooperates with the precise garbage collector as described in §3.1 “Cooperating

with 3m”.

In either your source in the in compiler command line, `#define MZ_PRECISE_GC` before including `"scheme.h"`. When using `mzc` with the `-cc` and `-3m` flags, `MZ_PRECISE_GC` is automatically defined.

In addition, some library details are different:

- Under Unix, the library is just `"libmzscheme3m.a"` (or `"libmzscheme3m.so"` for a dynamic-library build, with `"libmzscheme3m.la"` for use with `libtool`). There is no separate library for 3m analogous to CGC's `"libmzgc.a"`.
- Under Windows, the stub library for use with Microsoft tools is `"libmzsch3m x .lib"` (where x represents the version number). This library identifies the bindings that are provided by `"libmzsch3m x .dll"`. There is no separate library for 3m analogous to CGC's `"libmzgc x .lib"`.
- Under Mac OS X, 3m dynamic libraries are provided by the `"PLT_MzScheme"` framework, just as for CGC, but as a version suffixed with `"_3m"`.

For `MzScheme3m`, an embedding application must call `scheme_main_setup` with a non-zero first argument.

The simple embedding program from the previous section can be processed by `mzc -xform`, then compiled and linked with `MzScheme3m`. Alternately, the source code can be extended to work with either CGC or 3m depending on whether `MZ_PRECISE_GC` is defined on the compiler's command line:

```
#include "scheme.h"

#include "base.c"

static int run(Scheme_Env *e, int argc, char *argv[])
{
    Scheme_Object *curout = NULL, *v = NULL, *a[2] = {NULL, NULL};
    Scheme_Config *config = NULL;
    int i;
    mz_jmp_buf * volatile save = NULL, fresh;

    MZ_GC_DECL_REG(8);
    MZ_GC_VAR_IN_REG(0, e);
    MZ_GC_VAR_IN_REG(1, curout);
    MZ_GC_VAR_IN_REG(2, save);
    MZ_GC_VAR_IN_REG(3, config);
    MZ_GC_VAR_IN_REG(4, v);
    MZ_GC_ARRAY_VAR_IN_REG(5, a, 2);
```

```

MZ_GC_REG();

declare_modules(e);

v = scheme_intern_symbol("scheme/base");
scheme_namespace_require(v);

config = scheme_current_config();
curout = scheme_get_param(config, MZCONFIG_OUTPUT_PORT);

for (i = 1; i < argc; i++) {
    save = scheme_current_thread->error_buf;
    scheme_current_thread->error_buf = &fresh;
    if (scheme_setjmp(scheme_error_buf)) {
        scheme_current_thread->error_buf = save;
        return -1; /* There was an error */
    } else {
        v = scheme_eval_string(argv[i], e);
        scheme_display(v, curout);
        v = scheme_make_char('\n');
        scheme_display(v, curout);
        /* read-eval-print loop, uses initial Scheme_Env: */
        a[0] = scheme_intern_symbol("scheme/base");
        a[1] = scheme_intern_symbol("read-eval-print-loop");
        v = scheme_dynamic_require(2, a);
        scheme_apply(v, 0, NULL);
        scheme_current_thread->error_buf = save;
    }
}

MZ_GC_UNREG();

return 0;
}

int main(int argc, char *argv[])
{
    return scheme_main_setup(1, run, argc, argv);
}

```

Strictly speaking, the `config` and `v` variables above need not be registered with the garbage collector, since their values are not needed across function calls that allocate. The code is much easier to maintain, however, when all local variables are registered and when all temporary values are put into variables.

1.4 MzScheme and Threads

MzScheme implements threads for Scheme programs without aid from the operating system, so that MzScheme threads are cooperative from the perspective of C code. Under Unix, stand-alone MzScheme uses a single OS-implemented thread. Under Windows and Mac OS X, stand-alone MzScheme uses a few private OS-implemented threads for background tasks, but these OS-implemented threads are never exposed by the MzScheme API.

In an embedding application, MzScheme can co-exist with additional OS-implemented threads, but the additional OS threads must not call any `scheme_` function. Only the OS thread that originally calls `scheme_basic_env` can call `scheme_` functions. (This restriction is stronger than saying all calls must be serialized across threads. MzScheme relies on properties of specific threads to avoid stack overflow and garbage collection.) When `scheme_basic_env` is called a second time to reset the interpreter, it can be called in an OS thread that is different from the original call to `scheme_basic_env`. Thereafter, all calls to `scheme_` functions must originate from the new thread.

See §8 “Threads” for more information about threads, including the possible effects of MzScheme’s thread implementation on extension and embedding C code.

1.5 MzScheme, Unicode, Characters, and Strings

A character in MzScheme is a Unicode code point. In C, a character value has type `mzchar`, which is an alias for `unsigned` — which is, in turn, 4 bytes for a properly compiled MzScheme. Thus, a `mzchar*` string is effectively a UCS-4 string.

Only a few MzScheme functions use `mzchar*`. Instead, most functions accept `char*` strings. When such byte strings are to be used as character strings, they are interpreted as UTF-8 encodings. A plain ASCII string is always acceptable in such cases, since the UTF-8 encoding of an ASCII string is itself.

See also §2.3 “Strings” and §11 “String Encodings”.

1.6 Integers

MzScheme expects to be compiled in a mode where `short` is a 16-bit integer, `int` is a 32-bit integer, and `long` has the same number of bits as `void*`. The `mzlonglong` type has 64 bits for compilers that support a 64-bit integer type, otherwise it is the same as `long`; thus, `mzlonglong` tends to match `long long`. The `umzlonglong` type is the unsigned version of `mzlonglong`.

2 Values and Types

A Scheme value is represented by a pointer-sized value. The low bit is a mark bit: a 1 in the low bit indicates an immediate integer, a 0 indicates a (word-aligned) pointer.

A pointer Scheme value references a structure that begins with a `Scheme_Object` substructure, which in turn starts with a tag that has the C type `Scheme_Type`. The rest of the structure, following the `Scheme_Object` header, is type-dependent. PLT Scheme's C interface gives Scheme values the type `Scheme_Object*`. (The "object" here does not refer to objects in the sense of the `scheme/class` library.)

Examples of `Scheme_Type` values include `scheme_pair_type` and `scheme_symbol_type`. Some of these are implemented as instances of `Scheme_Simple_Object`, which is defined in "scheme.h", but extension or embedding code should never access this structure directly. Instead, the code should use macros, such as `SCHEME_CAR`, that provide access to the data of common Scheme types.

For most Scheme types, a constructor is provided for creating values of the type. For example, `scheme_make_pair` takes two `Scheme_Object*` values and returns the `cons` of the values.

The macro `SCHEME_TYPE` takes a `Scheme_Object *` and returns the type of the object. This macro performs the tag-bit check, and returns `scheme_integer_type` when the value is an immediate integer; otherwise, `SCHEME_TYPE` follows the pointer to get the type tag. Macros are provided to test for common Scheme types; for example, `SCHEME_PAIRP` returns 1 if the value is a cons cell, 0 otherwise.

In addition to providing constructors, PLT Scheme defines six global constant Scheme values: `scheme_true`, `scheme_false`, `scheme_null`, `scheme_eof`, `scheme_void`, and `scheme_undefined`. Each of these has a type tag, but each is normally recognized via its constant address.

An extension or embedding application can create new a primitive data type by calling `scheme_make_type`, which returns a fresh `Scheme_Type` value. To create a collectable instance of this type, allocate memory for the instance with `scheme_malloc`. From PLT Scheme's perspective, the main constraint on the data format of such an instance is that the first `sizeof(Scheme_Object)` bytes must correspond to a `Scheme_Object` record; furthermore, the first `sizeof(Scheme_Type)` bytes must contain the value returned by `scheme_make_type`. Extensions with modest needs can use `scheme_make_cptr`, instead of creating an entirely new type.

Scheme values should never be allocated on the stack, and they should never contain pointers to values on the stack. Besides the problem of restricting the value's lifetime to that of the stack frame, allocating values on the stack creates problems for continuations and threads, both of which copy into and out of the stack.

2.1 Standard Types

The following are the Scheme_Type values for the standard types:

- `scheme_bool_type` — the constants `scheme_true` and `scheme_false` are the only values of this type; use `SCHEME_FALSEP` to recognize `scheme_false` and use `SCHEME_TRUEP` to recognize anything except `scheme_false`; test for this type with `SCHEME_BOOLP`
- `scheme_char_type` — `SCHEME_CHAR_VAL` extracts the character (of type `mzchar`); test for this type with `SCHEME_CHARP`
- `scheme_integer_type` — fixnum integers, which are identified via the tag bit rather than following a pointer to this Scheme_Type value; `SCHEME_INT_VAL` extracts the integer; test for this type with `SCHEME_INTP`
- `scheme_double_type` — flonum inexact numbers; `SCHEME_FLOAT_VAL` or `SCHEME_DBL_VAL` extracts the floating-point value; test for this type with `SCHEME_DBLP`
- `scheme_float_type` — single-precision flonum inexact numbers, when specifically enabled when compiling PLT Scheme; `SCHEME_FLOAT_VAL` or `SCHEME_FLT_VAL` extracts the floating-point value; test for this type with `SCHEME_FLTP`
- `scheme_bignum_type` — test for this type with `SCHEME_BIGNUMP`
- `scheme_rational_type` — test for this type with `SCHEME_RATIONALP`
- `scheme_complex_type` — test for this type or `scheme_complex_izi_type` with `SCHEME_COMPLEXP`
- `scheme_complex_izi_type` — complex number with an inexact zero imaginary part (so it counts as a real number); test for this type specifically with `SCHEME_COMPLEX_IZIP`
- `scheme_char_string_type` — `SCHEME_CHAR_STR_VAL` extracts the string as a `mzchar*`; the string is always nul-terminated, but may also contain embedded nul characters, and the Scheme string is modified if this string is modified; `SCHEME_CHAR_STRLEN_VAL` extracts the string length (in characters, not counting the nul terminator); test for this type with `SCHEME_CHAR_STRINGP`
- `scheme_byte_string_type` — `SCHEME_BYTE_STR_VAL` extracts the string as a `char*`; the string is always nul-terminated, but may also contain embedded nul characters, and the Scheme string is modified if this string is modified; `SCHEME_BYTE_STRLEN_VAL` extracts the string length (in bytes, not counting the nul terminator); test for this type with `SCHEME_BYTE_STRINGP`
- `scheme_path_type` — `SCHEME_PATH_VAL` extracts the path as a `char*`; the string is always nul-terminated; `SCHEME_PATH_LEN` extracts the path length (in bytes, not counting the nul terminator); test for this type with `SCHEME_PATHP`

- `scheme_symbol_type` — `SCHEME_SYM_VAL` extracts the symbol’s string as a `char*` UTF-8 encoding (do not modify this string); `SCHEME_SYM_LEN` extracts the number of bytes in the symbol name (not counting the nul terminator); test for this type with `SCHEME_SYMBOLP`; 3m: see §3.1 “Cooperating with 3m” for a caution about `SCHEME_SYM_VAL`
- `scheme_keyword_type` — `SCHEME_KEYWORD_VAL` extracts the keyword’s string (without the leading hash colon) as a `char*` UTF-8 encoding (do not modify this string); `SCHEME_KEYWORD_LEN` extracts the number of bytes in the keyword name (not counting the nul terminator); test for this type with `SCHEME_KEYWORDP`; 3m: see §3.1 “Cooperating with 3m” for a caution about `SCHEME_KEYWORD_VAL`
- `scheme_box_type` — `SCHEME_BOX_VAL` extracts/sets the boxed value; test for this type with `SCHEME_BOXP`
- `scheme_pair_type` — `SCHEME_CAR` extracts/sets the `car` and `SCHEME_CDR` extracts/sets the `cdr`; test for this type with `SCHEME_PAIRP`
- `scheme_mutable_pair_type` — `SCHEME_MCAR` extracts/sets the `mcar` and `SCHEME_MCDR` extracts/sets the `mcdm`; test for this type with `SCHEME_MPAIRP`
- `scheme_vector_type` — `SCHEME_VEC_SIZE` extracts the length and `SCHEME_VEC_ELS` extracts the array of Scheme values (the Scheme vector is modified when this array is modified); test for this type with `SCHEME_VECTORP`; 3m: see §3.1 “Cooperating with 3m” for a caution about `SCHEME_VEC_ELS`
- `scheme_structure_type` — structure instances; test for this type with `SCHEME_STRUCTP`
- `scheme_struct_type_type` — structure types; test for this type with `SCHEME_STRUCT_TYPEP`
- `scheme_struct_property_type` — structure type properties
- `scheme_input_port_type` — `SCHEME_INPORT_VAL` extracts/sets the user data pointer; test for just this type with `SCHEME_INPORTP`, but use `SCHEME_INPUT_PORTP` to recognize all input ports (including structures with the `prop:input-port` property)
- `scheme_output_port_type` — `SCHEME_OUTPORT_VAL` extracts/sets the user data pointer; test for just this type with `SCHEME_OUTPORTP`, but use `SCHEME_OUTPUT_PORTP` to recognize all output ports (including structures with the `prop:output-port` property)
- `scheme_thread_type` — thread descriptors; test for this type with `SCHEME_THREADP`
- `scheme_sema_type` — semaphores; test for this type with `SCHEME_SEMAP`
- `scheme_hash_table_type` — test for this type with `SCHEME_HASHTP`

- `scheme_bucket_table_type` — test for this type with `SCHEME_BUCKTP`
- `scheme_weak_box_type` — test for this type with `SCHEME_WEAKP`; `SCHEME_WEAK_PTR` extracts the contained object, or `NULL` after the content is collected; do not set the content of a weak box
- `scheme_namespace_type` — namespaces; test for this type with `SCHEME_NAMESPACEP`
- `scheme_cpointer_type` — `#<void>` pointer with a type-describing `Scheme_Object`; `SCHEME_CPTR_VAL` extracts the pointer and `SCHEME_CPTR_TYPE` extracts the type tag object; test for this type with `SCHEME_CPTRP`. The tag is used when printing such objects when it's a symbol, a byte string, a string, or a pair holding one of these in its car.

The following are the procedure types:

- `scheme_prim_type` — a primitive procedure, possibly with data elements
- `scheme_closed_prim_type` — an old-style primitive procedure with a data pointer
- `scheme_compiled_closure_type` — a Scheme procedure
- `scheme_cont_type` — a continuation
- `scheme_escaping_cont_type` — an escape continuation
- `scheme_case_closure_type` — a case-lambda procedure
- `scheme_native_closure_type` — a procedure with native code generated by the just-in-time compiler

The predicate `SCHEME_PROCP` returns 1 for all procedure types and 0 for anything else.

The following are additional number predicates:

- `SCHEME_NUMBERP` — all numerical types
- `SCHEME_REALP` — all non-complex numerical types, plus `scheme_complex_izi_type`
- `SCHEME_EXACT_INTEGERP` — fixnums and bignums
- `SCHEME_EXACT_REALP` — fixnums, bignums, and rationals
- `SCHEME_FLOATP` — both single-precision (when enabled) and double-precision floats

2.2 Global Constants

There are six global constants:

- `scheme_null` — test for this value with `SCHEME_NULLP`
- `scheme_eof` — test for this value with `SCHEME_EOFP`
- `scheme_true`
- `scheme_false` — test for this value with `SCHEME_FALSEP`; test *against* it with `SCHEME_TRUEP`
- `scheme_void` — test for this value with `SCHEME_VOIDP`
- `scheme_undefined`

2.3 Strings

As noted in §1.5 “MzScheme, Unicode, Characters, and Strings”, a Scheme character is a Unicode code point represented by a `mzchar` value, and character strings are `mzchar` arrays. PLT Scheme also supplies byte strings, which are `char` arrays.

For a character string `s`, `SCHEME_CHAR_STR_VAL(s)` produces a pointer to `mzchars`, not `chars`. Convert a character string to its UTF-8 encoding as byte string with `scheme_char_string_to_byte_string`. For a byte string `bs`, `SCHEME_BYTE_STR_VAL(bs)` produces a pointer to `chars`. The function `scheme_byte_string_to_char_string` decodes a byte string as UTF-8 and produces a character string. The functions `scheme_char_string_to_byte_string_locale` and `scheme_byte_string_to_char_string_locale` are similar, but they use the current locale’s encoding instead of UTF-8.

For more fine-grained control over UTF-8 encoding, use the `scheme_utf8_decode` and `scheme_utf8_encode` functions, which are described in §11 “String Encodings”.

2.4 Value Functions

`Scheme_Object*` `scheme_make_char(mzchar ch)`

Returns the character value. The `ch` value must be a legal Unicode code point (and not a surrogate, for example). The first 256 characters are represented by constant Scheme values, and others are allocated.

`Scheme_Object* scheme_make_char_or_null(mzchar ch)`

Like `scheme_make_char`, but the result is NULL if `ch` is not a legal Unicode code point.

`Scheme_Object* scheme_make_character(mzchar ch)`

Returns the character value. This is a macro that directly accesses the array of constant characters when `ch` is less than 256.

`Scheme_Object* scheme_make_ascii_character(mzchar ch)`

Returns the character value, assuming that `ch` is less than 256. (This is a macro.)

`Scheme_Object* scheme_make_integer(long i)`

Returns the integer value; `i` must fit in a fixnum. (This is a macro.)

`Scheme_Object* scheme_make_integer_value(long i)`

Returns the integer value. If `i` does not fit in a fixnum, a bignum is returned.

`Scheme_Object*`
`scheme_make_integer_value_from_unsigned(unsigned long i)`

Like `scheme_make_integer_value`, but for unsigned integers.

`Scheme_Object*`
`scheme_make_integer_value_from_long_long(mzlonglong i)`

Like `scheme_make_integer_value`, but for `mzlonglong` values (see §1.6 “Integers”).

`Scheme_Object*`
`scheme_make_integer_value_from_unsigned_long_long(umzlonglong i)`

Like `scheme_make_integer_value_from_long_long`, but for unsigned integers.

`Scheme_Object*`
`scheme_make_integer_value_from_long_halves(unsigned long hi,
unsigned long lo)`

Creates an integer given the high and low longs of a signed integer. Note that on 64-bit

platforms where `long long` is the same as `long`, the resulting integer has 128 bits. (See also §1.6 “Integers”).

```
Scheme_Object*
scheme_make_integer_value_from_unsigned_long_halves(unsigned long hi,
                                                    unsigned long lo)
```

Creates an integer given the high and low longs of an unsigned integer. Note that on 64-bit platforms where `long long` is the same as `long`, the resulting integer has 128 bits.

```
int scheme_get_int_val(Scheme_Object* o,
                      long* i)
```

Extracts the integer value. Unlike the `SCHEME_INT_VAL` macro, this procedure will extract an integer that fits in a long from a Scheme bignum. If *o* fits in a long, the extracted integer is placed in **i* and 1 is returned; otherwise, 0 is returned and **i* is unmodified.

```
int scheme_get_unsigned_int_val(Scheme_Object* o,
                               unsigned long* i)
```

Like `scheme_get_int_val`, but for unsigned integers.

```
int scheme_get_long_long_val(Scheme_Object* o,
                             mzlonglong* i)
```

Like `scheme_get_int_val`, but for `mzlonglong` values (see §1.6 “Integers”).

```
int scheme_get_unsigned_long_long_val(Scheme_Object* o,
                                       umzlonglong* i)
```

Like `scheme_get_int_val`, but for unsigned `mzlonglong` values (see §1.6 “Integers”).

```
Scheme_Object* scheme_make_double(double d)
```

Creates a new floating-point value.

```
Scheme_Object* scheme_make_float(float d)
```

Creates a new single-precision floating-point value. The procedure is available only when PLT Scheme is compiled with single-precision numbers enabled.

```
double scheme_real_to_double(Scheme_Object* o)
```

Converts a Scheme real number to a double-precision floating-point value.

```
Scheme_Object* scheme_make_pair(Scheme_Object* carv,  
                                Scheme_Object* cdrv)
```

Makes a `cons` pair.

```
Scheme_Object* scheme_make_byte_string(char* bytes)
```

Makes a Scheme byte string from a nul-terminated C string. The *bytes* string is copied.

```
Scheme_Object*  
scheme_make_byte_string_without_copying(char* bytes)
```

Like `scheme_make_byte_string`, but the string is not copied.

```
Scheme_Object* scheme_make_sized_byte_string(char* bytes,  
                                             long len,  
                                             int copy)
```

Makes a byte string value with size *len*. A copy of *bytes* is made if *copy* is not 0. The string *bytes* should contain *len* bytes; *bytes* can contain the nul byte at any position, and need not be nul-terminated if *copy* is non-zero. However, if *len* is negative, then the nul-terminated length of *bytes* is used for the length, and if *copy* is zero, then *bytes* must be nul-terminated.

```
Scheme_Object* scheme_make_sized_offset_byte_string(char* bytes,  
                                                    long d,  
                                                    long len,  
                                                    int copy)
```

Like `scheme_make_sized_byte_string`, except the *len* characters start from position *d* in *bytes*. If *d* is non-zero, then *copy* must be non-zero.

```
Scheme_Object* scheme_alloc_byte_string(long size,  
                                         char fill)
```

Allocates a new Scheme byte string.

```
Scheme_Object* scheme_append_byte_string(Scheme_Object* a,  
                                         Scheme_Object* b)
```

Creates a new byte string by appending the two given byte strings.

```
Scheme_Object* scheme_make_locale_string(char* bytes)
```

Makes a Scheme string from a nul-terminated byte string that is a locale-specific encoding of a character string; a new string is allocated during decoding. The “locale” in the name of this function thus refers to *bytes*, and not the resulting string (which is internally stored as UCS-4).

```
Scheme_Object* scheme_make_utf8_string(char* bytes)
```

Makes a Scheme string from a nul-terminated byte string that is a UTF-8 encoding. A new string is allocated during decoding. The “utf8” in the name of this function thus refers to *bytes*, and not the resulting string (which is internally stored as UCS-4).

```
Scheme_Object* scheme_make_sized_utf8_string(char* bytes,  
                                             long len)
```

Makes a string value, based on *len* UTF-8-encoding bytes (so the resulting string is *len* characters or less). The string *bytes* should contain at least *len* bytes; *bytes* can contain the nul byte at any position, and need not be null-terminated. However, if *len* is negative, then the nul-terminated length of *bytes* is used for the length.

```
Scheme_Object* scheme_make_sized_offset_utf8_string(char* bytes,  
                                                    long d,  
                                                    long len)
```

Like `scheme_make_sized_char_string`, except the *len* characters start from position *d* in *bytes*.

```
Scheme_Object* scheme_make_char_string(mzchar* chars)
```

Makes a Scheme string from a nul-terminated UCS-4 string. The *chars* string is copied.

```
Scheme_Object*  
scheme_make_char_string_without_copying(mzchar* chars)
```

Like `scheme_make_char_string`, but the string is not copied.

```
Scheme_Object* scheme_make_sized_char_string(mzchar* chars,  
                                              long len,  
                                              int copy)
```

Makes a string value with size *len*. A copy of *chars* is made if *copy* is not 0. The string *chars* should contain *len* characters; *chars* can contain the nul character at any position, and need not be nul-terminated if *copy* is non-zero. However, if *len* is negative, then the nul-terminated length of *chars* is used for the length, and if *copy* is zero, then the *chars* must be nul-terminated.

```
Scheme_Object*  
scheme_make_sized_offset_char_string(mzchar* chars,  
                                     long d,  
                                     long len,  
                                     int copy)
```

Like `scheme_make_sized_char_string`, except the *len* characters start from position *d* in *chars*. If *d* is non-zero, then *copy* must be non-zero.

```
Scheme_Object* scheme_alloc_char_string(long size,  
                                         mzchar fill)
```

Allocates a new Scheme string.

```
Scheme_Object* scheme_append_char_string(Scheme_Object* a,  
                                         Scheme_Object* b)
```

Creates a new string by appending the two given strings.

```
Scheme_Object*  
scheme_char_string_to_byte_string(Scheme_Object* s)
```

Converts a Scheme character string into a Scheme byte string via UTF-8.

```
Scheme_Object*  
scheme_byte_string_to_char_string(Scheme_Object* s)
```

Converts a Scheme byte string into a Scheme character string via UTF-8.

```
Scheme_Object*  
scheme_char_string_to_byte_string_locale(Scheme_Object* s)
```

Converts a Scheme character string into a Scheme byte string via the locale's encoding.

```
Scheme_Object*  
scheme_byte_string_to_char_string_locale(Scheme_Object* s)
```

Converts a Scheme byte string into a Scheme character string via the locale's encoding.

```
Scheme_Object* scheme_intern_symbol(char* name)
```

Finds (or creates) the symbol matching the given nul-terminated, ASCII string (not UTF-8). The case of *name* is (non-destructively) normalized before interning if *scheme_case_sensitive* is 0.

```
Scheme_Object* scheme_intern_exact_symbol(char* name,  
                                           int len)
```

Creates or finds a symbol given the symbol's length in UTF-8-encoding bytes. The case of *name* is not normalized.

```
Scheme_Object* scheme_intern_exact_char_symbol(mzchar* name,  
                                               int len)
```

Like *scheme_intern_exact_symbol*, but given a character array instead of a UTF-8-encoding byte array.

```
Scheme_Object* scheme_make_symbol(char* name)
```

Creates an uninterned symbol from a nul-terminated, UTF-8-encoding string. The case is not normalized.

```
Scheme_Object* scheme_make_exact_symbol(char* name,  
                                         int len)
```

Creates an uninterned symbol given the symbol's length in UTF-8-encoded bytes.

```
Scheme_Object* scheme_intern_exact_keyword(char* name,  
                                           int len)
```

Creates or finds a keyword given the keywords length in UTF-8-encoding bytes. The case of *name* is not normalized, and it should not include the leading hash and colon of the keyword's printed form.

```
Scheme_Object* scheme_intern_exact_char_keyword(mzchar* name,  
                                                int len)
```

Like `scheme_intern_exact_keyword`, but given a character array instead of a UTF-8-encoding byte array.

```
Scheme_Object* scheme_make_vector(long size,  
                                  Scheme_Object* fill)
```

Allocates a new vector.

```
Scheme_Object* scheme_box(Scheme_Object* v)
```

Creates a new box containing the value *v*.

```
Scheme_Object* scheme_make_weak_box(Scheme_Object* v)
```

Creates a new weak box containing the value *v*.

```
Scheme_Type scheme_make_type(char* name)
```

Creates a new type (not a Scheme value).

```
Scheme_Object* scheme_make_cptr(void* ptr,  
                                const Scheme_Object* typetag)
```

Creates a C-pointer object that encapsulates *ptr* and uses *typetag* to identify the type of the pointer. The `SCHEME_CPTRP` macro recognizes objects created by `scheme_make_cptr`. The `SCHEME_CPTR_VAL` macro extracts the original *ptr* from the Scheme object, and `SCHEME_CPTR_TYPE` extracts the type tag. The `SCHEME_CPTR_OFFSETVAL` macro returns 0 for the result Scheme object.

