CHICKEN User's Manual - The User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

Chicken User's Manual

1 The User's Manual

(This document describes version 3.0.0)

CHICKEN is a compiler that translates Scheme source files into C, which in turn can be fed to a C-compiler to generate a standalone executable. An interpreter is also available and can be used as a scripting environment or for testing programs before compilation.

This package is distributed under the **BSD license** and as such is free to use and modify.

The method of compilation and the design of the runtime-system follow closely Henry Baker's <u>CONS Should Not CONS Its Arguments, Part II: Cheney on the M.T.A.</u> paper and expose a number of interesting properties:

- Consing (creation of data on the heap) is relatively inexpensive, because a generational garbage collection scheme is used, in which short-lived data structures are reclaimed extremely quickly.
- Moreover, `call-with-current-continuation` is practically for free and CHICKEN does not suffer under any performance penalties if first-class continuations are used in complex ways.

The generated C code is fully tail-recursive.

Some of the features supported by CHICKEN:

- SRFIs 0, 1, 2, 4, 6-19, 23, 25-31, 37-40, 42, 43, 45, 47, 55, 57, 60-63, 66, 69, 72, 78, 85 and 95.
- Lightweight threads based on first-class continuations
- Pattern matching with Andrew Wright's `match` package
- Record structures
- Extended comment- and string-literal syntaxes
- Libraries for regular expressions, string handling
- UNIX system calls and extended data structures
- Create interpreted or compiled shell scripts written in Scheme for UNIX or Windows
- Compiled C files can be easily distributed
- Allows the creation of fully self-contained statically linked executables
- On systems that support it, compiled code can be loaded dynamically

This manual is merely a reference for the CHICKEN system and assumes a working knowledge of Scheme.

The manual is split in the following sections:

<u>Basic mode of operation</u>
        Compiling Scheme files.
<u>Using the compiler</u>
        Explains how to use CHICKEN to compile programs and execute them.
<u>Using the interpreter</u>
        Invocation and usage of `csi`, the CHICKEN interpreter
<u>Supported language</u>
        The language implemented by CHICKEN (deviations from the standard and extensions).
<u>Interface to external functions and variables</u>
        Accessing C and C++ code and data.
<u>chicken-setup</u>
        Packaging and installing extension libraries.
<u>Data representation</u>
        How Scheme data is internally represented.
<u>Bugs and limitations</u>
        Yes, there are some.

FAQ
    A list of Frequently Asked Questions about CHICKEN (and their answers!).
Acknowledgements
    A list of some of the people that have contributed to make CHICKEN what it is.
Bibliography
    Links to documents that may be of interest.

2 Basic mode of operation

The compiler translates Scheme source code into fairly portable C that can be compiled and linked with most available C compilers. CHICKEN supports the generation of executables and libraries, linked either statically or dynamically. Compiled Scheme code can be loaded dynamically, or can be embedded in applications written in other languages. Separate compilation of modules is fully supported.

The most portable way of creating separately linkable entities is supported by so-called *unit*s. A unit is a single compiled object module that contains a number of toplevel expressions that are executed either when the unit is the *main* unit or if the unit is *used*. To use a unit, the unit has to be *declare*d as used, like this:

```
(declare (uses UNITNAME))
```

The toplevel expressions of used units are executed in the order in which the units appear in the `uses` declaration. Units may be used multiple times and `uses` declarations may be circular (the unit is initialized at most once). To compile a file as a unit, add a `unit` declaration:

```
(declare (unit UNITNAME))
```

When compiling different object modules, make sure to have one main unit. This unit is called initially and initializes all used units before executing its toplevel expressions. The main-unit has no `unit` declaration.

Another method of using definitions in separate source files is to *include* them. This simply inserts the code in a given file into the current file:

```
(include "FILENAME")
```

Macro definitions are only available when processed by `include` or `require-for-syntax`. Macro definitions in separate units are not available, since they are defined at compile time, i.e the time when that other unit was compiled (macros can optionally be available at runtime, see `define-macro` in Substitution forms and macros).

On platforms that support dynamic loading of compiled code (Windows, most ELF based systems like Linux or BSD, MacOS X, and others) code can be compiled into a shared object `.dll`, `.so`, `.dylib`) and loaded dynamically into a running application.

Previous: index.html Next: Using the compiler

3 Using the compiler

The interface to `chicken` is intentionally simple. System dependent makefiles, shell-scripts or batch-files should perform any necessary steps before and after invocation of `chicken`. A program named `csc` provides a much simpler interface to the Scheme- and C-compilers and linker. Enter

`csc -help`

on the command line for more information.

3.1 Compiler command line format

`chicken FILENAME {OPTION}`

`FILENAME` is the complete pathname of the source file that is to be translated into C. A filename argument of `-` specifies that the source text should be read from standard input. Note that the filename has to be the first argument to `chicken`.

Possible options are:

-analyze-only
> Stop compilation after first analysis pass.

-benchmark-mode
> Equivalent to `-no-trace -no-lambda-info -optimize-level 3 -fixnum-arithmetic -disable-interrupts -block -lambda-lift`.

-block
> Enable block-compilation. When this option is specified, the compiler assumes that global variables are not modified outside this compilation-unit. Specifically, toplevel bindings are not seen by `eval` and unused toplevel bindings are removed.

-case-insensitive
> Enables the reader to read symbols case insensitive. The default is to read case sensitive (in violation of R5RS). This option registers the `case-insensitive` feature identifier.

-check-imports
> Search for references to undefined global variables. For each library unit accessed via `(declare (uses ...))`, the compiler will search a file named `UNITNAME.exports` in the current include path and load its contents into the *import-table* (if found). Also, export-information for extensions (accessed through `(require-extension ...)`) will be searched and stored in the import-table. If a required extension does not provide explicit export-information a `.exports` file is searched (as with used units). After the analysis phase of the compiler, referenced toplevel variables for which no assignment was found will generate a warning. Also, re-assignments of imported variables will trigger a warning.

-check-syntax
> Aborts compilation process after macro-expansion and syntax checks.

-debug MODES
> Enables one or more compiler debugging modes. `MODES` is a string of characters that select debugging information about the compiler that will be printed to standard output.

```
t        show time needed for compilation
b        show breakdown of time needed for each compiler pass
o        show performed optimizations
r        show invocation parameters
```

```
    s          show program-size information and other statistics
    a          show node-matching during simplification
    p          show execution of compiler sub-passes
    l          show lambda-lifting information
    m          show GC statistics during compilation
    n          print the line-number database
    c          print every expression before macro-expansion
    u          lists all unassigned global variable references
    x          display information about experimental features
    D          when printing nodes, use node-tree output
    N          show the real-name mapping table
    U          show expressions after the secondary user pass
    0          show database before lambda-lifting pass
    L          show expressions after lambda-lifting
    M          show unit-information and syntax-/runtime-requirements
    1          show source expressions
    2          show canonicalized expressions
    3          show expressions converted into CPS
    4          show database after each analysis pass
    5          show expressions after each optimization pass
    6          show expressions after each inlining pass
    7          show expressions after complete optimization
    8          show database after final analysis
    9          show expressions after closure conversion
```

-debug-level LEVEL

Selects amount of debug-information. LEVEL should be an integer.

```
    -debug-level 0             is equivalent to -no-trace -no-lambda-info
    -debug-level 1             is equivalent to -no-trace
    -debug-level 2             does nothing (the default)
```

-disable-interrupts

Equivalent to the (disable-interrupts) declaration. No interrupt-checks are generated for compiled programs.

-disable-compiler-macros

disable expansion of compiler macros.

-disable-stack-overflow-checks

Disables detection of stack overflows. This is equivalent to running the compiled executable with the -:o runtime option.

-disable-warning CLASS : Disables specific class of warnings, may be given multiple times. The following classes are defined

```
    usage          warnings related to command-line arguments
    type           warnings related to type-conversion
    ext            warnings related to extension libraries
    var            warnings related to variable- and syntax-definitions and
    const          warnings related to constant-definitions
    syntax         syntax-related warnings
    redef          warnings about redefinitions of standard- or extended-bi
    call           warnings related to known procedure calls
    ffi            warnings related to the foreign function interface
```

-dynamic

This option should be used when compiling files intended to be loaded dynamically into a running Scheme program.

-epilogue FILENAME

Includes the file named `FILENAME` at the end of the compiled source file. The include-path is not searched. This option may be given multiple times.

-emit-exports FILENAME

Write exported toplevel variables to `FILENAME`.

-emit-external-prototypes-first

Emit prototypes for callbacks defined with `define-external` before any other foreign declarations. This is sometimes useful, when C/C++ code embedded into the a Scheme program has to access the callbacks. By default the prototypes are emitted after foreign declarations.

-explicit-use

Disables automatic use of the units `library, eval` and `extras`. Use this option if compiling a library unit instead of an application unit.

-extend FILENAME

Loads a Scheme source file or compiled Scheme program (on systems that support it) before compilation commences. This feature can be used to extend the compiler. This option may be given multiple times. The file is also searched in the current include path and in the extension-repository.

-extension

Mostly equivalent to `-prelude '(define-extension <NAME>)'`, where `<NAME>` is the basename of the currently compiled file. Note that if you want to compile a file as a normal (dynamically loadable) extension library, you should also pass the `-shared` option.

-feature SYMBOL

Registers `SYMBOL` to be a valid feature identifier for `cond-expand`. Multiple symbols may be given, if comma-separated.

-fixnum-arithmetic

Equivalent to `(fixnum-arithmetic)` declaration. Assume all mathematical operations use small integer arguments.

-heap-size NUMBER

Sets a fixed heap size of the generated executable to `NUMBER` bytes. The parameter may be followed by a `M` (`m`) or `K` (`k`) suffix which stand for mega- and kilobytes, respectively. The default heap size is 5 kilobytes. Note that only half of it is in use at every given time.

-heap-initial-size NUMBER

Sets the size that the heap of the compiled application should have at startup time.

-heap-growth PERCENTAGE

Sets the heap-growth rate for the compiled program at compile time (see: `-:hg`).

-heap-shrinkage PERCENTAGE

Sets the heap-shrinkage rate for the compiled program at compile time (see: `-:hs`).

-help

Print a summary of available options and the format of the command line parameters and exit the compiler.

-import FILENAME

Read exports from linked or loaded libraries from given file. See also `-check-imports`. This is equivalent to declaring `(declare (import FILENAME))`. Implies `-check-imports`.

-include-path PATHNAME

Specifies an additional search path for files included via the `include` special form. This option may be given multiple times. If the environment variable `CHICKEN_INCLUDE_PATH` is set, it should contain a list of alternative include pathnames separated by `;`.

-inline

Enable procedure inlining for known procedures of a size below the threshold (which can be set through the `-inline-limit` option).

-inline-limit THRESHOLD

Sets the maximum size of a potentially inlinable procedure. The default threshold is `10`.

-keyword-style STYLE

Enables alternative keyword syntax, where `STYLE` may be either `prefix` (as in Common Lisp), `suffix` (as in DSSSL) or `none`. Any other value is ignored. The default is `suffix`.

-keep-shadowed-macros

Do not remove macro definitions with the same name as assigned toplevel variables (the default is to remove the macro definition).

-lambda-lift

Enable the optimization known as lambda-lifting.

-no-lambda-info

Don't emit additional information for each `lambda` expression (currently the argument-list, after alpha-conversion/renaming).

-no-trace

Disable generation of tracing information. If a compiled executable should halt due to a runtime error, then a list of the name and the line-number (if available) of the last procedure calls is printed, unless `-no-trace` is specified. With this option the generated code is slightly faster.

-no-warnings

Disable generation of compiler warnings.

-nursery NUMBER

-stack-size NUMBER

Sets the size of the first heap-generation of the generated executable to `NUMBER` bytes. The parameter may be followed by a `M` (`m`) or `K` (`k`) suffix. The default stack-size depends on the target platform.

-optimize-leaf-routines

Enable leaf routine optimization.

-optimize-level LEVEL

Enables certain sets of optimization options. `LEVEL` should be an integer.

```
-optimize-level 0          does nothing.
-optimize-level 1          is equivalent to -optimize-leaf-routines
-optimize-level 2          is currently the same as -optimize-level 1
-optimize-level 3          is equivalent to -optimize-leaf-routines -unsafe
```

-output-file FILENAME

Specifies the pathname of the generated C file. Default is `FILENAME.c`.

-postlude EXPRESSIONS

Add `EXPRESSIONS` after all other toplevel expressions in the compiled file. This option may be given multiple times. Processing of this option takes place after processing of `-epilogue`.

-prelude EXPRESSIONS

Add `EXPRESSIONS` before all other toplevel expressions in the compiled file. This option may be given multiple times. Processing of this option takes place before processing of `-prologue`.

-profile

-accumulate-profile

Instruments the source code to count procedure calls and execution times. After the program terminates (either via an explicit `exit` or implicitly), profiling statistics are written to a file named `PROFILE`. Each line of the generated file contains a list with the procedure name, the number of calls and the time spent executing it. Use the `chicken-profile` program to display the profiling information in a more user-friendly form. Enter `chicken-profile` with no arguments at the command line to get a list of available options. The `-accumulate-profile` option is similar to `-profile`, but the resulting profile information will be appended to any existing `PROFILE` file. `chicken-profile` will merge and sum up the accumulated timing information, if several entries for the same procedure calls exist.

-profile-name FILENAME

Specifies name of the generated profile information (which defaults to `PROFILE`. Implies `-profile`.

-prologue FILENAME

Includes the file named `FILENAME` at the start of the compiled source file. The include-path is not

searched. This option may be given multiple times.

-quiet

Disables output of compile information.

-raw

Disables the generation of any implicit code that uses the Scheme libraries (that is all runtime system files besides `runtime.c` and `chicken.h`).

-require-extension NAME

Loads the extension `NAME` before the compilation process commences. This is identical to adding (`require-extension NAME`) at the start of the compiled program. If `-uses NAME` is also given on the command line, then any occurrences of `-require-extension NAME` are replaced with (`declare (uses NAME)`). Multiple names may be given and should be separated by `,`.

-run-time-macros

Makes macros also available at run-time. By default macros are not available at run-time.

-to-stdout

Write compiled code to standard output instead of creating a `.c` file.

-unit NAME

Compile this file as a library unit. Equivalent to `-prelude "(declare (unit NAME))"`

-unsafe

Disable runtime safety checks.

-unsafe-libraries

Marks the generated file for being linked with the unsafe runtime system. This should be used when generating shared object files that are to be loaded dynamically. If the marker is present, any attempt to load code compiled with this option will signal an error.

-uses NAME

Use definitions from the library unit `NAME`. This is equivalent to `-prelude "(declare (uses NAME))"`. Multiple arguments may be given, separated by `,`.

-no-usual-integrations

Specifies that standard procedures and certain internal procedures may be redefined, and can not be inlined. This is equivalent to declaring (`not usual-integrations`).

-version

Prints the version and some copyright information and exit the compiler.

-verbose

Prints progress information to standard output during compilation.

The environment variable `CHICKEN_OPTIONS` can be set to a string with default command-line options for the compiler.

3.2 Runtime options

After successful compilation a C source file is generated and can be compiled with a C compiler. Executables generated with CHICKEN (and the compiler itself) accept a small set of runtime options:

`-:?`

Shows a list of the available runtime options and exits the program.

`-:aNUMBER`

Specifies the length of the buffer for recording a trace of the last invoked procedures. Defaults to 16.

`-:b`

Enter a read-eval-print-loop when an error is encountered.

`-:B`

Sounds a bell (ASCII 7) on every major garbage collection.

`-:c`

Forces console mode. Currently this is only used in the interpreter (`csi`) to force output of the `#;N>` prompt even if stdin is not a terminal (for example if running in an `emacs` buffer under Windows).

`-:d`

Prints some debug-information at runtime.

`-:D`

Prints some more debug-information at runtime.

`-:fNUMBER`

Specifies the maximal number of currently pending finalizers before finalization is forced.

`-:hNUMBER`

Specifies fixed heap size

`-:hgPERCENTAGE`

Sets the growth rate of the heap in percent. If the heap is exhausted, then it will grow by `PERCENTAGE`. The default is 200.

`-:hiNUMBER`

Specifies the initial heap size

`-:hmNUMBER`

Specifies a maximal heap size. The default is (2GB - 15).

`-:hsPERCENTAGE`

Sets the shrink rate of the heap in percent. If no more than a quarter of `PERCENTAGE` of the heap is used, then it will shrink to `PERCENTAGE`. The default is 50. Note: If you want to make sure that the heap never shrinks, specify a value of `0`. (this can be useful in situations where an optimal heap-size is known in advance).

`-:o`

Disables detection of stack overflows at run-time.

`-:r`

Writes trace output to stderr. This option has no effect with in files compiled with the `-no-trace` options.

`-:sNUMBER`

Specifies stack size.

`-:tNUMBER`

Specifies symbol table size.

`-:w`

Enables garbage collection of unused symbols. By default unused and unbound symbols are not garbage collected.

`-:x`

Raises uncaught exceptions of separately spawned threads in primordial thread. By default uncaught exceptions in separate threads are not handled, unless the primordial one explicitly joins them. When warnings are enabled (the default) and `-:x` is not given, a warning will be shown, though.

The argument values may be given in bytes, in kilobytes (suffixed with `K` or `k`), in megabytes (suffixed with `M` or `m`), or in gigabytes (suffixed with `G` or `g`). Runtime options may be combined, like `-:dc`, but everything following a `NUMBER` argument is ignored. So `-:wh64m` is OK, but `-:h64mw` will not enable GC of unused symbols.

3.3 Examples

3.3.1 A simple example (with one source file)

To compile a Scheme program (assuming a UNIX-like environment) consisting of a single source file, perform the following steps.

3.3.1.1 Writing your source file

In this example we will assume your source file is called `foo.scm`:

```scheme
;;; foo.scm

(define (fac n)
  (if (zero? n)
      1
      (* n (fac (- n 1))) ) )

(write (fac 10))
(newline)
```

3.3.1.2 Compiling your program

Compile the file `foo.scm`:

```
% csc foo.scm
```

This will produce the `foo` executable:

```
% ls
foo  foo.scm
```

3.3.1.3 Running your program

To run your newly compiled executable use:

```
% foo
3628800
```

If you get a `foo: command not found` error, you might want to try with `./foo` instead (or, in Unix machines, modify your `PATH` environment variable to include your current directory).

3.3.2 An example with multiple files

If multiple bodies of Scheme code are to be combined into a single executable, then we have to compile each file and link the resulting object files together with the runtime system.

Let's consider an example where your program consists of multiple source files.

3.3.2.1 Writing your source files

The declarations in these files specify which of the compiled files is the main module, and which is the library module. An executable can only have one main module, since a program has only a single entry-point. In this case `foo.scm` is the main module, because it doesn't have a `unit` declaration:

```scheme
;;; foo.scm

; The declaration marks this source file as dependant on the symbols provided
; by the bar unit:
(declare (uses bar))

(write (fac 10)) (newline)
```

`bar.scm` will be our library:

```scheme
;;; bar.scm

; The declaration marks this source file as the bar unit.  The names of the
; units and your files don't need to match.
(declare (unit bar))

(define (fac n)
  (if (zero? n)
      1
      (* n (fac (- n 1))) ) )
```

3.3.2.2 Compiling and running your program

You should compile your two files with the following commands:

```
% csc -c bar.scm
% csc -c foo.scm
```

That should produce two files, `bar.o` and `foo.o`. They contain the code from your source files in compiled form.

To link your compiled files use the following command:

```
% csc foo.o bar.o -o foo
```

This should produce the `foo` executable, which you can run just as in the previous example. At this point you

can also erase the `*.o` files.

You could avoid one step and link the two files just as `foo.scm` is compiled:

```
% csc -c bar.scm
% csc foo.scm bar.o -o foo
```

Note that if you want to distribute your program, you might want it to follow the GNU Coding Standards. One relatively easy way to achieve this is to use Autoconf and Automake, two tools made for this specific purpose.