The *ptr* can refer to either memory managed by the garbage collector or by some other memory manager. Beware, however, of retaining a *ptr* that refers to memory released by another memory manager, since the enclosing memory range might later become managed by the garbage collector (in which case *ptr* might become an invalid pointer that can crash the garbage collector).

```
Scheme_Object*  
scheme_make_external_cptr(void* ptr,  
                          const Scheme_Object* typetag)
```

Like `scheme_make_cptr`, but `ptr` is never treated as referencing memory managed by the garbage collector.

```
Scheme_Object*
scheme_make_offset_cptr(void* ptr,
                       long offset,
                       const Scheme_Object* typetag)
```

Creates a C-pointer object that encapsulates both `ptr` and `offset`. The `SCHEME_CPTR_OFFSETVAL` macro returns `offset` for the result Scheme object (and the macro be used to change the offset, since it also works on objects with no offset).

The `ptr` can refer to either memory managed by the garbage collector or by some other memory manager; see also `scheme_make_cptr`.

```
Scheme_Object*
scheme_make_offset_external_cptr(void* ptr,
                                long offset,
                                const Scheme_Object* typetag)
```

Like `scheme_make_offset_cptr`, but `ptr` is never treated as referencing memory managed by the garbage collector.

```
void scheme_set_type_printer(Scheme_Type type,
                             Scheme_Type_Printer printer)
```

Installs a printer to be used for printing (or writing or displaying) values that have the type tag `type`.

The type of `printer` is defined as follows:

```
typedef void (*Scheme_Type_Printer)(Scheme_Object *v, int dis,
                                    Scheme_Print_Params *pp);
```

Such a printer must print a representation of the value using `scheme_print_bytes` and `scheme_print_string`. The first argument to the printer, `v`, is the value to be printed. The second argument indicates whether `v` is printed via `write` or `display`. The last argument is to be passed on to `scheme_print_bytes` or `scheme_print_string` to identify the printing context.

```
void scheme_print_bytes(Scheme_Print_Params* pp,
                       const char* str,
                       int offset,
                       int len)
```

Writes the content of *str* — starting from *offset* and running *len* bytes — into a printing context determined by *pp*. This function is for use by a printer that is installed with `scheme_set_type_printer`.

```
void scheme_print_string(Scheme_Print_Params* pp,
                        const mzchar* str,
                        int offset,
                        int len)
```

Writes the content of *str* — starting from *offset* and running *len* characters — into a printing context determined by *pp*. This function is for use by a printer that is installed with `scheme_set_type_printer`.

```
void scheme_set_type_equality(Scheme_Type type,
                              Scheme_Equal_Proc equalp,
                              Scheme_Primary_Hash_Proc hash1,
                              Scheme_Secondary_Hash_Proc hash2)
```

Installs an equality predicate and associated hash functions for values that have the type tag *type*. The *equalp* predicate is only applied to values that both have tag *type*.

The type of *equalp*, *hash1*, and *hash2* are defined as follows:

```
typedef int (*Scheme_Equal_Proc)(Scheme_Object* obj1,
                                 Scheme_Object* obj2,
                                 void* cycle_data);
typedef long (*Scheme_Primary_Hash_Proc)(Scheme_Object* obj,
                                          long base,
                                          void* cycle_data);
typedef long (*Scheme_Secondary_Hash_Proc)(Scheme_Object* obj,
                                           void* cycle_data);
```

The two hash functions are used to generate primary and secondary keys for double hashing in an `equal?`-based hash table. The result of the primary-key function should depend on both *obj* and *base*.

The *cycle_data* argument in each case allows checking and hashing on cyclic values. It is intended for use in recursive checking or hashing via `scheme_recur_equal`, `scheme_recur_equal_hash_key`, and `scheme_recur_equal_hash_key`. That is, do not call plain `scheme_equal`, `scheme_equal_hash_key`, or `scheme_equal_hash_key` for recursive checking or hashing on sub-elements of the given value(s).

3 Memory Allocation

PLT Scheme uses both `malloc` and allocation functions provided by a garbage collector. Embedding/extension C/C++ code may use either allocation method, keeping in mind that pointers to garbage-collectable blocks in `malloc`d memory are invisible (i.e., such pointers will not prevent the block from being garbage-collected).

PLT Scheme CGC uses a conservative garbage collector. This garbage collector normally only recognizes pointers to the beginning of allocated objects. Thus, a pointer into the middle of a GC-allocated string will normally not keep the string from being collected. The exception to this rule is that pointers saved on the stack or in registers may point to the middle of a collectable object. Thus, it is safe to loop over an array by incrementing a local pointer variable.

PLT Scheme 3m uses a precise garbage collector that moves objects during collection, in which case the C code must be instrumented to expose local pointer bindings to the collector, and to provide tracing procedures for (tagged) records containing pointers. This instrumentation is described further in §3.1 “Cooperating with 3m”.

The basic collector allocation functions are:

- `scheme_malloc` — Allocates collectable memory that may contain pointers to collectable objects; for 3m, the memory must be an array of pointers (though not necessarily to collectable objects). The newly allocated memory is initially zeroed.
- `scheme_malloc_atomic` — Allocates collectable memory that does not contain pointers to collectable objects. If the memory does contain pointers, they are invisible to the collector and will not prevent an object from being collected. Newly allocated atomic memory is not necessary zeroed.
Atomic memory is used for strings or other blocks of memory which do not contain pointers. Atomic memory can also be used to store intentionally-hidden pointers.
- `scheme_malloc_tagged` — Allocates collectable memory that contains a mixture of pointers and atomic data. With the conservative collector, this function is the same as `scheme_malloc`, but under 3m, the type tag stored at the start of the block is used to determine the size and shape of the object for future garbage collection (as described in §3.1 “Cooperating with 3m”).
- `scheme_malloc_allow_interior` — Allocates a large array of pointers such that references are allowed into the middle of the block under 3m, and such pointers prevent the block from being collected. This procedure is the same as `scheme_malloc` with the conservative collector, but in the that case, having *only* a pointer into the interior will not prevent the array from being collected.
- `scheme_malloc_atomic_allow_interior` — Like `scheme_malloc_allow_interior` for memory that does not contain pointers.

- `scheme_malloc_uncollectable` — Allocates uncollectable memory that may contain pointers to collectable objects. There is no way to free the memory. The newly allocated memory is initially zeroed. This function is not available in 3m.

If a PLT Scheme extension stores Scheme pointers in a global or static variable, then that variable must be registered with `scheme_register_extension_global`; this makes the pointer visible to the garbage collector. Registered variables need not contain a collectable pointer at all times (even with 3m, but the variable must contain some pointer, possibly uncollectable, at all times).

With conservative collection, no registration is needed for the global or static variables of an embedding program, unless it calls `scheme_main_setup` or `scheme_set_stack_base` with a non-zero first or second (respectively) argument. In that case, global and static variables containing collectable pointers must be registered with `scheme_register_static`. The `MZ_REGISTER_STATIC` macro takes any variable name and registers it with `scheme_register_static`. The `scheme_register_static` function can be safely called even when it's not needed, but it must not be called multiple times for a single memory address.

Collectable memory can be temporarily locked from collection by using the reference-counting function `scheme_dont_gc_ptr`. Under 3m, such locking does not prevent the object from being moved.

Garbage collection can occur during any call into Scheme or its allocator, on anytime that Scheme has control, except during functions that are documented otherwise. The predicate and accessor macros listed in §2.1 “Standard Types” never trigger a collection.

3.1 Cooperating with 3m

To allow 3m's precise collector to detect and update pointers during garbage collection, all pointer values must be registered with the collector, at least during the times that a collection may occur. The content of a word registered as a pointer must contain either NULL, a pointer to the start of a collectable object, a pointer into an object allocated by `scheme_malloc_allow_interior`, a pointer to an object currently allocated by another memory manager (and therefore not into a block that is currently managed by the collector), or a pointer to an odd-numbered address (e.g., a Scheme fixnum).

Pointers are registered in three different ways:

- Pointers in static variables should be registered with `scheme_register_static` or `MZ_REGISTER_STATIC`.
- Pointers in allocated memory are registered automatically when they are in an array allocated with `scheme_malloc`, etc. When a pointer resides in an object allocated

with `scheme_malloc_tagged`, etc.~the tag at the start of the object identifies the object's size and shape. Handling of tags is described in §3.1.1 “Tagged Objects”.

- Local pointers (i.e., pointers on the stack or in registers) must be registered through the `MZ_GC_DECL_REG`, etc. macros that are described in §3.1.2 “Local Pointers”.

A pointer must never refer to the interior of an allocated object (when a garbage collection is possible), unless the object was allocated with `scheme_malloc_allow_interior`. For this reason, pointer arithmetic must usually be avoided, unless the variable holding the generated pointer is NULLed before a collection.

IMPORTANT: The `SCHEME_SYM_VAL`, `SCHEME_KEYWORD_VAL`, `SCHEME_VEC_ELS`, and `SCHEME_PRIM_CLOSURE_ELS` macros produce pointers into the middle of their respective objects, so the results of these macros must not be held during the time that a collection can occur. Incorrectly retaining such a pointer can lead to a crash.

3.1.1 Tagged Objects

As explained in §2 “Values and Types”, the `scheme_make_type` function can be used to obtain a new tag for a new type of object. These new types are in relatively short supply for 3m; the maximum tag is 255, and Scheme itself uses nearly 200.

After allocating a new tag in 3m (and before creating instances of the tag), a *size procedure*, a *mark procedure*, and a *fixup procedure* must be installed for the tag using `GC_register_traversers`.

A size procedure simply takes a pointer to an object with the tag and returns its size in words (not bytes). The `gcBYTES_TO_WORDS` macro converts a byte count to a word count.

A mark procedure is used to trace references among objects without moving any objects. The procedure takes a pointer to an object, and it should apply the `gcMARK` macro to every pointer within the object. The mark procedure should return the same result as the size procedure.

A fixup procedure is used to update references to objects after or while they are moved. The procedure takes a pointer to an object, and it should apply the `gcFIXUP` macro to every pointer within the object; the expansion of this macro takes the address of its argument. The fixup procedure should return the same result as the size procedure.

Depending on the collector's implementation, the mark or fixup procedure might not be used. For example, the collector may only use the mark procedure and not actually move the object. Or it may use the fixup procedure to mark and move objects at the same time. To dereference an object pointer during a fixup procedure, use `GC_fixup_self` to convert the address passed to the procedure to refer to the potentially moved object, and use `GC_resolve` to convert an address that is not yet fixed up to determine the object's current

location.

When allocating a tagged object in 3m, the tag must be installed immediately after the object is allocated—or, at least, before the next possible collection.

3.1.2 Local Pointers

The 3m collector needs to know the address of every local or temporary pointer within a function call at any point when a collection can be triggered. Beware that nested function calls can hide temporary pointers; for example, in

```
scheme_make_pair(scheme_make_pair(scheme_true, scheme_false),
                 scheme_make_pair(scheme_false, scheme_true))
```

the result from one `scheme_make_pair` call is on the stack or in a register during the other call to `scheme_make_pair`; this pointer must be exposed to the garbage collection and made subject to update. Simply changing the code to

```
tmp = scheme_make_pair(scheme_true, scheme_false);
scheme_make_pair(tmp,
                 scheme_make_pair(scheme_false, scheme_true))
```

does not expose all pointers, since `tmp` must be evaluated before the second call to `scheme_make_pair`. In general, the above code must be converted to the form

```
tmp1 = scheme_make_pair(scheme_true, scheme_false);
tmp2 = scheme_make_pair(scheme_true, scheme_false);
scheme_make_pair(tmp1, tmp2);
```

and this is converted form must be instrumented to register `tmp1` and `tmp2`. The final result might be

```
{
  Scheme_Object *tmp1 = NULL, *tmp2 = NULL, *result;
  MZ_GC_DECL_REG(2);

  MZ_GC_VAR_IN_REG(0, tmp1);
  MZ_GC_VAR_IN_REG(1, tmp2);
  MZ_GC_REG();

  tmp1 = scheme_make_pair(scheme_true, scheme_false);
  tmp2 = scheme_make_pair(scheme_true, scheme_false);
  result = scheme_make_pair(tmp1, tmp2);

  MZ_GC_UNREG();
```

```

    return result;
}

```

Notice that `result` is not registered above. The `MZ_GC_UNREG` macro cannot trigger a garbage collection, so the `result` variable is never live during a potential collection. Note also that `tmp1` and `tmp2` are initialized with `NULL`, so that they always contain a pointer whenever a collection is possible.

The `MZ_GC_DECL_REG` macro expands to a local-variable declaration to hold information for the garbage collector. The argument is the number of slots to provide for registration. Registering a simple pointer requires a single slot, whereas registering an array of pointers requires three slots. For example, to register a pointer `tmp` and an array of 10 `char*s`:

```

{
    Scheme_Object *tmp1 = NULL;
    char *a[10];
    int i;
    MZ_GC_DECL_REG(4);

    MZ_GC_ARRAY_VAR_IN_REG(0, a, 10);
    MZ_GC_VAR_IN_REG(3, tmp1);
    /* Clear a before a potential GC: */
    for (i = 0; i < 10; i++) a[i] = NULL;
    ...
    f(a);
    ...
}

```

The `MZ_GC_ARRAY_VAR_IN_REG` macro registers a local array given a starting slot, the array variable, and an array size. The `MZ_GC_VAR_IN_REG` macro takes a slot and simple pointer variable. A local variable or array must not be registered multiple times.

In the above example, the first argument to `MZ_GC_VAR_IN_REG` is 3 because the information for `a` uses the first three slots. Even if `a` is not used after the call to `f`, `a` must be registered with the collector during the entire call to `f`, because `f` presumably uses `a` until it returns.

The name used for a variable need not be immediate. Structure members can be supplied as well:

```

{
    struct { void *s; int v; void *t; } x = {NULL, 0, NULL};
    MZ_GC_DECL_REG(2);

    MZ_GC_VAR_IN_REG(0, x.s);
    MZ_GC_VAR_IN_REG(0, x.t);
    ...
}

```

```
}
```

In general, the only constraint on the second argument to `MZ_GC_VAR_IN_REG` or `MZ_GC_ARRAY_VAR_IN_REG` is that `&` must produce the relevant address, and that address must be on the stack.

Pointer information is not actually registered with the collector until the `MZ_GC_REG` macro is used. The `MZ_GC_UNREG` macro de-registers the information. Each call to `MZ_GC_REG` must be balanced by one call to `MZ_GC_UNREG`.

Pointer information need not be initialized with `MZ_GC_VAR_IN_REG` and `MZ_GC_ARRAY_VAR_IN_REG` before calling `MZ_GC_REG`, and the set of registered pointers can change at any time—as long as all relevant pointers are registered when a collection might occur. The following example recycles slots and completely de-registers information when no pointers are relevant. The example also illustrates how `MZ_GC_UNREG` is not needed when control escapes from the function, such as when `scheme_signal_error` escapes.

```
{
  Scheme_Object *tmp1 = NULL, *tmp2 = NULL;
  mzchar *a, *b;
  MZ_GC_DECL_REG(2);

  MZ_GC_VAR_IN_REG(0, tmp1);
  MZ_GC_VAR_IN_REG(1, tmp2);

  tmp1 = scheme_make_utf8_string("foo");
  MZ_GC_REG();
  tmp2 = scheme_make_utf8_string("bar");
  tmp1 = scheme_append_char_string(tmp1, tmp2);

  if (SCHEME_FALSEP(tmp1))
    scheme_signal_error("shouldn't happen!");

  a = SCHEME_CHAR_VAL(tmp1);

  MZ_GC_VAR_IN_REG(0, a);

  tmp2 = scheme_make_pair(scheme_read_bignum(a, 0, 10), tmp2);

  MZ_GC_UNREG();

  if (SCHEME_INTP(tmp2)) {
    return 0;
  }

  MZ_GC_REG();
}
```

```

    tmp1 = scheme_make_pair(scheme_read_bignum(a, 0, 8), tmp2);
    MZ_GC_UNREG();

    return tmp1;
}

```

A `MZ_GC_DECL_REG` can be used in a nested block to hold declarations for the block's variables. In that case, the nested `MZ_GC_DECL_REG` must have its own `MZ_GC_REG` and `MZ_GC_UNREG` calls.

```

{
    Scheme_Object *accum = NULL;
    MZ_GC_DECL_REG(1);
    MZ_GC_VAR_IN_REG(0, accum);
    MZ_GC_REG();

    accum = scheme_make_pair(scheme_true, scheme_null);
    {
        Scheme_Object *tmp = NULL;
        MZ_GC_DECL_REG(1);
        MZ_GC_VAR_IN_REG(0, tmp);
        MZ_GC_REG();

        tmp = scheme_make_pair(scheme_true, scheme_false);
        accum = scheme_make_pair(tmp, accum);

        MZ_GC_UNREG();
    }
    accum = scheme_make_pair(scheme_true, accum);

    MZ_GC_UNREG();
    return accum;
}

```

Variables declared in a local block can also be registered together with variables from an enclosing block, but the local-block variable must be unregistered before it goes out of scope. The `MZ_GC_NO_VAR_IN_REG` macro can be used to unregister a variable or to initialize a slot as having no variable.

```

{
    Scheme_Object *accum = NULL;
    MZ_GC_DECL_REG(2);
    MZ_GC_VAR_IN_REG(0, accum);
    MZ_GC_NO_VAR_IN_REG(1);
    MZ_GC_REG();
}

```

```

accum = scheme_make_pair(scheme_true, scheme_null);
{
  Scheme_Object *tmp = NULL;
  MZ_GC_VAR_IN_REG(1, tmp);

  tmp = scheme_make_pair(scheme_true, scheme_false);
  accum = scheme_make_pair(tmp, accum);

  MZ_GC_NO_VAR_IN_REG(1);
}
accum = scheme_make_pair(scheme_true, accum);

MZ_GC_UNREG();
return accum;
}

```

The `MZ_GC_` macros all expand to nothing when `MZ_PRECISE_GC` is not defined, so the macros can be placed into code to be compiled for both conservative and precise collection.

The `MZ_GC_REG` and `MZ_GC_UNREG` macros must never be used in an OS thread other than Scheme's thread.

3.1.3 Local Pointers and `mzc -xform`

When `mzc` is run with the `-xform` flag and a source C program, it produces a C program that is instrumented in the way described in the previous section (but with a slightly different set of macros). For each input file "*name*.c", the transformed output is "*name*.3m.c".

The `-xform` mode for `mzc` does not change allocation calls, nor does it generate size, mark, or fixup predocures. It merely converts the code to register local pointers.

Furthermore, the `-xform` mode for `mzc` does not handle all of C. It's ability to rearrange compound expressions is particularly limited, because `-xform` merely converts expression text heuristically instead of parsing C. A future version of the tool will correct such problems. For now, `mzc` in `-xform` mode attempts to provide reasonable error messages when it is unable to convert a program, but beware that it can miss cases. To an even more limited degree, `-xform` can work on C++ code. Inspect the output of `-xform` mode to ensure that your code is correctly instrumented.

Some specific limitations:

- The body of a `for`, `while`, or `do` loop must be surrounded with curly braces. (A conversion error is normally reported, otherwise.)
- Function calls may not appear on the right-hand side of an assignment within a dec-

laration block. (A conversion error is normally reported if such an assignment is discovered.)

- Multiple function calls in `... ? ... : ...` cannot be lifted. (A conversion error is normally reported, otherwise.)
- In an assignment, the left-hand side must be a local or static variable, not a field selection, pointer dereference, etc. (A conversion error is normally reported, otherwise.)
- The conversion assumes that all function calls use an immediate name for a function, as opposed to a compound expression as in `s->f()`. The function name need not be a top-level function name, but it must be bound either as an argument or local variable with the form `type id`; the syntax `ret_type (*id)(...)` is not recognized, so bind the function type to a simple name with typedef, first: `typedef ret_type (*type)(...); ... type id`.
- Arrays and structs must be passed by address, only.
- GC-triggering code must not appear in system headers.
- Pointer-comparison expressions are not handled correctly when either of the compared expressions includes a function call. For example, `a() == b()` is not converted correctly when a and b produce pointer values.
- Passing the address of a local pointer to a function works only when the pointer variable remains live after the function call.
- A `return;` form can get converted to `{ stmt; return; };`, which can break an `if (...) return; else ...` pattern.
- Local instances of union types are generally not supported.
- Pointer arithmetic cannot be converted away, and is instead reported as an error.

3.1.4 Guiding `mzc -xform`

The following macros can be used (with care!) to navigate `-xform` around code that it cannot handle:

- `XFORM_START_SKIP` and `XFORM_END_SKIP`: code between these two statements is ignored by the transform tool, except to tokenize it.

Example:

```
int foo(int c, ...) {
    int r = 0;
    XFORM_START_SKIP;
    {
        /* va plays strange tricks that confuse xform */
    }
}
```

```

    va_list args;
    va_start(args, c);
    while (c-) {
        r += va_arg(args, int);
    }
}
XFORM_END_SKIP;
return r;
}

```

These macros can also be used at the top level, outside of any function. Since they have to be terminated by a semi-colon, however, top-level uses usually must be wrapped with `#ifdef MZ_PRECISE_GC` and `#endif`; a semi-colon by itself at the top level is not legal in C.

- `XFORM_SKIP_PROC`: annotate a function so that its body is skipped in the same way as bracketing it with `XFORM_START_SKIP` and `XFORM_END_SKIP`.

Example:

```

int foo(int c, ...) XFORM_END_SKIP {
}

```

- `XFORM_HIDE_EXPR`: a macro that takes wraps an expression to disable processing of the expression.

Example:

```

int foo(int c, ...) {
    int r = 0;
    {
        /* va plays strange tricks that confuse xform */
        XFORM_CAN_IGNORE va_list args; /* See below */
        XFORM_HIDE_EXPR(va_start(args, c));
        while (c-) {
            r += XFORM_HIDE_EXPR(va_arg(args, int));
        }
    }
    return r;
}

```

- `XFORM_CAN_IGNORE`: a macro that acts like a type modifier (must appear first) to indicate that a declared variable can be treated as atomic. See above for an example.
- `XFORM_START_SUSPEND` and `XFORM_END_SUSPEND`: for use at the top level (outside of any function definition), and similar to `XFORM_START_SKIP` and `XFORM_END_SKIP` in that function and class bodies are not transformed. Type and prototype information is still collected for use by later transformations, however. These forms must be terminated by a semi-colon.

- `XFORM_START_TRUST_ARITH` and `XFORM_END_TRUST_ARITH`: for use at the top level (outside of any function definition) to disable warnings about pointer arithmetic. Use only when you're absolutely certain that the garbage collector cannot be pointers offset into the middle of a collectable object. These forms must be terminated by a semicolon.
- `XFORM_TRUST_PLUS`: a replacement for `+` that does not trigger pointer-arithmetic warnings. Use with care.
- `XFORM_TRUST_MINUS`: a replacement for `-` that does not trigger pointer-arithmetic warnings. Use with care.

3.2 Memory Functions

```
void* scheme_malloc(size_t n)
```

Allocates n bytes of collectable memory, initially filled with zeros. In 3m, the allocated object is treated as an array of pointers.

```
void* scheme_malloc_atomic(size_t n)
```

Allocates n bytes of collectable memory containing no pointers visible to the garbage collector. The object is *not* initialized to zeros.

```
void* scheme_malloc_uncollectable(size_t n)
```

Non-3m, only. Allocates n bytes of uncollectable memory.

```
void* scheme_malloc_eternal(size_t n)
```

Allocates uncollectable atomic memory. This function is equivalent to `malloc`, except that the memory cannot be freed.

```
void* scheme_calloc(size_t num,
                   size_t size)
```

Allocates $num * size$ bytes of memory using `scheme_malloc`.

```
void* scheme_malloc_tagged(size_t n)
```

Like `scheme_malloc`, but in 3m, the type tag determines how the garbage collector traverses the object; see §3 “Memory Allocation”.

```
void* scheme_malloc_allow_interior(size_t n)
```

Like `scheme_malloc`, but in 3m, pointers are allowed to reference the middle of the object; see §3 “Memory Allocation”.

```
char* scheme_strdup(char* str)
```

Copies the null-terminated string `str`; the copy is collectable.

```
char* scheme_strdup_eternal(char* str)
```

Copies the null-terminated string `str`; the copy will never be freed.

```
void* scheme_malloc_fail_ok(size_t size,  
                           size_t size)
```

Attempts to allocate `size` bytes using `malloc`. If the allocation fails, the `exn:misc:out-of-memory` exception is raised.

```
void** scheme_malloc_immobile_box(void* p)
```

Allocates memory that is not garbage-collected and that does not move (even with 3m), but whose first word contains a pointer to a collectable object. The box is initialized with `p`, but the value can be changed at any time. An immobile box must be explicitly freed using `scheme_free_immobile_box`.

```
void scheme_free_immobile_box(void** b)
```

Frees an immobile box allocated with `scheme_malloc_immobile_box`.

```
void* scheme_malloc_code(long size)
```

Allocates non-collectable memory to hold executable machine code. Use this function instead of `malloc` to ensure that the allocated memory has “execute” permissions. Use `scheme_free_code` to free memory allocated by this function.

```
void scheme_free_code(void* p)
```

Frees memory allocated with `scheme_malloc_code`.

```
void scheme_register_extension_global(void* ptr,
                                     long size)
```

Registers an extension's global variable that can contain Scheme pointers. The address of the global is given in *ptr*, and its size in bytes in *size*. In addition to global variables, this function can be used to register any permanent memory that the collector would otherwise treat as atomic. A garbage collection can occur during the registration.

```
int scheme_main_setup(int no_auto_statics,
                     Scheme_Env_Main main,
                     int argc,
                     char** argv)
```

Initializes the GC stack base, creates the initial namespace by calling `scheme_basic_env`, and then calls *main* with the namespace, *argc*, and *argv*. (The *argc* and *argv* are just passed on to *main*, and are not inspected in any way.)

The `Scheme_Env_Main` type is defined as follows:

```
typedef int (*Scheme_Env_Main)(Scheme_Env *env,
                               int argc, char **argv);
```

The result of *main* is the result of `scheme_main_setup`.

If *no_auto_statics* is non-zero, then static variables must be explicitly registered with the garbage collector; see §3 “Memory Allocation” for more information.

```
int scheme_main_stack_setup(int no_auto_statics,
                            Scheme_Nested_Main main,
                            void* data)
```

A more primitive variant of `scheme_main_setup` that initializes the GC stack base but does not create the initial namespace (so an embedding application can perform other operations that involve garbage-collected data before creating a namespace).

The *data* argument is passed through to *main*, where the `Scheme_Nested_Main` type is defined as follows:

```
typedef int (*Scheme_Nested_Main)(void *data);
```

```
void scheme_set_stack_base(void* stack_addr,
                          int no_auto_statics)
```

Overrides the GC's auto-determined stack base, and/or disables the GC's automatic traversal

of global and static variables. If *stack_addr* is NULL, the stack base determined by the GC is used. Otherwise, it should be the “deepest” memory address on the stack where a collectable pointer might be stored. This function should be called only once, and before any other *scheme_function* is called. It never triggers a garbage collection.

The following example shows a typical use for setting the stack base for CGC:

```
int main(int argc, char **argv) {
    int dummy;
    scheme_set_stack_base(&dummy, 0);
    real_main(argc, argv); /* calls scheme_basic_env(), etc. */
}
```

Under 3m, the above code does not quite work, because *stack_addr* must be the beginning or end of a local-frame registration. Worse, in CGC or 3m, if *real_main* is declared *static*, the compiler may inline it and place variables containing collectable values deeper in the stack than *dummy*. To avoid these problems, use *scheme_main_setup* or *scheme_main_stack_setup*, instead.

```
void scheme_set_stack_bounds(void* stack_addr,
                             void* stack_end,
                             int no_auto_statics)
```

Like *scheme_set_stack_base*, except for the extra *stack_end* argument. If *stack_end* is non-NULL, then it corresponds to a point of C-stack growth after which Scheme should attempt to handle stack overflow. The *stack_end* argument should not correspond to the actual stack end, since detecting stack overflow may take a few frames, and since handling stack overflow requires a few frames.

If *stack_end* is NULL, then the stack end is computed automatically: the stack size assumed to be the limit reported by *getrlimit* under Unix and Mac OS X, or it is assumed to be 1 MB under Windows; if this size is greater than 8 MB, then 8 MB is assumed, instead; the size is decremented by 50000 bytes to cover a large margin of error; finally, the size is subtracted from (for stacks that grow down) or added to (for stacks that grow up) the stack base in *stack_addr* or the automatically computed stack base. Note that the 50000-byte margin of error is assumed to cover the difference between the actual stack start and the reported stack base, in addition to the margin needed for detecting and handling stack overflow.

```
void scheme_register_static(void* ptr,
                            long size)
```

Like *scheme_register_extension_global*, for use in embedding applications in situations where the collector does not automatically find static variables (i.e., when *scheme_set_stack_base* has been called with a non-zero second argument).

The macro `MZ_REGISTER_STATIC` can be used directly on a static variable. It expands to a comment if statics need not be registered, and a call to `scheme_register_static` (with the address of the static variable) otherwise.

```
void scheme_weak_reference(void** p)
```

Registers the pointer `*p` as a weak pointer; when no other (non-weak) pointers reference the same memory as `*p` references, then `*p` will be set to `NULL` by the garbage collector. The value in `*p` may change, but the pointer remains weak with respect to the value of `*p` at the time `p` was registered.

This function is not available in 3m.

```
void scheme_weak_reference_indirect(void** p,  
                                   void* v)
```

Like `scheme_weak_reference`, but `*p` is set to `NULL` (regardless of its prior value) when there are no references to `v`.

This function is not available in 3m.

```
void scheme_register_finalizer(void* p,  
                              fnl_proc f,  
                              void* data,  
                              fnl_proc* oldf,  
                              void** olddata)
```

Registers a callback function to be invoked when the memory `p` would otherwise be garbage-collected, and when no “will”-like finalizers are registered for `p`.

The `fnl_proc` type is not actually defined, but it is equivalent to

```
typedef void (*fnl_proc)(void *p, void *data)
```

The `f` argument is the callback function; when it is called, it will be passed the value `p` and the data pointer `data`; `data` can be anything — it is only passed on to the callback function. If `oldf` and `olddata` are not `NULL`, then `*oldf` and `*olddata` are filled with the old callback information (`f` and `data` will override this old callback).

To remove a registered finalizer, pass `NULL` for `f` and `data`.

Note: registering a callback not only keeps `p` from collection until the callback is invoked, but it also keeps `data` reachable until the callback is invoked.

```
void scheme_add_finalizer(void* p,  
                          fnl_proc f,  
                          void* data)
```

Adds a finalizer to a chain of primitive finalizers. This chain is separate from the single finalizer installed with `scheme_register_finalizer`; all finalizers in the chain are called immediately after a finalizer that is installed with `scheme_register_finalizer`.

See `scheme_register_finalizer`, above, for information about the arguments.

To remove an added finalizer, use `scheme_subtract_finalizer`.

```
void scheme_add_scheme_finalizer(void* p,  
                                 fnl_proc f,  
                                 void* data)
```

Installs a “will”-like finalizer, similar to [will-register](#). Scheme finalizers are called one at a time, requiring the collector to prove that a value has become inaccessible again before calling the next Scheme finalizer. Finalizers registered with `scheme_register_finalizer` or `scheme_add_finalizer` are not called until all Scheme finalizers have been exhausted.

See `scheme_register_finalizer`, above, for information about the arguments.

There is currently no facility to remove a “will”-like finalizer.

```
void scheme_add_finalizer_once(void* p,  
                              fnl_proc f,  
                              void* data)
```

Like `scheme_add_finalizer`, but if the combination `f` and `data` is already registered as a (non-“will”-like) finalizer for `p`, it is not added a second time.

```
void scheme_add_scheme_finalizer_once(void* p,  
                                      fnl_proc f,  
                                      void* data)
```

Like `scheme_add_scheme_finalizer`, but if the combination of `f` and `data` is already registered as a “will”-like finalizer for `p`, it is not added a second time.

```
void scheme_subtract_finalizer(void* p,  
                              fnl_proc f,  
                              void* data)
```

Removes a finalizer that was installed with `scheme_add_finalizer`.

```
void scheme_remove_all_finalization(void* p)
```

Removes all finalization (“will”-like or not) for *p*, including wills added in Scheme with `will-register` and finalizers used by custodians.

```
void scheme_dont_gc_ptr(void* p)
```

Keeps the collectable block *p* from garbage collection. Use this procedure when a reference to *p* is to be stored somewhere inaccessible to the collector. Once the reference is no longer used from the inaccessible region, de-register the lock with `scheme_gc_ptr_ok`. A garbage collection can occur during the registration.

This function keeps a reference count on the pointers it registers, so two calls to `scheme_dont_gc_ptr` for the same *p* should be balanced with two calls to `scheme_gc_ptr_ok`.

```
void scheme_gc_ptr_ok(void* p)
```

See `scheme_dont_gc_ptr`.

```
void scheme_collect_garbage()
```

Forces an immediate garbage-collection.

```
void GC_register_traversers(short tag,  
                           Size_Proc s,  
                           Mark_Proc m,  
                           Fixup_Proc f,  
                           int is_const_size,  
                           int is_atomic)
```

3m only. Registers a size, mark, and fixup procedure for a given type tag; see §3.1.1 “Tagged Objects” for more information.

Each of the three procedures takes a pointer and returns an integer:

```
typedef int (*Size_Proc)(void *obj);  
typedef int (*Mark_Proc)(void *obj);  
typedef int (*Fixup_Proc)(void *obj);
```

If the result of the size procedure is a constant, then pass a non-zero value for `is_const_size`. If the mark and fixup procedures are no-ops, then pass a non-zero value for `is_atomic`.

```
void* GC_resolve(void* p)
```

3m only. Can be called by a size, mark, or fixup procedure that is registered with `GC_register_traversers`. It returns the current address of an object *p* that might have been moved already, where *p* corresponds to an object that is referenced directly by the object being sized, marked, or fixed. This translation is necessary, for example, if the size or structure of an object depends on the content of an object it references. For example, the size of a class instance usually depends on a field count that is stored in the class. A fixup procedure should call this function on a reference *before* fixing it.

```
void* GC_fixup_self(void* p)
```

3m only. Can be called by a fixup procedure that is registered with `GC_register_traversers`. It returns the final address of *p*, which must be the pointer passed to the fixup procedure. For some implementations of the memory manager, the result is the same as *p*, either because objects are not moved or because the object is moved before it is fixed. With other implementations, an object might be moved after the fixup process, and the result is the location that the object will have after garbage collection finished.

4 Namespaces and Modules

A Scheme namespace (a top-level environment) is represented by a value of type `Scheme_Env*` — which is also a Scheme value, castable to `Scheme_Object*`. Calling `scheme_basic_env` returns a namespace that includes all of Scheme’s standard global procedures and syntax.

The `scheme_basic_env` function must be called once by an embedding program, before any other PLT Scheme function is called (except `scheme_make_param`), but `scheme_main_setup` automatically calls `scheme_basic_env`. The returned namespace is the initial current namespace for the main Scheme thread. Scheme extensions cannot call `scheme_basic_env`.

The current thread’s current namespace is available from `scheme_get_env`, given the current parameterization (see §9 “Parameterizations”): `scheme_get_env(scheme_config)`.

New values can be added as globals in a namespace using `scheme_add_global`. The `scheme_lookup_global` function takes a Scheme symbol and returns the global value for that name, or `NULL` if the symbol is undefined.

A module’s set of top-level bindings is implemented using the same machinery as a namespace. Use `scheme_primitive_module` to create a new `Scheme_Env*` that represents a primitive module. The name provided to `scheme_primitive_module` is subject to change through the `current-module-declare-name` parameter (which is normally set by the module name resolver when auto-loading module files). After installing variables into the module with `scheme_add_global`, etc., call `scheme_finish_primitive_module` on the `Scheme_Env*` value to make the module declaration available. All defined variables are exported from the primitive module.

The Scheme `#!/variable-reference` form produces a value that is opaque to Scheme code. Use `SCHEME_PTR_VAL` on the result of `#!/variable-reference` to obtain the same kind of value as returned by `scheme_global_bucket` (i.e., a bucket containing the variable’s value, or `NULL` if the variable is not yet defined).

```
void scheme_add_global(char* name,
                     Scheme_Object* val,
                     Scheme_Env* env)
```

Adds a value to the table of globals for the namespace `env`, where `name` is a null-terminated string. (The string’s case will be normalized in the same way as for interning a symbol.)

```
void scheme_add_global_symbol(Scheme_Object* name,
                             Scheme_Object* val,
                             Scheme_Env* env)
```

Adds a value to the table of globals by symbol name instead of string name.

```
Scheme_Object* scheme_lookup_global(Scheme_Object* symbol,  
                                   Scheme_Env* env)
```

Given a global variable name (as a symbol) in *sym*, returns the current value.

```
Scheme_Bucket* scheme_global_bucket(Scheme_Object* symbol,  
                                   Scheme_Env* env)
```

Given a global variable name (as a symbol) in *sym*, returns the bucket where the value is stored. When the value in this bucket is NULL, then the global variable is undefined.

The `Scheme_Bucket` structure is defined as:

```
typedef struct Scheme_Bucket {  
    Scheme_Object so; /* so.type = scheme_variable_type */  
    void *key;  
    void *val;  
} Scheme_Bucket;
```

```
Scheme_Bucket* scheme_module_bucket(Scheme_Object* mod,  
                                   Scheme_Object* symbol,  
                                   int pos,  
                                   Scheme_Env* env)
```

Like `scheme_global_bucket`, but finds a variable in a module. The *mod* and *symbol* arguments are as for [dynamic-require](#) in Scheme. The *pos* argument should be -1 always. The *env* argument represents the namespace in which the module is declared.

```
void scheme_set_global_bucket(char* procname,  
                              Scheme_Bucket* var,  
                              Scheme_Object* val,  
                              int set_undef)
```

Changes the value of a global variable. The *procname* argument is used to report errors (in case the global variable is constant, not yet bound, or bound as syntax). If *set_undef* is not 1, then the global variable must already have a binding. (For example, `set!` cannot set unbound variables, while `define` can.)

```
Scheme_Object* scheme_builtin_value(const char* name)
```

Gets the binding of a name as it would be defined in the initial namespace.

`Scheme_Env* scheme_get_env(Scheme_Config* config)`

Returns the current namespace for the given parameterization (see §9 “Parameterizations”). The current thread’s current parameterization is available as `scheme_config`.

`Scheme_Env* scheme_primitive_module(Scheme_Object* name,
Scheme_Env* for_env)`

Prepares a new primitive module whose name is the symbol *name* (or an alternative that is active via [current-module-declare-name](#)). The module will be declared within the namespace *for_env*. The result is a `Scheme_Env *` value that can be used with `scheme_add_global`, etc., but it represents a module instead of a namespace. The module is not fully declared until `scheme_finish_primitive_module` is called, at which point all variables defined in the module become exported.

`void scheme_finish_primitive_module(Scheme_Env* env)`

Finalizes a primitive module and makes it available for use within the module’s namespace.

5 Procedures

A *primitive procedure* is a Scheme-callable procedure that is implemented in C. Primitive procedures are created in Scheme with the function `scheme_make_prim_w_arity`, which takes a C function pointer, the name of the primitive, and information about the number of Scheme arguments that it takes; it returns a Scheme procedure value.

The C function implementing the procedure must take two arguments: an integer that specifies the number of arguments passed to the procedure, and an array of `Scheme_Object*` arguments. The number of arguments passed to the function will be checked using the arity information. (The arity information provided to `scheme_make_prim_w_arity` is also used for the Scheme `arity` procedure.) The procedure implementation is not allowed to mutate the input array of arguments, although it may mutate the arguments themselves when appropriate (e.g., a fill in a vector argument).

The function `scheme_make_prim_closure_w_arity` is similar to `scheme_make_prim_w_arity`, but it takes an additional count and `Scheme_Object*` array that is copied into the created procedure; the procedure is passed back to the C function when the closure is invoked. In this way, closure-like data from the C world can be associated with the primitive procedure.

The function `scheme_make_closed_prim_w_arity` is similar to `scheme_make_prim_closure_w_arity`, but it uses an older calling convention for passing closure data.

To work well with Scheme threads, a C function that performs substantial or unbounded work should occasionally call `SCHEME_USE_FUEL`; see §8.2 “Allowing Thread Switches” for details.

```
Scheme_Object* scheme_make_prim_w_arity(Scheme_Prim* prim,
                                       char* name,
                                       int mina,
                                       int maxa)
```

Creates a primitive procedure value, given the C function pointer *prim*. The form of *prim* is defined by:

```
typedef Scheme_Object *(Scheme_Prim)(int argc,
                                     Scheme_Object **argv);
```

The value *mina* should be the minimum number of arguments that must be supplied to the procedure. The value *maxa* should be the maximum number of arguments that can be supplied to the procedure, or -1 if the procedure can take arbitrarily many arguments. The *mina* and *maxa* values are used for automatically checking the argument count before the primitive is invoked, and also for the Scheme `arity` procedure. The *name* argument is used to

report application arity errors at run-time.

```
Scheme_Object* scheme_make_folding_prim(Scheme_Prim* prim,
                                         char* name,
                                         int mina,
                                         int maxa,
                                         short folding)
```

Like `scheme_make_prim_w_arity`, but if `folding` is non-zero, the compiler assumes that an application of the procedure to constant values can be folded to a constant. For example, `+`, `zero?`, and `string-length` are folding primitives, but `display` and `cons` are not.

```
Scheme_Object* scheme_make_prim(Scheme_Prim* prim)
```

Same as `scheme_make_prim_w_arity`, but the arity (0, -1) and the name “UNKNOWN” is assumed. This function is provided for backward compatibility only.

```
Scheme_Object*
scheme_make_prim_closure_w_arity(Scheme_Prim_Closure_Proc* prim,
                                 int c,
                                 Scheme_Object* vals,
                                 char* name,
                                 int mina,
                                 int maxa)
```

Creates a primitive procedure value that includes the `c` values in `vals`; when the C function `prim` is invoked, the generated primitive is passed as the last parameter. The form of `prim` is defined by:

```
typedef
Scheme_Object *(Scheme_Prim_Closure_Proc)(int argc,
                                           Scheme_Object **argv,
                                           Scheme_Object *prim);
```

The macro `SCHEME_PRIM_CLOSURE_ELS` takes a primitive-closure object and returns an array with the same length and content as `vals`. (3m: see §3.1 “Cooperating with 3m” for a caution about `SCHEME_PRIM_CLOSURE_ELS`.)

```
Scheme_Object*
scheme_make_closed_prim_w_arity(Scheme_Closed_Prim* prim,
                                void* data,
                                char* name,
                                int mina,
                                int maxa)
```

Creates an old-style primitive procedure value; when the C function *prim* is invoked, *data* is passed as the first parameter. The form of *prim* is defined by:

```
typedef  
Scheme_Object *(Scheme_Closed_Prim)(void *data, int argc,  
                                     Scheme_Object **argv);
```

```
Scheme_Object* scheme_make_closed_prim(Scheme_Closed_Prim* prim,  
                                       void* data)
```

Creates a closed primitive procedure value without arity information. This function is provided for backward compatibility only.

6 Evaluation

A Scheme S-expression is evaluated by calling `scheme_eval`. This function takes an S-expression (as a `Scheme_Object*`) and a namespace and returns the value of the expression in that namespace.

The function `scheme_apply` takes a `Scheme_Object*` that is a procedure, the number of arguments to pass to the procedure, and an array of `Scheme_Object *` arguments. The return value is the result of the application. There is also a function `scheme_apply_to_list`, which takes a procedure and a list (constructed with `scheme_make_pair`) and performs the Scheme `apply` operation.

The `scheme_eval` function actually calls `scheme_compile` followed by `scheme_eval_compiled`.

6.1 Top-level Evaluation Functions

The functions `scheme_eval`, `scheme_apply`, etc., are *top-level evaluation functions*. Continuation invocations are confined to jumps within a top-level evaluation.

The functions `_scheme_eval_compiled`, `_scheme_apply`, etc. (with a leading underscore) provide the same functionality without starting a new top-level evaluation; these functions should only be used within new primitive procedures. Since these functions allow full continuation hops, calls to non-top-level evaluation functions can return zero or multiple times.

Currently, escape continuations and primitive error escapes can jump out of all evaluation and application functions. For more information, see §7 “Exceptions and Escape Continuations”.

6.2 Tail Evaluation

All of Scheme’s built-in functions and syntax support proper tail-recursion. When a new primitive procedure or syntax is added to Scheme, special care must be taken to ensure that tail recursion is handled properly. Specifically, when the final return value of a function is the result of an application, then `scheme_tail_apply` should be used instead of `scheme_apply`. When `scheme_tail_apply` is called, it postpones the procedure application until control returns to the Scheme evaluation loop.

For example, consider the following implementation of a `thunk-or` primitive, which takes any number of thunks and performs `or` on the results of the thunks, evaluating only as many thunks as necessary.

```

static Scheme_Object *
thunk_or (int argc, Scheme_Object **argv)
{
    int i;
    Scheme_Object *v;

    if (!argc)
        return scheme_false;

    for (i = 0; i < argc - 1; i++)
        if (SCHEME_FALSEP((v = _scheme_apply(argv[i], 0, NULL))))
            return v;

    return scheme_tail_apply(argv[argc - 1], 0, NULL);
}

```

This `thunk-or` properly implements tail-recursion: if the final thunk is applied, then the result of `thunk-or` is the result of that application, so `scheme_tail_apply` is used for the final application.

6.3 Multiple Values

A primitive procedure can return multiple values by returning the result of calling `scheme_values`. The functions `scheme_eval_compiled_multi`, `scheme_apply_multi`, `_scheme_eval_compiled_multi`, and `_scheme_apply_multi` potentially return multiple values; all other evaluation and applications procedures return a single value or raise an exception.

Multiple return values are represented by the `scheme_multiple_values` “value.” This quasi-value has the type `Scheme_Object*`, but it is not a pointer or a fixnum. When the result of an evaluation or application is `scheme_multiple_values`, the number of actual values can be obtained as `scheme_multiple_count`, and the array of `Scheme_Object*` values as `scheme_multiple_array`. (Both of those identifiers are actually macros.)

A garbage collection must not occur between the return of a `scheme_multiple_values` “value” and the receipt of the values through `scheme_multiple_count` `scheme_multiple_array`. Furthermore, if `scheme_multiple_array` is to be used across a potential garbage collection, then it must be specifically received by calling `scheme_detach_multiple_array`; otherwise, a garbage collection or further evaluation may change the content of the array. Otherwise, if any application or evaluation procedure is called, the `scheme_multiple_count` and `scheme_multiple_array` variables may be modified (but the array previously referenced by `scheme_multiple_array` is never re-used if `scheme_detach_multiple_array` is called).

The `scheme_multiple_count` and `scheme_multiple_array` variables only contain

meaningful values when `scheme_multiple_values` is returned.

6.4 Evaluation Functions

```
Scheme_Object* scheme_eval(Scheme_Object* expr,  
                           Scheme_Env* env)
```

Evaluates the (uncompiled) S-expression *expr* in the namespace *env*.

```
Scheme_Object* scheme_eval_compiled(Scheme_Object* obj,  
                                    Scheme_Env* env)
```

Evaluates the compiled expression *obj*, which was previously returned from `scheme_compile`, first linking to the namespace *env*.

```
Scheme_Object* scheme_eval_compiled_multi(Scheme_Object* obj,  
                                          Scheme_Env* env)
```

Evaluates the compiled expression *obj*, possibly returning multiple values (see §6.3 “Multiple Values”).

```
Scheme_Object* _scheme_eval_compiled(Scheme_Object* obj,  
                                     Scheme_Env* env)
```

Non-top-level version of `scheme_eval_compiled`. (See §6.1 “Top-level Evaluation Functions”.)

```
Scheme_Object* _scheme_eval_compiled_multi(Scheme_Object* obj,  
                                           Scheme_Env* env)
```

Non-top-level version of `scheme_eval_compiled_multi`. (See §6.1 “Top-level Evaluation Functions”.)

```
Scheme_Env* scheme_basic_env()
```

Creates the main namespace for an embedded PLT Scheme. This procedure must be called before other Scheme library function (except `scheme_make_param`). Extensions to Scheme cannot call this function.

If it is called more than once, this function resets all threads (replacing the main thread), parameters, ports, namespaces, and finalizations.

```
Scheme_Object* scheme_make_namespace(int argc,  
                                     Scheme_Object** argv)
```

Creates and returns a new namespace. This values can be cast to `Scheme_Env *`. It can also be installed in a parameterization using `scheme_set_param` with `MZCONFIG_ENV`.

When PLT Scheme is embedded in an application, create the initial namespace with `scheme_basic_env` before calling this procedure to create new namespaces.

```
Scheme_Object* scheme_apply(Scheme_Object* f,  
                             int c,  
                             Scheme_Object** args)
```

Applies the procedure *f* to the given arguments.

```
Scheme_Object* scheme_apply_multi(Scheme_Object* f,  
                                   int c,  
                                   Scheme_Object** args)
```

Applies the procedure *f* to the given arguments, possibly returning multiple values (see §6.3 “Multiple Values”).

```
Scheme_Object* _scheme_apply(Scheme_Object* f,  
                              int c,  
                              Scheme_Object** args)
```

Non-top-level version of `scheme_apply`. (See §6.1 “Top-level Evaluation Functions”.)

```
Scheme_Object* _scheme_apply_multi(Scheme_Object* f,  
                                    int c,  
                                    Scheme_Object** args)
```

Non-top-level version of `scheme_apply_multi`. (See §6.1 “Top-level Evaluation Functions”.)

```
Scheme_Object* scheme_apply_to_list(Scheme_Object* f,  
                                     Scheme_Object* args)
```

Applies the procedure *f* to the list of arguments in *args*.

```
Scheme_Object* scheme_eval_string(char* str,  
                                 Scheme_Env* env)
```

Reads a single S-expression from *str* and evaluates it in the given namespace; the expression must return a single value, otherwise an exception is raised. The *str* argument is parsed as a UTF-8-encoded string of Unicode characters (so plain ASCII is fine).

```
Scheme_Object* scheme_eval_string_multi(char* str,  
                                       Scheme_Env* env)
```

Like `scheme_eval_string`, but returns `scheme_multiple_values` when the expression returns multiple values.

```
Scheme_Object* scheme_eval_string_all(char* str,  
                                     Scheme_Env* env,  
                                     int all)
```

Like `scheme_eval_string`, but if *all* is not 0, then expressions are read and evaluated from *str* until the end of the string is reached.

```
Scheme_Object* scheme_tail_apply(Scheme_Object* f,  
                                 int n,  
                                 Scheme_Object** args)
```

Applies the procedure as a tail-call. Actually, this function just registers the given application to be invoked when control returns to the evaluation loop. (Hence, this function is only useful within a primitive procedure that is returning to its caller.)

```
Scheme_Object* scheme_tail_apply_no_copy(Scheme_Object* f,  
                                         int n,  
                                         Scheme_Object** args)
```

Like `scheme_tail_apply`, but the array *args* is not copied. Use this only when *args* has infinite extent and will not be used again, or when *args* will certainly not be used again until the called procedure has returned.

```
Scheme_Object* scheme_tail_apply_to_list(Scheme_Object* f,  
                                         Scheme_Object* l)
```

Applies the procedure as a tail-call.

```
Scheme_Object* scheme_compile(Scheme_Object* form,  
                               Scheme_Env* env,  
                               int writable)
```

Compiles the S-expression *form* in the given namespace. The returned value can be used with `scheme_eval_compiled` et al. Provide a non-zero value for *writable* if the resulting compiled object will be marshalled via `write` instead of evaluated.

```
Scheme_Object* scheme_expand(Scheme_Object* form,  
                              Scheme_Env* env)
```

Expands all macros in the S-expression *form* using the given namespace.

```
Scheme_Object* scheme_values(int n,  
                              Scheme_Object** args)
```

Returns the given values together as multiple return values. Unless *n* is 1, the result will always be `scheme_multiple_values`.

```
void scheme_detach_multiple_array(Scheme_Object** args)
```

Called to receive multiple-value results; see §6.3 “Multiple Values”.