3.4 Extending the compiler

The compiler supplies a couple of hooks to add user-level passes to the compilation process. Before compilation commences any Scheme source files or compiled code specified using the `-extend` option are loaded and evaluated. The parameters `user-options-pass`, `user-read-pass`, `user-preprocessor-pass`, `user-pass`, `user-pass-2` and `user-post-analysis-pass` can be set to procedures that are called to perform certain compilation passes instead of the usual processing (for more information about parameters see: Supported language.

[parameter] user-options-pass
>   Holds a procedure that will be called with a list of command-line arguments and should return two values: the source filename and the actual list of options, where compiler switches have their leading `-` (hyphen) removed and are converted to symbols. Note that this parameter is invoked **before** processing of the `-extend` option, and so can only be changed in compiled user passes.

[parameter] user-read-pass
>   Holds a procedure of three arguments. The first argument is a list of strings with the code passed to the compiler via `-prelude` options. The second argument is a list of source files including any files specified by `-prologue` and `-epilogue`. The third argument is a list of strings specified using `-postlude` options. The procedure should return a list of toplevel Scheme expressions.

[parameter] user-preprocessor-pass
>   Holds a procedure of one argument. This procedure is applied to each toplevel expression in the source file **before** macro-expansion. The result is macro-expanded and compiled in place of the original expression.

[parameter] user-pass
>   Holds a procedure of one argument. This procedure is applied to each toplevel expression **after** macro-expansion. The result of the procedure is then compiled in place of the original expression.

[parameter] user-pass-2
>   Holds a procedure of three arguments, which is called with the canonicalized node-graph as its sole argument. The result is ignored, so this pass has to mutate the node-structure to cause any effect.

[parameter] user-post-analysis-pass
>   Holds a procedure that will be called after every performed program analysis pass. The procedure (when defined) will be called with seven arguments: a symbol indicating the analysis pass, the program database, the current node graph, a getter and a setter-procedure which can be used to access and manipulate the program database, which holds various information about the compiled program, a pass iteration count, and an analysis continuation flag. The getter procedure should be called with two arguments: a symbol representing the binding for which information should be retrieved, and a symbol that specifies the database-entry. The current value of the database entry will be returned or `#f`, if no such entry is available. The setter procedure is called with three arguments: the symbol and key and the new value. The pass iteration count currently is meaningful only for the 'opt pass. The analysis continuation flag will be `#f` for the last 'opt pass. For information about the contents of the program database contact the author.

Loaded code (via the `-extend` option) has access to the library units `extras, srfi-1, srfi-4, utils, regex` and the pattern matching macros. Multithreading is not available.

Note that the macroexpansion/canonicalization phase of the compiler adds certain forms to the source program. These extra expressions are not seen by `user-preprocessor-pass` but by `user-pass`.

3.5 Distributing compiled C files

It is relatively easy to create distributions of Scheme projects that have been compiled to C. The runtime system of CHICKEN consists of only two handcoded C files (`runtime.c` and `chicken.h`), plus the file `chicken-config.h`, which is generated by the build process. All other modules of the runtime system and the extension libraries are just compiled Scheme code. The following example shows a minimal application, which should run without changes on the most frequent operating systems, like Windows, Linux or FreeBSD:

Let's take a simple example.

```
; hello.scm

(print "Hello, world!")
```

```
 % chicken hello.scm -optimize-level 3 -output-file hello.c
```

Compiled to C, we get `hello.c`. We need the files `chicken.h` and `runtime.c`, which contain the basic runtime system, plus the three basic library files `library.c`, `eval.c` and `extras.c` which contain the same functionality as the library linked into a plain CHICKEN-compiled application, or which is available by default in the interpreter, `csi`:

```
 % cd /tmp
 %echo '(print "Hello World.")' > hello.scm
 % cp $CHICKEN_BUILD/runtime.c .
 % cp $CHICKEN_BUILD/library.c .
 % cp $CHICKEN_BUILD/eval.c    .
 % cp $CHICKEN_BUILD/extras.c  .
 % gcc -static -Os -fomit-frame-pointer runtime.c library.c eval.c \
   extras.c hello.c -o hello -lm
```

Now we have all files together, and can create an tarball containing all the files:

```
% tar cf hello.tar Makefile hello.c runtime.c library.c eval.c extras.c chicken
% gzip hello.tar
```

This is naturally rather simplistic. Things like enabling dynamic loading, estimating the optimal stack-size and selecting supported features of the host system would need more configuration- and build-time support. All this can be addressed using more elaborate build-scripts, makefiles or by using autoconf/automake.

Note also that the size of the application can still be reduced by removing `extras` and `eval` and compiling `hello.scm` with the `-explicit-use` option.

For more information, study the CHICKEN source code and/or get in contact with the author.

Next: <u>Using the interpreter</u>

4 Using the interpreter

CHICKEN provides an interpreter named `csi` for evaluating Scheme programs and expressions interactively.

4.1 Interpreter command line format

`csi {FILENAME|OPTION}`

where `FILENAME` specifies a file with Scheme source-code. If the extension of the source file is `.scm`, it may be omitted. The runtime options described in Compiler command line format are also available for the interpreter. If the environment variable `CSI_OPTIONS` is set to a list of options, then these options are additionally passed to every direct or indirect invocation of `csi`. Please note that runtime options (like `-:...`) can not be passed using this method. The options recognized by the interpreter are:

--
> Ignore everything on the command-line following this marker. Runtime options (`-:...`) are still recognized.

-i -case-insensitive
> Enables the reader to read symbols case insensitive. The default is to read case sensitive (in violation of R5RS). This option registers the `case-insensitive` feature identifier.

-b -batch
> Quit the interpreter after processing all command line options.

-e -eval EXPRESSIONS
> Evaluate `EXPRESSIONS`. This option implies `-batch` and `-quiet`, so no startup message will be printed and the interpreter exits after processing all `-eval` options and/or loading files given on the command-line.

-p -print EXPRESSIONS
> Evaluate `EXPRESSIONS` and print the results of each expression using `print`. Implies `-batch` and `-quiet`.

-P -pretty-print EXPRESSIONS
> Evaluate `EXPRESSIONS` and print the results of each expression using `pretty-print`. Implies `-batch` and `-quiet`.

-D -feature SYMBOL
> Registers `SYMBOL` to be a valid feature identifier for `cond-expand` and `feature?`.

-h -help
> Write a summary of the available command line options to standard output and exit.

-I -include-path PATHNAME
> Specifies an alternative search-path for files included via the `include` special form. This option may be given multiple times. If the environment variable `CHICKEN_INCLUDE_PATH` is set, it should contain a list of alternative include pathnames separated by `;`.

-k -keyword-style STYLE
> Enables alternative keyword syntax, where `STYLE` may be either `prefix` (as in Common Lisp) or `suffix` (as in DSSSL). Any other value is ignored.

-n -no-init
> Do not load initialization-file. If this option is not given and the file `./.csirc` or `$HOME/.csirc` exists, then it is loaded before the read-eval-print loop commences.

-w -no-warnings
> Disables any warnings that might be issued by the reader or evaluated code.

-q -quiet
> Do not print a startup message. Also disables generation of call-trace information for interpreted code.

-s -script PATHNAME

This is equivalent to `-batch -quiet -no-init` PATHNAME. Arguments following PATHNAME are available by using `command-line-arguments` and are not processed as interpreter options. Extra options in the environment variable `CSI_OPTIONS` are ignored.

-ss PATHNAME

The same as `-s` PATHNAME but invokes the procedure `main` with the value of `(command-line-arguments)` as its single argument. If the main procedure returns an integer result, then the interpreter is terminated, returning the integer as the status code back to the invoking process. Any other result terminates the interpreter with a zero exit status.

-R -require-extension NAME

Equivalent to evaluating `(require-extension NAME)`.

-v -version

Write the banner with version information to standard output and exit.

4.2 Writing Scheme scripts

Since UNIX shells use the `#!` notation for starting scripts, anything following the characters `#!` is ignored, with the exception of the special symbols `#!optional, #!key, #!rest` and `#!eof`.

The easiest way is to use the `-script` option like this:

```
% cat foo
#! /usr/local/bin/csi -script
(print (eval (with-input-from-string
              (car (command-line-arguments))
               read)))

% chmod +x foo
% foo "(+ 3 4)"
7
```

The parameter `command-line-arguments` is set to a list of the parameters that were passed to the Scheme script. Scripts can be compiled to standalone executables (don't forget to declare used library units).

CHICKEN supports writing shell scripts in Scheme for these platforms as well, using a slightly different approach. The first example would look like this on Windows:

```
C:>type foo.bat
@;csibatch %0 %1 %2 %3 %4 %5 %6 %7 %8 %9
(print (eval (with-input-from-string
              (car (command-line-arguments))
              read)))

C:>foo "(+ 3 4)"
7
```

Like UNIX scripts, batch files can be compiled. Windows batch scripts do not accept more than 8 arguments.

Since it is sometimes useful to run a script into the interpreter without actually running it (for example to test specific parts of it), the option `-ss` can be used as an alternative to `-script`. `-ss` PATHNAME is equivalent to `-script` PATHNAME but invokes `(main (command-line-arguments))` after loading all top-level forms of the script file. The result of `main` is returned as the exit status to the shell. Any non-numeric result exits with status zero:

```
% cat hi.scm
(define (main args)
  (print "Hi, " (car args))
  0)
% csi -ss hi.scm you
Hi, you
% csi -q
#;1> ,l hi.scm
#;2> (main (list "ye all"))
Hi, ye all
0
#;3>
```

4.3 Toplevel commands

The toplevel loop understands a number of special commands:

,?

> Show summary of available toplevel commands.

,l FILENAME ...

> Load files with given FILENAMEs

,ln FILENAME ...

> Load files and print result(s) of each top-level expression.

,p EXP

> Pretty-print evaluated expression EXP.

,d EXP

> Describe result of evaluated expression EXP.

,du EXP

> Dump contents of the result of evaluated expression EXP.

,dur EXP N

> Dump N bytes of the result of evaluated expression EXP.

,exn

> Describes the last exception that occurred and adds it to the result history (it can be accessed using the # notation).

,q

> Quit the interpreter.

,r

> Show system information.

,s TEXT ...

> Execute shell-command.

,t EXP

> Evaluate form and print elapsed time.

,x EXP

> Pretty-print macroexpanded expression EXP (the expression is not evaluated).

,tr SYMBOL ...

> Enables tracing of the toplevel procedures with the given names.

```
#;1> (fac 10)                              ==> 3628800
#;2> ,tr fac
#;3> (fac 3)
|(fac 3)
| (fac 2)
```

```
|   (fac 1)
|    (fac 0)
|    fac -> 1
|   fac -> 1
| fac -> 2
|fac -> 6                              ==> 6
#;4> ,utr fac
#;5> (fac 3)                          ==> 6
```

k

,utr SYMBOL ...
>   Disables tracing of the given toplevel procedures.

,br SYMBOL ...
>   Sets a breakpoint at the procedures named SYMBOL .... Breakpoint can also be trigged using the breakpoint procedure.

,ubr SYMBOL ...
>   Removes breakpoints.

,c
>   Continues execution from the last invoked breakpoint.

,breakall
>   Enable breakpoints for all threads (this is the default).

,breakonly THREAD
>   Enable breakpoints only for the thread returned by the expression THREAD.

,info
>   Lists traced procedures and breakpoints.

,step EXPR
>   Evaluates EXPR in single-stepping mode. On each procedure call you will be presented with a menu that allows stepping to the next call, leaving single-stepping mode or triggering a breakpoint. Note that you will see some internal calls, and unsafe or heavily optimized compiled code might not be stepped at all. Single-stepping mode is also possible by invoking the singlestep procedure.

You can define your own toplevel commands using the toplevel-command procedure:

4.4 toplevel-command

[procedure] (toplevel-command SYMBOL PROC [HELPSTRING])

Defines or redefines a toplevel interpreter command which can be invoked by entering ,SYMBOL. PROC will be invoked when the command is entered and may read any required argument via read (or read-line). If the optional argument HELPSTRING is given, it will be listed by the ,? command.

4.5 History access

The interpreter toplevel accepts the special object #[INDEX] which returns the result of entry number INDEX in the history list. If the expression for that entry resulted in multiple values, the first result (or an unspecified value for no values) is returned. If no INDEX is given (and if a whitespace or closing paranthesis character follows the #, then the result of the last expression is returned. Note that the value returned is implicitly quoted.

4.6 set-describer!

[procedure] (set-describer! TAG PROC)

Sets a custom description handler that invokes PROC when the ,d command is invoked with a record-type object that has the type TAG (a symbol). PROC is called with two arguments: the object to be described and an output-port. It should write a possibly useful textual description of the object to the passed output-port. For example:

```
#;1> (define-record point x y)
#;2> (set-describer! 'point
        (lambda (pt o)
          (print "a point with x=" (point-x pt) " and y=" (point-y pt))))
#;3> ,d (make-point 1 2)
a point with x=1 and y=2
```

4.7 Auto-completion and edition

On platforms that support it, it is possible to get auto-completion of symbols, history (over different csi sessions) and a more feature-full editor for the expressions you type using the http://www.call-with-current-continuation.org/eggs/readline.html egg by Tony Garnock Jones. It is very useful for interactive use of csi.

To enable it install the egg and put this in your ~/.csirc file:

```
(use readline regex)
(current-input-port (make-gnu-readline-port))
(gnu-history-install-file-manager
  (string-append (or (getenv "HOME") ".") "/.csi.history"))
```

More details are available in the egg's documentation.

4.8 Accessing documentation

You can access the manual directly from csi using the man extension by Mario Domenech Goulart.

To enable it install the egg and put this in your ~/.csirc file:

```
(use man)
(man:load)
```

Then, in csi, you can search for definitions using man:search as in:

```
(man:search "case")
```

Note that the search uses regular expressions. To view the documentation for one entry from the manual, use `man:help` as in:

```
(man:help "case-lambda")
```

Note: Currently the documentation provided by the `man` extension corresponds to Chicken's 2.429, one of the last releases whose documentation was in the texinfo format (the format the `man` extension parses).

Previous: Using the compiler

Next: Supported language

5 Supported language

- Deviations from the standard
- Extensions to the standard
- Non-standard read syntax
- Non-standard macros and special forms
- Pattern matching
- Declarations
- Parameters
- Unit library basic Scheme definitions
- Unit eval evaluation and macro-handling
- Unit extras useful utility definitions
- Unit srfi-1 List Library
- Unit srfi-4 Homogeneous numeric vectors
- Unit srfi-13 String library
- Unit srfi-14 character set library
- Unit match pattern matching runtime-support
- Unit regex regular expressions
- Unit srfi-18 multithreading
- Unit posix Unix-like services
- Unit utils Shell scripting and file operations
- Unit tcp basic TCP-sockets
- Unit lolevel low-level operations

Previous: Using the interpreter

Next: Interface to external functions and variables

6 Deviations from the standard

Identifiers are by default case-sensitive (see <u>Compiler command line format</u>).

[4.1.3] The maximal number of arguments that may be passed to a compiled procedure or macro is 120. A macro-definition that has a single rest-parameter can have any number of arguments. If the `libffi` library is available on this platform, and if it is installed, then CHICKEN can take advantage of this. See the <u>README</u> file for more details.

[4.2.2] `letrec` does evaluate the initial values for the bound variables sequentially and not in parallel, that is:

```
(letrec ((x 1) (y 2)) (cons x y))
```

is equivalent to

```
(let ((x (void)) (y (void)))
  (set! x 1)
  (set! y 2)
  (cons x y) )
```

where R5RS requires

```
(let ((x (void)) (y (void)))
  (let ((tmp1 1) (tmp2 2))
    (set! x tmp1)
    (set! y tmp2)
    (cons x y) ) )
```

[4.3] `syntax-rules` macros are not provided but available separately.

[6.1] `equal?` compares all structured data recursively, while R5RS specifies that `eqv?` is used for data other than pairs, strings and vectors.

[6.2.4] The runtime system uses the numerical string-conversion routines of the underlying C library and so does only understand standard (C-library) syntax for floating-point constants.

[6.2.5] There is no built-in support for rationals, complex numbers or extended-precision integers (bignums). The routines `complex?`, `real?` and `rational?` are identical to the standard procedure `number?`. The procedures `numerator`, `denominator`, `rationalize`, `make-rectangular` and `make-polar` are not implemented. Fixnums are limited to $±2^{30}$ (or $±2^{62}$ on 64-bit hardware). Support for extended numbers is available as a separate package, provided the GNU multiprecision library is installed.

[6.2.6] The procedure `string->number` does not obey read/write invariance on inexact numbers.

[6.4] The maximum number of values that can be passed to continuations captured using `call-with-current-continuation` is 120.

[6.5] Code evaluated in `scheme-report-environment` or `null-environment` still sees non-standard syntax.

[6.6.2] The procedure `char-ready?` always returns `#t` for terminal ports. The procedure `read` does not obey read/write invariance on inexact numbers.

[6.6.3] The procedures `write` and `display` do not obey read/write invariance to inexact numbers.

[6.6.4] The `transcript-on` and `transcript-off` procedures are not implemented.

Previous: Supported language

Next: Extensions to the standard

7 Extensions to the standard

[2.1] Identifiers may contain special characters if delimited with `|  ...  |`.

[2.3] The brackets `[  ...  ]` and the braces `{  ...  }` are provided as an alternative syntax for `(  ...  )`. A number of reader extensions is provided. See <u>Non-standard read syntax</u>.

[4] Numerous non-standard macros are provided. See <u>Non-standard macros and special forms</u> for more information.

[4.1.4] Extended DSSSL style lambda lists are supported. DSSSL parameter lists are defined by the following grammar:

```
<parameter-list> ==> <required-parameter>*
                     [(#!optional <optional-parameter>*)]
                     [(#!rest <rest-parameter>)]
                     [(#!key <keyword-parameter>*)]
<required-parameter> ==> <ident>
<optional-parameter> ==> <ident>
                         | (<ident> <initializer>)
<rest-parameter> ==> <ident>
<keyword-parameter> ==> <ident>
                        | (<ident> <initializer>)
<initializer> ==> <expr>
```

When a procedure is applied to a list of arguments, the parameters and arguments are processed from left to right as follows:

- Required-parameters are bound to successive arguments starting with the first argument. It shall be an error if there are fewer arguments than required-parameters.
- Next, the optional-parameters are bound with the remaining arguments. If there are fewer arguments than optional-parameters, then the remaining optional-parameters are bound to the result of the evaluation of their corresponding <initializer>, if one was specified, otherwise `#f`. The corresponding <initializer> is evaluated in an environment in which all previous parameters have been bound.
- If there is a rest-parameter, then it is bound to a list containing all the remaining arguments left over after the argument bindings with required-parameters and optional-parameters have been made.
- If `#!key` was specified in the parameter-list, there should be an even number of remaining arguments. These are interpreted as a series of pairs, where the first member of each pair is a keyword specifying the parameter name, and the second member is the corresponding value. If the same keyword occurs more than once in the list of arguments, then the corresponding value of the first keyword is the binding value. If there is no argument for a particular keyword-parameter, then the variable is bound to the result of evaluating <initializer>, if one was specified, otherwise `#f`. The corresponding <initializer> is evaluated in an environment in which all previous parameters have been bound.

Needing a special mention is the close relationship between the rest-parameter and possible keyword-parameters. Declaring a rest-parameter binds up all remaining arguments in a list, as described above. These same remaining arguments are also used for attempted matches with declared keyword-parameters, as described above, in which case a matching keyword-parameter binds to the corresponding value argument at the same time that both the keyword and value arguments are added to the rest parameter list. Note that for efficiency reasons, the keyword-parameter matching does nothing more than simply attempt to match with pairs that may exist in the remaining arguments. Extra arguments that don't match are simply unused and forgotten if no rest-parameter has been declared. Because of this, the caller of a procedure containing one or more keyword-parameters cannot rely on any kind of system error to report

wrong keywords being passed in.

It shall be an error for an `<ident>` to appear more than once in a parameter-list.

If there is no rest-parameter and no keyword-parameters in the parameter-list, then it shall be an error for any extra arguments to be passed to the procedure.