7 Exceptions and Escape Continuations

When Scheme encounters an error, it raises an exception. The default exception handler invokes the error display handler and then the error escape handler. The default error escape handler escapes via a *primitive error escape*, which is implemented by calling `scheme_longjmp(*scheme_current_thread->error_buf)`.

An embedding program should install a fresh buffer into `scheme_current_thread->error_buf` and call `scheme_setjmp(*scheme_current_thread->error_buf)` before any top-level entry into Scheme evaluation to catch primitive error escapes. When the new buffer goes out of scope, restore the original in `scheme_current_thread->error_buf`. The macro `scheme_error_buf` is a shorthand for `*scheme_current_thread->error_buf`.

```
mz_jump_buf * volatile save, fresh;
...
save = scheme_current_thread->error_buf;
scheme_current_thread->error_buf = &fresh;
if (scheme_setjmp(scheme_error_buf)) {
    /* There was an error */
    ...
} else {
    v = scheme_eval_string(s, env);
}
scheme_current_thread->error_buf = save;
...
```

3m: when `scheme_setjmp` is used, the enclosing context must provide a local-variable registration record via `MZ_GC_DECL_REG`. Use `MZ_GC_DECL_REG(0)` if the context has no local variables to register. Unfortunately, when using `-xform` with `mzc` instead of `MZ_GC_DECL_REG`, etc., you may need to declare a dummy pointer and use it after `scheme_setjmp` to ensure that a local-variable registration is generated.

New primitive procedures can raise a generic exception by calling `scheme_signal_error`. The arguments for `scheme_signal_error` are roughly the same as for the standard C function `printf`. A specific primitive exception can be raised by calling `scheme_raise_exn`.

Full continuations are implemented in Scheme by copying the C stack and using `scheme_setjmp` and `scheme_longjmp`. As long a C/C++ application invokes Scheme evaluation through the top-level evaluation functions (`scheme_eval`, `scheme_apply`, etc., as opposed to `_scheme_apply`, `_scheme_eval_compiled`, etc.), the code is protected against any unusual behavior from Scheme evaluations (such as returning twice from a function) because continuation invocations are confined to jumps within a single top-level evaluation. However, escape continuation jumps are still allowed; as explained in the following subsection, special care must be taken in extension that is sensitive to escapes.

7.1 Temporarily Catching Error Escapes

When implementing new primitive procedure, it is sometimes useful to catch and handle errors that occur in evaluating subexpressions. One way to do this is the following: save `scheme_current_thread->error_buf` to a temporary variable, set `scheme_current_thread->error_buf` to the address of a stack-allocated `mz_jump_buf`, invoke `scheme_setjmp(scheme_error_buf)`, perform the function's work, and then restore `scheme_current_thread->error_buf` before returning a value. (3m: A stack-allocated `mz_jump_buf` instance need not be registered with the garbage collector, and a heap-allocated `mz_jump_buf` should be allocated as atomic.)

However, beware that a prompt abort or the invocation of an escaping continuation looks like a primitive error escape. In that case, the special indicator flag `scheme_jumping_to_continuation` is non-zero (instead of its normal zero value); this situation is only visible when implementing a new primitive procedure. When `scheme_jumping_to_continuation` is non-zero, honor the escape request by chaining to the previously saved error buffer; otherwise, call `scheme_clear_escape`.

```
mz_jump_buf * volatile save, fresh;
save = scheme_current_thread->error_buf;
scheme_current_thread->error_buf = &fresh;
if (scheme_setjmp(scheme_error_buf)) {
    /* There was an error or continuation invocation */
    if (scheme_jumping_to_continuation) {
        /* It was a continuation jump */
        scheme_longjmp(*save, 1);
        /* To block the jump, instead: scheme_clear_escape(); */
    } else {
        /* It was a primitive error escape */
    }
} else {
    scheme_eval_string("x", scheme_env);
}
scheme_current_thread->error_buf = save;
```

This solution works fine as long as the procedure implementation only calls top-level evaluation functions (`scheme_eval`, `scheme_apply`, etc., as opposed to `_scheme_apply`, `_scheme_eval_compiled`, etc.). Otherwise, use `scheme_dynamic_wind` to protect your code against full continuation jumps in the same way that `dynamic-wind` is used in Scheme.

The above solution simply traps the escape; it doesn't report the reason that the escape occurred. To catch exceptions and obtain information about the exception, the simplest route is to mix Scheme code with C-implemented thunks. The code below can be used to catch exceptions in a variety of situations. It implements the function `_apply_catch_exceptions`, which catches exceptions during the application of a thunk. (This code is in "collects/mzscheme/examples/catch.c" in the distribution.)

```

static Scheme_Object *exn_catching_apply, *exn_p, *exn_message;

static void init_exn_catching_apply()
{
    if (!exn_catching_apply) {
        char *e =
            "(lambda (thunk) "
            "  (with-handlers ([void (lambda (exn) (cons #f exn))]) "
            "    (cons #t (thunk))))";
        /* make sure we have a namespace with the standard bindings: */
        Scheme_Env *env = (Scheme_Env *)scheme_make_namespace(0, NULL);

        scheme_register_extension_global(&exn_catching_apply,
                                         sizeof(Scheme_Object *));
        scheme_register_extension_global(&exn_p,
                                         sizeof(Scheme_Object *));
        scheme_register_extension_global(&exn_message,
                                         sizeof(Scheme_Object *));

        exn_catching_apply = scheme_eval_string(e, env);
        exn_p = scheme_lookup_global(scheme_intern_symbol("exn?"),
env);
        exn_message
            = scheme_lookup_global(scheme_intern_symbol("exn-message"),
env);
    }
}

/* This function applies a thunk, returning the Scheme value if
   there's no exception, otherwise returning NULL and setting *exn
   to the raised value (usually an exn structure). */
Scheme_Object *_apply_thunk_catch_exceptions(Scheme_Object *f,
                                             Scheme_Object **exn)
{
    Scheme_Object *v;

    init_exn_catching_apply();

    v = _scheme_apply(exn_catching_apply, 1, &f);
    /* v is a pair: (cons #t value) or (cons #f exn) */

    if (SCHEME_TRUEP(SCHEME_CAR(v)))
        return SCHEME_CDR(v);
    else {
        *exn = SCHEME_CDR(v);
        return NULL;
    }
}

```

```

    }
}

Scheme_Object *extract_exn_message(Scheme_Object *v)
{
    init_exn_catching_apply();

    if (SCHEME_TRUEP(_scheme_apply(exn_p, 1, &v)))
        return _scheme_apply(exn_message, 1, &v);
    else
        return NULL; /* Not an exn structure */
}

```

In the following example, the above code is used to catch exceptions that occur during while evaluating source code from a string.

```

    static Scheme_Object *do_eval(void *s, int noargc, Scheme_Object
**noargv)
    {
        return scheme_eval_string((char *)s, scheme_get_env(scheme_config));
    }

    static Scheme_Object *eval_string_or_get_exn_message(char *s)
    {
        Scheme_Object *v, *exn;

        v = scheme_make_closed_prim(do_eval, s);
        v = _apply_thunk_catch_exceptions(v, &exn);
        /* Got a value? */
        if (v)
            return v;

        v = extract_exn_message(exn);
        /* Got an exn? */
        if (v)
            return v;

        /* 'raise' was called on some arbitrary value */
        return exn;
    }
}

```

7.2 Enabling and Disabling Breaks

When embedding PLT Scheme, asynchronous break exceptions are disabled by default. Call `scheme_set_can_break` (which is the same as calling the Scheme function `break-enabled`) to enable or disable breaks. To enable or disable breaks during the dynamic extent of another evaluation (where you would use `with-break-parameterization` in Scheme), use `scheme_push_break_enable` before and `scheme_pop_break_enable` after, instead.

7.3 Exception Functions

```
void scheme_signal_error(char* msg,  
                        ... ...)
```

Raises a generic primitive exception. The parameters are roughly as for `printf`, but with the following format directives:

- `%c` : a Unicode character (of type `mzchar`)
- `%d` : an integer
- `%ld` : a long integer
- `%f` : a floating-point double
- `%s` : a nul-terminated char string
- `%5` : a nul-terminated `mzchar` string
- `%S` : a Scheme symbol (a `Scheme_Object*`)
- `%t` : a char string with a long size (two arguments), possibly containing a non-terminating nul byte, and possibly without a nul-terminator
- `%u` : a `mzchar` string with a long size (two arguments), possibly containing a non-terminating nul character, and possibly without a nul-terminator
- `%T` : a Scheme string (a `Scheme_Object*`)
- `%q` : a string, truncated to 253 characters, with ellipses printed if the string is truncated
- `%Q` : a Scheme string (a `Scheme_Object*`), truncated to 253 characters, with ellipses printed if the string is truncated
- `%V` : a Scheme value (a `Scheme_Object*`), truncated according to the current error print width.
- `%e` : an `errno` value, to be printed as a text message.

- `%E` : a platform-specific error value, to be printed as a text message.
- `%Z` : a potential platform-specific error value and a char string; if the string is non-NULL, then the error value is ignored, otherwise the error value is used as for `%E`.
- `%%` : a percent sign

The arguments following the format string must include no more than 25 strings and Scheme values, 25 integers, and 25 floating-point numbers. (This restriction simplifies the implementation with precise garbage collection.)

```
void scheme_raise_exn(int exnid,
                    ... ..)
```

Raises a specific primitive exception. The *exnid* argument specifies the exception to be raised. If an instance of that exception has *n* fields, then the next *n-2* arguments are values for those fields (skipping the `message` and `debug-info` fields). The remaining arguments start with an error string and proceed roughly as for `printf`; see `scheme_signal_error` above for more details.

Exception ids are `#defined` using the same names as in Scheme, but prefixed with “MZ”, all letters are capitalized, and all “:”s, “-”s, and “/”s are replaced with underscores. For example, `MZEXN_FAIL_FILESYSTEM` is the exception id for a filesystem exception.

```
void scheme_warning(char* msg,
                  ... ..)
```

Signals a warning. The parameters are roughly as for `printf`; see `scheme_signal_error` above for more details.

```
void scheme_wrong_count(char* name,
                       int minc,
                       int maxc,
                       int argc,
                       Scheme_Object** argv)
```

This function is automatically invoked when the wrong number of arguments are given to a primitive procedure. It signals that the wrong number of parameters was received and escapes (like `scheme_signal_error`). The *name* argument is the name of the procedure that was given the wrong number of arguments; *minc* is the minimum number of expected arguments; *maxc* is the maximum number of expected arguments, or -1 if there is no maximum; *argv* and *argc* contain all of the received arguments.

```
void scheme_wrong_type(char* name,
                       char* expected,
                       int which,
                       int argc,
                       Scheme_Object** argv)
```

Signals that an argument of the wrong type was received, and escapes (like `scheme_signal_error`). The *name* argument is the name of the procedure that was given the wrong type of argument; *expected* is the name of the expected type; *which* is the offending argument in the *argv* array; *argc* and *argv* contain all of the received arguments. If the original *argc* and *argv* are not available, provide -1 for *which* and a pointer to the bad value in *argv*; *argc* is ignored in this case.

```
void scheme_wrong_return_arity(char* name,
                                int expected,
                                int got,
                                Scheme_Object** argv,
                                const char* detail)
```

Signals that the wrong number of values were returned to a multiple-values context. The *expected* argument indicates how many values were expected, *got* indicates the number received, and *argv* are the received values. The *detail* string can be NULL or it can contain a `printf`-style string (with additional arguments) to describe the context of the error; see `scheme_signal_error` above for more details about the `printf`-style string.

```
void scheme_unbound_global(char* name)
```

Signals an unbound-variable error, where *name* is the name of the variable.

```
char* scheme_make_provided_string(Scheme_Object* o,
                                   int count,
                                   int* len)
```

Converts a Scheme value into a string for the purposes of reporting an error message. The *count* argument specifies how many Scheme values total will appear in the error message (so the string for this value can be scaled appropriately). If *len* is not NULL, it is filled with the length of the returned string.

```
char* scheme_make_args_string(char* s,
                               int which,
                               int argc,
                               Scheme_Object** argv,
                               long* len)
```

Converts an array of Scheme values into a byte string, skipping the array element indicated by *which*. This function is used to specify the “other” arguments to a function when one argument is bad (thus giving the user more information about the state of the program when the error occurred). If *len* is not NULL, it is filled with the length of the returned string.

```
void scheme_check_proc_arity(char* where,
                             int a,
                             int which,
                             int argc,
                             Scheme_Object** argv)
```

Checks the *which*th argument in *argv* to make sure it is a procedure that can take *a* arguments. If there is an error, the *where*, *which*, *argc*, and *argv* arguments are passed on to *scheme_wrong_type*. As in *scheme_wrong_type*, *which* can be -1, in which case **argv* is checked.

```
Scheme_Object* scheme_dynamic_wind(Pre_Post_Proc pre,
                                   Action_Proc action,
                                   Pre_Post_Proc post,
                                   Action_Proc jmp_handler,
                                   void* data)
```

Evaluates calls the function *action* to get a value for the *scheme_dynamic_wind* call. The *Pre_Post_Proc* and *Action_Proc* types are not actually defined; instead the types are inlined as if they were defined as follows:

```
typedef void (*Pre_Post_Proc)(void *data);
typedef Scheme_Object* (*Action_Proc)(void *data);
```

The functions *pre* and *post* are invoked when jumping into and out of *action*, respectively.

The function *jmp_handler* is called when an error is signaled (or an escaping continuation is invoked) during the call to *action*; if *jmp_handler* returns NULL, then the error is passed on to the next error handler, otherwise the return value is used as the return value for the *scheme_dynamic_wind* call.

The pointer *data* can be anything; it is passed along in calls to *action*, *pre*, *post*, and *jmp_handler*.

```
void scheme_clear_escape()
```

Clears the “jumping to escape continuation” flag associated with a thread. Call this function when blocking escape continuation hops (see the first example in §7.1 “Temporarily Catching Error Escapes”).

```
void scheme_set_can_break(int on)
```

Enables or disables breaks in the same way as calling [break-enabled](#).

```
void scheme_push_break_enable(Scheme_Cont_Frame_Data* cframe,  
                             int on,  
                             int pre_check)
```

Use this function with `scheme_pop_break_enable` to enable or disable breaks in the same way as [with-break-parameterization](#); this function writes to `cframe` to initialize it, and `scheme_pop_break_enable` reads from `cframe`. If `pre_check` is non-zero and breaks are currently enabled, any pending break exception is raised.

```
void scheme_pop_break_enable(Scheme_Cont_Frame_Data* cframe,  
                             int post_check)
```

Use this function with `scheme_push_break_enable`. If `post_check` is non-zero and breaks are enabled after restoring the previous state, then any pending break exception is raised.

```
Scheme_Object*  
scheme_current_continuation_marks(Scheme_Object* prompt_tag)
```

Like [current-continuation-marks](#). Passing NULL as `prompt_tag` is the same as providing the default continuation prompt tag.

8 Threads

The initializer function `scheme_basic_env` creates the main Scheme thread; all other threads are created through calls to `scheme_thread`.

Information about each internal Scheme thread is kept in a `Scheme_Thread` structure. A pointer to the current thread's structure is available as `scheme_current_thread`. A `Scheme_Thread` structure includes the following fields:

- `error_buf` — the `mz_jmp_buf` value used to escape from errors. The `error_buf` value of the current thread is available as `scheme_error_buf`.
- `cjs.jumping_to_continuation` — a flag that distinguishes escaping-continuation invocations from error escapes. The `cjs.jumping_to_continuation` value of the current thread is available as `scheme_jumping_to_continuation`.
- `init_config` — the thread's initial parameterization. See also §9 “Parameterizations”.
- `cell_values` — The thread's values for thread cells (see also §9 “Parameterizations”).
- `next` — The next thread in the linked list of threads; this is `NULL` for the main thread.

The list of all scheduled threads is kept in a linked list; `scheme_first_thread` points to the first thread in the list. The last thread in the list is always the main thread.

8.1 Integration with Threads

Scheme's threads can break external C code under two circumstances:

- *Pointers to stack-based values can be communicated between threads.* For example, if thread A stores a pointer to a stack-based variable in a global variable, if thread B uses the pointer in the global variable, it may point to data that is not currently on the stack.
- *C functions that can invoke Scheme (and also be invoked by Scheme) depend on strict function-call nesting.* For example, suppose a function F uses an internal stack, pushing items on to the stack on entry and popping the same items on exit. Suppose also that F invokes Scheme to evaluate an expression. If the evaluation of this expression invokes F again in a new thread, but then returns to the first thread before completing the second F, then F's internal stack will be corrupted.

If either of these circumstances occurs, Scheme will probably crash.

8.2 Allowing Thread Switches

C code that performs substantial or unbounded work should occasionally call `SCHEME_USE_FUEL`—actually a macro—which allows Scheme to swap in another Scheme thread to run, and to check for breaks on the current thread. In particular, if breaks are enabled, then `SCHEME_USE_FUEL` may trigger an exception.

The macro consumes an integer argument. On most platforms, where thread scheduling is based on timer interrupts, the argument is ignored. On some platforms, however, the integer represents the amount of “fuel” that has been consumed since the last call to `SCHEME_USE_FUEL`. For example, the implementation of `vector->list` consumes a unit of fuel for each created cons cell:

```
Scheme_Object *scheme_vector_to_list(Scheme_Object *vec)
{
    int i;
    Scheme_Object *pair = scheme_null;

    i = SCHEME_VEC_SIZE(vec);

    for (; i--;) {
        SCHEME_USE_FUEL(1);
        pair = scheme_make_pair(SCHEME_VEC_ELS(vec)[i], pair);
    }

    return pair;
}
```

The `SCHEME_USE_FUEL` macro expands to a C block, not an expression.

8.3 Blocking the Current Thread

Embedding or extension code sometimes needs to block, but blocking should allow other Scheme threads to execute. To allow other threads to run, block using `scheme_block_until`. This procedure takes two functions: a polling function that tests whether the blocking operation can be completed, and a prepare-to-sleep function that sets bits in `fd_sets` when Scheme decides to sleep (because all Scheme threads are blocked). Under Windows, an “`fd_set`” can also accommodate OS-level semaphores or other handles via `scheme_add_fd_handle`.

Since the functions passed to `scheme_block_until` are called by the Scheme thread scheduler, they must never raise exceptions, call `scheme_apply`, or trigger the evaluation of Scheme code in any way. The `scheme_block_until` function itself may call the current exception handler, however, in reaction to a break (if breaks are enabled).

When a blocking operation is associated with an object, then the object might make sense as an argument to `sync`. To extend the set of objects accepted by `sync`, either register polling and sleeping functions with `scheme_add_evt`, or register a semaphore accessor with `scheme_add_evt_through_sema`.

The `scheme_signal_received` function can be called to wake up Scheme when it is sleeping. In particular, calling `scheme_signal_received` ensures that Scheme will poll all blocking synchronizations soon afterward. Furthermore, `scheme_signal_received` can be called from any OS-level thread. Thus, when no adequate prepare-to-sleep function can be implemented for `scheme_block_until` in terms of file descriptors or Windows handles, calling `scheme_signal_received` when the poll result changes will ensure that a poll is issued.

8.4 Threads in Embedded Scheme with Event Loops

When Scheme is embedded in an application with an event-based model (i.e., the execution of Scheme code in the main thread is repeatedly triggered by external events until the application exits) special hooks must be set to ensure that non-main threads execute correctly. For example, during the execution in the main thread, a new thread may be created; the new thread may still be running when the main thread returns to the event loop, and it may be arbitrarily long before the main thread continues from the event loop. Under such circumstances, the embedding program must explicitly allow Scheme to execute the non-main threads; this can be done by periodically calling the function `scheme_check_threads`.

Thread-checking only needs to be performed when non-main threads exist (or when there are active callback triggers). The embedding application can set the global function pointer `scheme_notify_multithread` to a function that takes an integer parameter and returns `void`. This function is called with 1 when thread-checking becomes necessary, and then with 0 when thread checking is no longer necessary. An embedding program can use this information to prevent unnecessary `scheme_check_threads` polling.

The below code illustrates how MrEd formerly set up `scheme_check_threads` polling using the wxWindows `wxTimer` class. (Any regular event-loop-based callback is appropriate.) The `scheme_notify_multithread` pointer is set to `MrEdInstallThreadTimer`. (MrEd no longer work this way, however.)

```
class MrEdThreadTimer : public wxTimer
{
public:
    void Notify(void); /* callback when timer expires */
};

static int threads_go;
static MrEdThreadTimer *theThreadTimer;
#define THREAD_WAIT_TIME 40
```

```

void MrEdThreadTimer::Notify()
{
    if (threads_go)
        Start(THREAD_WAIT_TIME, TRUE);

    scheme_check_threads();
}

static void MrEdInstallThreadTimer(int on)
{
    if (!theThreadTimer)
        theThreadTimer = new MrEdThreadTimer;

    if (on)
        theThreadTimer->Start(THREAD_WAIT_TIME, TRUE);
    else
        theThreadTimer->Stop();

    threads_go = on;
    if (on)
        do_this_time = 1;
}

```

An alternate architecture, which MrEd now uses, is to send the main thread into a loop, which blocks until an event is ready to handle. Scheme automatically takes care of running all threads, and it does so efficiently because the main thread blocks on a file descriptor, as explained in §8.3 “Blocking the Current Thread”.

8.4.1 Callbacks for Blocked Threads

Scheme threads are sometimes blocked on file descriptors, such as an input file or the X event socket. Blocked non-main threads do not block the main thread, and therefore do not affect the event loop, so `scheme_check_threads` is sufficient to implement this case correctly. However, it is wasteful to poll these descriptors with `scheme_check_threads` when nothing else is happening in the application and when a lower-level poll on the file descriptors can be installed. If the global function pointer `scheme_wakeup_on_input` is set, then this case is handled more efficiently by turning off thread checking and issuing a “wakeup” request on the blocking file descriptors through `scheme_wakeup_on_input`.

A `scheme_wakeup_on_input` procedure takes a pointer to an array of three `fd_sets` (use `MZ_FD_SET` instead of `FD_SET`, etc.) and returns `void`. The `scheme_wakeup_on_input` function does not sleep immediately; it just sets up callbacks on the specified file descriptors. When input is ready on any of those file descriptors, the callbacks are removed and

scheme_wake_up is called.

For example, the X Windows version of MrEd formerly set scheme_wakeup_on_input to this MrEdNeedWakeup:

```
static XtInputId *scheme_cb_ids = NULL;
static int num_cbs;

static void MrEdNeedWakeup(void *fds)
{
    int limit, count, i, p;
    fd_set *rd, *wr, *ex;

    rd = (fd_set *)fds;
    wr = ((fd_set *)fds) + 1;
    ex = ((fd_set *)fds) + 2;

    limit = getdtablesize();

    /* See if we need to do any work, really: */
    count = 0;
    for (i = 0; i < limit; i++) {
        if (MZ_FD_ISSET(i, rd))
            count++;
        if (MZ_FD_ISSET(i, wr))
            count++;
        if (MZ_FD_ISSET(i, ex))
            count++;
    }

    if (!count)
        return;

    /* Remove old callbacks: */
    if (scheme_cb_ids)
        for (i = 0; i < num_cbs; i++)
            notify_set_input_func((Notify_client)NULL, (Notify_func)NULL,
                                  scheme_cb_ids[i]);

    num_cbs = count;
    scheme_cb_ids = new int[num_cbs];

    /* Install callbacks */
    p = 0;
    for (i = 0; i < limit; i++) {
        if (MZ_FD_ISSET(i, rd))
```

```

        scheme_cb_ids[p++] = XtAppAddInput(wxAPP_CONTEXT, i,
                                           (XtPointer *)XtInputRead-
Mask,
                                           (XtInputCallbackProc)MrEdWakeUp,
NULL);
    if (MZ_FD_ISSET(i, wr))
        scheme_cb_ids[p++] = XtAppAddInput(wxAPP_CONTEXT, i,
                                           (XtPointer *)XtIn-
putWriteMask,
                                           (XtInputCallbackProc)MrEdWakeUp,
NULL);
    if (MZ_FD_ISSET(i, ex))
        scheme_cb_ids[p++] = XtAppAddInput(wxAPP_CONTEXT, i,
                                           (XtPointer *)XtInputEx-
ceptMask,
                                           (XtInputCallbackProc)MrEdWakeUp,
                                           NULL);
    }
}

/* callback function when input/exception is detected: */
Bool MrEdWakeUp(XtPointer, int *, XtInputId *)
{
    int i;

    if (scheme_cb_ids) {
        /* Remove all callbacks: */
        for (i = 0; i < num_cbs; i++)
            XtRemoveInput(scheme_cb_ids[i]);

        scheme_cb_ids = NULL;

        /* “wake up” */
        scheme_wake_up();
    }

    return FALSE;
}

```

8.5 Sleeping by Embedded Scheme

When all Scheme threads are blocked, Scheme must “sleep” for a certain number of seconds or until external input appears on some file descriptor. Generally, sleeping should block the main event loop of the entire application. However, the way in which sleeping is performed

may depend on the embedding application. The global function pointer `scheme_sleep` can be set by an embedding application to implement a blocking sleep, although Scheme implements this function for you.

A `scheme_sleep` function takes two arguments: a `float` and a `void*`. The latter is really points to an array of three “`fd_set`” records (one for read, one for write, and one for exceptions); these records are described further below. If the `float` argument is non-zero, then the `scheme_sleep` function blocks for the specified number of seconds, at most. The `scheme_sleep` function should block until there is input one of the file descriptors specified in the “`fd_set`,” indefinitely if the `float` argument is zero.

The second argument to `scheme_sleep` is conceptually an array of three `fd_set` records, but always use `scheme_get_fdset` to get anything other than the zeroth element of this array, and manipulate each “`fd_set`” with `MZ_FD_SET`, `MZ_FD_CLR`, etc. instead of `FD_SET`, `FD_CLR`, etc.

The following function `mzsleep` is an appropriate `scheme_sleep` function for most any Unix or Windows application. (This is approximately the built-in sleep used by Scheme.)

```
void mzsleep(float v, void *fds)
{
    if (v) {
        sleep(v);
    } else {
        int limit;
        fd_set *rd, *wr, *ex;

        # ifdef WIN32
            limit = 0;
        # else
            limit = getdtablesize();
        # endif

        rd = (fd_set *)fds;
        wr = (fd_set *)scheme_get_fdset(fds, 1);
        ex = (fd_set *)scheme_get_fdset(fds, 2);

        select(limit, rd, wr, ex, NULL);
    }
}
```

8.6 Thread Functions

`Scheme_Object*` `scheme_thread(Scheme_Object* thunk)`

Creates a new thread, just like [thread](#).

```
Scheme_Object*
scheme_thread_w_details(Scheme_Object* thunk,
                        Scheme_Config* config,
                        Scheme_Thread_Cell_Table* cells,
                        Scheme_Custodian* cust,
                        int suspend_to_kill)
```

Like `scheme_thread`, except that the created thread belongs to `cust` instead of the current custodian, it uses the given `config` for its initial configuration, it uses `cells` for its thread-cell table, and if `suspend_to_kill` is non-zero, then the thread is merely suspended when it would otherwise be killed (through either [kill-thread](#) or [custodian-shutdown-all](#)).

The `config` argument is typically obtained through `scheme_current_config` or `scheme_extend_config`. A `config` is immutable, so different threads can safely use the same value. The `cells` argument should be obtained from `scheme_inherit_cells`; it is mutable, and a particular cell table should be used by only one thread.

```
Scheme_Object* scheme_make_sema(long v)
```

Creates a new semaphore.

```
void scheme_post_sema(Scheme_Object* sema)
```

Posts to `sema`.

```
int scheme_wait_sema(Scheme_Object* sema,
                    int try)
```

Waits on `sema`. If `try` is not 0, the wait can fail and 0 is returned for failure, otherwise 1 is returned.

```
void scheme_thread_block(float sleep_time)
```

Allows the current thread to be swapped out in favor of other threads. If `sleep_time` positive, then the current thread will sleep for at least `sleep_time` seconds.

After calling this function, a program should almost always call `scheme_making_progress` next. The exception is when `scheme_thread_block` is called in a polling loop that performs no work that affects the progress of other threads. In that case, `scheme_making_progress` should be called immediately after exiting the loop.

See also `scheme_block_until`, and see also the `SCHEME_USE_FUEL` macro in §8.2 “Al-

lowing Thread Switches”.

```
void scheme_thread_block_enable_break(float sleep_time,
                                     int break_on)
```

Like `scheme_thread_block`, but breaks are enabled while blocking if `break_on` is true.

```
void scheme_swap_thread(Scheme_Thread* thread)
```

Swaps out the current thread in favor of `thread`.

```
void scheme_break_thread(Scheme_Thread* thread)
```

Sends a break signal to the given thread.

```
int scheme_break_waiting(Scheme_Thread* thread)
```

Returns 1 if a break from `break-thread` or `scheme_break_thread` has occurred in the specified thread but has not yet been handled.

```
int scheme_block_until(Scheme_Ready_Fun f,
                      Scheme_Needs_Wakeup_Fun fdf,
                      Scheme_Object* data,
                      float sleep)
```

The `Scheme_Ready_Fun` and `Scheme_Needs_Wakeup_Fun` types are defined as follows:

```
typedef int (*Scheme_Ready_Fun)(Scheme_Object *data);
typedef void (*Scheme_Needs_Wakeup_Fun)(Scheme_Object *data,
                                       void *fds);
```

Blocks the current thread until `f` with `data` returns a true value. The `f` function is called periodically—at least once per potential swap-in of the blocked thread—and it may be called multiple times even after it returns a true value. If `f` with `data` ever returns a true value, it must continue to return a true value until `scheme_block_until` returns. The argument to `f` is the same `data` as provided to `scheme_block_until`, and `data` is ignored otherwise. (The `data` argument is not actually required to be a `Scheme_Object*` value, because it is only used by `f` and `fdf`.)

If Scheme decides to sleep, then the `fdf` function is called to sets bits in `fds`, conceptually an array of three `fd_sets`: one for reading, one for writing, and one for exceptions. Use `scheme_get_fdset` to get elements of this array, and manipulate an “`fd_set`” with `MZ_FD_SET` instead of `FD_SET`, etc. Under Windows, an “`fd_set`” can also accommodate

OS-level semaphores or other handles via `scheme_add_fd_handle`.

The *fdf* argument can be NULL, which implies that the thread becomes unblocked (i.e., *ready* changes its result to true) only through Scheme actions, and never through external processes (e.g., through a socket or OS-level semaphore)—with the exception that `scheme_signal_received` may be called to indicate an external change.

If *sleep* is a positive number, then `scheme_block_until` polls *f* at least every *sleep* seconds, but `scheme_block_until` does not return until *f* returns a true value. The call to `scheme_block_until` can return before *sleep* seconds if *f* returns a true value.

The return value from `scheme_block_until` is the return value of its most recent call to *f*, which enables *f* to return some information to the `scheme_block_until` caller.

See §8.3 “Blocking the Current Thread” for information about restrictions on the *f* and *fdf* functions.

```
int scheme_block_until_enable_break(Scheme_Ready_Fun f,
                                   Scheme_Needs_Wakeup_Fun fdf,
                                   Scheme_Object* data,
                                   float sleep,
                                   int break_on)
```

Like `scheme_block_until`, but breaks are enabled while blocking if *break_on* is true.

```
int scheme_block_until_unless(Scheme_Ready_Fun f,
                              Scheme_Needs_Wakeup_Fun fdf,
                              Scheme_Object* data,
                              float sleep,
                              Scheme_Object* unless_evt,
                              int break_on)
```

Like `scheme_block_until_enable_break`, but the function returns if *unless_evt* becomes ready, where *unless_evt* is a port progress event implemented by `scheme_progress_evt_via_get`. See `scheme_make_input_port` for more information.

```
void scheme_signal_received()
```

Indicates that an external event may have caused the result of a synchronization poll to have a different result. Unlike most other Scheme functions, this one can be called from any OS-level thread, and it wakes up if the Scheme thread if it is sleeping.

```
void scheme_check_threads()
```

This function is periodically called by the embedding program to give background processes time to execute. See §8.4 “Threads in Embedded Scheme with Event Loops” for more information.

```
void scheme_wake_up()
```

This function is called by the embedding program when there is input on an external file descriptor. See §8.5 “Sleeping by Embedded Scheme” for more information.

```
void* scheme_get_fdset(void* fds,  
                      int pos)
```

Extracts an “fd_set” from an array passed to `scheme_sleep`, a callback for `scheme_block_until`, or an input port callback for `scheme_make_input_port`.

```
void scheme_add_fd_handle(void* h,  
                         void* fds,  
                         int repost)
```

Adds an OS-level semaphore (Windows) or other waitable handle (Windows) to the “fd_set” `fds`. When Scheme performs a “select” to sleep on `fds`, it also waits on the given semaphore or handle. This feature makes it possible for Scheme to sleep until it is awakened by an external process.

Scheme does not attempt to deallocate the given semaphore or handle, and the “select” call using `fds` may be unblocked due to some other file descriptor or handle in `fds`. If `repost` is a true value, then `h` must be an OS-level semaphore, and if the “select” unblocks due to a post on `h`, then `h` is reposted; this allows clients to treat `fds`-installed semaphores uniformly, whether or not a post on the semaphore was consumed by “select”.

The `scheme_add_fd_handle` function is useful for implementing the second procedure passed to `scheme_wait_until`, or for implementing a custom input port.

Under Unix and Mac OS X, this function has no effect.

```
void scheme_add_fd_eventmask(void* fds,  
                            int mask)
```

Adds an OS-level event type (Windows) to the set of types in the “fd_set” `fds`. When Scheme performs a “select” to sleep on `fds`, it also waits on events of them specified type. This feature makes it possible for Scheme to sleep until it is awakened by an external process.

The event mask is only used when some handle is installed with `scheme_add_fd_handle`. This awkward restriction may force you to create a dummy semaphore that is never posted.

Under Unix, and Mac OS X, this function has no effect.

```
void scheme_add_evt(Scheme_Type type,
                   Scheme_Ready_Fun ready,
                   Scheme_Needs_Wakeup_Fun wakeup,
                   Scheme_Wait_Filter_Fun filter,
                   int can_redirect)
```

The argument types are defined as follows:

```
typedef int (*Scheme_Ready_Fun)(Scheme_Object *data);
typedef void (*Scheme_Needs_Wakeup_Fun)(Scheme_Object *data,
                                         void *fds);
typedef int (*Scheme_Wait_Filter_Fun)(Scheme_Object *data);
```

Extends the set of waitable objects for `sync` to those with the type tag `type`. If `filter` is non-NULL, it constrains the new waitable set to those objects for which `filter` returns a non-zero value.

The `ready` and `wakeup` functions are used in the same way was the arguments to `scheme_block_until`.

The `can_redirect` argument should be 0.

```
void scheme_add_evt_through_sema(Scheme_Type type,
                                Scheme_Wait_Sema_Fun getsema,
                                Scheme_Wait_Filter_Fun filter)
```

Like `scheme_add_evt`, but for objects where waiting is based on a semaphore. Instead of `ready` and `wakeup` functions, the `getsema` function extracts a semaphore for a given object:

```
typedef
Scheme_Object *(*Scheme_Wait_Sema_Fun)(Scheme_Object *data,
                                       int *repost);
```

If a successful wait should leave the semaphore waited, then `getsema` should set `*repost` to 0. Otherwise, the given semaphore will be re-posted after a successful wait. A `getsema` function should almost always set `*repost` to 1.

```
void scheme_making_progress()
```

Notifies the scheduler that the current thread is not simply calling `scheme_thread_block` in a loop, but that it is actually making progress.

```
int scheme_tls_allocate()
```

Allocates a thread local storage index to be used with `scheme_tls_set` and `scheme_tls_get`.

```
void scheme_tls_set(int index,  
                   void* v)
```

Stores a thread-specific value using an index allocated with `scheme_tls_allocate`.

```
void* scheme_tls_get(int index)
```

Retrieves a thread-specific value installed with `scheme_tls_set`. If no thread-specific value is available for the given index, NULL is returned.

```
Scheme_Object* scheme_call_enable_break(Scheme_Prim* prim,  
                                       int argc,  
                                       Scheme_Object** argv)
```

Calls *prim* with the given *argc* and *argv* with breaks enabled. The *prim* function can block, in which case it might be interrupted by a break. The *prim* function should not block, yield, or check for breaks after it succeeds, where “succeeds” depends on the operation. For example, [tcp-accept/enable-break](#) is implemented by wrapping this function around the implementation of `tcp-accept`; the `tcp-accept` implementation does not block or yield after it accepts a connection.

```
Scheme_Object*  
scheme_make_thread_cell(Scheme_Object* def_val,  
                       int preserved,  
                       Scheme_Object* cell,  
                       Scheme_Thread_Cell_Table* cells,  
                       Scheme_Object* cell,  
                       Scheme_Thread_Cell_Table* cells,  
                       Scheme_Object* v)
```

Prevents Scheme thread swaps until `scheme_end_atomic` or `scheme_end_atomic_no_swap` is called. Start-atomic and end-atomic pairs can be nested.

```
void scheme_end_atomic()
```

Ends an atomic region with respect to Scheme threads. The current thread may be swapped

out immediately (i.e., the call to `scheme_end_atomic` is assumed to be a safe point for thread swaps).

```
void scheme_end_atomic_no_swap()
```

Ends an atomic region with respect to Scheme threads, and also prevents an immediate thread swap. (In other words, no Scheme thread swaps will occur until a future safe point.)

```
void scheme_add_swap_callback(Scheme_Closure_Func f,  
                             Scheme_Object* data)
```

Registers a callback to be invoked just after a Scheme thread is swapped in. The *data* is provided back to *f* when it is called, where `Closure_Func` is defined as follows:

```
typedef Scheme_Object *(*Scheme_Closure_Func)(Scheme_Object *);
```

```
void scheme_add_swap_out_callback(Scheme_Closure_Func f,  
                                 Scheme_Object* data)
```

Like `scheme_add_swap_callback`, but registers a callback to be invoked just before a Scheme thread is swapped out.

9 Parameterizations

A *parameterization* is a set of parameter values. Each thread has its own initial parameterization, which is extended functionally and superseded by parameterizations that are attached to a particular continuation mark.

Parameterization information is stored in a `Scheme_Config` record. For the currently executing thread, `scheme_current_config` returns the current parameterization.

To obtain parameter values, a `Scheme_Config` is combined with the current threads `Scheme_Thread_Cell_Table`, as stored in the thread record's `cell_values` field.

Parameter values for built-in parameters are obtained and modified (for the current thread) using `scheme_get_param` and `scheme_set_param`. Each parameter is stored as a `Scheme_Object * value`, and the built-in parameters are accessed through the following indices:

- `MZCONFIG_ENV` — `current-namespace` (use `scheme_get_env`)
- `MZCONFIG_INPUT_PORT` — `current-input-port`
- `MZCONFIG_OUTPUT_PORT` — `current-output-port`
- `MZCONFIG_ERROR_PORT` — `current-error-port`
- `MZCONFIG_ERROR_DISPLAY_HANDLER` — `error-display-handler`
- `MZCONFIG_ERROR_PRINT_VALUE_HANDLER` — `error-value->string-handler`
- `MZCONFIG_EXIT_HANDLER` — `exit-handler`
- `MZCONFIG_INIT_EXN_HANDLER` — `uncaught-exception-handler`
- `MZCONFIG_EVAL_HANDLER` — `current-eval`
- `MZCONFIG_LOAD_HANDLER` — `current-load`
- `MZCONFIG_PRINT_HANDLER` — `current-print`
- `MZCONFIG_PROMPT_READ_HANDLER` — `current-prompt-read`
- `MZCONFIG_CAN_READ_GRAPH` — `read-accept-graph`
- `MZCONFIG_CAN_READ_COMPILED` — `read-accept-compiled`
- `MZCONFIG_CAN_READ_BOX` — `read-accept-box`
- `MZCONFIG_CAN_READ_PIPE_QUOTE` — `read-accept-bar-quote`
- `MZCONFIG_PRINT_GRAPH` — `print-graph`

- MZCONFIG_PRINT_STRUCT — [print-struct](#)
- MZCONFIG_PRINT_BOX — [print-box](#)
- MZCONFIG_CASE_SENS — [read-case-sensitive](#)
- MZCONFIG_SQUARE_BRACKETS_ARE_PARENS — [read-square-brackets-as-parens](#)
- MZCONFIG_CURLY_BRACES_ARE_PARENS — [read-curly-braces-as-parens](#)
- MZCONFIG_ERROR_PRINT_WIDTH — [error-print-width](#)
- MZCONFIG_ALLOW_SET_UNDEFINED — [allow-compile-set!-undefined](#)
- MZCONFIG_CUSTODIAN — [current-custodian](#)
- MZCONFIG_USE_COMPILED_KIND — [use-compiled-file-paths](#)
- MZCONFIG_LOAD_DIRECTORY — [current-load-relative-directory](#)
- MZCONFIG_COLLECTION_PATHS — [current-library-collection-paths](#)
- MZCONFIG_PORT_PRINT_HANDLER — [global-port-print-handler](#)
- MZCONFIG_LOAD_EXTENSION_HANDLER — [current-load-extension](#)

To get or set a parameter value for a thread other than the current one, use `scheme_get_thread_param` and `scheme_set_thread_param`, each of which takes a `Scheme_Thread_Cell_Table` to use in resolving or setting a parameter value.

When installing a new parameter with `scheme_set_param`, no check is performed on the supplied value to ensure that it is a legal value for the parameter; this is the responsibility of the caller of `scheme_set_param`. Note that Boolean parameters should only be set to the values `#t` and `#f`.

New primitive parameter indices are created with `scheme_new_param` and implemented with `scheme_make_parameter` and `scheme_param_config`.

```
Scheme_Object* scheme_get_param(Scheme_Config* config,
                               int param_id)
```

Gets the current value (for the current thread) of the parameter specified by `param_id`.

```
Scheme_Object* scheme_set_param(Scheme_Config* config,
                               int param_id,
                               Scheme_Object* v)
```

Sets the current value (for the current thread) of the parameter specified by `param_id`.

```
Scheme_Object*
scheme_get_thread_param(Scheme_Config* config,
                       Scheme_Thread_Cell_Table* cells,
                       int param_id)
```

Like `scheme_get_param`, but using an arbitrary thread's cell-value table.

```
Scheme_Object*
scheme_set_thread_param(Scheme_Config* config,
                       Scheme_Thread_Cell_Table* cells,
                       int param_id,
                       Scheme_Object* v)
```

Like `scheme_set_param`, but using an arbitrary thread's cell-value table.

```
Scheme_Object* scheme_extend_config(Scheme_Config* base,
                                    int param_id,
                                    Scheme_Object* v)
```

Creates and returns a parameterization that extends *base* with a new value *v* (in all threads) for the parameter *param_id*. Use `scheme_install_config` to make this configuration active in the current thread.

```
void scheme_install_config(Scheme_Config* config)
```

Adjusts the current thread's continuation marks to make *config* the current parameterization. Typically, this function is called after `scheme_push_continuation_frame` to establish a new continuation frame, and then `scheme_pop_continuation_frame` is called later to remove the frame (and thus the parameterization).

```
Scheme_Thread_Cell_Table*
scheme_inherit_cells(Scheme_Thread_Cell_Table* cells)
```

Creates a new thread-cell-value table, copying values for preserved thread cells from *cells*.

```
int scheme_new_param()
```

Allocates a new primitive parameter index. This function must be called *before* `scheme_basic_env`, so it is only available to embedding applications (i.e., not extensions).

```
Scheme_Object* scheme_register_parameter(Scheme_Prim* function,  
                                         char* name,  
                                         int exnid)
```

Use this function instead of the other primitive-constructing functions, like `scheme_make_prim`, to create a primitive parameter procedure. See also `scheme_param_config`, below. This function is only available to embedding applications (i.e., not extensions).

```
Scheme_Object* scheme_param_config(char* name,  
                                   Scheme_Object* param,  
                                   int argc,  
                                   Scheme_Object** argv,  
                                   int arity,  
                                   Scheme_Prim* check,  
                                   char* expected,  
                                   int isbool)
```

Call this procedure in a primitive parameter procedure to implement the work of getting or setting the parameter. The *name* argument should be the parameter procedure name; it is used to report errors. The *param* argument is a fixnum corresponding to the primitive parameter index returned by `scheme_new_param`. The *argc* and *argv* arguments should be the un-touched and un-tested arguments that were passed to the primitive parameter. Argument-checking is performed within `scheme_param_config` using *arity*, *check*, *expected*, and *isbool*:

- If *arity* is non-negative, potential parameter values must be able to accept the specified number of arguments. The *check* and *expected* arguments should be NULL.
- If *check* is not NULL, it is called to check a potential parameter value. The arguments passed to *check* are always 1 and an array that contains the potential parameter value. If *isbool* is 0 and *check* returns `scheme_false`, then a type error is reported using *name* and *expected*. If *isbool* is 1, then a type error is reported only when *check* returns NULL and any non-NULL return value is used as the actual value to be stored for the parameter.
- Otherwise, *isbool* should be 1. A potential procedure argument is then treated as a Boolean value.

This function is only available to embedding applications (i.e., not extensions).

10 Continuation Marks

A mark can be attached to the current continuation frame using `scheme_set_cont_mark`. To force the creation of a new frame (e.g., during a nested function call within your function), use `scheme_push_continuation_frame`, and then remove the frame with `scheme_pop_continuation_frame`.

```
void scheme_set_cont_mark(Scheme_Object* key,  
                          Scheme_Object* val)
```

Add/sets a continuation mark in the current continuation.

```
void scheme_push_continuation_frame(Scheme_Cont_Frame_Data* data)
```

Creates a new continuation frame. The *data* record need not be initialized, and it can be allocated on the C stack. Supply *data* to `scheme_pop_continuation_frame` to remove the continuation frame.

```
void scheme_pop_continuation_frame(Scheme_Cont_Frame_Data* data)
```

Removes a continuation frame created by `scheme_pop_continuation_frame`.

11 String Encodings

The `scheme_utf8_decode` function decodes a char array as UTF-8 into either a UCS-4 `mzchar` array or a UTF-16 short array. The `scheme_utf8_encode` function encodes either a UCS-4 `mzchar` array or a UTF-16 short array into a UTF-8 char array.

These functions can be used to check or measure an encoding or decoding without actually producing the result decoding or encoding, and variations of the function provide control over the handling of decoding errors.

```
int scheme_utf8_decode(const unsigned char* s,
                      int start,
                      int end,
                      mzchar* us,
                      int dstart,
                      int dend,
                      long* ipos,
                      char utf16,
                      int permissive)
```

Decodes a byte array as UTF-8 to produce either Unicode code points into `us` (when `utf16` is zero) or UTF-16 code units into `us` cast to `short*` (when `utf16` is non-zero). No nul terminator is added to `us`.

The result is non-negative when all of the given bytes are decoded, and the result is the length of the decoding (in `mzchars` or `shorts`). A -2 result indicates an invalid encoding sequence in the given bytes (possibly because the range to decode ended mid-encoding), and a -3 result indicates that decoding stopped because not enough room was available in the result string.

The `start` and `end` arguments specify a range of `s` to be decoded. If `end` is negative, `strlen(s)` is used as the end.

If `us` is NULL, then decoded bytes are not produced, but the result is valid as if decoded bytes were written. The `dstart` and `dend` arguments specify a target range in `us` (in `mzchar` or `short` units) for the decoding; a negative value for `dend` indicates that any number of bytes can be written to `us`, which is normally sensible only when `us` is NULL for measuring the length of the decoding.

If `ipos` is non-NULL, it is filled with the first undecoded index within `s`. If the function result is non-negative, then `*ipos` is set to the ending index (with is `end` if non-negative, `strlen(s)` otherwise). If the result is -1 or -2, then `*ipos` effectively indicates how many bytes were decoded before decoding stopped.

If `permissive` is non-zero, it is used as the decoding of bytes that are not part of a valid UTF-8

encoding. Thus, the function result can be -2 only if *permissive* is 0.

This function does not allocate or trigger garbage collection.

```
int scheme_utf8_decode_as_prefix(const unsigned char* s,
                                int start,
                                int end,
                                mzchar* us,
                                int dstart,
                                int dend,
                                long* ipos,
                                char utf16,
                                int permissive)
```

Like `scheme_utf8_decode`, but the result is always the number of the decoded `mzchars` or shorts. If a decoding error is encountered, the result is still the size of the decoding up until the error.

```
int scheme_utf8_decode_all(const unsigned char* s,
                           int len,
                           mzchar* us,
                           int permissive)
```

Like `scheme_utf8_decode`, but with fewer arguments. The decoding produces UCS-4 `mzchars`. If the buffer `us` is non-NULL, it is assumed to be long enough to hold the decoding (which cannot be longer than the length of the input, though it may be shorter). If `len` is negative, `strlen(s)` is used as the input length.

```
int scheme_utf8_decode_prefix(const unsigned char* s,
                              int len,
                              mzchar* us,
                              int permissive)
```

Like `scheme_utf8_decode`, but with fewer arguments. The decoding produces UCS-4 `mzchars`. If the buffer `us` **must** be non-NULL, and it is assumed to be long enough to hold the decoding (which cannot be longer than the length of the input, though it may be shorter). If `len` is negative, `strlen(s)` is used as the input length.

In addition to the result of `scheme_utf8_decode`, the result can be -1 to indicate that the input ended with a partial (valid) encoding. A -1 result is possible even when *permissive* is non-zero.

```
mzchar* scheme_utf8_decode_to_buffer(const unsigned char* s,
                                     int len,
                                     mzchar* buf,
                                     int blen)
```

Like `scheme_utf8_decode_all` with *permissive* as 0, but if *buf* is not large enough (as indicated by *blen*) to hold the result, a new buffer is allocated. Unlike other functions, this one adds a nul terminator to the decoding result. The function result is either *buf* (if it was big enough) or a buffer allocated with `scheme_malloc_atomic`.

```
mzchar* scheme_utf8_decode_to_buffer_len(const unsigned char* s,
                                         int len,
                                         mzchar* buf,
                                         int blen,
                                         long* ulen)
```

Like `scheme_utf8_decode_to_buffer`, but the length of the result (not including the terminator) is placed into *ulen* if *ulen* is non-NULL.

```
int scheme_utf8_decode_count(const unsigned char* s,
                             int start,
                             int end,
                             int* state,
                             int might_continue,
                             int permissive)
```

Like `scheme_utf8_decode`, but without producing the decoded `mzchars`, and always returning the number of decoded `mzchars` up until a decoding error (if any). If *might_continue* is non-zero, the a partial valid encoding at the end of the input is not decoded when *permissive* is also non-zero.

If *state* is non-NULL, it holds information about partial encodings; it should be set to zero for an initial call, and then passed back to `scheme_utf8_decode` along with bytes that extend the given input (i.e., without any unused partial encodings). Typically, this mode makes sense only when *might_continue* and *permissive* are non-zero.

```
int scheme_utf8_encode(const mzchar* us,
                      int start,
                      int end,
                      unsigned char* s,
                      int dstart,
                      char utf16)
```

Encodes the given UCS-4 array of `mzchars` (if `utf16` is zero) or UTF-16 array of shorts (if `utf16` is non-zero) into `s`. The `end` argument must be no less than `start`.

The array `s` is assumed to be long enough to contain the encoding, but no encoding is written if `s` is NULL. The `dstart` argument indicates a starting place in `s` to hold the encoding. No nul terminator is added to `s`.

The result is the number of bytes produced for the encoding (or that would be produced if `s` was non-NULL). Encoding never fails.

This function does not allocate or trigger garbage collection.

```
int scheme_utf8_encode_all(const mzchar* us,
                          int len,
                          unsigned char* s)
```

Like `scheme_utf8_encode` with 0 for `start`, `len` for `end`, 0 for `dstart` and 0 for `utf16`.

```
char* scheme_utf8_encode_to_buffer(const mzchar* s,
                                  int len,
                                  char* buf,
                                  int blen)
```

Like `scheme_utf8_encode_all`, but the length of `buf` is given, and if it is not long enough to hold the encoding, a buffer is allocated. A nul terminator is added to the encoded array. The result is either `buf` or an array allocated with `scheme_malloc_atomic`.

```
char* scheme_utf8_encode_to_buffer_len(const mzchar* s,
                                       int len,
                                       char* buf,
                                       int blen,
                                       long* rlen)
```

Like `scheme_utf8_encode_to_buffer`, but the length of the resulting encoding (not including a nul terminator) is reported in `rlen` if it is non-NULL.

```
unsigned-short* scheme_ucs4_to_utf16(const mzchar* text,
                                     int start,
                                     int end,
                                     unsigned short* buf,
                                     int bufsize,
                                     long* ulen,
                                     int term_size)
```

Converts a UCS-4 encoding (the indicated range of *text*) to a UTF-16 encoding. The *end* argument must be no less than *start*.