Example:

```
((lambda x x) 3 4 5 6)         => (3 4 5 6)
((lambda (x y #!rest z) z)
 3 4 5 6)                      => (5 6)
((lambda (x y #!optional z #!rest r #!key i (j 1))
    (list x y z i: i j: j))
 3 4 5 i: 6 i: 7)             => (3 4 5 i: 6 j: 1)
```

[4.1.6] `set!` for unbound toplevel variables is allowed. `set! (PROCEDURE ...) ...)` is supported, as CHICKEN implements SRFI-17. [4.2.1] The `cond` form supports SRFI-61.

[4.2.2] It is allowed for initialization values of bindings in a `letrec` construct to refer to previous variables in the same set of bindings, so

```
(letrec ((foo 123)
         (bar foo) )
  bar)
```

is allowed and returns `123`.

[4.2.3] `(begin)` is allowed in non-toplevel contexts and evaluates to an unspecified value.

[4.2.5] Delayed expressions may return multiple values.

[5.2.2] CHICKEN extends standard semantics by allowing internal definitions everywhere, and not only at the beginning of a body. A set of internal definitions is equivalent to a `letrec` form enclosing all following expressions in the body:

```
(let ((foo 123))
  (bar)
  (define foo 456)
  (baz foo) )
```

expands into

```
(let ((foo 123))
  (bar)
  (letrec ((foo 456))
    (baz foo) ) )
```

[5.2] `define` with a single argument is allowed and initializes the toplevel or local binding to an unspecified value. CHICKEN supports *curried* definitions, where the variable name may also be a list specifying a name and a nested lambda list. So

```
(define ((make-adder x) y) (+ x y))
```

is equivalent to

```
(define (make-adder x) (lambda (y) (+ x y)))
```

[6] CHICKEN provides numerous non-standard procedures. See the manual sections on library units for more information.

[6.2.4] The special IEEE floating-point numbers *+nan*, *+inf* and *-inf* are supported, as is negative zero.

[6.3.4] User defined character names are supported. See `char-name`. Characters can be given in hexadecimal notation using the *#\xXX* syntax where *XX* specifies the character code. Character codes above 255 are supported and can be read (and are written) using the *#\uXXXX* and *#\UXXXXXXXX* notations.

Non-standard characters names supported are `#\tab`, `#\linefeed`, `#\return`, `#\alarm`, `#\vtab`, `#\nul`, `#\page`, `#\esc`, `#\delete` and `#\backspace`.

[6.3.5] CHICKEN supports special characters preceded with a backslash \ in quoted string constants. *\n* denotes the newline-character, *\r* carriage return, *\b* backspace, *\t* TAB, *\v* vertical TAB, *\a* alarm, *\f* formfeed, *\xXX* a character with the code XX in hex and *\uXXXX* (and *\UXXXXXXXX*) a unicode character with the code `XXXX`. The latter is encoded in UTF-8 format.

The third argument to `substring` is optional and defaults to the length of the string.

[6.4] `force` called with an argument that is not a promise returns that object unchanged. Captured continuations can be safely invoked inside before- and after-thunks of a `dynamic-wind` form and execute in the outer dynamic context of the `dynamic-wind` form.

Implicit non-multival continuations accept multiple values by discarding all but the first result. Zero values result in the continuation receiving an unspecified value. Note that this slight relaxation of the behaviour of returning mulitple values to non-multival continuations does not apply to explicit continuations (created with `call-with-current-continuation`).

[6.5] The second argument to `eval` is optional and defaults to the value of `(interaction-environment)`. `scheme-report-environment` and `null-environment` accept an optional 2nd parameter: if not `#f` (which is the default), toplevel bindings to standard procedures are mutable and new toplevel bindings may be introduced.

[6.6] The *tilde* character (~) is automatically expanded in pathnames. Additionally, if a pathname starts with `$VARIABLE...`, then the prefix is replaced by the value of the given environment variable.

[6.6.1] if the procedures `current-input-port` and `current-output-port` are called with an argument (which should be a port), then that argument is selected as the new current input- and output-port, respectively. The procedures `open-input-file`, `open-output-file`, `with-input-from-file`, `with-output-to-file`, `call-with-input-file` and `call-with-output-file` accept an optional second (or third) argument which should be one or more keywords, if supplied. These arguments specify the mode in which the file is opened. Possible values are the keywords `#:text`, `#:binary` or `#:append`.

[6.7] The `exit` procedure exits a program right away and does *not* invoke pending `dynamic-wind` thunks.

Previous:

Next:

8 Non-standard read syntax

8.1 Multiline Block Comment

`#| ... |#`

A multiline *block* comment. May be nested. Implements <u>SRFI-30</u>

8.2 Expression Comment

`#;EXPRESSION`

Treats `EXPRESSION` as a comment.

8.3 External Representation

`#,(CONSTRUCTORNAME DATUM ...)`

Allows user-defined extension of external representations. (For more information see the documentation for <u>SRFI-10</u>)

8.4 Syntax Expression

`#'EXPRESSION`

An abbreviation for `(syntax EXPRESSION)`.

8.5 Location Expression

`#$EXPRESSION`

An abbreviation for `(location EXPRESSION)`.

8.6 Keyword

```
#:SYMBOL
```

Syntax for keywords. Keywords are symbols that evaluate to themselves, and as such don't have to be quoted.

8.7 Multiline String Constant

```
#<<TAG
```

Specifies a multiline string constant. Anything up to a line equal to `TAG` (or end of file) will be returned as a single string:

```
(define msg #<<END
 "Hello, world!", she said.
END
)
```

is equivalent to

```
(define msg "\"Hello, world!\", she said.")
```

8.8 Multiline String Constant with Embedded Expressions

```
#<#TAG
```

Similar to `#<<`, but allows substitution of embedded Scheme expressions prefixed with `#` and optionally enclosed in curly brackets. Two consecutive `#`s are translated to a single `#`:

```
(define three 3)
(display #<#EOF
This is a simple string with an embedded `##' character
and substituted expressions: (+ three 99) ==> #(+ three 99)
(three is "#{three}")
EOF
)
```

prints

```
This is a simple string with an embedded `#' character
and substituted expressions: (+ three 99) ==> 102
(three is "3")
```

8.9 Foreign Declare

`#> ... <#`

Abbreviation for `foreign-declare " ... ")`.

8.10 Sharp Prefixed Symbol

`#%...`

Reads like a normal symbol.

8.11 Bang

`#!...`

Interpretation depends on the directly following characters. Only the following are recognized. Any other case results in a read error.

8.11.1 Line Comment

- If followed by whitespace or a slash, then everything up the end of the current line is ignored

8.11.2 Eof Object

- If followed by the character sequence `eof`, then the (self-evaluating) end-of-file object is returned

8.11.3 DSSSL Formal Parameter List Annotation

- If followed by any of the character sequences `optional`, `rest` or `key`, then a symbol with the same name (and prefixed with `#!`) is returned

8.11.4 Read Mark Invocation

- If a *read mark* with the same name as the token is registered, then its procedure is called and the result of the read-mark procedure will be returned

8.12 Case Sensitive Expression

```
#cs...
```

Read the next expression in case-sensitive mode (regardless of the current global setting).

8.13 Case Insensitive Expression

```
#ci...
```

Read the next expression in case-insensitive mode (regardless of the current global setting).

8.14 Conditional Expansion

```
#+FEATURE EXPR
```

Equivalent to

```
(cond-expand (FEATURE EXPR) (else))
```

Previous: Extensions to the standard

Next: Non-standard macros and special forms

9 Non-standard macros and special forms

9.1 Making extra libraries and extensions available

9.1.1 require-extension

```
[syntax] (require-extension ID ...)
[syntax] (use ID ...)
```

This form does all the necessary steps to make the libraries or extensions given in `ID ...` available. It loads syntactic extensions, if needed and generates code for loading/linking with core library modules or separately installed extensions. `use` is just a shorter alias for `require-extension`. This implementation of `require-extension` is compliant with SRFI-55 (see the SRFI-55 document for more information).

During interpretation/evaluation `require-extension` performs one of the following:

- If `ID` names a built-in feature `chicken srfi-0 srfi-2 srfi-6 srfi-8 srfi-9 srfi-10 srfi-17 srfi-23 srfi-30 srfi-39 srfi-55`, then nothing is done.
- If `ID` names one of the syntactic extensions `chicken-more-macros chicken-ffi-macros`, then this extension will be loaded.
- If `ID` names one of the core library units shipped with CHICKEN, then a `(load-library 'ID)` will be performed.
- If `ID` names an installed extension with the `syntax` or `require-at-runtime` attribute, then the equivalent of `(require-for-syntax 'ID)` is performed, probably followed by `(require ...)` for any run-time requirements.
- Otherwise, `(require-extension ID)` is equivalent to `(require 'ID)`.

During compilation, one of the following happens instead:

- If `ID` names a built-in feature `chicken srfi-0 srfi-2 srfi-6 srfi-8 srfi-9 srfi-10 srfi-17 srfi-23 srfi-30 srfi-39 srfi-55`, then nothing is done.
- If `ID` names one of the syntactic extensions `chicken-more-macros chicken-ffi-macros`, then this extension will be loaded at compile-time, making the syntactic extensions available in compiled code.
- If `ID` names one of the core library units shipped with CHICKEN, or if the option `-uses ID` has been passed to the compiler, then a `(declare (uses ID))` is generated.
- If `ID` names an installed extension with the `syntax` or `require-at-runtime` attribute, then the equivalent of `(require-for-syntax 'ID)` is performed, and code is emitted to `(require ...)` any needed run-time requirements.
- Otherwise `(require-extension ID)` is equivalent to `(require 'ID)`.

To make long matters short - just use `require-extension` and it will normally figure everything out for dynamically loadable extensions and core library units.

`ID` should be a pure extension name and should not contain any path prefixes (for example `dir/lib...`) is illegal).

`ID` may also be a list that designates an extension-specifier. Currently the following extension specifiers are defined:

- `(srfi NUMBER ...)` is required for SRFI-55 compatibility and is fully implemented
- `(version ID NUMBER)` is equivalent to `ID`, but checks at compile-time whether the extension named `ID` is installed and whether its version is equal or higher than `NUMBER`. `NUMBER` may be a string or a number, the comparison is done lexicographically (using `string>=?`).

See also: `set-extension-specifier!`

When syntax extensions are loaded that redefine the global toplevel macro-expander (for example the <u>syntax-case</u> extension), then all remaining expression *in the same toplevel form* are still expanded with the old toplevel macro-expander.

9.1.2 define-extension

`[syntax] (define-extension NAME CLAUSE ...)`

This macro simplifies the task of writing extensions that can be linked both statically and dynamically. If encountered in interpreted code or code that is compiled into a shared object (specifically if compiled with the feature `chicken-compile-shared`, done automatically by `csc` when compiling with the `-shared` or `-dynamic` option) then the code given by clauses of the form

`(dynamic EXPRESSION ...)`

are inserted into the output as a `begin` form.

If compiled statically (specifically if the feature `chicken-compile-shared` has not been given), then this form expands into the following:

```
(declare (unit NAME))
(provide 'NAME)
```

and all clauses of the form

`(static EXPRESSION ...)`

all additionally inserted into the expansion.

As a convenience, the clause

`(export IDENTIFIER ...)`

is also allowed and is identical to `(declare (export IDENTIFIER ...))` (unless the `define-extension` form occurs in interpreted code, in with it is simply ignored).

Note that the compiler option `-extension NAME` is equivalent to prefixing the compiled file with

`(define-extension NAME)`

9.2 Binding forms for optional arguments

9.2.1 optional

[syntax] (optional ARGS DEFAULT)

Use this form for procedures that take a single optional argument. If `ARGS` is the empty list `DEFAULT` is evaluated and returned, otherwise the first element of the list `ARGS`. It is an error if `ARGS` contains more than one value.

```
(define (incr x . i) (+ x (optional i 1)))
(incr 10)                                ==> 11
(incr 12 5)                              ==> 17
```

9.2.2 case-lambda

[syntax] (case-lambda (LAMBDA-LIST1 EXP1 ...) ...)

Expands into a lambda that invokes the body following the first matching lambda-list.

```
(define plus
  (case-lambda
    (() 0)
    ((x) x)
    ((x y) (+ x y))
    ((x y z) (+ (+ x y) z))
    (args (apply + args))))

(plus)               ==> 9
(plus 1)             ==> 1
(plus 1 2 3)         ==> 6
```

For more information see the documentation for SRFI-16

9.2.3 let-optionals

[syntax] (let-optionals ARGS ((VAR1 DEFAULT1) ...) BODY ...)

Binding constructs for optional procedure arguments. `ARGS` should be a rest-parameter taken from a lambda-list. `let-optionals` binds `VAR1 ...` to available arguments in parallel, or to `DEFAULT1 ...` if not enough arguments were provided. `let-optionals*` binds `VAR1 ...` sequentially, so every variable sees the previous ones. it is an error if any excess arguments are provided.

```
(let-optionals '(one two) ((a 1) (b 2) (c 3))
  (list a b c) )                                ==> (one two 3)
```

9.2.4 let-optionals*

[syntax] (let-optionals* ARGS ((VAR1 DEFAULT1) ... [RESTVAR]) BODY ...)

Binding constructs for optional procedure arguments. `ARGS` should be a rest-parameter taken from a lambda-list. `let-optionals` binds `VAR1 ...` to available arguments in parallel, or to `DEFAULT1 ...` if not enough arguments were provided. `let-optionals*` binds `VAR1 ...` sequentially, so every variable sees the previous ones. If a single variable `RESTVAR` is given, then it is bound to any remaining arguments, otherwise it is an error if any excess arguments are provided.

```
(let-optionals* '(one two) ((a 1) (b 2) (c a))
  (list a b c) )                                    ==> (one two one)
```

9.3 Other binding forms

9.3.1 and-let*

[syntax] (and-let* (BINDING ...) EXP1 EXP2 ...)

SRFI-2. Bind sequentially and execute body. `BINDING` can be a list of a variable and an expression, a list with a single expression, or a single variable. If the value of an expression bound to a variable is `#f`, the `and-let*` form evaluates to `#f` (and the subsequent bindings and the body are not executed). Otherwise the next binding is performed. If all bindings/expressions evaluate to a true result, the body is executed normally and the result of the last expression is the result of the `and-let*` form. See also the documentation for SRFI-2.

9.3.2 rec

[syntax] (rec NAME EXPRESSION)
[syntax] (rec (NAME VARIABLE ...) BODY ...)

Allows simple definition of recursive definitions. `(rec NAME EXPRESSION)` is equivalent to `(letrec ((NAME EXPRESSION)) NAME)` and `(rec (NAME VARIABLE ...) BODY ...)` is the same as `(letrec ((NAME (lambda (VARIABLE ...) BODY ...))) NAME)`.

9.3.3 cut

[syntax] (cut SLOT ...)
[syntax] (cute SLOT ...)

Syntactic sugar for specializing parameters.

9.3.4 define-values

`[syntax] (define-values (NAME ...) EXP)`

Defines several variables at once, with the result values of expression EXP.

9.3.5 fluid-let

`[syntax] (fluid-let ((VAR1 X1) ...) BODY ...)`

Binds the variables VAR1 ... dynamically to the values X1 ... during execution of BODY ....

9.3.6 let-values

`[syntax] (let-values (((NAME ...) EXP) ...) BODY ...)`

Binds multiple variables to the result values of EXP .... All variables are bound simultaneously.

9.3.7 let*-values

`[syntax] (let*-values (((NAME ...) EXP) ...) BODY ...)`

Binds multiple variables to the result values of EXP .... The variables are bound sequentially.

```
(let*-values (((a b) (values 2 3))
              ((p) (+ a b)) )
  p)                              ==> 5
```

9.3.8 letrec-values

`[syntax] (letrec-values (((NAME ...) EXP) ...) BODY ...)`

Binds the result values of EXP ... to multiple variables at once. All variables are mutually recursive.

```
(letrec-values (((odd even)
                 (values
                   (lambda (n) (if (zero? n) #f (even (sub1 n))))
                   (lambda (n) (if (zero? n) #t (odd (sub1 n)))) ) ) )
  (odd 17) )                              ==> #t
```

9.3.9 parameterize

[syntax] (parameterize ((PARAMETER1 X1) ...) BODY ...)

Binds the parameters PARAMETER1 ... dynamically to the values X1 ... during execution of BODY
.... (see also: `make-parameter` in Parameters). Note that PARAMETER may be any expression that
evaluates to a parameter procedure.

9.3.10 receive

[syntax] (receive (NAME1 ... [. NAMEn]) VALUEEXP BODY ...)
[syntax] (receive VALUEEXP)

SRFI-8. Syntactic sugar for `call-with-values`. Binds variables to the result values of VALUEEXP and
evaluates BODY ....

The syntax

(receive VALUEEXP)

is equivalent to

(receive _ VALUEEXP _)

9.3.11 set!-values

[syntax] (set!-values (NAME ...) EXP)

Assigns the result values of expression EXP to multiple variables.

9.4 Substitution forms and macros

9.4.1 define-constant

[syntax] (define-constant NAME CONST)

Define a variable with a constant value, evaluated at compile-time. Any reference to such a constant should
appear textually **after** its definition. This construct is equivalent to `define` when evaluated or interpreted.
Constant definitions should only appear at toplevel. Note that constants are local to the current compilation
unit and are not available outside of the source file in which they are defined. Names of constants still exist in
the Scheme namespace and can be lexically shadowed. If the value is mutable, then the compiler is careful to
preserve its identity. CONST may be any constant expression, and may also refer to constants defined via
`define-constant` previously. This for should only be used at top-level.

9.4.2 define-inline

```
[syntax] (define-inline (NAME VAR ... [. VAR]) BODY ...)
[syntax] (define-inline NAME EXP)
```

Defines an inline procedure. Any occurrence of `NAME` will be replaced by `EXP` or `(lambda (VAR ... [. VAR]) BODY ...)`. This is similar to a macro, but variable-names and -scope will be correctly handled. Inline substitutions take place **after** macro-expansion. `EXP` should be a lambda-expression. Any reference to `NAME` should appear textually **after** its definition. Note that inline procedures are local to the current compilation unit and are not available outside of the source file in which they are defined. Names of inline procedures still exist in the Scheme namespace and can be lexically shadowed. This construct is equivalent to `define` when evaluated or interpreted. Inline definitions should only appear at toplevel.

9.4.3 define-macro

```
[syntax] (define-macro (NAME VAR ... [. VAR]) EXP1 ...)
[syntax] (define-macro NAME (lambda (VAR ... [. VAR]) EXP1 ...))
[syntax] (define-macro NAME1 NAME2)
```

Define a globally visible macro special form. The macro is available as soon as it is defined, i.e. it is registered at compile-time. If the file containing this definition invokes `eval` and the declaration `run-time-macros` (or the command line option `-run-time-macros`) has been used, then the macro is visible in evaluated expressions during runtime. The second possible syntax for `define-macro` is allowed for portability purposes only. In this case the second argument **must** be a lambda-expression or a macro name. Only global macros can be defined using this form. `(define-macro NAME1 NAME2)` simply copies the macro definition from `NAME2` to `NAME1`, creating an alias.

Extended lambda list syntax (`#!optional`, etc.) can be used but note that arguments are source expressions and thus default values for optional or keyword arguments should take this into consideration.

9.4.4 define-for-syntax

```
[syntax] (define-for-syntax (NAME VAR ... [. VAR]) EXP1 ...)
[syntax] (define-for-syntax NAME [VALUE])
```

Defines the toplevel variable `NAME` at macro-expansion time. This can be helpful when you want to define support procedures for use in macro-transformers, for example.

9.5 Conditional forms

9.5.1 select

[syntax] (select EXP ((KEY ...) EXP1 ...) ... [(else EXPn ...)])

This is similar to `case`, but the keys are evaluated.

9.5.2 unless

[syntax] (unless TEST EXP1 EXP2 ...)

Equivalent to:

```
(if (not TEST) (begin EXP1 EXP2 ...))
```

9.5.3 when

[syntax] (when TEST EXP1 EXP2 ...)

Equivalent to:

```
(if TEST (begin EXP1 EXP2 ...))
```

9.6 Record structures

9.6.1 define-record

[syntax] (define-record NAME SLOTNAME ...)