A result buffer is allocated if *buf* is not long enough (as indicated by *bufsize*). If *ulen* is non-NULL, it is filled with the length of the UTF-16 encoding. The *term_size* argument indicates a number of shorts to reserve at the end of the result buffer for a terminator (but no terminator is actually written).

```
mzchar* scheme_utf16_to_ucs4(const unsigned short* text,
                             int start,
                             int end,
                             mzchar* buf,
                             int bufsize,
                             long* ulen,
                             int term_size)
```

Converts a UTF-16 encoding (the indicated range of *text*) to a UCS-4 encoding. The *end* argument must be no less than *start*.

A result buffer is allocated if *buf* is not long enough (as indicated by *bufsize*). If *ulen* is non-NULL, it is filled with the length of the UCS-4 encoding. The *term_size* argument indicates a number of mzchars to reserve at the end of the result buffer for a terminator (but no terminator is actually written).

12 Bignums, Rationals, and Complex Numbers

Scheme supports integers of an arbitrary magnitude; when an integer cannot be represented as a fixnum (i.e., 30 or 62 bits plus a sign bit), then it is represented by the Scheme type `scheme_bignum_type`. There is no overlap in integer values represented by fixnums and bignums.

Rationals are implemented by the type `scheme_rational_type`, composed of a numerator and a denominator. The numerator and denominator fixnums or bignums (possibly mixed).

Complex numbers are implemented by the types `scheme_complex_type` and `scheme_complex_izi_type`, composed of a real and imaginary part. The real and imaginary parts will either be both flonums, both exact numbers (fixnums, bignums, and rationals can be mixed in any way), or one part will be exact 0 and the other part will be a flonum. If the inexact part is inexact 0, the type is `scheme_complex_izi_type`, otherwise the type is `scheme_complex_type`; this distinction make it easy to test whether a complex number should be treated as a real number.

```
int scheme_is_exact(Scheme_Object* n)
```

Returns 1 if n is an exact number, 0 otherwise (n need not be a number).

```
int scheme_is_inexact(Scheme_Object* n)
```

Returns 1 if n is an inexact number, 0 otherwise (n need not be a number).

```
Scheme_Object* scheme_make_bignum(long v)
```

Creates a bignum representing the integer v . This can create a bignum that otherwise fits into a fixnum. This must only be used to create temporary values for use with the bignum functions. Final results can be normalized with `scheme_bignum_normalize`. Only normalized numbers can be used with procedures that are not specific to bignums.

```
Scheme_Object* scheme_make_bignum_from_unsigned(unsigned long v)
```

Like `scheme_make_bignum`, but works on unsigned integers.

```
double scheme_bignum_to_double(Scheme_Object* n)
```

Converts a bignum to a floating-point number, with reasonable but unspecified accuracy.

```
float scheme_bignum_to_float(Scheme_Object* n)
```

If PLT Scheme is not compiled with single-precision floats, this procedure is actually a macro alias for `scheme_bignum_to_double`.

```
Scheme_Object* scheme_bignum_from_double(double d)
```

Creates a bignum that is close in magnitude to the floating-point number *d*. The conversion accuracy is reasonable but unspecified.

```
Scheme_Object* scheme_bignum_from_float(float f)
```

If PLT Scheme is not compiled with single-precision floats, this procedure is actually a macro alias for `scheme_bignum_from_double`.

```
char* scheme_bignum_to_string(Scheme_Object* n,  
                             int radix)
```

Writes a bignum into a newly allocated byte string.

```
Scheme_Object* scheme_read_bignum(mzchar* str,  
                                  int offset,  
                                  int radix)
```

Reads a bignum from a `mzchar` string, starting from position *offset* in *str*. If the string does not represent an integer, then `NULL` will be returned. If the string represents a number that fits in 31 bits, then a `scheme_integer_type` object will be returned.

```
Scheme_Object* scheme_read_bignum_bytes(char* str,  
                                       int offset,  
                                       int radix)
```

Like `scheme_read_bignum`, but from a UTF-8-encoding byte string.

```
Scheme_Object* scheme_bignum_normalize(Scheme_Object* n)
```

If *n* fits in 31 bits, then a `scheme_integer_type` object will be returned. Otherwise, *n* is returned.

```
Scheme_Object* scheme_make_rational(Scheme_Object* n,  
                                   Scheme_Object* d)
```

Creates a rational from a numerator and denominator. The *n* and *d* parameters must be fixnums or bignums (possibly mixed). The resulting will be normalized (thus, a bignum or

fixnum might be returned).

```
double scheme_rational_to_double(Scheme_Object* n)
```

Converts the rational *n* to a double.

```
float scheme_rational_to_float(Scheme_Object* n)
```

If PLT Scheme is not compiled with single-precision floats, this procedure is actually a macro alias for `scheme_rational_to_double`.

```
Scheme_Object* scheme_rational_numerator(Scheme_Object* n)
```

Returns the numerator of the rational *n*.

```
Scheme_Object* scheme_rational_denominator(Scheme_Object* n)
```

Returns the denominator of the rational *n*.

```
Scheme_Object* scheme_rational_from_double(double d)
```

Converts the given double into a maximally-precise rational.

```
Scheme_Object* scheme_rational_from_float(float d)
```

If PLT Scheme is not compiled with single-precision floats, this procedure is actually a macro alias for `scheme_rational_from_double`.

```
Scheme_Object* scheme_make_complex(Scheme_Object* r,  
                                   Scheme_Object* i)
```

Creates a complex number from real and imaginary parts. The *r* and *i* arguments must be fixnums, bignums, flonums, or rationals (possibly mixed). The resulting number will be normalized (thus, a real number might be returned).

```
Scheme_Object* scheme_complex_real_part(Scheme_Object* n)
```

Returns the real part of the complex number *n*.

```
Scheme_Object* scheme_complex_imaginary_part(Scheme_Object* n)
```

Returns the imaginary part of the complex number n .

13 Ports and the Filesystem

Ports are represented as Scheme values with the types `scheme_input_port_type` and `scheme_output_port_type`. The function `scheme_read` takes an input port value and returns the next S-expression from the port. The function `scheme_write` takes an output port and a value and writes the value to the port. Other standard low-level port functions are also provided, such as `scheme_getc`.

File ports are created with `scheme_make_file_input_port` and `scheme_make_file_output_port`; these functions take a `FILE *` file pointer and return a Scheme port. Strings are read or written with `scheme_make_byte_string_input_port`, which takes a nul-terminated byte string, and `scheme_make_byte_string_output_port`, which takes no arguments. The contents of a string output port are obtained with `scheme_get_byte_string_output`.

Custom ports, with arbitrary read/write handlers, are created with `scheme_make_input_port` and `scheme_make_output_port`.

When opening a file for any reason using a name provided from Scheme, use `scheme_expand_filename` to normalize the filename and resolve relative paths.

```
Scheme_Object* scheme_read(Scheme_Object* port)
```

`reads` the next S-expression from the given input port.

```
void scheme_write(Scheme_Object* obj,  
                 Scheme_Object* port)
```

`writes` the Scheme value `obj` to the given output port.

```
void scheme_write_w_max(Scheme_Object* obj,  
                       Scheme_Object* port,  
                       int n)
```

Like `scheme_write`, but the printing is truncated to `n` bytes. (If printing is truncated, the last bytes are printed as “.”.)

```
void scheme_display(Scheme_Object* obj,  
                  Scheme_Object* port)
```

`displays` the Scheme value `obj` to the given output port.

```
void scheme_display_w_max(Scheme_Object* obj,
                          Scheme_Object* port,
                          int n)
```

Like `scheme_display`, but the printing is truncated to n bytes. (If printing is truncated, the last three bytes are printed as “?”.)

```
void scheme_write_byte_string(char* str,
                              long len,
                              Scheme_Object* port)
```

Writes len bytes of str to the given output port.

```
void scheme_write_char_string(mzchar* str,
                              long len,
                              Scheme_Object* port)
```

Writes len characters of str to the given output port.

```
long scheme_put_byte_string(const char* who,
                            Scheme_Object* port,
                            char* str,
                            long d,
                            long len,
                            int rarely_block)
```

Writes len bytes of str , starting with the d th character. Bytes are written to the given output port, and errors are reported as from who .

If $rarely_block$ is 0, the write blocks until all len bytes are written, possibly to an internal buffer. If $rarely_block$ is 2, the write never blocks, and written bytes are not buffered. If $rarely_block$ is 1, the write blocks only until at least one byte is written (without buffering) or until part of an internal buffer is flushed.

Supplying 0 for len corresponds to a buffer-flush request. If $rarely_block$ is 2, the flush request is non-blocking, and if $rarely_block$ is 0, it is blocking. (A $rarely_block$ of 1 is the same as 0 in this case.)

The result is -1 if no bytes are written from str and unflushed bytes remain in the internal buffer. Otherwise, the return value is the number of written characters.

```
long scheme_put_char_string(const char* who,
                           Scheme_Object* port,
                           char* str,
                           long d,
                           long len)
```

Like `scheme_put_byte_string`, but for a `mzchar` string, and without the non-blocking option.

```
char* scheme_write_to_string(Scheme_Object* obj,
                             long* len)
```

Prints the Scheme value `obj` using `write` to a newly allocated string. If `len` is not NULL, `*len` is set to the length of the bytes string.

```
void scheme_write_to_string_w_max(Scheme_Object* obj,
                                  long* len,
                                  int n)
```

Like `scheme_write_to_string`, but the string is truncated to `n` bytes. (If the string is truncated, the last three bytes are “.”.)

```
char* scheme_display_to_string(Scheme_Object* obj,
                               long* len)
```

Prints the Scheme value `obj` using `display` to a newly allocated string. If `len` is not NULL, `*len` is set to the length of the string.

```
void scheme_display_to_string_w_max(Scheme_Object* obj,
                                    long* len,
                                    int n)
```

Like `scheme_display_to_string`, but the string is truncated to `n` bytes. (If the string is truncated, the last three bytes are “.”.)

```
void scheme_debug_print(Scheme_Object* obj)
```

Prints the Scheme value `obj` using `write` to the main thread’s output port.

```
void scheme_flush_output(Scheme_Object* port)
```

If `port` is a file port, a buffered data is written to the file. Otherwise, there is no effect. `port`

must be an output port.

```
int scheme_get_byte(Scheme_Object* port)
```

Get the next byte from the given input port. The result can be EOF.

```
int scheme_getc(Scheme_Object* port)
```

Get the next character from the given input port (by decoding bytes as UTF-8). The result can be EOF.

```
int scheme_peek_byte(Scheme_Object* port)
```

Peeks the next byte from the given input port. The result can be EOF.

```
int scheme_peekc(Scheme_Object* port)
```

Peeks the next character from the given input port (by decoding bytes as UTF-8). The result can be EOF.

```
int scheme_peek_byte_skip(Scheme_Object* port,  
                          Scheme_Object* skip)
```

Like `scheme_peek_byte`, but with a skip count. The result can be EOF.

```
int scheme_peekc_skip(Scheme_Object* port,  
                     Scheme_Object* skip)
```

Like `scheme_peekc`, but with a skip count. The result can be EOF.

```
long scheme_get_byte_string(const char* who,  
                           Scheme_Object* port,  
                           char* buffer,  
                           int offset,  
                           long size,  
                           int only_avail,  
                           int peek,  
                           Scheme_Object* peek_skip)
```

Gets multiple bytes at once from a port, reporting errors with the name *who*. The *size* argument indicates the number of requested bytes, to be put into the *buffer* array starting at *offset*. The return value is the number of bytes actually read, or EOF if an end-of-file is

encountered without reading any bytes.

If *only_avail* is 0, then the function blocks until *size* bytes are read or an end-of-file is reached. If *only_avail* is 1, the function blocks only until at least one byte is read. If *only_avail* is 2, the function never blocks. If *only_avail* is -1, the function blocks only until at least one byte is read but also allows breaks (with the guarantee that bytes are read or a break is raised, but not both).

If *peek* is non-zero, then the port is peeked instead of read. The *peek_skip* argument indicates a portion of the input stream to skip as a non-negative, exact integer (fixnum or bignum). In this case, an *only_avail* value of 1 means to continue the skip until at least one byte can be returned, even if it means multiple blocking reads to skip bytes.

If *peek* is zero, then *peek_skip* should be either NULL (which means zero) or the fixnum zero.

```
long scheme_get_char_string(const char* who,
                           Scheme_Object* port,
                           char* buffer,
                           int offset,
                           long size,
                           int peek,
                           Scheme_Object* peek_skip)
```

Like `scheme_get_byte_string`, but for characters (by decoding bytes as UTF-8), and without the non-blocking option.

```
long scheme_get_bytes(Scheme_Object* port,
                     long size,
                     char* buffer,
                     int offset)
```

For backward compatibility: calls `scheme_get_byte_string` in essentially the obvious way with *only_avail* as 0; if *size* is negative, then it reads *-size* bytes with *only_avail* as 1.

```
void scheme_ungetc(int ch,
                  Scheme_Object* port)
```

Puts the byte *ch* back as the next character to be read from the given input port. The character need not have been read from *port*, and `scheme_ungetc` can be called to insert up to five characters at the start of *port*.

Use `scheme_get_byte` followed by `scheme_ungetc` only when your program will certainly call `scheme_get_byte` again to consume the byte. Otherwise, use `scheme_peek_byte`, because some a port may implement peeking and getting differently.

```
int scheme_byte_ready(Scheme_Object* port)
```

Returns 1 if a call to `scheme_get_byte` is guaranteed not to block for the given input port.

```
int scheme_char_ready(Scheme_Object* port)
```

Returns 1 if a call to `scheme_getc` is guaranteed not to block for the given input port.

```
void scheme_need_wakeup(Scheme_Object* port,  
                        void* fds)
```

Requests that appropriate bits are set in *fds* to specify which file descriptors(s) the given input port reads from. (*fds* is sortof a pointer to an `fd_set` struct; see §8.4.1 “Callbacks for Blocked Threads”.)

```
long scheme_tell(Scheme_Object* port)
```

Returns the current read position of the given input port, or the current file position of the given output port.

```
long scheme_tell_line(Scheme_Object* port)
```

Returns the current read line of the given input port. If lines are not counted, -1 is returned.

```
void scheme_count_lines(Scheme_Object* port)
```

Turns on line-counting for the given input port. To get accurate line counts, call this function immediately after creating a port.

```
long scheme_set_file_position(Scheme_Object* port,  
                             long pos)
```

Sets the file position of the given input or output port (from the start of the file). If the port does not support position setting, an exception is raised.

```
void scheme_close_input_port(Scheme_Object* port)
```

Closes the given input port.

```
void scheme_close_output_port(Scheme_Object* port)
```

Closes the given output port.

```
int scheme_get_port_file_descriptor(Scheme_Object* port,
                                   long* fd)
```

Fills **fd* with a file-descriptor value for *port* if one is available (i.e., the port is a file-stream port and it is not closed). The result is non-zero if the file-descriptor value is available, zero otherwise. Under Windows, a “file dscriptor” is a file HANDLE.

```
long scheme_get_port_fd(Scheme_Object* port)
```

Like `scheme_get_port_file_descriptor`, but a file descriptor or HANDLE is returned directly, and the result is -1 if no file descriptor or HANDLE is available.

```
long scheme_get_port_socket(Scheme_Object* port,
                            long* s)
```

Fills **s* with a socket value for *port* if one is available (i.e., the port is a TCP port and it is not closed). The result is non-zero if the socket value is available, zero otherwise. Under Windows, a socket value has type SOCKET.

```
Scheme_Object* scheme_make_port_type(char* name)
```

Creates a new port subtype.

```
Scheme_Input_Port*
scheme_make_input_port(Scheme_Object* subtype,
                      void* data,
                      Scheme_Object* name,
                      Scheme_Get_String_Fun get_bytes_fun,
                      Scheme_Peek_String_Fun peek_bytes_fun,
                      Scheme_Progress_Evt_Fun progress_evt_fun,
                      Scheme_Peeked_Read_Fun peeked_read_fun,
                      Scheme_In_Ready_Fun char_ready_fun,
                      Scheme_Close_Input_Fun close_fun,
                      Scheme_Need_Wakeup_Input_Fun need_wakeup_fun,
                      int must_close)
```

Creates a new input port with arbitrary control functions. The *subtype* is an arbitrary value to distinguish the port’s class. The pointer *data* will be installed as the port’s user data, which can be extracted/set with the `SCHEME_INPORT_VAL` macro. The *name* object is used as the port’s name (for [object-name](#) and as the default source name for [read-syntax](#)).

If *must_close* is non-zero, the new port will be registered with the current custodian, and *close_fun* is guaranteed to be called before the port is garbage-collected.

Although the return type of *scheme_make_input_port* is *Scheme_Input_Port**, it can be cast into a *Scheme_Object**.

The functions are as follows.

```
long get_bytes_fun(Scheme_Input_Port* port,
                  char* buffer,
                  long offset,
                  long size,
                  int nonblock,
                  Scheme_Object* unless)
```

Reads bytes into *buffer*, starting from *offset*, up to *size* bytes (i.e., *buffer* is at least *offset* plus *size* long). If *nonblock* is 0, then the function can block indefinitely, but it should return when at least one byte of data is available. If *nonblock* is 1, the function should never block. If *nonblock* is 2, a port in unbuffered mode should return only bytes previously forced to be buffered; other ports should treat a *nonblock* of 2 like 1. If *nonblock* is -1, the function can block, but should enable breaks while blocking. The function should return 0 if no bytes are ready in non-blocking mode. It should return EOF if an end-of-file is reached (and no bytes were read into *buffer*). Otherwise, the function should return the number of read bytes. The function can raise an exception to report an error.

The *unless* argument will be non-NULL only when *nonblocking* is non-zero (except as noted below), and only if the port supports progress events. If *unless* is non-NULL and *SCHEME_CDR(unless)* is non-NULL, the latter is a progress event specific to the port. The *get_bytes_fun* function should return *SCHEME_UNLESS_READY* instead of reading bytes if the event in *unless* becomes ready before bytes can be read. In particular, *get_bytes_fun* should check the event in *unless* before taking any action, and it should check the event in *unless* after any operation that may allow Scheme thread swaps. If the read must block, then it should unblock if the event in *unless* becomes ready.

If *scheme_progress_evt_via_get* is used for *progress_evt_fun*, then *unless* can be non-NULL even when *nonblocking* is 0. In all modes, *get_bytes_fun* must call *scheme_unless_ready* to check *unless_evt*. Furthermore, after any potentially thread-swapping operation, *get_bytes_fun* must call *scheme_wait_input_allowed*, because another thread may be attempting to commit, and *unless_evt* must be checked after *scheme_wait_input_allowed* returns. To block, the port should use *scheme_block_until_unless* instead of *scheme_block_until*. Finally, in blocking mode, *get_bytes_fun* must return after immediately reading data, without allowing a Scheme thread swap.

```
long peek_bytes_fun(Scheme_Input_Port* port,
                   char* buffer,
                   long offset,
                   long size,
                   Scheme_Object* skip,
                   int nonblock,
                   Scheme_Object* unless_evt)
```

Can be NULL to use a default implementation of peeking that uses *get_bytes_fun*. Otherwise, the protocol is the same as for *get_bytes_fun*, except that an extra *skip* argument indicates the number of input elements to skip (but *skip* does not apply to *buffer*). The *skip* value will be a non-negative exact integer, either a fixnum or a bignum.

```
Scheme_Object* progress_evt_fun(Scheme_Input_Port* port)
```

Called to obtain a progress event for the port, such as for [port-progress-evt](#). This function can be NULL if the port does not support progress events. Use `scheme_progress_evt_via_get` to obtain a default implementation, in which case *peeked_read_fun* should be `scheme_peeked_read_via_get`, and *get_bytes_fun* and *peek_bytes_fun* should handle *unless* as described above.

```
int peeked_read_fun(Scheme_Input_Port* port,
                   long amount,
                   Scheme_Object* unless_evt,
                   Scheme_Object* target_ch)
```

Called to commit previously peeked bytes, just like the sixth argument to [make-input-port](#). Use `scheme_peeked_read_via_get` for the default implementation of commits when *progress_evt_fun* is `scheme_progress_evt_via_get`.

```
int char_ready_fun(Scheme_Input_Port* port)
```

Returns 1 when a non-blocking *get_bytes_fun* will return bytes or an EOF.

```
void close_fun(Scheme_Input_Port* port)
```

Called to close the port. The port is not considered closed until the function returns.

```
void need_wakeup_fun(Scheme_Input_Port* port,
                    void* fds)
```

Called when the port is blocked on a read; *need_wakeup_fun* should set appropriate bits in *fds* to specify which file descriptor(s) it is blocked on. The *fds* argument is conceptually an array of three *fd_set* structs (one for read, one for write, one for exceptions), but manipulate this array using *scheme_get_fdset* to get a particular element of the array, and use *MZ_FD_XXX* instead of *FD_XXX* to manipulate a single “*fd_set*”. Under Windows, the first “*fd_set*” can also contain OS-level semaphores or other handles via *scheme_add_fd_handle*.

```
Scheme_Output_Port*
scheme_make_output_port(Scheme_Object* subtype,
                       void* data,
                       Scheme_Object* name,
                       Scheme_Write_String_Evt_Fun write_bytes_evt_fun,
                       Scheme_Write_String_Fun write_bytes_fun,
                       Scheme_Out_Ready_Fun char_ready_fun,
                       Scheme_Close_Output_Fun close_fun,
                       Scheme_Need_Wakeup_Output_Fun need_wakeup_fun,
                       Scheme_Write_Special_Fun write_special_fun,
                       Scheme_Write_Special_Evt_Fun write_special_evt_fun,
                       Scheme_Write_Special_Fun write_special_fun,
                       int must_close)
```

Creates a new output port with arbitrary control functions. The *subtype* is an arbitrary value to distinguish the port’s class. The pointer *data* will be installed as the port’s user data, which can be extracted/set with the *SCHEME_OUTPORT_VAL* macro. The *name* object is used as the port’s name.

If *must_close* is non-zero, the new port will be registered with the current custodian, and *close_fun* is guaranteed to be called before the port is garbage-collected.

Although the return type of *scheme_make_output_port* is *Scheme_Output_Port**, it can be cast into a *Scheme_Object**.

The functions are as follows.

```
long write_bytes_evt_fun(Scheme_Output_Port* port,
                        const char* buffer,
                        long offset,
                        long size)
```

Returns an event that writes up to *size* bytes atomically when event is chosen in a synchronization. Supply *NULL* if bytes cannot be written atomically, or supply *scheme_write_evt_via_write* to use the default implementation in terms of *write_bytes_fun* (with *rarely_block* as 2).

```
long write_bytes_fun(Scheme_Output_Port* port,
                    const char* buffer,
                    long offset,
                    long size,
                    int rarely_block,
                    int enable_break)
```

Write bytes from *buffer*, starting from *offset*, up to *size* bytes (i.e., *buffer* is at least *offset* plus *size* long). If *rarely_block* is 0, then the function can block indefinitely, and it can buffer output. If *rarely_block* is 2, the function should never block, and it should not buffer output. If *rarely_block* is 1, the function should not buffer data, and it should block only until writing at least one byte, either from *buffer* or an internal buffer. The function should return the number of bytes from *buffer* that were written; when *rarely_block* is non-zero and bytes remain in an internal buffer, it should return -1. The *size* argument can be 0 when *rarely_block* is 0 for a blocking flush, and it can be 0 if *rarely_block* is 2 for a non-blocking flush. If *enable_break* is true, then it should enable breaks while blocking. The function can raise an exception to report an error.

```
int char_ready_fun(Scheme_Output_Port* port)
```

Returns 1 when a non-blocking *write_bytes_fun* will write at least one byte or flush at least one byte from the port's internal buffer.

```
void close_fun(Scheme_Output_Port* port)
```

Called to close the port. The port is not considered closed until the function returns. This function is allowed to block (usually to flush a buffer) unless *scheme_close_should_force_port_closed* returns a non-zero result, in which case the function must return without blocking.

```
void need_wakeup_fun(Scheme_Output_Port* port,
                    void* fds)
```

Called when the port is blocked on a write; *need_wakeup_fun* should set appropriate bits in *fds* to specify which file descriptor(s) it is blocked on. The *fds* argument is conceptually an array of three *fd_set* structs (one for read, one for write, one for exceptions), but manipulate this array using *scheme_get_fdset* to get a particular element of the array, and use *MZ_FD_XXX* instead of *FD_XXX* to manipulate a single "fd_set". Under Windows, the first "fd_set" can also contain OS-level semaphores or other handles via *scheme_add_fd_handle*.

```
int write_special_evt_fun(Scheme_Output_Port* port,
                        Scheme_Object* v)
```

Returns an event that writes *v* atomically when event is chosen in a synchronization. Supply NULL if specials cannot be written atomically (or at all), or supply `scheme_write_special_evt_via_write_special` to use the default implementation in terms of `write_special_fun` (with *non_block* as 1).

```
int write_special_fun(Scheme_Output_Port* port,
                    Scheme_Object* v,
                    int non_block)
```

Called to write the special value *v* for `write-special` (when *non_block* is 0) or `write-special-avail` (when *non_block* is 1). If NULL is supplied instead of a function pointer, then `write-special` and `write-special-avail` produce an error for this port.

```
Scheme_Object* scheme_make_file_input_port(FILE* fp)
```

Creates a Scheme input file port from an ANSI C file pointer. The file must never block on reads.

```
Scheme_Object* scheme_open_input_file(const char* filename,
                                     const char* who)
```

Opens *filename* for reading. In an exception is raised, the exception message uses *who* as the name of procedure that raised the exception.

```
Scheme_Object*
scheme_make_named_file_input_port(FILE* fp,
                                 Scheme_Object* name)
```

Creates a Scheme input file port from an ANSI C file pointer. The file must never block on reads. The *name* argument is used as the port's name.

```
Scheme_Object* scheme_open_output_file(const char* filename,
                                       const char* who)
```

Opens *filename* for writing in `'truncate/replace` mode. If an exception is raised, the exception message uses *who* as the name of procedure that raised the exception.

```
Scheme_Object* scheme_make_file_output_port(FILE* fp)
```

Creates a Scheme output file port from an ANSI C file pointer. The file must never block on writes.

```
Scheme_Object* scheme_make_fd_input_port(int fd,
                                         Scheme_Object* name,
                                         int regfile,
                                         int win_textmode)
```

Creates a Scheme input port for a file descriptor *fd*. Under Windows, *fd* can be a HANDLE for a stream, and it should never be a file descriptor from the C library or a WinSock socket.

The *name* object is used for the port's name. Specify a non-zero value for *regfile* only if the file descriptor corresponds to a regular file (which implies that reading never blocks, for example).

Under Windows, *win_textmode* can be non-zero to make trigger auto-conversion (at the byte level) of CRLF combinations to LF.

Closing the resulting port closes the file descriptor.

Instead of calling both `scheme_make_fd_input_port` and `scheme_make_fd_output_port` on the same file descriptor, call `scheme_make_fd_output_port` with a non-zero last argument. Otherwise, closing one of the ports causes the file descriptor used by the other to be closed as well.

```
Scheme_Object* scheme_make_fd_output_port(int fd,
                                           Scheme_Object* name,
                                           int regfile,
                                           int win_textmode,
                                           int read_too)
```

Creates a Scheme output port for a file descriptor *fd*. Under Windows, *fd* can be a HANDLE for a stream, and it should never be a file descriptor from the C library or a WinSock socket.

The *name* object is used for the port's name. Specify a non-zero value for *regfile* only if the file descriptor corresponds to a regular file (which implies that reading never blocks, for example).

Under Windows, *win_textmode* can be non-zero to make trigger auto-conversion (at the byte level) of CRLF combinations to LF.

Closing the resulting port closes the file descriptor.

If *read_too* is non-zero, the function produces multiple values (see §6.3 “Multiple Values”) instead of a single port. The first result is an input port for *fd*, and the second is an output port for *fd*. These ports are connected in that the file descriptor is closed only when both of

the ports are closed.

```
void scheme_socket_to_ports(long s,
                           const char* name,
                           int close,
                           Scheme_Object** inp,
                           Scheme_Object** outp)
```

Creates Scheme input and output ports for a TCP socket *s*. The *name* argument supplies the name for the ports. If *close* is non-zero, then the ports assume responsibility for closing the socket. The resulting ports are written to *inp* and *outp*.

```
Scheme_Object* scheme_make_byte_string_input_port(char* str)
```

Creates a Scheme input port from a byte string; successive [read-chars](#) on the port return successive bytes in the string.

```
Scheme_Object* scheme_make_byte_string_output_port()
```

Creates a Scheme output port; all writes to the port are kept in a byte string, which can be obtained with `scheme_get_byte_string_output`.

```
char* scheme_get_byte_string_output(Scheme_Object* port)
```

Returns (in a newly allocated byte string) all data that has been written to the given string output port so far. (The returned string is nul-terminated.)

```
char* scheme_get_sized_byte_string_output(Scheme_Object* port,
                                          long* len)
```

Returns (in a newly allocated byte string) all data that has been written to the given string output port so far and fills in **len* with the length of the string in bytes (not including the nul terminator).

```
void scheme_pipe(Scheme_Object** read,
                 Scheme_Object** write)
```

Creates a pair of ports, setting **read* and **write*; data written to **write* can be read back out of **read*. The pipe can store arbitrarily many unread characters,

```
void scheme_pipe_with_limit(Scheme_Object** read,
                           Scheme_Object** write,
                           int limit)
```

Like `scheme_pipe` if `limit` is 0. If `limit` is positive, creates a pipe that stores at most `limit` unread characters, blocking writes when the pipe is full.

```
Scheme_Input_Port* scheme_input_port_record(Scheme_Object* port)
```

Returns the input-port record for `port`, which may be either a raw-port object with type `scheme_input_port_type` or a structure with the `prop:input-port` property.

```
Scheme_Output_Port*
scheme_output_port_record(Scheme_Object* port)
```

Returns the output-port record for `port`, which may be either a raw-port object with type `scheme_output_port_type` or a structure with the `prop:output-port` property.

```
int scheme_file_exists(char* name)
```

Returns 1 if a file by the given name exists, 0 otherwise. If `name` specifies a directory, FALSE is returned. The `name` should be already expanded.

```
int scheme_directory_exists(char* name)
```

Returns 1 if a directory by the given name exists, 0 otherwise. The `name` should be already expanded.

```
char* scheme_expand_filename(const char* name,
                             int len,
                             const char* where,
                             int* expanded,
                             int checks)
```

Cleanses the pathname `name` (see [cleanse-path](#)) and resolves relative paths with respect to the current directory parameter. The `len` argument is the length of the input string; if it is -1, the string is assumed to be null-terminated. The `where` argument is used to raise an exception if there is an error in the filename; if this is NULL, an error is not reported and NULL is returned instead. If `expanded` is not NULL, `*expanded` is set to 1 if some expansion takes place, or 0 if the input name is simply returned.

If `guards` is not 0, then `scheme_security_check_file` (see §15 “Security Guards”) is

called with *name*, *where*, and *checks* (which implies that *where* should never be NULL unless *guards* is 0). Normally, *guards* should be SCHEME_GUARD_FILE_EXISTS at a minimum. Note that a failed access check will result in an exception.

```
char* scheme_expand_string_filename(Scheme_Object* name,
                                   const char* where,
                                   int* expanded,
                                   int checks)
```

Like `scheme_expand_string`, but given a *name* that can be a character string or a path value.

```
Scheme_Object* scheme_char_string_to_path(Scheme_Object* s)
```

Converts a Scheme character string into a Scheme path value.

```
Scheme_Object* scheme_path_to_char_string(Scheme_Object* s)
```

Converts a Scheme path value into a Scheme character string.

```
Scheme_Object* scheme_make_path(char* bytes)
```

Makes a path value given a byte string. The *bytes* string is copied.

```
Scheme_Object* scheme_make_path_without_copying(char* bytes)
```

Like `scheme_make_path`, but the string is not copied.

```
Scheme_Object* scheme_make_sized_path(char* bytes,
                                     long len,
                                     int copy)
```

Makes a path whose byte form has size *len*. A copy of *bytes* is made if *copy* is not 0. The string *bytes* should contain *len* bytes, and if *copy* is zero, *bytes* must have a nul terminator in addition. If *len* is negative, then the nul-terminated length of *bytes* is used for the length.

```
Scheme_Object* scheme_make_sized_offset_path(char* bytes,
                                             long d,
                                             long len,
                                             int copy)
```

Like `scheme_make_sized_path`, except the *len* bytes start from position *d* in *bytes*. If *d* is

non-zero, then *copy* must be non-zero.

```
char* scheme_build_mac_filename(FSSpec* spec,  
                                int isdir)
```

Mac OS X only: Converts an FSSpec record (defined by Mac OS X) into a pathname string. If *spec* contains only directory information (via the *vRefNum* and *parID* fields), *isdir* should be 1, otherwise it should be 0.

```
int scheme_mac_path_to_spec(const char* filename,  
                           FSSpec* spec,  
                           long* type)
```

Mac OS X only: Converts a pathname into an FSSpec record (defined by Mac OS X), returning 1 if successful and 0 otherwise. If *type* is not NULL and *filename* is a file that exists, *type* is filled with the file's four-character Mac OS X type. If *type* is not NULL and *filename* is not a file that exists, *type* is filled with 0.

```
char* scheme_os_getcwd(char* buf,  
                       int buflen,  
                       int* actlen,  
                       int noexn)
```

Gets the current working directory according to the operating system. This is separate from MzScheme's current directory parameter.

The directory path is written into *buf*, of length *buflen*, if it fits. Otherwise, a new (collectable) string is allocated for the directory path. If *actlen* is not NULL, **actlen* is set to the length of the current directory path. If *noexn* is not 0, then an exception is raised if the operation fails.

```
int scheme_os_setcwd(char* buf,  
                    int noexn)
```

Sets the current working directory according to the operating system. This is separate from MzScheme's current directory parameter.

If *noexn* is not 0, then an exception is raised if the operation fails.

```
char* scheme_format(mzchar* format,
                    int flen,
                    int argc,
                    Scheme_Object** argv,
                    long* rlen)
```

Creates a string like MzScheme's `format` procedure, using the format string *format* (of length *flen*) and the extra arguments specified in *argc* and *argv*. If *rlen* is not NULL, **rlen* is filled with the length of the resulting string.

```
void scheme_printf(char* format,
                  int flen,
                  int argc,
                  Scheme_Object** argv)
```

Writes to the current output port like MzScheme's `printf` procedure, using the format string *format* (of length *flen*) and the extra arguments specified in *argc* and *argv*.

```
char* scheme_format_utf8(char* format,
                         int flen,
                         int argc,
                         Scheme_Object** argv,
                         long* rlen)
```

Like `scheme_format`, but takes a UTF-8-encoding byte string.

```
void scheme_printf_utf8(char* format,
                       int flen,
                       int argc,
                       Scheme_Object** argv)
```

Like `scheme_printf`, but takes a UTF-8-encoding byte string.

```
int scheme_close_should_force_port_closed()
```

This function must be called by the close function for a port created with `scheme_make_output_port`.

14 Structures

A new Scheme structure type is created with `scheme_make_struct_type`. This creates the structure type, but does not generate the constructor, etc. procedures. The `scheme_make_struct_values` function takes a structure type and creates these procedures. The `scheme_make_struct_names` function generates the standard structure procedure names given the structure type's name. Instances of a structure type are created with `scheme_make_struct_instance` and the function `scheme_is_struct_instance` tests a structure's type. The `scheme_struct_ref` and `scheme_struct_set` functions access or modify a field of a structure.

The structure procedure values and names generated by `scheme_make_struct_values` and `scheme_make_struct_names` can be restricted by passing any combination of these flags:

- `SCHEME_STRUCT_NO_TYPE` — the structure type value/name is not returned.
- `SCHEME_STRUCT_NO_CONSTR` — the constructor procedure value/name is not returned.
- `SCHEME_STRUCT_NO_PRED` — the predicate procedure value/name is not returned.
- `SCHEME_STRUCT_NO_GET` — the selector procedure values/names are not returned.
- `SCHEME_STRUCT_NO_SET` — the mutator procedure values/names are not returned.
- `SCHEME_STRUCT_GEN_GET` — the field-independent selector procedure value/name is returned.
- `SCHEME_STRUCT_GEN_SET` — the field-independent mutator procedure value/name is returned.

When all values or names are returned, they are returned as an array with the following order: structure type, constructor, predicate, first selector, first mutator, second selector, etc., field-independent select, field-independent mutator. When particular values/names are omitted, the array is compressed accordingly.

```
Scheme_Object* scheme_make_struct_type(Scheme_Object* base_name,
                                       Scheme_Object* super_type,
                                       Scheme_Object* inspector,
                                       int num_init_fields,
                                       int num_auto_fields,
                                       Scheme_Object* auto_val,
                                       Scheme_Object* properties,
                                       Scheme_Object* guard)
```

Creates and returns a new structure type. The *base_name* argument is used as the name of the new structure type; it must be a symbol. The *super_type* argument should be NULL or an existing structure type to use as the super-type. The *inspector* argument should be NULL or an inspector to manage the type. The *num_init_fields* argument specifies the number of fields for instances of this structure type that have corresponding constructor arguments. (If a super-type is used, this is the number of additional fields, rather than the total number.) The *num_auto_fields* argument specifies the number of additional fields that have no corresponding constructor arguments, and they are initialized to *auto_val*. The *properties* argument is a list of property-value pairs. The *guard* argument is either NULL or a procedure to use as a constructor guard.

```
Scheme_Object**
scheme_make_struct_names(Scheme_Object* base_name,
                        Scheme_Object* field_names,
                        int flags,
                        int* count_out)
```

Creates and returns an array of standard structure value name symbols. The *base_name* argument is used as the name of the structure type; it should be the same symbol passed to the associated call to `scheme_make_struct_type`. The *field_names* argument is a (Scheme) list of field name symbols. The *flags* argument specifies which names should be generated, and if *count_out* is not NULL, *count_out* is filled with the number of names returned in the array.

```
Scheme_Object**
scheme_make_struct_values(Scheme_Object* struct_type,
                        Scheme_Object** names,
                        int count,
                        int flags)
```

Creates and returns an array of the standard structure value and procedure values for *struct_type*. The *struct_type* argument must be a structure type value created by `scheme_make_struct_type`. The *names* procedure must be an array of name symbols, generally the array returned by `scheme_make_struct_names`. The *count* argument specifies the length of the *names* array (and therefore the number of expected return values) and the *flags* argument specifies which values should be generated.

```
Scheme_Object*
scheme_make_struct_instance(Scheme_Object* struct_type,
                          int argc,
                          Scheme_Object** argv)
```

Creates an instance of the structure type *struct_type*. The *argc* and *argv* arguments provide the field values for the new instance.

```
int scheme_is_struct_instance(Scheme_Object* struct_type,  
                             Scheme_Object* v)
```

Returns 1 if *v* is an instance of *struct_type* or 0 otherwise.

```
Scheme_Object* scheme_struct_ref(Scheme_Object* s,  
                                 int n)
```

Returns the *n*th field (counting from 0) in the structure *s*.

```
void scheme_struct_set(Scheme_Object* s,  
                      int n,  
                      Scheme_Object* v)
```

Sets the *n*th field (counting from 0) in the structure *s* to *v*.

15 Security Guards

Before a primitive procedure accesses the filesystem or creates a network connection, it should first consult the current security guard to determine whether such access is allowed for the current thread.

File access is normally preceded by a call to `scheme_expand_filename`, which accepts flags to indicate the kind of filesystem access needed, so that the security guard is consulted automatically.

An explicit filesystem-access check can be made by calling `scheme_security_check_file`. Similarly, an explicit network-access check is performed by calling `scheme_security_check_network`.

```
void scheme_security_check_file(const char* who,
                               char* filename,
                               int guards)
```

Consults the current security manager to determine whether access is allowed to *filename*. The *guards* argument should be a bitwise combination of the following:

- SCHEME_GUARD_FILE_READ
- SCHEME_GUARD_FILE_WRITE
- SCHEME_GUARD_FILE_EXECUTE
- SCHEME_GUARD_FILE_DELETE
- SCHEME_GUARD_FILE_EXISTS (do not combine with other values)

The *filename* argument can be NULL (in which case `#f` is sent to the security manager's procedure), and *guards* should be SCHEME_GUARD_FILE_EXISTS in that case.

If access is denied, an exception is raised.

```
void scheme_security_check_network(const char* who,
                                  char* host,
                                  int portno)
```

Consults the current security manager to determine whether access is allowed for creating a client connection to *host* on port number *portno*. If *host* is NULL, the security manager is consulted for creating a server at port number *portno*.

If access is denied, an exception is raised.

16 Custodians

When an extension allocates resources that must be explicitly freed (in the same way that a port must be explicitly closed), a Scheme object associated with the resource should be placed into the management of the current custodian with `scheme_add_managed`.

Before allocating the resource, call `scheme_custodian_check_available` to ensure that the relevant custodian is not already shut down. If it is, `scheme_custodian_check_available` will raise an exception. If the custodian is shut down when `scheme_add_managed` is called, the close function provided to `scheme_add_managed` will be called immediately, and no exception will be reported.

```
Scheme_Custodian* scheme_make_custodian(Scheme_Custodian* m)
```

Creates a new custodian as a subordinate of *m*. If *m* is NULL, then the main custodian is used as the new custodian's supervisor. Do not use NULL for *m* unless you intend to create an especially privileged custodian.

```
Scheme_Custodian_Reference*  
scheme_add_managed(Scheme_Custodian* m,  
                  Scheme_Object* o,  
                  Scheme_Close_Custodian_Client* f,  
                  void* data,  
                  int strong)
```

Places the value *o* into the management of the custodian *m*. If *m* is NULL, the current custodian is used.

The *f* function is called by the custodian if it is ever asked to “shutdown” its values; *o* and *data* are passed on to *f*, which has the type

```
typedef void (*Scheme_Close_Custodian_Client)(Scheme_Object *o,  
                                              void *data);
```

If *strong* is non-zero, then the newly managed value will be remembered until either the custodian shuts it down or `scheme_remove_managed` is called. If *strong* is zero, the value is allowed to be garbage collected (and automatically removed from the custodian).

Independent of whether *strong* is zero, the value *o* is initially weakly held. A value associated with a custodian can therefore be finalized via will executors.

The return value from `scheme_add_managed` can be used to refer to the value's custodian later in a call to `scheme_remove_managed`. A value can be registered with at most one custodian.

If *m* (or the current custodian if *m* is NULL) is shut down, then *f* is called immediately, and the result is NULL.

```
void scheme_custodian_check_available(Scheme_Custodian* m,
                                     const char* name,
                                     const char* rename)
```

Checks whether *m* is already shut down, and raises an error if so. If *m* is NULL, the current custodian is used. The *name* argument is used for error reporting. The *rename* argument will likely be used for checking pre-set limits in the future; pre-set limits will have symbolic names, and the *rename* string will be compared to the symbols.

```
void scheme_remove_managed(Scheme_Custodian_Reference* mref,
                           Scheme_Object* o)
```

Removes *o* from the management of its custodian. The *mref* argument must be a value returned by `scheme_add_managed` or NULL.

```
void scheme_close_managed(Scheme_Custodian* m)
```

Instructs the custodian *m* to shutdown all of its managed values.

```
void scheme_add_atexit_closer(Scheme_Exit_Closer_Func f)
```

Installs a function to be called on each custodian-registered item and its closer when MzScheme is about to exit. The registered function has the type

```
typedef
void (*Scheme_Exit_Closer_Func)(Scheme_Object *o,
                                Scheme_Close_Custodian_Client *f,
                                void *d);
```

where *d* is the second argument for *f*.

17 Miscellaneous Utilities

The `MZSCHEME_VERSION` preprocessor macro is defined as a string describing the version of Scheme. The `MZSCHEME_VERSION_MAJOR` and `MZSCHEME_VERSION_MINOR` macros are defined as the major and minor version numbers, respectively.

```
int scheme_eq(Scheme_Object* obj1,
              Scheme_Object* obj2)
```

Returns 1 if the Scheme values are `eq?`.

```
int scheme_eqv(Scheme_Object* obj1,
               Scheme_Object* obj2)
```

Returns 1 if the Scheme values are `eqv?`.

```
int scheme_equal(Scheme_Object* obj1,
                 Scheme_Object* obj2)
```

Returns 1 if the Scheme values are `equal?`.

```
int scheme_recur_equal(Scheme_Object* obj1,
                       Scheme_Object* obj2,
                       void* cycle_data)
```

Like `scheme_equal`, but accepts an extra value for cycle tracking. This procedure is meant to be called by a procedure installed with `scheme_set_type_equality`.

Returns 1 if the Scheme values are `equal?.`}

```
long scheme_equal_hash_key(Scheme_Object* obj)
```

Returns the primary `equal?`-hash key for `obj`.

```
long scheme_equal_hash_key2(Scheme_Object* obj)
```

Returns the secondary `equal?`-hash key for `obj`.

```
long scheme_recur_equal_hash_key(Scheme_Object* obj,
                                 void* cycle_data)
```

Like `scheme_equal_hash_key`, but accepts an extra value for cycle tracking. This procedure is meant to be called by a hashing procedure installed with `scheme_set_type_equality`.

Returns the primary `equal?`-hash key for *obj*.)

```
long scheme_recur_equal_hash_key2(Scheme_Object* obj,  
                                void* cycle_data)
```

Like `scheme_equal_hash_key2`, but accepts an extra value for cycle tracking. This procedure is meant to be called by a secondary hashing procedure installed with `scheme_set_type_equality`.

```
Scheme_Object* scheme_build_list(int c,  
                                Scheme_Object** elems)
```

Creates and returns a list of length *c* with the elements *elems*.

```
int scheme_list_length(Scheme_Object* list)
```

Returns the length of the list. If *list* is not a proper list, then the last `cdr` counts as an item. If there is a cycle in *list* (involving only `cdrs`), this procedure will not terminate.

```
int scheme_proper_list_length(Scheme_Object* list)
```

Returns the length of the list, or -1 if it is not a proper list. If there is a cycle in *list* (involving only `cdrs`), this procedure returns -1.

```
Scheme_Object* scheme_car(Scheme_Object* pair)
```

Returns the `car` of the pair.

```
Scheme_Object* scheme_cdr(Scheme_Object* pair)
```

Returns the `cdr` of the pair.

```
Scheme_Object* scheme_cadr(Scheme_Object* pair)
```

Returns the `cadr` of the pair.

```
Scheme_Object* scheme_caddr(Scheme_Object* pair)
```

Returns the [caddr](#) of the pair.

```
Scheme_Object* scheme_vector_to_list(Scheme_Object* vec)
```

Creates a list with the same elements as the given vector.

```
Scheme_Object* scheme_list_to_vector(Scheme_Object* list)
```

Creates a vector with the same elements as the given list.

```
Scheme_Object* scheme_append(Scheme_Object* lstx,  
                             Scheme_Object* lsty)
```

Non-destructively appends the given lists.

```
Scheme_Object* scheme_unbox(Scheme_Object* obj)
```

Returns the contents of the given box.

```
void scheme_set_box(Scheme_Object* b,  
                  Scheme_Object* v)
```

Sets the contents of the given box.

```
Scheme_Object* scheme_dynamic_require(int argc,  
                                     Scheme_Object** argv)
```

The same as [dynamic-require](#). The *argc* argument must be 2, and *argv* contains the arguments.

```
Scheme_Object*  
scheme_namespace_require(Scheme_Object* prim_req_spec)
```

The same as [namespace-require](#).

```
Scheme_Object* scheme_load(char* file)
```

Loads the specified Scheme file, returning the value of the last expression loaded, or NULL if the load fails.

```
Scheme_Object* scheme_load_extension(char* filename)
```

Loads the specified Scheme extension file, returning the value provided by the extension's initialization function.

`Scheme_Hash_Table* scheme_make_hash_table(int type)`

Creates a hash table. The *type* argument must be either `SCHEME_hash_ptr` or `SCHEME_hash_string`, which determines how keys are compared (unless the hash and compare functions are modified in the hash table record; see below). A `SCHEME_hash_ptr` table hashes on a key's pointer address, while `SCHEME_hash_string` uses a key as a `char*` and hashes on the null-terminated string content. Since a hash table created with `SCHEME_hash_string` (instead of `SCHEME_hash_ptr`) does not use a key as a Scheme value, it cannot be used from Scheme code.

Although the hash table interface uses the type `Scheme_Object*` for both keys and values, the table functions never inspect values, and they inspect keys only for `SCHEME_hash_string` hashing. Thus, the actual types of the values (and keys, for `SCHEME_hash_ptr` tables) can be anything.

The public portion of the `Scheme_Hash_Table` type is defined roughly as follows:

```
typedef struct Scheme_Hash_Table {
    Scheme_Object so; /* so.type == scheme_hash_table_type */
    /* ... */
    int size; /* size of keys and vals arrays */
    int count; /* number of mapped keys */
    Scheme_Object **keys;
    Scheme_Object **vals;
    void (*make_hash_indices)(void *v, long *h1, long *h2);
    int (*compare)(void *v1, void *v2);
    /* ... */
} Scheme_Hash_Table;
```

The `make_hash_indices` and `compare` function pointers can be set to arbitrary hashing and comparison functions (before any mapping is installed into the table). A hash function should fill *h1* with a primary hash value and *h2* with a secondary hash value; the values are for double-hashing, where the caller takes appropriate modulus. Either *h1* or *h2* can be `NULL` if the corresponding hash code is not needed.

To traverse the hash table content, iterate over *keys* and *vals* in parallel from 0 to `size-1`, and ignore *keys* where the corresponding *vals* entry is `NULL`. The `count` field indicates the number of non-`NULL` values that will be encountered.

`Scheme_Hash_Table* scheme_make_hash_table_equal()`

Like `scheme_make_hash_table`, except that keys are treated as Scheme values and hashed based on `equal?` instead of `eq?`.

```
void scheme_hash_set(Scheme_Hash_Table* table,
                    Scheme_Object* key,
                    Scheme_Object* val)
```

Sets the current value for *key* in *table* to *val*. If *val* is NULL, the *key* is unmapped in *table*.

```
Scheme_Object* scheme_hash_get(Scheme_Hash_Table* table,
                               Scheme_Object* key)
```

Returns the current value for *key* in *table*, or NULL if *key* has no value.

```
Scheme_Bucket_Table* scheme_make_bucket_table(int size_hint,
                                              int type)
```

Like `make_hash_table`, but bucket tables are somewhat more flexible, in that hash buckets are accessible and weak keys are supported. (They also consume more space than hash tables.)

The *type* argument must be either `SCHEME_hash_ptr`, `SCHEME_hash_string`, or `SCHEME_hash_weak_ptr`. The first two are the same as for hash tables. The last is like `SCHEME_hash_ptr`, but the keys are weakly held.

The public portion of the `Scheme_Bucket_Table` type is defined roughly as follows:

```
typedef struct Scheme_Bucket_Table {
    Scheme_Object so; /* so.type == scheme_variable_type */
    /* ... */
    int size; /* size of buckets array */
    int count; /* number of buckets, >= number of mapped keys */
    Scheme_Bucket **buckets;
    void (*make_hash_indices)(void *v, long *h1, long *h2);
    int (*compare)(void *v1, void *v2);
    /* ... */
} Scheme_Bucket_Table;
```

The `make_hash_indices` and `compare` functions are used as for hash tables. Note that `SCHEME_hash_weak_ptr` supplied as the initial type makes keys weak even if the hash and comparison functions are changed.

See `scheme_bucket_from_table` for information on buckets.

```
void scheme_add_to_table(Scheme_Bucket_Table* table,
                        const char* key,
                        void* val,
                        int const)
```

Sets the current value for *key* in *table* to *val*. If *const* is non-zero, the value for *key* must never be changed.

```
void scheme_change_in_table(Scheme_Bucket_Table* table,
                            const char* key,
                            void* val)
```

Sets the current value for *key* in *table* to *val*, but only if *key* is already mapped in the table.

```
void* scheme_lookup_in_table(Scheme_Bucket_Table* table,
                             const char* key)
```

Returns the current value for *key* in *table*, or NULL if *key* has no value.

```
Scheme_Bucket*
scheme_bucket_from_table(Scheme_Bucket_Table* table,
                        const char* key)
```

Returns the bucket for *key* in *table*. The `Scheme_Bucket` structure is defined as:

```
typedef struct Scheme_Bucket {
    Scheme_Object so; /* so.type == scheme_bucket_type */
    /* ... */
    void *key;
    void *val;
} Scheme_Bucket;
```

Setting *val* to NULL unmaps the bucket's key, and *key* can be NULL in that case as well. If the table holds keys weakly, then *key* points to a (weak) pointer to the actual key, and the weak pointer's value can be NULL.

```
long scheme_double_to_int(char* where,
                          double d)
```

Returns a fixnum value for the given floating-point number *d*. If *d* is not an integer or if it is too large, then an error message is reported; *name* is used for error-reporting.