Defines a record type. Call `make-NAME` to create an instance of the structure (with one initialization-argument for each slot). `(NAME? STRUCT)` tests any object for being an instance of this structure. Slots are accessed via `(NAME-SLOTNAME STRUCT)` and updated using `(NAME-SLOTNAME-set! STRUCT VALUE)`.

```
(define-record point x y)
(define p1 (make-point 123 456))
(point? p1)                       ==> #t
(point-x p1)                      ==> 123
(point-y-set! p1 99)
(point-y p1)                      ==> 99
```

9.6.2 define-record-printer

```
[syntax] (define-record-printer (NAME RECORDVAR PORTVAR) BODY ...)
[syntax] (define-record-printer NAME PROCEDURE)
```

Defines a printing method for record of the type `NAME` by associating a procedure with the record type. When a record of this type is written using `display, write` or `print`, then the procedure is called with two arguments: the record to be printed and an output-port.

```
(define-record foo x y z)
(define f (make-foo 1 2 3))
(define-record-printer (foo x out)
  (fprintf out "#,(foo ~S ~S ~S)"
           (foo-x x) (foo-y x) (foo-z x)) )
(define-reader-ctor 'foo make-foo)
(define s (with-output-to-string
              (lambda () (write f))))
s                                      ==> "#,(foo 1 2 3)"
(equal? f (with-input-from-string
              s read)))                ==> #t
```

`define-record-printer` works also with SRFI-9 record types.

9.6.3 define-record-type

```
[syntax] (define-record-type NAME
                             (CONSTRUCTOR TAG ...)
                             PREDICATE
                             (FIELD ACCESSOR [MODIFIER]) ...)
```

SRFI-9 record types. For more information see the documentation for SRFI-9.

9.7 Other forms

9.7.1 assert

```
[syntax] (assert EXP [STRING ARG ...])
```

Signals an error if `EXP` evaluates to false. An optional message `STRING` and arguments `ARG ...` may be supplied to give a more informative error-message. If compiled in *unsafe* mode (either by specifying the `-unsafe` compiler option or by declaring (`unsafe`)), then this expression expands to an unspecified value. The result is the value of `EXP`.

9.7.2 cond-expand

`[syntax] (cond-expand FEATURE-CLAUSE ...)`

Expands by selecting feature clauses. This form is allowed to appear in non-toplevel expressions.

Predefined feature-identifiers are "situation" specific:

compile
>   `eval`, `library`, `match`, `compiling`, `srfi-11`, `srfi-15`, `srfi-31`, `srfi-26`, `srfi-16`, `utils`, `regex`, `srfi-4`, `match`, `srfi-1`, `srfi-69`, `srfi-28`, `extras`, `srfi-8`, `srfi-6`, `srfi-2`, `srfi-0`, `srfi-10`, `srfi-9`, `srfi-55`, `srfi-61` `chicken`, `srfi-23`, `srfi-30`, `srfi-39`, `srfi-62`, `srfi-17`, `srfi-12`.

load
>   `srfi-69`, `srfi-28`, `extras`, `srfi-8`, `srfi-6`, `srfi-2`, `srfi-0`, `srfi-10`, `srfi-9`, `srfi-55`, `srfi-61`, `chicken`, `srfi-23`, `srfi-30`, `srfi-39`, `srfi-62`, `srfi-17`, `srfi-12`. `library` is implicit.

eval
>   `match`, `csi`, `srfi-11`, `srfi-15`, `srfi-31`, `srfi-26`, `srfi-16`, `srfi-69`, `srfi-28`, `extras`, `srfi-8`, `srfi-6`, `srfi-2`, `srfi-0`, `srfi-10`, `srfi-9`, `srfi-55`, `srfi-61`, `chicken`, `srfi-23`, `srfi-30`, `srfi-39`, `srfi-62`, `srfi-17`, `srfi-12`. `library` is implicit.

The following feature-identifiers are available in all situations: `(machine-byte-order)`, `(machine-type)`, `(software-type)`, `(software-version)`, where the actual feature-identifier is platform dependent.

In addition the following feature-identifiers may exist: `applyhook`, `extraslot`, `ptables`, `dload`.

For further information, see the documentation for SRFI-0.

9.7.3 ensure

`[syntax] (ensure PREDICATE EXP [ARGUMENTS ...])`

Evaluates the expression `EXP` and applies the one-argument procedure `PREDICATE` to the result. If the predicate returns `#f` an error is signaled, otherwise the result of `EXP` is returned. If compiled in *unsafe* mode (either by specifying the `-unsafe` compiler option or by declaring `(unsafe)`), then this expression expands to an unspecified value. If specified, the optional `ARGUMENTS` are used as arguments to the invocation of the error-signalling code, as in `(error ARGUMENTS ...)`. If no `ARGUMENTS` are given, a generic error message is displayed with the offending value and `PREDICATE` expression.

9.7.4 eval-when

`[syntax] (eval-when (SITUATION ...) EXP ...)`

Controls evaluation/compilation of subforms. `SITUATION` should be one of the symbols `eval`, `compile` or `load`. When encountered in the evaluator, and the situation specifier `eval` is not given, then this form is not evaluated and an unspecified value is returned. When encountered while compiling code, and the situation

specifier `compile` is given, then this form is evaluated at compile-time. When encountered while compiling code, and the situation specifier `load` is not given, then this form is ignored and an expression resulting into an unspecified value is compiled instead.

The following table should make this clearer:

	In compiled code	In interpreted code
`eval`	ignore	evaluate
`compile`	evaluate at compile time	ignore
`load`	compile as normal	ignore

The situation specifiers `compile-time` and `run-time` are also defined and have the same meaning as `compile` and `load`, respectively.

9.7.5 include

```
[syntax] (include STRING)
```

Include toplevel-expressions from the given source file in the currently compiled/interpreted program. If the included file has the extension `.scm`, then it may be omitted. The file is searched in the current directory and, if not found, in all directories specified in the `-include-path` option.

9.7.6 nth-value

```
[syntax] (nth-value N EXP)
```

Returns the `N`th value (counting from zero) of the values returned by expression `EXP`.

9.7.7 time

```
[syntax] (time EXP1 ...)
```

Evaluates `EXP1 ...` and prints elapsed time and some values about GC use, like time spent in major GCs, number of minor and major GCs.

Previous: <u>Non-standard read syntax</u>

Next: <u>Pattern matching</u>

10 Pattern matching

(This description has been taken mostly from Andrew Wright's postscript document)

Pattern matching allows complicated control decisions based on data structure to be expressed in a concise manner. Pattern matching is found in several modern languages, notably Standard ML, Haskell and Miranda. These syntactic extensions internally use the `match` library unit.

The basic form of pattern matching expression is:

```
(match exp [pat body] ...)
```

where `exp` is an expression, `pat` is a pattern, and `body` is one or more expressions (like the body of a lambda-expression). The `match` form matches its first subexpression against a sequence of patterns, and branches to the `body` corresponding to the first pattern successfully matched. For example, the following code defines the usual `map` function:

```
(define map
  (lambda (f l)
    (match l
      [() '()]
      [(x . y) (cons (f x) (map f y))])))
```

The first pattern `()` matches the empty list. The second pattern `(x . y)` matches a pair, binding `x` to the first component of the pair and `y` to the second component of the pair.

10.1 Pattern Matching Expressions

The complete syntax of the pattern matching expressions follows:

```
exp ::= (match exp clause ...)
      | (match-lambda clause ...)
      | (match-lambda* clause ...)
      | (match-let ([pat exp] ...) body)
      | (match-let* ([pat exp] ...) body)
      | (match-letrec ([pat exp] ...) body)
      | (match-let var ([pat exp] ...) body)
      | (match-define pat exp)

clause ::= [pat body]
         | [pat (=> identifier) body]

pat ::= identifier         matches anything, and binds identifier as a variab
      | _                  anything
      | ()                 itself (the empty list)
      | #t                 itself
      | #f                 itself
      | string             an `equal?' string
      | number             an `equal?' number
      | character          an `equal?' character
      | 's-expression      an `equal?' s-expression
```

```
      | (pat-1 ... pat-n)    a proper list of n elements
      | (pat-1 ... pat-n . pat-n+1)
                              a list of n or more elements
      | (pat-1 ... pat-n pat-n+1 ..k)
                              a proper list of n+k or more elements [1]
      | #(pat-1 ... pat-n)   a vector of n elements
      | #(pat-1 ... pat-n pat-n+1 ..k)
                              a vector of n+k or more elements
      | ($ struct pat-1 ... pat-n)
                              a structure
      | (= field pat)        a field of a structure
      | (and pat-1 ... pat-n)
                              if all of pat-1 through pat-n match
      | (or pat-1 ... pat-n)
                              if any of pat-1 through pat-n match
      | (not pat-1 ... pat-n)
                              if none of pat-1 through pat-n match
      | (? predicate pat-1 ... pat-n)
                              if predicate true and pat-1 through pat-n all matc
      | (set! identifier)    anything, and binds identifier as a setter
      | (get! identifier)    anything, and binds identifier as a getter
      | `qp                  a quasipattern

qp ::= ()                    itself (the empty list)
      | #t                   itself
      | #f                   itself
      | string               an `equal?' string
      | number               an `equal?' number
      | character            an `equal?' character
      | symbol               an `equal?' symbol
      | (qp-1 ... qp-n)       a proper list of n elements
      | (qp-1 ... qp-n . qp-n+1)
                              a list of n or more elements
      | (qp-1 ... qp-n qp-n+1 ..k)
                              a proper list of n+k or more elements
      | #(qp-1 ... qp-n)      a vector of n elements
      | #(qp-1 ... qp-n qp-n+1 ..k)
                              a vector of n+k or more elements
      | ,pat                 a pattern
      | ,@pat                a pattern, spliced
```

The notation `..k` denotes a keyword consisting of three consecutive dots (ie., `...`), or two dots and an non-negative integer (eg., `..1`, `..2`), or three consecutive underscores (ie., `___`), or two underscores and a non-negative integer. The keywords `..k` and `__k` are equivalent. The keywords `...`, `___`, `..0`, and `__0` are equivalent.

The next subsection describes the various patterns.

The `match-lambda` and `match-lambda*` forms are convenient combinations of `match` and `lambda`, and can be explained as follows:

```
(match-lambda [pat body] ...)   =  (lambda (x) (match x [pat body] ...))
(match-lambda* [pat body] ...)  =  (lambda x (match x [pat body] ...))
```

where `x` is a unique variable. The `match-lambda` form is convenient when defining a single argument function that immediately destructures its argument. The `match-lambda*` form constructs a function that accepts any number of arguments; the patterns of `match-lambda*` should be lists.

The `match-let`, `match-let*`, `match-letrec`, and `match-define` forms generalize Scheme's `let`, `let*`, `letrec`, and `define` expressions to allow patterns in the binding position rather than just variables. For example, the following expression:

```
(match-let ([(x y z) (list 1 2 3)]) body ...)
```

binds `x` to 1, `y` to 2, and `z` to 3 in `body ...`. These forms are convenient for destructuring the result of a function that returns multiple values as a list or vector. As usual for `letrec` and `define`, pattern variables bound by `match-letrec` and `match-define` should not be used in computing the bound value.

The `match`, `match-lambda`, and `match-lambda*` forms allow the optional syntax `(=> identifier)` between the pattern and the body of a clause. When the pattern match for such a clause succeeds, the `identifier` is bound to a `failure procedure' of zero arguments within the `body`. If this procedure is invoked, it jumps back to the pattern matching expression, and resumes the matching process as if the pattern had failed to match. The `body` must not mutate the object being matched, otherwise unpredictable behavior may result.

10.2 Patterns

`identifier`: (excluding the reserved names `?`, `,`, `=`, `_`, `and`, `or`, `not`, `set!`, `get!`, `...`, and `..k` for non-negative integers `k`) matches anything, and binds a variable of this name to the matching value in the `body`.

`_`: matches anything, without binding any variables.

`()`, `#t`, `#f`, `string`, `number`, `character`, `'s-expression`: These constant patterns match themselves, i.e., the corresponding value must be `equal?` to the pattern.

`(pat-1 ... pat-n)`: matches a proper list of `n` elements that match `pat-1` through `pat-n`.

`(pat-1 ... pat-n . pat-n+1)`: matches a (possibly improper) list of at least `n` elements that ends in something matching `pat-n+1`.

`(pat-1 ... pat-n pat-n+1 ...)`: matches a proper list of `n` or more elements, where each element of the tail matches `pat-n+1`. Each pattern variable in `pat-n+1` is bound to a list of the matching values. For example, the expression:

```
(match '(let ([x 1][y 2]) z)
  [('let ((binding values) ...) exp)  body])
```

binds `binding` to the list `'(x y)`, `values` to the list `\'(1 2)`, and `exp` to `'z` in the body of the `match`-expression. For the special case where `pat-n+1` is a pattern variable, the list bound to that variable may share with the matched value.

`(pat-1 ... pat-n pat-n+1 ___ )`: This pattern means the same thing as the previous pattern.

`(pat-1 ... pat-n pat-n+1 ..k)`: This pattern is similar to the previous pattern, but the tail must be at least `k` elements long. The pattern keywords `..0` and `...` are equivalent.

`(pat-1 ... pat-n ~ pat-n+1 __k)`: This pattern means the same thing as the previous pattern.

`#(pat-1 ... pat-n)`: matches a vector of length `n`, whose elements match `pat-1` through `pat-n`.

`#(pat-1 ... pat-n pat-n+1 ...)`: matches a vector of length `n` or more, where each element beyond `n` matches `pat-n+1`.

`#(pat-1 ... pat-n pat-n+1 ..k)`: matches a vector of length `n+k` or more, where each element beyond `n` matches `pat-n+1`.

`($ struct pat-1 ... pat-n)`: matches a structure declared with `define-record` or `define-record-type`.

`(= field pat)`: is intended for selecting a field from a structure. *field* may be any expression; it is applied to the value being matched, and the result of this application is matched against `pat`.

`(and pat-1 ... pat-n)`: matches if all of the subpatterns match. At least one subpattern must be present. This pattern is often used as `(and x pat)` to bind `x` to to the entire value that matches `pat` (cf. *as-patterns* in ML or Haskell).

`(or pat-1 ... pat-n)`: matches if any of the subpatterns match. At least one subpattern must be present. All subpatterns must bind the same set of pattern variables.

`(not pat-1 ... pat-n)`: matches if none of the subpatterns match. At least one subpattern must be present. The subpatterns may not bind any pattern variables.

`(? predicate pat-1 ... pat-n)`: In this pattern, `predicate` must be an expression evaluating to a single argument function. This pattern matches if `predicate` applied to the corresponding value is true, and the subpatterns `pat-1 ... pat-n` all match. The `predicate` should not have side effects, as the code generated by the pattern matcher may invoke predicates repeatedly in any order. The `predicate` expression is bound in the same scope as the match expression, i.e., free variables in `predicate` are not bound by pattern variables.

`(set! identifier)`: matches anything, and binds `identifier` to a procedure of one argument that mutates the corresponding field of the matching value. This pattern must be nested within a pair, vector, box, or structure pattern. For example, the expression:

```
(define x (list 1 (list 2 3)))
(match x [(_ (_ (set! setit)))  (setit 4)])
```

mutates the `cadadr` of `x` to 4, so that `x` is `'(1 (2 4))`.

`(get! identifier)`: matches anything, and binds `identifier` to a procedure of zero arguments that accesses the corresponding field of the matching value. This pattern is the complement to `set!`. As with `set!`, this pattern must be nested within a pair, vector, box, or structure pattern.

Quasipatterns: Quasiquote introduces a quasipattern, in which identifiers are considered to be symbolic constants. Like Scheme's quasiquote for data, `unquote` (,) and `unquote-splicing` (,@) escape back to normal patterns.

10.3 Match Failure

If no clause matches the value, the default action is to invoke the procedure `(match-error-procedure)` with the value that did not match. The default definition of `(match-error-procedure)` calls `error` with an appropriate message:

```
#;1> (match 1 (2 2))

Failed match:
Error: no matching clause for : 1
```

For most situations, this behavior is adequate, but it can be changed by altering the value of the parameter `match-error-control`:

{procedure} match-error-control

```
(match-error-control [MODE])
```

Selects a mode that specifies how `match...` macro forms are to be expanded. With no argument this procedure returns the current mode. A single argument specifies the new mode that decides what should happen if no match-clause applies. The following modes are supported:

> Signal an error. This is the default.
1. :error

> Signal an error and output the offending form.
1. :match

> Omits `pair?` tests when the consequence is to fail in `car` or `cdr` rather than to signal an error.
1. :fail

unspecified Non-matching expressions will either fail in `car` or `cdr` or return an unspecified value. This mode applies to files compiled with the `unsafe` option or declaration.

When an error is signalled, the raised exception will be of kind `(exn match)`.

[procedure] match-error-procedure

```
(match-error-procedure [PROCEDURE])
```

Sets or returns the procedure called upon a match error. The procedure takes one argument, the value which failed to match. When the error control mode is `#:match` a second argument, the source form of the match expression is available.

10.4 Record Structures Pattern

The `$` pattern handles native record structures and SRFI-9 records transparently. Currently it is required that SRFI-9 record predicates are named exactly like the record type name, followed by a `?` (question mark) character.

10.5 Code Generation

Pattern matching macros are compiled into `if`-expressions that decompose the value being matched with standard Scheme procedures, and test the components with standard predicates. Rebinding or lexically shadowing the names of any of these procedures will change the semantics of the `match` macros. The names that should not be rebound or shadowed are:

```
null? pair? number? string? symbol? boolean? char? procedure? vector? list?
equal?
car cdr cadr cdddr ...
vector-length vector-ref
reverse length call/cc
```

Additionally, the code generated to match a structure pattern like `($ Foo pat-1 ... pat-n)` refers to the name `Foo?`. This name also should not be shadowed.

Previous: Non-standard macros and special forms

Next: Declarations

11 Declarations

11.1 declare

[syntax] (declare DECLSPEC ...)

Process declaration specifiers. Declarations always override any command-line settings. Declarations are valid for the whole compilation-unit (source file), the position of the declaration in the source file can be arbitrary. Declarations are ignored in the interpreter but not in code evaluated at compile-time (by `eval-when` or in syntax extensions loaded via `require-extension` or `require-for-syntax`. DECLSPEC may be any of the following:

11.2 always-bound

[declaration specifier] (always-bound SYMBOL ...)

Declares that the given variables are always bound and accesses to those have not to be checked.

11.3 block

[declaration specifier] (block)

Assume global variables are never redefined. This is the same as specifying the `-block` option.

11.4 block-global

11.5 hide

[declaration specifier] (block-global SYMBOL ...)
[declaration specifier] (hide SYMBOL ...)

Declares that the toplevel bindings for SYMBOL ... should not be accessible from code in other compilation units or by `eval`. Access to toplevel bindings declared as block global is also more efficient.

11.6 bound-to-procedure

```
[declaration specifier] (bound-to-procedure SYMBOL ...)
```

Declares that the given identifiers are always bound to procedure values.

11.7 c-options

```
[declaration specifier] (c-options STRING ...)
```

Declares additional C/C++ compiler options that are to be passed to the subsequent compilation pass that translates C to machine code. This declaration will only work if the source file is compiled with the `csc` compiler driver.

11.8 check-c-syntax

```
[declaration specifier] (check-c-syntax)
[declaration specifier] (not check-c-syntax)
```

Enables or disables syntax-checking of embedded C/C++ code fragments. Checking C syntax is the default.

11.9 constant

```
[declaration specifier] (constant SYMBOL ...)
```

Declares the procedures with the names `SYMBOL ...` as constant, that is, as not having any side effects. This can help the compiler to remove non-side-effecting expressions.

11.10 export

```
[declaration specifier] (export SYMBOL ...)
```

The opposite of `hide`. All given identifiers will be exported and all toplevel variables not listed will be hidden and not be accessible outside of this compilation unit.

11.11 emit-exports

```
[declaration specifier] (emit-exports STRING)
```

Write exported toplevel variables to file with name `STRING`.

11.12 emit-external-prototypes-first

`[declaration specifier] (emit-external-prototypes-first)`

Emit prototypes for callbacks defined with `define-external` before any other foreign declarations. Equivalent to giving the `-emit-external-prototypes-first` option to the compiler.