`long scheme_get_milliseconds()`

Returns the current “time” in milliseconds, just like [current-milliseconds](#).

`long scheme_get_process_milliseconds()`

Returns the current process “time” in milliseconds, just like [current-process-milliseconds](#).

`char* scheme_banner()`

Returns the string that is used as the Scheme startup banner.

`char* scheme_version()`

Returns a string for the executing version of Scheme.

18 Flags and Hooks

The following flags and hooks are available when PLT Scheme is embedded:

- `scheme_exit` — This pointer can be set to a function that takes an integer argument and returns `void`; the function will be used as the default exit handler. The default is `NULL`.
- `scheme_make_stdin`, `scheme_make_stdout`, `scheme_make_stderr`, — These pointers can be set to a function that takes no arguments and returns a Scheme port `Scheme_Object *` to be used as the starting standard input, output, and/or error port. The defaults are `NULL`. Setting the initial error port is particularly important for seeing unexpected error messages if `stderr` output goes nowhere.
- `scheme_console_output` — This pointer can be set to a function that takes a string and a long string length; the function will be called to display internal MzScheme warnings and messages that possibly contain non-terminating nuls. The default is `NULL`.
- `scheme_check_for_break` — This points to a function of no arguments that returns an integer. It is used as the default user-break polling procedure in the main thread. A non-zero return value indicates a user break, and each time the function returns a non-zero value, it counts as a new break signal (though the break signal may be ignored if a previous signal is still pending). The default is `NULL`.
- `scheme_case_sensitive` — If this flag is set to a non-zero value before `scheme_basic_env` is called, then MzScheme will not ignore capitalization for symbols and global variable names. The value of this flag should not change once it is set. The default is zero.
- `scheme_allow_set_undefined` — This flag determines the initial value of `compile-allow-set!-undefined`. The default is zero.
- `scheme_console_printf` — This function pointer was left for backward compatibility. The default builds a string and calls `scheme_console_output`.

```
void scheme_set_collects_path(Scheme_Object* path)
```

Sets the path to be returned by `(find-system-path 'collects-dir)`.

```
void scheme_set_addon_path(Scheme_Object* path)
```

Sets the path to be returned by `(find-system-path 'addon-dir)`.

```
void
scheme_init_collection_paths_post(Scheme_Env* env,
                                  Scheme_Object* pre_extra_paths,
                                  Scheme_Object* post_extra_paths)
```

Initializes the `current-library-collection-paths` parameter using `find-library-collection-paths`. The `pre_extra_paths` and `post_extra-paths` arguments are propagated to `find-library-collection-paths`.

The function calls `scheme_seal_parameters` automatically.

```
void scheme_init_collection_paths(Scheme_Env* env,
                                  Scheme_Object* pre_extra_paths)
```

Like `scheme_init_collection_paths_post`, but with `null` as the last argument.

```
void scheme_seal_parameters()
```

Takes a snapshot of the current values of built-in parameters. These values are used for privileged actions, such as installing a PLaneT package.

Index

- "escheme.h", 5
- "libmzgc.a", 9
- "libmzgc.la", 9
- "libmzgc.so", 9
- "libmzscheme.a", 9
- "libmzscheme.la", 9
- "libmzscheme.so", 9
- "libmzscheme3m.a", 13
- "libmzscheme3m.la", 13
- "libmzscheme3m.so", 13
- "mzdyn.o", 6
- "mzdyn.obj", 6
- "mzdyn3m.o", 8
- "mzdyn3m.obj", 8
- "scheme.h", 10
- #!/variable-reference, 48
- 3m, 8
- cc, 6
- cgc, 6
- ld, 6
- xform, 7
- 3m, 5
- 3m Embedding, 12
- 3m Extensions, 7
- _scheme_apply, 57
- _scheme_apply_multi, 57
- _scheme_eval_compiled, 56
- _scheme_eval_compiled_multi, 56
- allocation, 7
- allocation, 11
- Allowing Thread Switches, 70
- arity, 51
- Bignums, Rationals, and Complex Numbers, 93
- Blocking the Current Thread, 70
- Callbacks for Blocked Threads, 72
- CGC, 5
- CGC Embedding, 9
- CGC Extensions, 5
- CGC versus 3m, 5
- Continuation Marks, 87
- continuations, 60
- Cooperating with 3m, 31
- current working directory, 113
- Custodians, 119
- Declaring a Module in an Extension, 8
- embedding MzScheme, 9
- Embedding MzScheme into a Program, 9
- Enabling and Disabling Breaks, 64
- Evaluation, 54
- Evaluation Functions, 56
- Exception Functions, 64
- Exceptions and Escape Continuations, 60
- extending MzScheme, 5
- Flags and Hooks, 128
- garbage collection, 30
- GC_fixup_self, 47
- GC_register_traversers, 46
- GC_resolve, 47
- Global Constants, 20
- globals, in extension code, 31
- globals, 48
- Guiding mzc -xform, 38
- header files, 5
- Inside**: PLT Scheme C API, 1
- Integers, 15
- Integration with Threads, 69
- Local Pointers, 33
- Local Pointers and mzc -xform, 37
- memory, 30
- Memory Allocation, 30
- Memory Functions, 40
- Miscellaneous Utilities, 121
- module, 48
- Multiple Values, 55
- MZ_GC_ARRAY_VAR_IN_REG, 34
- MZ_GC_DECL_REG, 34
- MZ_GC_NO_VAR_IN_REG, 36
- MZ_GC_REG, 35
- MZ_GC_UNREG, 34
- MZ_GC_VAR_IN_REG, 34
- mzc, 6

MZCONFIG_ALLOW_SET_UNDEFINED, 84
MZCONFIG_CAN_READ_BOX, 83
MZCONFIG_CAN_READ_COMPILED, 83
MZCONFIG_CAN_READ_GRAPH, 83
MZCONFIG_CAN_READ_PIPE_QUOTE, 83
MZCONFIG_CASE_SENS, 84
MZCONFIG_COLLECTION_PATHS, 84
MZCONFIG_CURLY_BRACES_ARE_PARENS, 84
MZCONFIG_CUSTODIAN, 84
MZCONFIG_ENV, 83
MZCONFIG_ERROR_DISPLAY_HANDLER, 83
MZCONFIG_ERROR_PORT, 83
MZCONFIG_ERROR_PRINT_VALUE_HANDLER, 83
MZCONFIG_ERROR_PRINT_WIDTH, 84
MZCONFIG_EVAL_HANDLER, 83
MZCONFIG_EXIT_HANDLER, 83
MZCONFIG_INIT_EXN_HANDLER, 83
MZCONFIG_INPUT_PORT, 83
MZCONFIG_LOAD_DIRECTORY, 84
MZCONFIG_LOAD_EXTENSION_HANDLER, 84
MZCONFIG_LOAD_HANDLER, 83
MZCONFIG_OUTPUT_PORT, 83
MZCONFIG_PORT_PRINT_HANDLER, 84
MZCONFIG_PRINT_BOX, 84
MZCONFIG_PRINT_GRAPH, 83
MZCONFIG_PRINT_HANDLER, 83
MZCONFIG_PRINT_STRUCT, 84
MZCONFIG_PROMPT_READ_HANDLER, 83
MZCONFIG_SQUARE_BRACKETS_ARE_PARENS, 84
MZCONFIG_USE_COMPILED_KIND, 84
MzScheme and Threads, 15
MzScheme, Unicode, Characters, and Strings, 15
Namespaces and Modules, 48
Overview, 5
Parameterizations, 83
Ports and the Filesystem, 97
Procedures, 51
scheme_add_atexit_closer, 120
scheme_add_evt, 80
scheme_add_evt_through_sema, 80
scheme_add_fd_eventmask, 79
scheme_add_fd_handle, 79
scheme_add_finalizer, 45
scheme_add_finalizer_once, 45
scheme_add_global, 48
scheme_add_global_symbol, 48
scheme_add_managed, 119
scheme_add_scheme_finalizer, 45
scheme_add_scheme_finalizer_once, 45
scheme_add_swap_callback, 82
scheme_add_swap_out_callback, 82
scheme_add_to_table, 126
scheme_alloc_byte_string, 23
scheme_alloc_char_string, 25
scheme_allow_set_undefined, 128
scheme_append, 123
scheme_append_byte_string, 24
scheme_append_char_string, 25
scheme_apply, 57
scheme_apply_multi, 57
scheme_apply_to_list, 57
scheme_banner, 127
scheme_basic_env, 56
scheme_bignum_from_double, 94
scheme_bignum_from_float, 94
scheme_bignum_normalize, 94
scheme_bignum_to_double, 93
scheme_bignum_to_float, 93
scheme_bignum_to_string, 94
scheme_bignum_type, 17
SCHEME_BIGNUMP, 17
scheme_block_until, 77
scheme_block_until_enable_break, 78
scheme_block_until_unless, 78
scheme_bool_type, 17
SCHEME_BOOLP, 17
scheme_box, 27
scheme_box_type, 18
SCHEME_BOX_VAL, 18

SCHEME_BOXP, 18
 scheme_break_thread, 77
 scheme_break_waiting, 77
 scheme_bucket_from_table, 126
 scheme_bucket_table_type, 19
 SCHEME_BUCKTP, 19
 scheme_build_list, 122
 scheme_build_mac_filename, 113
 scheme_builtin_value, 49
 scheme_byte_ready, 102
 SCHEME_BYTE_STR_VAL, 17
 scheme_byte_string_to_char_string, 25
 scheme_byte_string_to_char_string_lossy, 26
 scheme_byte_string_type, 17
 SCHEME_BYTE_STRINGP, 17
 SCHEME_BYTE_STRLEN_VAL, 17
 scheme_caddr, 122
 scheme_cadr, 122
 scheme_call_enable_break, 81
 scheme_calloc, 40
 SCHEME_CAR, 18
 scheme_car, 122
 scheme_case_closure_type, 19
 scheme_case_sensitive, 128
 SCHEME_CDR, 18
 scheme_cdr, 122
 scheme_change_in_table, 126
 scheme_char_ready, 102
 SCHEME_CHAR_STR_VAL, 17
 scheme_char_string_to_byte_string, 25
 scheme_char_string_to_byte_string_lossy, 25
 scheme_char_string_to_path, 112
 scheme_char_string_type, 17
 SCHEME_CHAR_STRINGP, 17
 SCHEME_CHAR_STRLEN_VAL, 17
 scheme_char_type, 17
 SCHEME_CHAR_VAL, 17
 SCHEME_CHARP, 17
 scheme_check_for_break, 128
 scheme_check_proc_arity, 67
 scheme_check_threads, 78
 scheme_clear_escape, 67
 scheme_close_input_port, 102
 scheme_close_managed, 120
 scheme_close_output_port, 102
 scheme_close_should_force_port_closed, 114
 scheme_closed_prim_type, 19
 scheme_collect_garbage, 46
 scheme_compile, 59
 scheme_compiled_closure_type, 19
 scheme_complex_imaginary_part, 95
 scheme_complex_izi_type, 17
 SCHEME_COMPLEX_IZIP, 17
 scheme_complex_real_part, 95
 scheme_complex_type, 17
 SCHEME_COMPLEXP, 17
 scheme_console_output, 128
 scheme_console_printf, 128
 scheme_cont_type, 19
 scheme_count_lines, 102
 scheme_cpointer_type, 19
 SCHEME_CPTR_TYPE, 19
 SCHEME_CPTR_VAL, 19
 SCHEME_CPTRP, 19
 scheme_current_continuation_marks, 68
 scheme_custodian_check_available, 120
 SCHEME_DBL_VAL, 17
 SCHEME_DBLP, 17
 scheme_debug_print, 99
 scheme_detach_multiple_array, 59
 scheme_directory_exists, 111
 scheme_display, 97
 scheme_display_to_string, 99
 scheme_display_to_string_w_max, 99
 scheme_display_w_max, 98
 scheme_dont_gc_ptr, 46
 scheme_double_to_int, 126

scheme_double_type, 17
 scheme_dynamic_require, 123
 scheme_dynamic_wind, 67
 scheme_end_atomic, 81
 scheme_end_atomic_no_swap, 82
 scheme_eof, 20
 SCHEME_EOF, 20
 scheme_eq, 121
 scheme_equal, 121
 scheme_equal_hash_key, 121
 scheme_equal_hash_key2, 121
 scheme_eqv, 121
 scheme_escaping_cont_type, 19
 scheme_eval, 56
 scheme_eval_compiled, 56
 scheme_eval_compiled_multi, 56
 scheme_eval_string, 58
 scheme_eval_string_all, 58
 scheme_eval_string_multi, 58
 SCHEME_EXACT_INTEGERP, 19
 SCHEME_EXACT_REALP, 19
 scheme_exit, 128
 scheme_expand, 59
 scheme_expand_filename, 111
 scheme_expand_string_filename, 112
 scheme_extend_config, 85
 scheme_false, 20
 SCHEME_FALSE, 20
 scheme_file_exists, 111
 scheme_finish_primitive_module, 50
 scheme_float_type, 17
 SCHEME_FLOAT_VAL, 17
 SCHEME_FLOATP, 19
 SCHEME_FLT_VAL, 17
 SCHEME_FLTP, 17
 scheme_flush_output, 99
 scheme_format, 114
 scheme_format_utf8, 114
 scheme_free_code, 41
 scheme_free_immobile_box, 41
 scheme_gc_ptr_ok, 46
 scheme_get_byte, 100
 scheme_get_byte_string, 100
 scheme_get_byte_string_output, 110
 scheme_get_bytes, 101
 scheme_get_char_string, 101
 scheme_get_env, 50
 scheme_get_fdset, 79
 scheme_get_int_val, 22
 scheme_get_long_long_val, 22
 scheme_get_milliseconds, 127
 scheme_get_param, 84
 scheme_get_port_fd, 103
 scheme_get_port_file_descriptor, 103
 scheme_get_port_socket, 103
 scheme_get_process_milliseconds, 127
 scheme_get_sized_byte_string_output, 110
 scheme_get_thread_param, 85
 scheme_get_unsigned_int_val, 22
 scheme_get_unsigned_long_long_val, 22
 scheme_getc, 100
 scheme_global_bucket, 49
 scheme_hash_get, 125
 scheme_hash_set, 125
 scheme_hash_table_type, 18
 SCHEME_HASHTP, 18
 scheme_inherit_cells, 85
 scheme_init_collection_paths, 129
 scheme_init_collection_paths_post, 129
 SCHEME_INPORT_VAL, 18
 SCHEME_INPORTP, 18
 scheme_input_port_record, 111
 scheme_input_port_type, 18
 SCHEME_INPUT_PORTP, 18
 scheme_install_config, 85
 SCHEME_INT_VAL, 17
 scheme_integer_type, 17
 scheme_intern_exact_char_keyword, 27
 scheme_intern_exact_char_symbol, 26

scheme_intern_exact_keyword, 26
 scheme_intern_exact_symbol, 26
 scheme_intern_symbol, 26
 SCHEME_INTP, 17
 scheme_is_exact, 93
 scheme_is_inexact, 93
 scheme_is_struct_instance, 117
 SCHEME_KEYWORD_LEN, 18
 scheme_keyword_type, 18
 SCHEME_KEYWORD_VAL, 18
 SCHEME_KEYWORDP, 18
 scheme_list_length, 122
 scheme_list_to_vector, 123
 scheme_load, 123
 scheme_load_extension, 123
 scheme_lookup_global, 49
 scheme_lookup_in_table, 126
 scheme_mac_path_to_spec, 113
 scheme_main_setup, 42
 scheme_main_stack_setup, 42
 scheme_make_args_string, 66
 scheme_make_ascii_character, 21
 scheme_make_bignum, 93
 scheme_make_bignum_from_unsigned,
 93
 scheme_make_bucket_table, 125
 scheme_make_byte_string, 23
 scheme_make_byte_string_input_port,
 110
 scheme_make_byte_string_output_port,
 110
 scheme_make_byte_string_without_copying,
 23
 scheme_make_char, 20
 scheme_make_char_or_null, 21
 scheme_make_char_string, 24
 scheme_make_char_string_without_copying,
 24
 scheme_make_character, 21
 scheme_make_closed_prim, 53
 scheme_make_closed_prim_w_arity, 52
 scheme_make_complex, 95
 scheme_make_cptr, 27
 scheme_make_custodian, 119
 scheme_make_double, 22
 scheme_make_exact_symbol, 26
 scheme_make_external_cptr, 27
 scheme_make_fd_input_port, 109
 scheme_make_fd_output_port, 109
 scheme_make_file_input_port, 108
 scheme_make_file_output_port, 108
 scheme_make_float, 22
 scheme_make_folding_prim, 52
 scheme_make_hash_table, 124
 scheme_make_hash_table_equal, 124
 scheme_make_input_port, 103
 scheme_make_integer, 21
 scheme_make_integer_value, 21
 scheme_make_integer_value_from_long_halves,
 21
 scheme_make_integer_value_from_long_long,
 21
 scheme_make_integer_value_from_unsigned,
 21
 scheme_make_integer_value_from_unsigned_long_halves,
 22
 scheme_make_integer_value_from_unsigned_long_long,
 21
 scheme_make_locale_string, 24
 scheme_make_named_file_input_port,
 108
 scheme_make_namespace, 57
 scheme_make_offset_cptr, 28
 scheme_make_offset_external_cptr,
 28
 scheme_make_output_port, 106
 scheme_make_pair, 23
 scheme_make_path, 112
 scheme_make_path_without_copying,
 112
 scheme_make_port_type, 103
 scheme_make_prim, 52
 scheme_make_prim_closure_w_arity,
 52
 scheme_make_prim_w_arity, 51

scheme_make_provided_string, 66
 scheme_make_rational, 94
 scheme_make_sema, 76
 scheme_make_sized_byte_string, 23
 scheme_make_sized_char_string, 25
 scheme_make_sized_offset_byte_strings
 23
 scheme_make_sized_offset_char_strings
 25
 scheme_make_sized_offset_path, 112
 scheme_make_sized_offset_utf8_strings
 24
 scheme_make_sized_path, 112
 scheme_make_sized_utf8_string, 24
 scheme_make_stderr, 128
 scheme_make_stdin, 128
 scheme_make_stdout, 128
 scheme_make_struct_instance, 116
 scheme_make_struct_names, 116
 scheme_make_struct_type, 115
 scheme_make_struct_values, 116
 scheme_make_symbol, 26
 scheme_make_thread_cell, 81
 scheme_make_type, 27
 scheme_make_utf8_string, 24
 scheme_make_vector, 27
 scheme_make_weak_box, 27
 scheme_making_progress, 80
 scheme_malloc, 40
 scheme_malloc_allow_interior, 41
 scheme_malloc_atomic, 40
 scheme_malloc_code, 41
 scheme_malloc_eternal, 40
 scheme_malloc_fail_ok, 41
 scheme_malloc_immobile_box, 41
 scheme_malloc_tagged, 40
 scheme_malloc_uncollectable, 40
 SCHEME_MCAR, 18
 SCHEME_MCDR, 18
 scheme_module_bucket, 49
 SCHEME_MPAIRP, 18
 scheme_mutable_pair_type, 18
 scheme_namespace_require, 123
 scheme_namespace_type, 19
 SCHEME_NAMESPACEP, 19
 scheme_native_closure_type, 19
 scheme_need_wakeup, 102
 scheme_new_param, 85
 scheme_null, 20
 SCHEME_NULLP, 20
 SCHEME_NUMBERP, 19
 scheme_open_input_file, 108
 scheme_open_output_file, 108
 scheme_os_getcwd, 113
 scheme_os_setcwd, 113
 SCHEME_OUTPORT_VAL, 18
 SCHEME_OUTPORTP, 18
 scheme_output_port_record, 111
 scheme_output_port_type, 18
 SCHEME_OUTPUT_PORTP, 18
 scheme_pair_type, 18
 SCHEME_PAIRP, 18
 scheme_param_config, 86
 SCHEME_PATH_LEN, 17
 scheme_path_to_char_string, 112
 scheme_path_type, 17
 SCHEME_PATH_VAL, 17
 SCHEME_PATHP, 17
 scheme_peek_byte, 100
 scheme_peek_byte_skip, 100
 scheme_peekc, 100
 scheme_peekc_skip, 100
 scheme_pipe, 110
 scheme_pipe_with_limit, 111
 scheme_pop_break_enable, 68
 scheme_pop_continuation_frame, 87
 scheme_post_sema, 76
 scheme_prim_type, 19
 scheme_primitive_module, 50
 scheme_print_bytes, 28
 scheme_print_string, 29
 scheme_printf, 114
 scheme_printf_utf8, 114
 SCHEME_PROCP, 19

scheme_proper_list_length, 122
 scheme_push_break_enable, 68
 scheme_push_continuation_frame, 87
 scheme_put_byte_string, 98
 scheme_put_char_string, 99
 scheme_raise_exn, 65
 scheme_rational_denominator, 95
 scheme_rational_from_double, 95
 scheme_rational_from_float, 95
 scheme_rational_numerator, 95
 scheme_rational_to_double, 95
 scheme_rational_to_float, 95
 scheme_rational_type, 17
 SCHEME_RATIONALP, 17
 scheme_read, 97
 scheme_read_bignum, 94
 scheme_read_bignum_bytes, 94
 scheme_real_to_double, 23
 SCHEME_REALP, 19
 scheme_recur_equal, 121
 scheme_recur_equal_hash_key, 121
 scheme_recur_equal_hash_key2, 122
 scheme_register_extension_global,
 42
 scheme_register_finalizer, 44
 scheme_register_parameter, 86
 scheme_register_static, 43
 scheme_remove_all_finalization, 46
 scheme_remove_managed, 120
 scheme_seal_parameters, 129
 scheme_security_check_file, 118
 scheme_security_check_network, 118
 scheme_sema_type, 18
 SCHEME_SEMAP, 18
 scheme_set_addon_path, 128
 scheme_set_box, 123
 scheme_set_can_break, 68
 scheme_set_collects_path, 128
 scheme_set_cont_mark, 87
 scheme_set_file_position, 102
 scheme_set_global_bucket, 49
 scheme_set_param, 84
 scheme_set_stack_base, 42
 scheme_set_stack_bounds, 43
 scheme_set_thread_param, 85
 scheme_set_type_equality, 29
 scheme_set_type_printer, 28
 scheme_signal_error, 64
 scheme_signal_received, 78
 scheme_socket_to_ports, 110
 scheme_strdup, 41
 scheme_strdup_eternal, 41
 scheme_struct_property_type, 18
 scheme_struct_ref, 117
 scheme_struct_set, 117
 scheme_struct_type_type, 18
 SCHEME_STRUCT_TYPEP, 18
 SCHEME_STRUCTP, 18
 scheme_structure_type, 18
 scheme_subtract_finalizer, 45
 scheme_swap_thread, 77
 SCHEME_SYM_LEN, 18
 SCHEME_SYM_VAL, 18
 scheme_symbol_type, 18
 SCHEME_SYMBOLP, 18
 scheme_tail_apply, 58
 scheme_tail_apply_no_copy, 58
 scheme_tail_apply_to_list, 58
 scheme_tell, 102
 scheme_tell_line, 102
 scheme_thread, 75
 scheme_thread_block, 76
 scheme_thread_block_enable_break,
 77
 scheme_thread_type, 18
 scheme_thread_w_details, 76
 SCHEME_THREADP, 18
 scheme_tls_allocate, 81
 scheme_tls_get, 81
 scheme_tls_set, 81
 scheme_true, 20
 SCHEME_TRUEP, 20
 SCHEME_TYPE, 16
 scheme_ucs4_to_utf16, 91

scheme_unbound_global, 66
 scheme_unbox, 123
 scheme_undefined, 20
 scheme_ungetc, 101
 scheme_utf16_to_ucs4, 92
 scheme_utf8_decode, 88
 scheme_utf8_decode_all, 89
 scheme_utf8_decode_as_prefix, 89
 scheme_utf8_decode_count, 90
 scheme_utf8_decode_prefix, 89
 scheme_utf8_decode_to_buffer, 90
 scheme_utf8_decode_to_buffer_len, 90
 scheme_utf8_encode, 90
 scheme_utf8_encode_all, 91
 scheme_utf8_encode_to_buffer, 91
 scheme_utf8_encode_to_buffer_len, 91
 scheme_values, 59
 SCHEME_VEC_ELS, 18
 SCHEME_VEC_SIZE, 18
 scheme_vector_to_list, 123
 scheme_vector_type, 18
 SCHEME_VECTORP, 18
 scheme_version, 127
 scheme_void, 20
 SCHEME_VOIDP, 20
 scheme_wait_sema, 76
 scheme_wake_up, 79
 scheme_warning, 65
 scheme_weak_box_type, 19
 SCHEME_WEAK_PTR, 19
 scheme_weak_reference, 44
 scheme_weak_reference_indirect, 44
 SCHEME_WEAKP, 19
 scheme_write, 97
 scheme_write_byte_string, 98
 scheme_write_char_string, 98
 scheme_write_to_string, 99
 scheme_write_to_string_w_max, 99
 scheme_write_w_max, 97
 scheme_wrong_count, 65
 scheme_wrong_return_arity, 66
 scheme_wrong_type, 66
 Security Guards, 118
 Sleeping by Embedded Scheme, 74
 Standard Types, 17
 String Encodings, 88
 strings, conversion to C, 17
 strings, conversion to C, 17
 Strings, 20
 Structures, 115
 sync, 71
 Tagged Objects, 32
 Tail Evaluation, 54
 tail recursion, 54
 Temporarily Catching Error Escapes, 61
 Thread Functions, 75
 Threads, 69
 Threads in Embedded Scheme with Event Loops, 71
 Top-level Evaluation Functions, 54
 types, creating, 16
 Value Functions, 20
 Values and Types, 16
 Writing MzScheme Extensions, 5
 XFORM_CAN_IGNORE, 39
 XFORM_END_SKIP, 38
 XFORM_END_SUSPEND, 39
 XFORM_END_TRUST_ARITH, 40
 XFORM_HIDE_EXPR, 39
 XFORM_SKIP_PROC, 39
 XFORM_START_SKIP, 38
 XFORM_START_SUSPEND, 39
 XFORM_START_TRUST_ARITH, 40
 XFORM_TRUST_MINUS, 40
 XFORM_TRUST_PLUS, 40