11.13 disable-interrupts

`[declaration specifier] (disable-interrupts)`
`[declaration specifier] (not interrupts-enabled)`

Disable timer-interrupts checks in the compiled program. Threads can not be preempted in main- or library-units that contain this declaration.

11.14 disable-warning

`[declaration specifier] (disable-warning CLASS ...)`

Disable warnings of type `CLASS ...` (equivalent to the `-disable-warning CLASS` compiler option).

11.15 import

`[declaration specifier] (import SYMBOL-OR-STRING ...)`

Adds new imports to the list of externally available toplevel variables. Arguments to this declaration may be either strings (designating `.exports` files, without the file-extension) or symbols which directly designate imported variables.

11.16 inline

`[declaration specifier] (inline)`
`[declaration specifier] (not inline)`
`[declaration specifier] (inline IDENTIFIER ...)`
`[declaration specifier] (not inline IDENTIFIER ...)`

If given without an identifier-list, inlining of known procedures is enabled (this is equivalent to the `-inline` compiler option). When an identifier-list is given, then inlining is enabled only for the specified global procedures. The negated forms `(not inline)` and `(not inline IDENTIFIER)` disable global inlining, or inlining for the given global procedures only, respectively.

11.17 inline-limit

[declaration specifier] (inline-limit THRESHOLD)

Sets the maximum size of procedures which may potentially be inlined. The default threshold is 10.

11.18 interrupts-enabled

[declaration specifier] (interrupts-enabled)

Enable timer-interrupts checks in the compiled program (the default).

11.19 keep-shadowed-macros

[declaration specifier] (keep-shadowed-macros)

Normally, when a toplevel variable is assigned or defined that has the same name as a macro, the macro-definition will be removed (in addition to showing a warning). This declaration will disable the removal of the macro.

11.20 lambda-lift

[declaration specifier] (lambda-lift)

Enables lambda-lifting (equivalent to the -lambda-lift option).

11.21 link-options

[declaration specifier] (link-options STRING ...)

Declares additional linker compiler options that are to be passed to the subsequent compilation pass that links the generated code into an executable or library. This declaration will only work if the source file is compiled with the csc compiler driver.

11.22 no-argc-checks

```
[declaration specifier] (no-argc-checks)
```

Disables argument count checking.

11.23 no-bound-checks

```
[declaration specifier] (no-bound-checks)
```

Disables the bound-checking of toplevel bindings.

11.24 no-procedure-checks

```
[declaration specifier] (no-procedure-checks)
```

Disables checking of values in operator position for being of procedure type.

11.25 post-process

```
[declaration specifier] (post-process STRING ...)
```

Arranges for the shell commands `STRING ...` to be invoked after the current file has been translated to C. Any occurrences of the substring `$@@` in the strings given for this declaration will be replaced by the pathname of the currently compiled file, without the file-extension. This declaration will only work if the source file is compiled with the `csc` compiler driver.

11.26 number-type

11.27 fixnum-arithmetic

```
[declaration specifier] ([number-type] TYPE)
[declaration specifier] (fixnum-arithmetic)
```

Declares that only numbers of the given type are used. `TYPE` may be `fixnum` or `generic` (which is the default).

11.28 run-time-macros

```
[declaration specifier] (run-time-macros)
```

Equivalent to the compiler option of the same name - macros defined in the compiled code are also made available at runtime.

11.29 standard-bindings

```
[declaration specifier] (standard-bindings SYMBOL ...)
[declaration specifier] (not standard-bindings SYMBOL ...)
```

Declares that all given standard procedures (or all if no symbols are specified) are never globally redefined. If `not` is specified, then all but the given standard bindings are assumed to be never redefined.

11.30 extended-bindings

```
[declaration specifier] (extended-bindings SYMBOL ...)
[declaration specifier] (not extended-bindings SYMBOL ...)
```

Declares that all given non-standard and CHICKEN-specific procedures (or all if no symbols are specified) are never globally redefined. If `not` is specified, then all but the given extended bindings are assumed to be never redefined.

11.31 usual-integrations

```
[declaration specifier] (usual-integrations SYMBOL ...)
[declaration specifier] (not usual-integrations SYMBOL ...)
```

Declares that all given standard and extended bindings (or all if no symbols are specified) are never globally redefined. If `not` is specified, then all but the given standard and extended bindings are assumed to be never redefined. Note that this is the default behaviour, unless the `-no-usual-integrations` option has been given.

11.32 unit

```
[declaration specifier] (unit SYMBOL)
```

Specify compilation unit-name (if this is a library)

11.33 unsafe

```
[declaration specifier] (unsafe)
[declaration specifier] (not safe)
```

Do not generate safety-checks. This is the same as specifying the `-unsafe` option. Also implies

```
(declare (no-bound-checks) (no-procedure-checks) (no-argc-checks))
```

11.34 unused

```
[declaration specifier] (unused SYMBOL ...)
```

Disables any warnings when the global variable `SYMBOL` is not defined but used, or defined but never used and not exported.

11.35 uses

```
[declaration specifier] (uses SYMBOL ...)
```

Gives a list of used library-units. Before the toplevel-expressions of the main-module are executed, all used units evaluate their toplevel-expressions in the order in which they appear in this declaration. If a library unit A uses another unit B, then B's toplevel expressions are evaluated before A's. Furthermore, the used symbols are registered as features during compile-time, so `cond-expand` knows about them.

Previous: Pattern matching

Next: Parameters

12 Parameters

Certain behavior of the interpreter and compiled programs can be customized via 'parameters', where a parameter is a procedure of zero or one arguments. To retrieve the value of a parameter call the parameter-procedure with zero arguments. To change the setting of the parameter, call the parameter-procedure with the new value as argument:

```
(define foo (make-parameter 123))
(foo)                              ==> 123
(foo 99)
(foo)                              ==> 99
```

Parameters are fully thread-local, each thread of execution owns a local copy of a parameters' value.

CHICKEN implements SRFI-39.

12.1 make-parameter

```
[procedure] (make-parameter VALUE [GUARD])
```

Returns a procedure that accepts zero or one argument. Invoking the procedure with zero arguments returns VALUE. Invoking the procedure with one argument changes its value to the value of that argument (subsequent invocations with zero parameters return the new value). GUARD should be a procedure of a single argument. Any new values of the parameter (even the initial value) are passed to this procedure. The guard procedure should check the value and/or convert it to an appropriate form.

12.2 case-sensitive

If true, then read reads symbols and identifiers in case-sensitive mode and uppercase characters in symbols are printed escaped. Defaults to #t.

12.3 dynamic-load-libraries

A list of strings containing shared libraries that should be checked for explicitly loaded library units (this facility is not available on all platforms). See load-library.

12.4 command-line-arguments

Contains the list of arguments passed to this program, with the name of the program and any runtime options (all options starting with -:) removed.

12.5 current-read-table

A read-table object that holds read-procedures for special non-standard read-syntax (see `set-read-syntax!` for more information).

12.6 exit-handler

A procedure of a single optional argument. When `exit` is called, then this procedure will be invoked with the exit-code as argument. The default behavior is to terminate the program.

12.7 eval-handler

A procedure of one or two arguments. When `eval` is invoked, it calls the value of this parameter with the same arguments. The default behavior is to evaluate the argument expression and to ignore the second parameter.

12.8 force-finalizers

If true, force and execute all pending finalizers before exiting the program (either explicitly by `exit` or implicitly when the last toplevel expression has been executed). Default is `#t`.

12.9 implicit-exit-handler

A procedure of no arguments. When the last toplevel expression of the program has executed, then the value of this parameter is called. The default behaviour is to invoke all pending finalizers.

12.10 keyword-style

Enables alternative keyword syntax, where `STYLE` may be either `#:prefix` (as in Common Lisp) or `#:suffix` (as in DSSSL). Any other value disables the alternative syntaxes.

12.11 load-verbose

A boolean indicating whether loading of source files, compiled code (if available) and compiled libraries should display a message.

12.12 program-name

The name of the currently executing program. This is equivalent to `(car (argv))` for compiled programs or the filename following the `-script` option in interpreted scripts.

12.13 repl-prompt

A procedure that should evaluate to a string that will be printed before reading interactive input from the user in a read-eval-print loop. Defaults to `(lambda () "#;N> ")`.

12.14 reset-handler

A procedure of zero arguments that is called via `reset`. The default behavior in compiled code is to invoke the value of `(exit-handler)`. The default behavior in the interpreter is to abort the current computation and to restart the read-eval-print loop.

Previous: <u>Declarations</u>

Next: <u>Unit library</u>

13 Unit library

This unit contains basic Scheme definitions. This unit is used by default, unless the program is compiled with the `-explicit-use` option.

13.1 Arithmetic

13.1.1 add1/sub1

[procedure] (add1 N)
[procedure] (sub1 N)

Adds/subtracts 1 from `N`.

13.1.2 Binary integer operations

Binary integer operations. `arithmetic-shift` shifts the argument `N1` by `N2` bits to the left. If `N2` is negative, than `N1` is shifted to the right. These operations only accept exact integers or inexact integers in word range (32 bit signed on 32-bit platforms, or 64 bit signed on 64-bit platforms).

[procedure] (bitwise-and N1 ...)
[procedure] (bitwise-ior N1 ...)
[procedure] (bitwise-xor N1 ...)
[procedure] (bitwise-not N)
[procedure] (arithmetic-shift N1 N2)

13.1.3 bit-set?

[procedure] (bit-set? N INDEX)

Returns `#t` if the bit at the position `INDEX` in the integer `N` is set, or `#f` otherwise. The rightmost/least-significant bit is bit 0.

13.1.4 fixnum?

[procedure] (fixnum? X)

Returns `#t` if `X` is a fixnum, or `#f` otherwise.

13.1.5 Arithmetic fixnum operations

These procedures do not check their arguments, so non-fixnum parameters will result in incorrect results. `fxneg` negates its argument.

On division by zero, `fx/` and `fxmod` signal a condition of kind `(exn arithmetic)`.

`fxshl` and `fxshr` perform arithmetic shift left and right, respectively.

[procedure] (fx+ N1 N2)
[procedure] (fx- N1 N2)
[procedure] (fx* N1 N2)
[procedure] (fx/ N1 N2)
[procedure] (fxmod N1 N2)
[procedure] (fxneg N)
[procedure] (fxmin N1 N2)
[procedure] (fxmax N1 N2)
[procedure] (fx= N1 N2)
[procedure] (fx> N1 N2)
[procedure] (fx< N1 N2)
[procedure] (fx>= N1 N2)
[procedure] (fx N1 N2)
[procedure] (fxand N1 N2)
[procedure] (fxior N1 N2)
[procedure] (fxxor N1 N2)
[procedure] (fxnot N)
[procedure] (fxshl N1 N2)
[procedure] (fxshr N1 N2)

13.1.6 Arithmetic floating-point operations

In safe mode, these procedures throw a type error with non-float arguments (except `flonum?`, which returns `#f`). In unsafe mode, these procedures do not check their arguments. A non-flonum argument in unsafe mode can crash the system.

[procedure] (flonum? X)
[procedure] (fp+ X Y)
[procedure] (fp- X Y)
[procedure] (fp* X Y)
[procedure] (fp/ X Y)
[procedure] (fpneg X)
[procedure] (fpmin X Y)
[procedure] (fpmax X Y)
[procedure] (fp= X Y)
[procedure] (fp> X Y)
[procedure] (fp< X Y)
[procedure] (fp>= X Y)
[procedure] (fp X Y)

13.1.7 signum

[procedure] (signum N)

Returns 1 if N is positive, -1 if N is negative or 0 if N is zero. signum is exactness preserving.

13.1.8 finite?

[procedure] (finite? N)

Returns #f if N is negative or positive infinity, and #t otherwise.

13.2 File Input/Output

13.2.1 current-output-port

[procedure] (current-output-port [PORT])

Returns default output port. If PORT is given, then that port is selected as the new current output port.

Note that the default output port is not buffered. Use [[http://galinha.ucpel.tche.br/Unit posix#Setting the file buffering mode|set-buffering-mode!]] if you need a different behavior.

13.2.2 current-error-port

[procedure] (current-error-port [PORT])

Returns default error output port. If PORT is given, then that port is selected as the new current error output port.

Note that the default error output port is not buffered. Use [[http://galinha.ucpel.tche.br/Unit posix#Setting the file buffering mode|set-buffering-mode!]] if you need a different behavior.

13.2.3 flush-output

[procedure] (flush-output [PORT])

Write buffered output to the given output-port. PORT defaults to the value of (current-output-port).

13.2.4 port-name

[procedure] (port-name [PORT])

Fetch filename from PORT. This returns the filename that was used to open this file. Returns a special tag string, enclosed into parentheses for non-file ports. PORT defaults to the value of (current-input-port).

13.2.5 port-position

[procedure] (port-position [PORT])

Returns the current position of PORT as two values: row and column number. If the port does not support such an operation an error is signaled. This procedure is currently only available for input ports. PORT defaults to the value of (current-input-port).

13.2.6 set-port-name!

[procedure] (set-port-name! PORT STRING)

Sets the name of PORT to STRING.

13.3 Files

13.3.1 delete-file

[procedure] (delete-file STRING)

Deletes the file with the pathname STRING. If the file does not exist, an error is signaled.

13.3.2 file-exists?

[procedure] (file-exists? STRING)

Returns STRING if a file with the given pathname exists, or #f otherwise.

13.3.3 rename-file

[procedure] (rename-file OLD NEW)

Renames the file or directory with the pathname OLD to NEW. If the operation does not succeed, an error is signaled.

13.4 String ports

13.4.1 get-output-string

[procedure] (get-output-string PORT)

Returns accumulated output of a port created with (open-output-string).

13.4.2 open-input-string

[procedure] (open-input-string STRING)

Returns a port for reading from STRING.

13.4.3 open-output-string

[procedure] (open-output-string)

Returns a port for accumulating output in a string.

13.5 Feature identifiers

CHICKEN maintains a global list of *features* naming functionality available in the current system. Additionally the cond-expand form accesses this feature list to infer what features are provided. Predefined features are chicken, and the SRFIs (Scheme Request For Implementation) provided by the base system: srfi-23, srfi-30, srfi-39. If the eval unit is used (the default), the features srfi-0, srfi-2, srfi-6, srfi-8, srfi-9 and srfi-10 are defined. When compiling code (during compile-time) the feature compiling is registered. When evaluating code in the interpreter (csi), the feature csi is registered.

13.5.1 features

```
[procedure] (features)
```

Returns a list of all registered features that will be accepted as valid feature-identifiers by `cond-expand`.

13.5.2 feature?

```
[procedure] (feature? ID ...)
```

Returns `#t` if all features with the given feature-identifiers `ID ...` are registered.

13.5.3 register-feature!

```
[procedure] (register-feature! FEATURE ...)
```

Register one or more features that will be accepted as valid feature-identifiers by `cond-expand`. `FEATURE ...` may be a keyword, string or symbol.

13.5.4 unregister-feature!

```
[procedure] (unregister-feature! FEATURE ...)
```

Unregisters the specified feature-identifiers. `FEATURE ...` may be a keyword, string or symbol.

13.6 Keywords

Keywords are special symbols prefixed with `#:` that evaluate to themselves. Procedures can use keywords to accept optional named parameters in addition to normal required parameters. Assignment to and bindings of keyword symbols is not allowed. The parameter `keyword-style` and the compiler/interpreter option `-keyword-style` can be used to allow an additional keyword syntax, either compatible to Common LISP, or to DSSSL.

13.6.1 get-keyword

```
[procedure] (get-keyword KEYWORD ARGLIST [THUNK])
```

Returns the argument from `ARGLIST` specified under the keyword `KEYWORD`. If the keyword is not found, then the zero-argument procedure `THUNK` is invoked and the result value is returned. If `THUNK` is not given, `#f` is returned.

```
(define (increase x . args)
  (+ x (get-keyword #:amount args (lambda () 1))) )
(increase 123)                                         ==> 124
(increase 123 #:amount 10)                             ==> 133
```

Note: the KEYWORD may actually be any kind of object.

13.6.2 keyword?

[procedure] (keyword? X)

Returns #t if X is a keyword symbol, or #f otherwise.

13.6.3 keywordstring

[procedure] (keyword->string KEYWORD)

Transforms KEYWORD into a string.

13.6.4 stringkeyword

[procedure] (string->keyword STRING)

Returns a keyword with the name STRING.

13.7 Exceptions

CHICKEN implements the (currently withdrawn) SRFI-12 exception system. For more information, see the SRFI-12 document.

13.7.1 condition-case

[syntax] (condition-case EXPRESSION CLAUSE ...)

Evaluates EXPRESSION and handles any exceptions that are covered by CLAUSE ..., where CLAUSE should be of the following form:

CLAUSE = ([VARIABLE] (KIND ...) BODY ...)

If provided, VARIABLE will be bound to the signaled exception object. BODY ... is executed when the exception is a property- or composite condition with the kinds given KIND ... (unevaluated). If no clause

applies, the exception is re-signaled in the same dynamic context as the `condition-case` form.

```scheme
(define (check thunk)
  (condition-case (thunk)
    [(exn file) (print "file error")]
    [(exn) (print "other error")]
    [var () (print "something else")] ) )

(check (lambda () (open-input-file "")))   ; -> "file error"
(check (lambda () some-unbound-variable))  ; -> "othererror"
(check (lambda () (signal 99)))            ; -> "something else"

(condition-case some-unbound-variable
  [(exn file) (print "ignored")] )         ; -> signals error
```

13.7.2 breakpoint

`[procedure] (breakpoint [NAME])`

Programmatically triggers a breakpoint (similar to the `,br` top-level csi command).

All error-conditions signaled by the system are of kind `exn`. The following composite conditions are additionally defined:

`(exn arity)`

Signaled when a procedure is called with the wrong number of arguments.
`(exn type)`

Signaled on type-mismatch errors, for example when an argument of the wrong type is passed to a built-in procedure.
`(exn arithmetic)`

Signaled on arithmetic errors, like division by zero.
`(exn i/o)`

Signaled on input/output errors.
`(exn i/o file)`

Signaled on file-related errors.
`(exn i/o net)`

Signaled on network errors.
`(exn bounds)`

Signaled on errors caused by accessing non-existent elements of a collection.
`(exn runtime)`

Signaled on low-level runtime-system error-situations.

13.7.1 condition-case

```
(exn runtime limit)
```

Signaled when an internal limit is exceeded (like running out of memory).

```
(exn match)
```

Signaled on errors raised by failed matches (see the section on `match`).

```
(exn syntax)
```

Signaled on syntax errors.

```
(exn breakpoint)
```

Signaled when a breakpoint is reached.

Notes:

- All error-exceptions (of the kind `exn`) are non-continuable.
- Error-exceptions of the `exn` kind have additional `arguments` and `location` properties that contain the arguments passed to the exception-handler and the name of the procedure where the error occurred (if available).
- When the `posix` unit is available and used, then a user-interrupt (`signal/int`) signals an exception of the kind `user-interrupt`.
- the procedure `condition-property-accessor` accepts an optional third argument. If the condition does not have a value for the desired property and if the optional argument is given, no error is signaled and the accessor returns the third argument.
- In composite conditions all properties are currently collected in a single property-list, so in the case that to conditions have the same named property, only one will be visible.

13.8 Environment information and system interface

13.8.1 argv

```
[procedure] (argv)
```

Return a list of all supplied command-line arguments. The first item in the list is a string containing the name of the executing program. The other items are the arguments passed to the application. This list is freshly created on every invocation of `(argv)`. It depends on the host-shell whether arguments are expanded ('globbed') or not.

13.8.2 exit

```
[procedure] (exit [CODE])
```

Exit the running process and return exit-code, which defaults to 0 (Invokes `exit-handler`).

Note that pending `dynamic-wind` thunks are *not* invoked when exiting your program in this way.

13.8.3 build-platform

```
[procedure] (build-platform)
```

Returns a symbol specifying the toolset which has been used for building the executing system, which is one of the following:

```
cygwin
mingw32
gnu
intel
unknown
```

13.8.4 chicken-version

```
[procedure] (chicken-version [FULL])
```

Returns a string containing the version number of the CHICKEN runtime system. If the optional argument FULL is given and true, then a full version string is returned.

13.8.5 errno

```
[procedure] (errno)
```

Returns the error code of the last system call.

13.8.6 getenv

```
[procedure] (getenv STRING)
```

Returns the value of the environment variable STRING or #f if that variable is not defined.

13.8.7 machine-byte-order

```
[procedure] (machine-byte-order)
```

Returns the symbol little-endian or big-endian, depending on the machine's byte-order.

13.8.8 machine-type

[procedure] (machine-type)

Returns a symbol specifying the processor on which this process is currently running, which is one of the following:

```
alpha
mips
hppa
ultrasparc
sparc
ppc
ppc64
ia64
x86
x86-64
unknown
```

13.8.9 on-exit

[procedure] (on-exit THUNK)

Schedules the zero-argument procedures THUNK to be executed before the process exits, either explicitly via exit or implicitly after execution of the last top-level form. Note that finalizers for unreferenced finalized data are run before exit procedures.

13.8.10 software-type

[procedure] (software-type)

Returns a symbol specifying the operating system on which this process is currently running, which is one of the following:

```
windows
unix
macos
ecos
unknown
```

13.8.11 software-version

[procedure] (software-version)

Returns a symbol specifying the operating system version on which this process is currently running, which is one of the following:

```
linux
freebsd
netbsd
openbsd
macosx
hpux
solaris
sunos
unknown
```

13.8.12 c-runtime

`[procedure] (c-runtime)`

Returns a symbol that designates what kind of C runtime library has been linked with this version of the Chicken libraries. Possible return values are `static`, `dynamic` or `unknown`. On systems not compiled with the Microsoft C compiler, `c-runtime` always returns `unknown`.

13.8.13 system

`[procedure] (system STRING)`

Execute shell command. The functionality offered by this procedure depends on the capabilities of the host shell. If the forking of a subprocess failed, an exception is raised. Otherwise the return status of the subprocess is returned unaltered.

13.9 Execution time

13.9.1 cpu-time

`[procedure] (cpu-time)`

Returns the used CPU time of the current process in milliseconds as two values: the time spent in user code, and the time spent in system code. On platforms where user and system time can not be differentiated, system time will be always be 0.

13.9.2 current-milliseconds

`[procedure] (current-milliseconds)`

Returns the number of milliseconds since process- or machine startup.

13.9.3 current-seconds

[procedure] (current-seconds)

Returns the number of seconds since midnight, Jan. 1, 1970.

13.9.4 current-gc-milliseconds

[procedure] (current-gc-milliseconds)

Returns the number of milliseconds spent in major garbage collections since the last call of
`current-gc-milliseconds` and returns an exact integer.

13.10 Interrupts and error-handling

13.10.1 enable-warnings

[procedure] (enable-warnings [BOOL])

Enables or disables warnings, depending on wether `BOOL` is true or false. If called with no arguments, this
procedure returns `#t` if warnings are currently enabled, or `#f` otherwise. Note that this is not a parameter. The
current state (whether warnings are enabled or disabled) is global and not thread-local.

13.10.2 error

[procedure] (error [LOCATION] [STRING] EXP ...)

Prints error message, writes all extra arguments to the value of `(current-error-port)` and invokes the
current exception-handler. This conforms to SRFI-23. If `LOCATION` is given and a symbol, it specifies the
location (the name of the procedure) where the error occurred.

13.10.3 get-call-chain

[procedure] (get-call-chain [START [THREAD]])

Returns a list with the call history. Backtrace information is only generated in code compiled without
`-no-trace` and evaluated code. If the optional argument `START` is given, the backtrace starts at this offset,
i.e. when `START` is 1, the next to last trace-entry is printed, and so on. If the optional argument `THREAD` is
given, then the call-chain will only be constructed for calls performed by this thread.

13.10.4 print-call-chain

[procedure] (print-call-chain [PORT [START [THREAD]]])

Prints a backtrace of the procedure call history to PORT, which defaults to (current-output-port).

13.10.5 print-error-message

[procedure] (print-error-message EXN [PORT [STRING]])

Prints an appropriate error message to PORT (which defaults to the value of (current-output-port) for the object EXN. EXN may be a condition, a string or any other object. If the optional argument STRING is given, it is printed before the error-message. STRING defaults to "Error:".

13.10.6 procedure-information

[procedure] (procedure-information PROC)

Returns an s-expression with debug information for the procedure PROC, or #f, if PROC has no associated debug information.

13.10.7 reset

[procedure] (reset)

Reset program (Invokes reset-handler).

13.10.8 warning

[procedure] (warning STRING EXP ...)

Displays a warning message (if warnings are enabled with enable-warnings) and continues execution.

13.10.9 singlestep

[procedure] (singlestep THUNK)

Executes the code in the zero-procedure THUNK in single-stepping mode.

13.11 Garbage collection

13.11.1 gc

[procedure] (gc [FLAG])

Invokes a garbage-collection and returns the number of free bytes in the heap. The flag specifies whether a minor (#f) or major (#t) GC is to be triggered. If no argument is given, #t is assumed. An explicit #t argument will cause all pending finalizers to be executed.

13.11.2 memory-statistics

[procedure] (memory-statistics)

Performs a major garbage collection and returns a three element vector containing the total heap size in bytes, the number of bytes currently used and the size of the nursery (the first heap generation). Note that the actual heap is actually twice the size given in the heap size, because CHICKEN uses a copying semi-space collector.

13.11.3 set-finalizer!

[procedure] (set-finalizer! X PROC)

Registers a procedure of one argument PROC, that will be called as soon as the non-immediate data object X is about to be garbage-collected (with that object as its argument). Note that the finalizer will **not** be called while interrupts are disabled. This procedure returns X.

13.11.4 set-gc-report!

[procedure] (set-gc-report! FLAG)

Print statistics after every GC, depending on FLAG. A value of #t shows statistics after every major GC. A true value different from #t shows statistics after every minor GC. #f switches statistics off.

13.12 Other control structures

13.12.1 promise?

`[procedure] (promise? X)`

Returns `#t` if X is a promise returned by `delay`, or `#f` otherwise.

13.13 String utilities

13.13.1 reverse-liststring

`[procedure] (reverse-list->string LIST)`

Returns a string with the characters in `LIST` in reverse order. This is equivalent to `(list->string (reverse LIST))`, but much more efficient.

13.14 Generating uninterned symbols

13.14.1 gensym

`[procedure] (gensym [STRING-OR-SYMBOL])`

Returns a newly created uninterned symbol. If an argument is provided, the new symbol is prefixed with that argument.

13.14.2 stringuninterned-symbol

`[procedure] (string->uninterned-symbol STRING)`

Returns a newly created, unique symbol with the name `STRING`.

13.15 Standard Input/Output

13.15.1 port?

[procedure] (port? X)

Returns #t if X is a port object or #f otherwise.

13.15.2 print

[procedure] (print [EXP1 ...])

Outputs the optional arguments EXP1 ... using display and writes a newline character to the port that is the value of (current-output-port). Returns (void).

13.15.3 print*

[procedure] (print* [EXP1 ...])

Similar to print, but does not output a terminating newline character and performs a flush-output after writing its arguments.

13.16 User-defined named characters

13.16.1 char-name

[procedure] (char-name SYMBOL-OR-CHAR [CHAR])

This procedure can be used to inquire about character names or to define new ones. With a single argument the behavior is as follows: If SYMBOL-OR-CHAR is a symbol, then char-name returns the character with this name, or #f if no character is defined under this name. If SYMBOL-OR-CHAR is a character, then the name of the character is returned as a symbol, or #f if the character has no associated name.

If the optional argument CHAR is provided, then SYMBOL-OR-CHAR should be a symbol that will be the new name of the given character. If multiple names designate the same character, then the write will use the character name that was defined last.

```
(char-name 'space)                 ==> #\space
(char-name #\space)                ==> space
(char-name 'bell)                  ==> #f
(char-name (integer->char 7))      ==> #f
(char-name 'bell (integer->char 7))
(char-name 'bell)                  ==> #\bell
(char->integer (char-name 'bell))  ==> 7
```

13.17 Blobs

"blobs" are collections of unstructured bytes. You can't do much with them, but allow conversion to and from SRFI-4 number vectors.

13.17.1 make-blob

[procedure] (make-blob SIZE)

Returns a blob object of SIZE bytes, aligned on an 8-byte boundary, uninitialized.

13.17.2 blob?

[procedure] (blob? X)

Returns #t if X is a blob object, or #f otherwise.

13.17.3 blob-size

[procedure] (blob-size BLOB)

Returns the number of bytes in BLOB.

13.17.4 blobstring

[procedure] (blob->string BLOB)

Returns a string with the contents of BLOB.

13.17.5 stringblob

[procedure] (string->blob STRING)

Returns a blob with the contents of STRING.

13.17.6 blob=?

[procedure] (blob=? BLOB1 BLOB2)

Returns #t if the two argument blobs are of the same size and have the same content.

13.18 Vectors

13.18.1 vector-copy!

[procedure] (vector-copy! VECTOR1 VECTOR2 [COUNT])

Copies contents of VECTOR1 into VECTOR2. If the argument COUNT is given, it specifies the maximal number of elements to be copied. If not given, the minimum of the lengths of the argument vectors is copied.

Exceptions: (exn bounds)

13.18.2 vector-resize

[procedure] (vector-resize VECTOR N [INIT])

Creates and returns a new vector with the contents of VECTOR and length N. If N is greater than the original length of VECTOR, then all additional items are initialized to INIT. If INIT is not specified, the contents are initialized to some unspecified value.

13.19 The unspecified value

13.19.1 void

[procedure] (void)

Returns an unspecified value.

13.20 Continuations

13.20.1 call/cc

[procedure] (call/cc PROCEDURE)

An alias for `call-with-current-continuation`.

13.20.2 continuation-capture

[procedure] (continuation-capture PROCEDURE)

Creates a continuation object representing the current continuation and tail-calls `PROCEDURE` with this continuation as the single argument.

More information about this continuation API can be found in the paper http://repository.readscheme.org/ftp/papers/sw2001/feeley.pdf *A Better API for first class Continuations* by Marc Feeley.

13.20.3 continuation?

[procedure] (continuation? X)

Returns `#t` if `X` is a continuation object, or `#f` otherwise. Please note that this applies only to continuations created by the Continuation API, but not by call/cc, i.e.: `(call-with-current-continuation continuation?)` returns `#f`, whereas `(continuation-capture continuation?)` returns `#t`.

13.20.4 continuation-graft

[procedure] (continuation-graft CONT THUNK)

Calls the procedure `THUNK` with no arguments and the implicit continuation `CONT`.

13.20.5 continuation-return

[procedure] (continuation-return CONT VALUE ...)

Returns the value(s) to the continuation `CONT`. `continuation-return` could be implemented like this:

```
(define (continuation-return k . vals)
  (continuation-graft
    k
    (lambda () (apply values vals)) ) )
```

13.21 Setters

SRFI-17 is fully implemented. For more information see: <u>SRFI-17</u>.

13.21.1 setter

[procedure] (setter PROCEDURE)

Returns the setter-procedure of PROCEDURE, or signals an error if PROCEDURE has no associated setter-procedure.

Note that (set! (setter PROC) ...) for a procedure that has no associated setter procedure yet is a very slow operation (the old procedure is replaced by a modified copy, which involves a garbage collection).

13.21.2 getter-with-setter

[procedure] (getter-with-setter GETTER SETTER)

Returns a copy of the procedure GETTER with the associated setter procedure SETTER. Contrary to the SRFI specification, the setter of the returned procedure may be changed.

13.22 Reader extensions

13.22.1 define-reader-ctor

[procedure] (define-reader-ctor SYMBOL PROC)

Define new read-time constructor for #, read syntax. For further information, see the documentation for <u>SRFI-10</u>.

13.22.2 set-read-syntax!

[procedure] (set-read-syntax! CHAR-OR-SYMBOL PROC)

When the reader encounters the non-whitespace character CHAR while reading an expression from a given port, then the procedure PROC will be called with that port as its argument. The procedure should return a value that will be returned to the reader:

```
; A simple RGB color syntax:

(set-read-syntax! #\%
```

```
  (lambda (port)
    (apply vector
      (map (cut string->number <> 16)
           (string-chop (read-string 6 port) 2) ) ) ) )

 (with-input-from-string "(1 2 %f0f0f0 3)" read)
 ; ==> (1 2 #(240 240 240) 3)
```

If `CHAR-OR-SYMBOL` is a symbol, then a so-called *read-mark* handler is defined. In that case the handler procedure will be called when a character-sequence of the form

`#!SYMBOL`

is encountered.

You can undo special handling of read-syntax by passing `#f` as the second argument (if the syntax was previously defined via `set-read-syntax!`).

Note that all of CHICKEN's special non-standard read-syntax is handled directly by the reader. To disable built-in read-syntax, define a handler that triggers an error (for example).

13.22.3 set-sharp-read-syntax!

[procedure] (set-sharp-read-syntax! CHAR-OR-SYMBOL PROC)

Similar to `set-read-syntax!`, but allows defining new `#<CHAR> ...` reader syntax. If the first argument is a symbol, then this procedure is equivalent to `set-read-syntax!`.

13.22.4 set-parameterized-read-syntax!

[procedure] (set-parameterized-read-syntax! CHAR-OR-SYMBOL PROC)

Similar to `set-sharp-read-syntax!`, but intended for defining reader syntax of the form `#<NUMBER><CHAR> ...`. The handler procedure `PROC` will be called with two arguments: the input port and the number preceding the dispatching character. If the first argument is a symbol, then this procedure is equivalent to `set-read-syntax!`.

13.22.5 copy-read-table

[procedure] (copy-read-table READ-TABLE)

Returns a copy of the given read-table. You can access the currently active read-table with `(current-read-table)`.

13.23 Property lists

As in other Lisp dialects, CHICKEN supports "property lists" associated with symbols. Properties are accessible via a key that can be any kind of value but which will be compared using `eq?`.

13.23.1 get

```
[procedure] (get SYMBOL PROPERTY [DEFAULT])
```

Returns the value stored under the key `PROPERTY` in the property list of `SYMBOL`. If no such property is stored, returns `DEFAULT`. The `DEFAULT` is optional and defaults to `#f`.

13.23.2 put!

```
[procedure] (put! SYMBOL PROPERTY VALUE)
[setter] (set! (get SYMBOL PROPERTY) VALUE)
```

Stores `VALUE` under the key `PROPERTY` in the property list of `SYMBOL` replacing any previously stored value.

13.23.3 remprop!

```
[procedure] (remprop! SYMBOL PROPERTY)
```

Deletes the first property matching the key `PROPERTY` in the property list of `SYMBOL`. Returns `#t` when a deletion performed, and `#f` otherwise.

13.23.4 symbol-plist

```
[procedure] (symbol-plist SYMBOL)
[setter] (set! (symbol-plist SYMBOL) LST)
```

Returns the property list of `SYMBOL` or sets it.

13.23.5 get-properties

```
[procedure] (get-properties SYMBOL PROPERTIES)
```

Searches the property list of `SYMBOL` for the first property with a key in the list `PROPERTIES`. Returns 3 values: the matching property key, value, and the tail of property list after the matching property. When no match found all values are `#f`.

`PROPERTIES` may also be an atom, in which case it is treated as a list of one element.

Previous: <u>Parameters</u>

Next: <u>Unit eval</u>

14 Unit eval

This unit has support for evaluation and macro-handling. This unit is used by default, unless the program is compiled with the `-explicit-use` option.

14.1 Loading code

14.1.1 load

[procedure] (load FILE [EVALPROC])

Loads and evaluates expressions from the given source file, which may be either a string or an input port. Each expression read is passed to `EVALPROC` (which defaults to `eval`). On platforms that support it (currently native Windows, Linux ELF and Solaris), `load` can be used to load compiled programs:

```
% cat x.scm
(define (hello) (print "Hello!"))
% csc -s x.scm
% csi -q
#;1> (load "x.so")
; loading x.so ...
#;2> (hello)
Hello!
#;3>
```

The second argument to `load` is ignored when loading compiled code. If source code is loaded from a port, then that port is closed after all expressions have been read.

Compiled code can be re-loaded, but care has to be taken, if code from the replaced dynamically loaded module is still executing (i.e. if an active continuation refers to compiled code in the old module).

Support for reloading compiled code dynamically is still experimental.

14.1.2 load-relative

[procedure] (load-relative FILE [EVALPROC])

Similar to `load`, but loads `FILE` relative to the path of the currently loaded file.

14.1.3 load-noisily

`[procedure] (load-noisily FILE #!key EVALUATOR TIME PRINTER)`

As `load` but the result(s) of each evaluated toplevel-expression is written to standard output. If `EVALUATOR` is given and not `#f`, then each expression is evaluated by calling this argument with the read expression as argument. If `TIME` is given and not false, then the execution time of each expression is shown (as with the `time` macro). If `PRINTER` is given and not false, then each expression is printed before evaluation by applying the expression to the value of this argument, which should be a one-argument procedure.

See also the <u>load-verbose</u> parameter.

14.1.4 load-library

`[procedure] (load-library UNIT [LIBRARYFILE])`

On platforms that support dynamic loading, `load-library` loads the compiled library unit `UNIT` (which should be a symbol). If the string `LIBRARYFILE` is given, then the given shared library will be loaded and the toplevel code of the contained unit will be executed. If no `LIBRARYFILE` argument is given, then the following libraries are checked for the required unit:

- a file named *<UNIT>.so*
- the files given in the parameter `dynamic-load-libraries`

If the unit is not found, an error is signaled. When the library unit can be successfully loaded, a feature-identifier named `UNIT` is registered. If the feature is already registered before loading, the `load-library` does nothing.

14.1.5 set-dynamic-load-mode!

`[procedure] (set-dynamic-load-mode! MODELIST)`

On systems that support dynamic loading of compiled code via the `dlopen(3)` interface (for example Linux and Solaris), some options can be specified to fine-tune the behaviour of the dynamic linker. `MODE` should be a list of symbols (or a single symbol) taken from the following set:

`local`
> If `local` is given, then any C/C++ symbols defined in the dynamically loaded file are not available for subsequently loaded files and libraries. Use this if you have linked foreign code into your dynamically loadable file and if you don't want to export them (for example because you want to load another file that defines the same symbols).

`global`
> The default is `global`, which means all C/C++ symbols are available to code loaded at a later stage.

`now`
> If `now` is specified, all symbols are resolved immediately.

`lazy`
> Unresolved symbols are resolved as code from the file is executed. This is the default.

Note that this procedure does not control the way Scheme variables are handled - this facility is mainly of interest when accessing foreign code.

14.2 Read-eval-print loop

14.2.1 repl

[procedure] (repl)

Start a new read-eval-print loop. Sets the `reset-handler` so that any invocation of `reset` restarts the read-eval-print loop. Also changes the current exception-handler to display a message, write any arguments to the value of (`current-error-port`) and reset.

14.3 Macros

14.3.1 get-line-number

[procedure] (get-line-number EXPR)

If EXPR is a pair with the car being a symbol, and line-number information is available for this expression, then this procedure returns the associated line number. If line-number information is not available, then `#f` is returned. Note that line-number information for expressions is only available in the compiler.

14.3.2 macro?

[procedure] (macro? SYMBOL)

Returns `#t` if there exists a macro-definition for SYMBOL.

14.3.3 macroexpand

[procedure] (macroexpand X)

If X is a macro-form, expand the macro (and repeat expansion until expression is a non-macro form). Returns the resulting expression.

14.3.4 macroexpand-1

[procedure] (macroexpand-1 X)

If X is a macro-form, expand the macro. Returns the resulting expression.

14.3.5 undefine-macro!

[procedure] (undefine-macro! SYMBOL)

Remove the current macro-definition of the macro named SYMBOL.

14.3.6 syntax-error

[procedure] (syntax-error [LOCATION] MESSAGE ARGUMENT ...)

Signals an exception of the kind (exn syntax). Otherwise identical to error.

14.4 Loading extension libraries

This functionality is only available on platforms that support dynamic loading of compiled code. Currently Linux, BSD, Solaris, Windows (with Cygwin) and HP/UX are supported.

14.4.1 repository-path

[parameter] repository-path

Contains a string naming the path to the extension repository, which defaults to either the value of the environment variable CHICKEN_REPOSITORY or the default library path (usually /usr/local/lib/chicken on UNIX systems).

14.4.2 extension-information

[procedure] (extension-information ID)

If an extension with the name ID is installed and if it has a setup-information list registered in the extension repository, then the info-list is returned. Otherwise extension-information returns #f.

14.4.3 provide

```
[procedure] (provide ID ...)
```

Registers the extension IDs `ID ...` as loaded. This is mainly intended to provide aliases for certain extension identifiers.

14.4.4 provided?

```
[procedure] (provided? ID ...)
```

Returns `#t` if the extension with the IDs `ID ...` are currently loaded, or `#f` otherwise.

14.4.5 require

```
[procedure] (require ID ...)
[procedure] (require-for-syntax ID ...)
```

If the extension library `ID` is not already loaded into the system, then `require` will lookup the location of the shared extension library and load it. If `ID` names a library-unit of the base system, then it is loaded via `load-library`. If no extension library is available for the given ID, then an attempt is made to load the file `ID.so` or `ID.scm` (in that order) from one of the following locations:

- the current include path, which defaults to the pathnames given in `CHICKEN_INCLUDE_PATH`.
- the current directory

`ID` should be a string or a symbol. The difference between `require` and `require-for-syntax` is the the latter loads the extension library at compile-time (the argument is still evaluated), while the former loads it at run-time.

14.4.6 set-extension-specifier!

```
[procedure] (set-extension-specifier! SYMBOL PROC)
```

Registers the handler-procedure `PROC` as a extension-specifier with the name `SYMBOL`. This facility allows extending the set of valid extension specifiers to be used with `require-extension`. When `register-extension` is called with an extension specifier of the form `(SPEC ...)` and `SPEC` has been registered with `set-extension-specifier!`, then `PROC` will be called with two arguments: the specifier and the previously installed handler (or `#f` if no such handler was defined). The handler should return a new specifier that will be processed recursively. If the handler returns a vector, then each element of the vector will be processed recursively. Alternatively the handler may return a string which specifies a file to be loaded:

```
(eval-when (compile eval)
  (set-extension-specifier!
    'my-package
    (lambda (spec old)
```

```
     (make-pathname my-package-directory (->string (cadr spec))) ) ) )
```

```
(require-extension (my-package stuff))     ; --> expands into '(load "my-packag
```

Note that the handler has to be registered at compile time, if it is to be visible in compiled code.

14.5 System information

14.5.1 chicken-home

```
[procedure] (chicken-home)
```

Returns a string given the installation directory (usually `/usr/local/share/chicken` on UNIX-like systems). As a last option, if the environment variable `CHICKEN_PREFIX` is set, then `chicken-home` will return `$CHICKEN_PREFIX/share`.

14.6 Eval

14.6.1 eval

```
[procedure] (eval EXP [ENVIRONMENT])
```

Evaluates `EXP` and returns the result of the evaluation. The second argument is optional and defaults to the value of (`interaction-environment`).

Previous: <u>Unit library</u>

Next: <u>Unit extras</u>

15 Unit extras

This unit contains a collection of useful utility definitions. This unit is used by default, unless the program is compiled with the `-explicit-use` option.

15.1 Lists

15.1.1 alist-ref

[procedure] (alist-ref KEY ALIST [TEST [DEFAULT]])

Looks up KEY in ALIST using TEST as the comparison function (or `eqv?` if no test was given) and returns the cdr of the found pair, or DEFAULT (which defaults to `#f`).

15.1.2 alist-update!

[procedure] (alist-update! KEY VALUE ALIST [TEST])

If the list ALIST contains a pair of the form (KEY . X), then this procedure replaces X with VALUE and returns ALIST. If ALIST contains no such item, then `alist-update!` returns ((KEY . VALUE) . ALIST). The optional argument TEST specifies the comparison procedure to search a matching pair in ALIST and defaults to `eqv?`.

15.1.3 atom?

[procedure] (atom? X)

Returns `#t` if X is not a pair. This is identical to `not-pair?` from <u>Unit srfi-1</u> but kept for historical reasons.

15.1.4 rassoc

[procedure] (rassoc KEY LIST [TEST])

Similar to `assoc`, but compares KEY with the `cdr` of each pair in LIST using TEST as the comparison procedures (which defaults to `eqv?`.

15.1.5 butlast

```
[procedure] (butlast LIST)
```

Returns a fresh list with all elements but the last of LIST.

15.1.6 chop

```
[procedure] (chop LIST N)
```

Returns a new list of sublists, where each sublist contains N elements of LIST. If LIST has a length that is not a multiple of N, then the last sublist contains the remaining elements.

```
(chop '(1 2 3 4 5 6) 2) ==> ((1 2) (3 4) (5 6))
(chop '(a b c d) 3)      ==> ((a b c) (d))
```

15.1.7 compress

```
[procedure] (compress BLIST LIST)
```

Returns a new list with elements taken from LIST with corresponding true values in the list BLIST.

```
(define nums '(99 100 110 401 1234))
(compress (map odd? nums) nums)      ==> (99 401)
```

15.1.8 flatten

```
[procedure] (flatten LIST1 ...)
```

Returns LIST1 ... concatenated together, with nested lists removed (flattened).

15.1.9 intersperse

```
[procedure] (intersperse LIST X)
```

Returns a new list with X placed between each element.

15.1.10 join

`[procedure] (join LISTOFLISTS [LIST])`

Concatenates the lists in `LISTOFLISTS` with `LIST` placed between each sublist. `LIST` defaults to the empty list.

```
(join '((a b) (c d) (e)) '(x y)) ==> (a b x y c d x y e)
(join '((p q) () (r (s) t)) '(-))  ==> (p q - - r (s) t)
```

`join` could be implemented as follows:

```
(define (join lstoflsts #!optional (lst '()))
  (apply append (intersperse lstoflists lst)) )
```

15.1.11 shuffle

`[procedure] (shuffle LIST)`

Returns `LIST` with its elements sorted in a random order.

15.1.12 tail?

`[procedure] (tail? X LIST)`

Returns true if `X` is one of the tails (cdr's) of `LIST`.

15.2 String-port extensions

15.2.1 call-with-input-string

`[procedure] (call-with-input-string STRING PROC)`

Calls the procedure `PROC` with a single argument that is a string-input-port with the contents of `STRING`.

15.2.2 call-with-output-string

`[procedure] (call-with-output-string PROC)`

Calls the procedure `PROC` with a single argument that is a string-output-port. Returns the accumulated output-string.

15.2.3 with-input-from-string

[procedure] (with-input-from-string STRING THUNK)

Call procedure THUNK with the current input-port temporarily bound to an input-string-port with the contents of STRING.

15.2.4 with-output-to-string

[procedure] (with-output-to-string THUNK)

Call procedure THUNK with the current output-port temporarily bound to a string-output-port and return the accumulated output string.

15.3 Formatted output

15.3.1 printf

15.3.2 fprintf

15.3.3 sprintf

[procedure] (fprintf PORT FORMATSTRING ARG ...)
[procedure] (printf FORMATSTRING ARG ...)
[procedure] (sprintf FORMATSTRING ARG ...)

Simple formatted output to a given port (fprintf), the value of (current-output-port) (printf), or a string (sprintf). The FORMATSTRING can contain any sequence of characters. There must be at least as many ARG arguments given as there are format directives that require an argument in FORMATSTRING. Extra ARG arguments are ignored. The character `~' prefixes special formatting directives:

~%

write newline character

~N

the same as ~%

~S

write the next argument

`~A`

display the next argument

`~\n`

skip all whitespace in the format-string until the next non-whitespace character

`~B`

write the next argument as a binary number

`~O`

write the next argument as an octal number

`~X`

write the next argument as a hexadecimal number

`~C`

write the next argument as a character

`~~`

display `~'

`~!`

flush all pending output

`~?`

invoke formatted output routine recursively with the next two arguments as format-string and list of parameters

15.3.4 format

`[procedure] (format [DESTINATION] FORMATSTRING ARG ...)`

The parameters `FORMATSTRING` and `ARG ...` are as for (`printf/sprintf/fprintf`).

The optional `DESTINATION`, when supplied, performs a (`sprintf`) for a `#f`, a (`printf`) for a `#t`, and a (`fprintf`) for an output-port. When missing a (`sprintf`) is performed.

15.4 Hash tables

CHICKEN implements SRFI-69. For more information, see SRFI-69.

A setter for `hash-table-ref` is defined, so

`(set! (hash-table-ref HT KEY) VAL)`

is equivalent to

15.3.3 sprintf   93

```
(hash-table-set! HT KEY VAL)
```

As an extension to SRFI-69, `hash-table-update!` and `hash-table-update!/default` return the new value (after applying the update procedure).

15.4.1 hash-table-remove!

[procedure] (hash-table-remove! HASHTABLE PROC)

Calls `PROC` for all entries in `HASHTABLE` with the key and value of each entry. If `PROC` returns true, then that entry is removed.

15.5 Queues

15.5.1 listqueue

[procedure] (list->queue LIST)

Returns `LIST` converted into a queue, where the first element of the list is the same as the first element of the queue. The resulting queue may share memory with the list and the list should not be modified after this operation.

15.5.2 make-queue

[procedure] (make-queue)

Returns a newly created queue.

15.5.3 queue?

[procedure] (queue? X)

Returns `#t` if X is a queue, or `#f` otherwise.

15.5.4 queuelist

[procedure] (queue->list QUEUE)

Returns QUEUE converted into a list, where the first element of the list is the same as the first element of the queue. The resulting list may share memory with the queue object and should not be modified.

15.5.5 queue-add!

[procedure] (queue-add! QUEUE X)

Adds X to the rear of QUEUE.

15.5.6 queue-empty?

[procedure] (queue-empty? QUEUE)

Returns #t if QUEUE is empty, or #f otherwise.

15.5.7 queue-first

[procedure] (queue-first QUEUE)

Returns the first element of QUEUE. If QUEUE is empty an error is signaled

15.5.8 queue-last

[procedure] (queue-last QUEUE)

Returns the last element of QUEUE. If QUEUE is empty an error is signaled

15.5.9 queue-remove!

[procedure] (queue-remove! QUEUE)

Removes and returns the first element of QUEUE. If QUEUE is empty an error is signaled

15.5.10 queue-push-back!

[procedure] (queue-push-back! QUEUE ITEM)

Pushes an item into the first position of a queue, i.e. the next queue-remove! will return ITEM.

15.5.11 queue-push-back-list!

[procedure] (queue-push-back-list! QUEUE LIST)

Pushes the items in item-list back onto the queue, so that `(car LIST)` becomes the next removable item.

15.6 Sorting

15.6.1 merge

[procedure] (merge LIST1 LIST2 LESS?)
[procedure] (merge! LIST1 LIST2 LESS?)

Joins two lists in sorted order. `merge!` is the destructive version of merge. `LESS?` should be a procedure of two arguments, that returns true if the first argument is to be ordered before the second argument.

15.6.2 sort

[procedure] (sort SEQUENCE LESS?)
[procedure] (sort! SEQUENCE LESS?)

Sort `SEQUENCE`, which should be a list or a vector. `sort!` is the destructive version of sort.

15.6.3 sorted?

[procedure] (sorted? SEQUENCE LESS?)

Returns true if the list or vector `SEQUENCE` is already sorted.

15.7 Random numbers

15.7.1 random

[procedure] (random N)

Returns an exact random integer from 0 to `N`-1.

15.7.2 randomize

`[procedure] (randomize [X])`

Set random-number seed. If X is not supplied, the current time is used. On startup (when the `extras` unit is initialized), the random number generator is initialized with the current time.

15.8 Input/Output extensions

15.8.1 make-input-port

`[procedure] (make-input-port READ READY? CLOSE [PEEK])`

Returns a custom input port. Common operations on this port are handled by the given parameters, which should be procedures of no arguments. READ is called when the next character is to be read and should return a character or `#!eof`. READY? is called when `char-ready?` is called on this port and should return `#t` or `#f`. CLOSE is called when the port is closed. PEEK is called when `peek-char` is called on this port and should return a character or `#!eof`. if the argument PEEK is not given, then READ is used instead and the created port object handles peeking automatically (by calling READ and buffering the character).

15.8.2 make-output-port

`[procedure] (make-output-port WRITE CLOSE [FLUSH])`

Returns a custom output port. Common operations on this port are handled by the given parameters, which should be procedures. WRITE is called when output is sent to the port and receives a single argument, a string. CLOSE is called when the port is closed and should be a procedure of no arguments. FLUSH (if provided) is called for flushing the output port.

15.8.3 pretty-print

`[procedure] (pretty-print EXP [PORT])`
`[procedure] (pp EXP [PORT])`

Print expression nicely formatted. PORT defaults to the value of `(current-output-port)`.

15.8.4 pretty-print-width

(Parameter) Specifies the maximal line-width for pretty printing, after which line wrap will occur.

15.8.5 read-byte

15.8.6 write-byte

```
[procedure] (read-byte [PORT])
[procedure] (write-byte BYTE [PORT])
```

Read/write a byte to the port given in `PORT`, which default to the values of `(current-input-port)` and `(current-output-port)`, respectively.

15.8.7 read-file

```
[procedure] (read-file [FILE-OR-PORT [READER [MAXCOUNT]]])
```

Returns a list containing all toplevel expressions read from the file or port `FILE-OR-PORT`. If no argument is given, input is read from the port that is the current value of `(current-input-port)`. After all expressions are read, and if the argument is a port, then the port will not be closed. The `READER` argument specifies the procedure used to read expressions from the given file or port and defaults to `read`. The reader procedure will be called with a single argument (an input port). If `MAXCOUNT` is given then only up to `MAXCOUNT` expressions will be read in.

15.8.8 read-line

15.8.9 write-line

```
[procedure] (read-line [PORT [LIMIT]])
[procedure] (write-line STRING [PORT])
```

Line-input and -output. `PORT` defaults to the value of `(current-input-port)` and `(current-output-port)`, respectively. If the optional argument `LIMIT` is given and not `#f`, then `read-line` reads at most `LIMIT` characters per line. `read-line` returns a string without the terminating newline and `write-line` adds a terminating newline before outputting.

15.8.10 read-lines

```
[procedure] (read-lines [PORT [MAX]])
```

Read MAX or fewer lines from PORT. PORT defaults to the value of `(current-input-port)`. PORT may optionally be a string naming a file. Returns a list of strings, each string representing a line read, not including any line separation character(s).

15.8.11 read-string

15.8.12 read-string!

15.8.13 write-string

```
[procedure] (read-string [NUM [PORT]])
[procedure] (read-string! NUM STRING [PORT [START]])
[procedure] (write-string STRING [NUM [PORT]]
```

Read or write NUM characters from/to PORT, which defaults to the value of `(current-input-port)` or `(current-output-port)`, respectively. If NUM is `#f` or not given, then all data up to the end-of-file is read, or, in the case of `write-string` the whole string is written. If no more input is available, `read-string` returns the empty string. `read-string!` reads destructively into the given STRING argument, but never more characters that would fit into STRING. If START is given, then the read characters are stored starting at that position. `read-string!` returns the actual number of characters read.

15.8.14 read-token

```
[procedure] (read-token PREDICATE [PORT])
```

Reads characters from PORT (which defaults to the value of `(current-input-port)`) and calls the procedure PREDICATE with each character until PREDICATE returns false. Returns a string with the accumulated characters.

15.8.15 with-error-output-to-port

```
[procedure] (with-error-output-to-port PORT THUNK)
```

Call procedure THUNK with the current error output-port temporarily bound to PORT.

15.8.16 with-input-from-port

[procedure] (with-input-from-port PORT THUNK)

Call procedure THUNK with the current input-port temporarily bound to PORT.

15.8.17 with-output-to-port

[procedure] (with-output-to-port PORT THUNK)

Call procedure THUNK with the current output-port temporarily bound to PORT.

15.9 Strings

15.9.1 conc

[procedure] (conc X ...)

Returns a string with the string-represenation of all arguments concatenated together. conc could be implemented as

```
(define (conc . args)
  (apply string-append (map ->string args)) )
```

15.9.2 string

[procedure] (->string X)

Returns a string-representation of X.

15.9.3 string-chop

[procedure] (string-chop STRING LENGTH)

Returns a list of substrings taken by *chopping* STRING every LENGTH characters:

```
(string-chop "one two three" 4)  ==>  ("one " "two " "thre" "e")
```

15.9.4 string-chomp

[procedure] (string-chomp STRING [SUFFIX])

If STRING ends with SUFFIX, then this procedure returns a copy of its first argument with the suffix removed, otherwise returns STRING unchanged. SUFFIX defaults to "\n".

15.9.5 string-compare3

[procedure] (string-compare3 STRING1 STRING2)
[procedure] (string-compare3-ci STRING1 STRING2)

Perform a three-way comparison between the STRING1 and STRING2, returning either -1 if STRING1 is lexicographically less than STRING2, 0 if it is equal, or 1 if it s greater. string-compare3-ci performs a case-insensitive comparison.

15.9.6 string-intersperse

[procedure] (string-intersperse LIST [STRING])

Returns a string that contains all strings in LIST concatenated together. STRING is placed between each concatenated string and defaults to " ".

(string-intersperse '("one" "two") "three")

is equivalent to

(apply string-append (intersperse '("one" "two") "three"))

15.9.7 string-split

[procedure] (string-split STRING [DELIMITER-STRING [KEEPEMPTY]])

Split string into substrings separated by the given delimiters. If no delimiters are specified, a string comprising the tab, newline and space characters is assumed. If the parameter KEEPEMPTY is given and not #f, then empty substrings are retained:

(string-split "one two three") ==> ("one" "two" "three")
(string-split "foo:bar::baz:" ":" #t) ==> ("foo" "bar" "" "baz" "")

15.9.8 string-translate

[procedure] (string-translate STRING FROM [TO])

Returns a fresh copy of STRING with characters matching FROM translated to TO. If TO is omitted, then matching characters are removed. FROM and TO may be a character, a string or a list. If both FROM and TO are strings, then the character at the same position in TO as the matching character in FROM is substituted.

15.9.9 string-translate*

[procedure] (string-translate* STRING SMAP)

Substitutes elements of STRING according to SMAP. SMAP should be an association-list where each element of the list is a pair of the form (MATCH \. REPLACEMENT). Every occurrence of the string MATCH in STRING will be replaced by the string REPLACEMENT:

```
(string-translate*
  "<h1>this is a \"string\"</h1>"
  '(("<" . "&lt;") (">" . "&gt;") ("\"" . "&quot;")) )
=>  "&lt;h1&gt;this is a &quot;string&quot;&lt;/ht&gt;"
```

15.9.10 substring=?

[procedure] (substring=? STRING1 STRING2 [START1 [START2 [LENGTH]]])
[procedure] (substring-ci=? STRING1 STRING2 [START1 [START2 [LENGTH]]])

Returns #t if the strings STRING1 and STRING2 are equal, or #f otherwise. The comparison starts at the positions START1 and START2 (which default to 0), comparing LENGTH characters (which defaults to the minimum of the remaining length of both strings).

15.9.11 substring-index

[procedure] (substring-index WHICH WHERE [START])
[procedure] (substring-index-ci WHICH WHERE [START])

Searches for first index in string WHERE where string WHICH occurs. If the optional argument START is given, then the search starts at that index. substring-index-ci is a case-insensitive version of substring-index.

15.10 Combinators

15.10.1 any?

[procedure] (any? X)

Ignores its argument and always returns `#t`. This is actually useful sometimes.

15.10.2 constantly

[procedure] (constantly X ...)

Returns a procedure that always returns the values `X ...` regardless of the number and value of its arguments.

```
(constantly X) <=> (lambda args X)
```

15.10.3 complement

[procedure] (complement PROC)

Returns a procedure that returns the boolean inverse of `PROC`.

```
(complement PROC) <=> (lambda (x) (not (PROC x)))
```

15.10.4 compose

[procedure] (compose PROC1 PROC2 ...)

Returns a procedure that represents the composition of the argument-procedures `PROC1 PROC2 ...`.

```
(compose F G) <=> (lambda args
                    (call-with-values
                      (lambda () (apply G args))
                      F))
```

`(compose)` is equivalent to `values`.

15.10.5 conjoin

```
[procedure] (conjoin PRED ...)
```

Returns a procedure that returns `#t` if its argument satisfies the predicates `PRED ....`

```
((conjoin odd? positive?) 33)   ==>  #t
((conjoin odd? positive?) -33)  ==>  #f
```

15.10.6 disjoin

```
[procedure] (disjoin PRED ...)
```

Returns a procedure that returns `#t` if its argument satisfies any predicate `PRED ....`

```
((disjoin odd? positive?) 32)    ==>  #t
((disjoin odd? positive?) -32)   ==>  #f
```

15.10.7 each

```
[procedure] (each PROC ...)
```

Returns a procedure that applies `PROC ...` to its arguments, and returns the result(s) of the last procedure application. For example

```
(each pp eval)
```

is equivalent to

```
(lambda args
  (apply pp args)
  (apply eval args) )
```

`(each PROC)` is equivalent to `PROC` and `(each)` is equivalent to `noop`.

15.10.8 flip

```
[procedure] (flip PROC)
```

Returns a two-argument procedure that calls `PROC` with its arguments swapped:

```
(flip PROC) <=> (lambda (x y) (PROC y x))
```

15.10.9 identity

```
[procedure] (identity X)
```

Returns its sole argument X.

15.10.10 project

```
[procedure] (project N)
```

Returns a procedure that returns its Nth argument (starting from 0).

15.10.11 list-of

```
[procedure] (list-of PRED)
```

Returns a procedure of one argument that returns #t when applied to a list of elements that all satisfy the predicate procedure PRED, or #f otherwise.

```
((list-of even?) '(1 2 3))   ==> #f
((list-of number?) '(1 2 3)) ==> #t
```

15.10.12 noop

```
[procedure] (noop X ...)
```

Ignores it's arguments, does nothing and returns an unspecified value.

15.10.13 o

```
[procedure] (o PROC ...)
```

A single value version of compose (slightly faster). (o) is equivalent to identity.

15.11 Binary searching

15.11.1 binary-search

```
[procedure] (binary-search SEQUENCE PROC)
```

Performs a binary search in SEQUENCE, which should be a sorted list or vector. PROC is called to compare items in the sequence, should accept a single argument and return an exact integer: zero if the searched value is equal to the current item, negative if the searched value is *less* than the current item, and positive otherwise. Returns the index of the found value or #f otherwise.

Previous: Unit eval

Next: Unit srfi-1

16 Unit srfi-1

List library, see the documentation for <u>SRFI-1</u>

Previous: <u>Unit extras</u>

Next: <u>Unit srfi-4</u>

17 Unit srfi-4

Homogeneous numeric vectors, see the documentation for <u>SRFI-4</u> 64-bit integer vectors (`u64vector` and `s64vector` are not supported.

The basic constructor procedures for number vectors are extended to allow allocating the storage in non garbage collected memory:

17.1 make-XXXvector

`[procedure] (make-XXXvector SIZE [INIT NONGC FINALIZE])`

Creates a SRFI-4 homogenous number vector of length `SIZE`. If `INIT` is given, it specifies the initial value for each slot in the vector. The optional arguments `NONGC` and `FINALIZE` define whether the vector should be allocated in a memory area not subject to garbage collection and whether the associated storage should be automatically freed (using finalization) when there are no references from Scheme variables and data. `NONGC` defaults to `#f` (the vector will be located in normal garbage collected memory) and `FINALIZE` defaults to `#t`. Note that the `FINALIZE` argument is only used when `NONGC` is true.

Additionally, the following procedures are provided:

17.2 u8vectorblob

17.3 s8vectorblob

17.4 u16vectorblob

17.5 s16vectorblob

17.6 u32vectorblob

17.7 s32vectorblob

17.8 f32vectorblob

17.9 f64vectorblob

17.10 u8vectorblob/shared

17.11 s8vectorblob/shared

17.12 u16vectorblob/shared

17.13 s16vectorblob/shared

17.14 u32vectorblob/shared

17.15 s32vectorblob/shared

17.16 f32vectorblob/shared

17.17 f64vectorblob/shared

```
[procedure] (u8vector->blob U8VECTOR)
[procedure] (s8vector->blob S8VECTOR)
[procedure] (u16vector->blob U16VECTOR)
[procedure] (s16vector->blob S16VECTOR)
```

```
[procedure] (u32vector->blob U32VECTOR)
[procedure] (s32vector->blob S32VECTOR)
[procedure] (f32vector->blob F32VECTOR)
[procedure] (f64vector->blob F64VECTOR)
[procedure] (u8vector->blob/shared U8VECTOR)
[procedure] (s8vector->blob/shared S8VECTOR)
[procedure] (u16vector->blob/shared U16VECTOR)
[procedure] (s16vector->blob/shared S16VECTOR)
[procedure] (u32vector->blob/shared U32VECTOR)
[procedure] (s32vector->blob/shared S32VECTOR)
[procedure] (f32vector->blob/shared F32VECTOR)
[procedure] (f64vector->blob/shared F64VECTOR)
```

Each of these procedures return the contents of the given vector as a 'packed' blob. The byte order in that vector is platform-dependent (for example little-endian on an **Intel** processor). The `/shared` variants return a blob that shares memory with the contents of the vector.

17.18 blobu8vector

17.19 blobs8vector

17.20 blobu16vector

17.21 blobs16vector

17.22 blobu32vector

17.23 blobs32vector

17.24 blobf32vector

17.25 blobf64vector

17.26 blobu8vector/shared

17.27 blobs8vector/shared

17.28 blobu16vector/shared

17.29 blobs16vector/shared

17.30 blobu32vector/shared

17.31 blobs32vector/shared

17.32 blobf32vector/shared

17.33 blobf64vector/shared

```
[procedure] (blob->u8vector BLOB)
[procedure] (blob->s8vector BLOB)
[procedure] (blob->u16vector BLOB)
[procedure] (blob->s16vector BLOB)
[procedure] (blob->u32vector BLOB)
[procedure] (blob->s32vector BLOB)
[procedure] (blob->f32vector BLOB)
[procedure] (blob->f64vector BLOB)
[procedure] (blob->u8vector/shared BLOB)
[procedure] (blob->s8vector/shared BLOB)
[procedure] (blob->u16vector/shared BLOB)
[procedure] (blob->s16vector/shared BLOB)
[procedure] (blob->u32vector/shared BLOB)
[procedure] (blob->s32vector/shared BLOB)
```

```
[procedure] (blob->f32vector/shared BLOB)
[procedure] (blob->f64vector/shared BLOB)
```

Each of these procedures return a vector where the argument `BLOB` is taken as a 'packed' representation of the contents of the vector. The `/shared` variants return a vector that shares memory with the contents of the blob.

17.34 subu8vector

17.35 subu16vector

17.36 subu32vector

17.37 subs8vector

17.38 subs16vector

17.39 subs32vector

17.40 subf32vector

17.41 subf64vector

```
[procedure] (subu8vector U8VECTOR FROM TO)
[procedure] (subu16vector U16VECTOR FROM TO)
[procedure] (subu32vector U32VECTOR FROM TO)
[procedure] (subs8vector S8VECTOR FROM TO)
[procedure] (subs16vector S16VECTOR FROM TO)
[procedure] (subs32vector S32VECTOR FROM TO)
[procedure] (subf32vector F32VECTOR FROM TO)
[procedure] (subf64vector F64VECTOR FROM TO)
```

Creates a number vector of the same type as the argument vector with the elements at the positions FROM up to but not including TO.

SRFI-17 Setters for XXXvector-ref are defined.

17.42 read-u8vector

[procedure] (read-u8vector LENGTH [PORT])

Reads LENGTH bytes from the PORT and returns a fresh u8vector or less if end-of-file is encountered. PORT defaults to the value of (current-input-port). If LENGTH is #f, the vector will be filled completely until end-of-file is reached.

17.43 read-u8vector!

[procedure] (read-u8vector! LENGTH U8VECTOR [PORT [START]])

Reads LENGTH bytes from the PORT writing the read input into U8VECTOR beginning at START (or 0 if not given). PORT defaults to the value of (current-input-port). If LENGTH is #f, the vector will be filled completely until end-of-file is reached. This procedure returns the number of bytes read.

17.44 write-u8vector

[procedure] (write-u8vector U8VECTOR [PORT [START [END]]])

Writes the bytes U8VECTOR between the indices START (inclusive) and END (exclusive) to PORT. PORT defaults to the value of (current-output-port).

Previous: <u>Unit srfi-1</u>

Next: <u>Unit srfi-13</u>

18 Unit srfi-13

String library, see the documentation for <u>SRFI-13</u>

On systems that support dynamic loading, the `srfi-13` unit can be made available in the interpreter (`csi`) by entering

```
(require-extension srfi-13)
```

Previous: <u>Unit srfi-4</u>

Next: <u>Unit srfi-14</u>

19 Unit srfi-14

Character set library, see the documentation for <u>SRFI-14</u>

On systems that support dynamic loading, the `srfi-14` unit can be made available in the interpreter (`csi`) by entering

```
(require-extension srfi-14)
```

This library provides only the Latin-1 character set.

Previous: <u>Unit srfi-13</u>

Next: <u>Unit match</u>

20 Unit match

The runtime-support code for the Pattern Matching extensions. Note that to use the macros in normal compiled code it is not required to declare this unit as used. It is only necessary to do so if forms containing these macros are to be expanded at runtime.

Previous: Unit srfi-14

Next: Unit regex

21 Unit regex

This library unit provides support for regular expressions. The regular expression package used is PCRE (*Perl Compatible Regular Expressions*) written by Philip Hazel. See http://www.pcre.org for information about the particular regexp flavor and extensions provided by this library.

To test that PCRE support has been built into Chicken properly, try:

```
(require 'regex)
(test-feature? 'pcre) => t
```

21.1 grep

[procedure] (grep REGEX LIST)

Returns all items of LIST that match the regular expression REGEX. This procedure could be defined as follows:

```
(define (grep regex lst)
  (filter (lambda (x) (string-search regex x)) lst) )
```

21.2 globregexp

[procedure] (glob->regexp PATTERN)

Converts the file-pattern PATTERN into a regular expression.

```
(glob->regexp "foo.*")
=> "foo\..*"
```

PATTERN should follow "glob" syntax. Allowed wildcards are

```
*
[C...]
[C1-C2]
[-C...]
?
```

21.3 glob?

[procedure] (glob? STRING)

Does the STRING have any "glob" wildcards?

A string without any "glob" wildcards does not meet the criteria, even though it technically is a valid "glob" file-pattern.

21.4 regexp

```
[procedure] (regexp STRING [IGNORECASE [IGNORESPACE [UTF8]]])
```

Returns a precompiled regular expression object for `string`. The optional arguments `IGNORECASE`, `IGNORESPACE` and `UTF8` specify whether the regular expression should be matched with case- or whitespace-differences ignored, or whether the string should be treated as containing UTF-8 encoded characters, respectively.

21.5 regexp*

```
[procedure] (regexp* STRING [OPTIONS [TABLES]])
```

Returns a precompiled regular expression object for `string`. The optional argument `OPTIONS` must be a list of option symbols. The optional argument `TABLES` must be a character definitions table (not defined here).

Option Symbols:

caseless
      Character case insensitive match
multiline
      Equivalent to Perl's /m option
dotall
      Equivalent to Perl's /s option
extended
      Ignore whitespace
anchored
      Anchor pattern match
dollar-endonly
      `$' metacharacter in the pattern matches only at the end of the subject string
extra
      Currently of very little use
notbol
      First character of the string is not the beginning of a line
noteol
      End of the string is not the end of a line
ungreedy
      Inverts the "greediness" of the quantifiers so that they are not greedy by default
notempty
      The empty string is not considered to be a valid match
utf8
      UTF-8 encoded characters
no-auto-capture
      Disables the use of numbered capturing parentheses
no-utf8-check
      Skip valid UTF-8 sequence check
auto-callout
      Automatically inserts callout items (not defined here)
partial

Partial match ok

firstline

An unanchored pattern is required to match before or at the first newline

dupnames

Names used to identify capturing subpatterns need not be unique

newline-cr

Newline definition is `\r'

newline-lf

Newline definition is `\n'

newline-crlf

Newline definition is `\r\n'

newline-anycrlf

Newline definition is any of `\r', `\n', or `\r\n'

newline-any

ewline definition is any Unicode newline sequence

bsr-anycrlf

`\R' escape sequence matches only CR, LF, or CRLF

bsr-unicode

`\R' escape sequence matches only Unicode newline sequence

dfa-shortest

Currently unused

dfa-restart

Currently unused

21.6 regexp?

```
[procedure] (regexp? X)
```

Returns `#t` if `X` is a precompiled regular expression, or `#f` otherwise.

21.7 regexp-optimize

```
[procedure] (regexp-optimize RX)
```

Perform available optimizations for the precompiled regular expression `RX`. Returns `#t` when optimization performed, and `#f` otherwise.

21.8 string-match

21.9 string-match-positions

[procedure] (string-match REGEXP STRING [START])
[procedure] (string-match-positions REGEXP STRING [START])

Matches the regular expression in REGEXP (a string or a precompiled regular expression) with STRING and returns either #f if the match failed, or a list of matching groups, where the first element is the complete match. If the optional argument START is supplied, it specifies the starting position in STRING. For each matching group the result-list contains either: #f for a non-matching but optional group; a list of start- and end-position of the match in STRING (in the case of string-match-positions); or the matching substring (in the case of string-match). Note that the exact string is matched. For searching a pattern inside a string, see below. Note also that string-match is implemented by calling string-search with the regular expression wrapped in ^ ... $. If invoked with a precompiled regular expression argument (by using regexp), string-match is identical to string-search.

21.10 string-search

21.11 string-search-positions

[procedure] (string-search REGEXP STRING [START [RANGE]])
[procedure] (string-search-positions REGEXP STRING [START [RANGE]])

Searches for the first match of the regular expression in REGEXP with STRING. The search can be limited to RANGE characters.

21.12 string-split-fields

[procedure] (string-split-fields REGEXP STRING [MODE [START]])

Splits STRING into a list of fields according to MODE, where MODE can be the keyword #:infix (REGEXP matches field separator), the keyword #:suffix (REGEXP matches field terminator) or #t (REGEXP matches field), which is the default.

```
(define s "this is a string 1, 2, 3,")

(string-split-fields "[^ ]+" s)

  => ("this" "is" "a" "string" "1," "2," "3,")

(string-split-fields " " s #:infix)

  => ("this" "is" "a" "string" "1," "2," "3,")

(string-split-fields "," s #:suffix)

  => ("this is a string 1" " 2" " 3")
```

21.13 string-substitute

[procedure] (string-substitute REGEXP SUBST STRING [MODE])

Searches substrings in STRING that match REGEXP and substitutes them with the string SUBST. The substitution can contain references to subexpressions in REGEXP with the \NUM notation, where NUM refers to the NUMth parenthesized expression. The optional argument MODE defaults to 1 and specifies the number of the match to be substituted. Any non-numeric index specifies that all matches are to be substituted.

```
(string-substitute "([0-9]+) (eggs|chicks)"
                    "\\2 (\\1)" "99 eggs or 99 chicks" 2)
=> "99 eggs or chicks (99)"
```

Note that a regular expression that matches an empty string will signal an error.

21.14 string-substitute*

[procedure] (string-substitute* STRING SMAP [MODE])

Substitutes elements of STRING with string-substitute according to SMAP. SMAP should be an association-list where each element of the list is a pair of the form (MATCH . REPLACEMENT). Every occurrence of the regular expression MATCH in STRING will be replaced by the string REPLACEMENT

```
(string-substitute* "<h1>Hello, world!</h1>"
                    '(("<[/A-Za-z0-9]+>" . "")))

=>  "Hello, world!"
```

21.15 regexp-escape

[procedure] (regexp-escape STRING)

Escapes all special characters in STRING with \, so that the string can be embedded into a regular expression.

```
(regexp-escape "^[0-9]+:.*$")
=>  "\\^\\[0-9\\]\\+:.\n.\\*\\$"
```

21.16 make-anchored-pattern

[procedure] (make-anchored-pattern REGEXP [WITHOUT-BOL [WITHOUT-EOL]])

Makes an anchored pattern from REGEXP (a string or a precompiled regular expression) and returns the updated pattern. When WITHOUT-BOL is #t the beginning-of-line anchor is not added. When

`WITHOUT-EOL` is `#t` the end-of-line anchor is not added.

The `WITHOUT-BOL` and {WITHOUT-EOL}} arguments are ignored for a precompiled regular expression.

22 Unit srfi-18

A simple multithreading package. This threading package follows largely the specification of SRFI-18. For more information see the documentation for SRFI-18.

Notes:

- `thread-start!` accepts a thunk (a zero argument procedure) as argument, which is equivalent to `(thread-start! (make-thread THUNK))`.
- `thread-sleep!` accepts a seconds real number value in addition to a time object.
- When an uncaught exception (i.e. an error) is signalled in a thread other than the primordial thread and warnings are enabled (see: `enable-warnings`, then a warning message is written to the port that is the value of `(current-error-port)`.
- Blocking I/O will block all threads, except for some socket operations (see the section about the `tcp` unit). An exception is the read-eval-print loop on UNIX platforms: waiting for input will not block other threads, provided the current input port reads input from a console.
- It is generally not a good idea for one thread to call a continuation created by another thread, if `dynamic-wind` is involved.
- When more than one thread compete for the current time-slice, the thread that was waiting first will become the next runnable thread.
- The dynamic environment of a thread consists of the following state:

    - The current input-, output- and error-port
    - The current exception handler
    - The values of all current parameters (created by `make-parameter`)
    - Any pending `dynamic-wind` thunks.

The following procedures are provided, in addition to the procedures defined in SRFI-18:

22.1 thread-signal!

[procedure] (thread-signal! THREAD X)

This will cause THREAD to signal the condition X once it is scheduled for execution. After signalling the condition, the thread continues with its normal execution.

22.2 thread-quantum

[procedure] (thread-quantum THREAD)

Returns the quantum of THREAD, which is an exact integer specifying the approximate time-slice of the thread in milliseconds.

22.3 thread-quantum-set!

[procedure] (thread-quantum-set! THREAD QUANTUM)

Sets the quantum of THREAD to QUANTUM.

22.4 thread-suspend!

[procedure] (thread-suspend! THREAD)

Suspends the execution of THREAD until resumed.

22.5 thread-resume!

[procedure] (thread-resume! THREAD)

Readies the suspended thread THREAD.

22.6 thread-wait-for-i/o!

[procedure] (thread-wait-for-i/o! FD [MODE])

Suspends the current thread until input (MODE is #:input), output (MODE is #:output) or both (MODE is #:all) is available. FD should be a file-descriptor (not a port!) open for input or output, respectively.

22.7 timemilliseconds

[procedure] (time->milliseconds TIME)

Converts a time object (as created via current-time) into an exact integer representing the number of milliseconds since process startup.

23 Unit posix

This unit provides services as used on many UNIX-like systems. Note that the following definitions are not all available on non-UNIX systems like Windows. See below for Windows specific notes.

This unit uses the `regex`, `scheduler`, `extras` and `utils` units.

All errors related to failing file-operations will signal a condition of kind `(exn i/o file)`.

23.1 Constants

23.1.1 File-control Commands

23.1.1.1 fcntl/dupfd

23.1.1.2 fcntl/getfd

23.1.1.3 fcntl/setfd

23.1.1.4 fcntl/getfl

23.1.1.5 fcntl/setfl

23.1.2 Standard I/O file-descriptors

23.1.2.1 fileno/stdin

23.1.2.2 fileno/stdout

23.1.2.3 fileno/stderr

23.1.3 Open flags

23.1.3.1 open/rdonly

23.1.3.2 open/wronly

23.1.3.3 open/rdwr

23.1.3.4 open/read

23.1.3.5 open/write

23.1.3.6 open/creat

23.1.3.7 open/append

23.1.3.8 open/excl

23.1.3.9 open/noctty

23.1.3.10 open/nonblock

23.1.3.11 open/trunc

23.1.3.12 open/sync

23.1.3.13 open/fsync

23.1.3.14 open/binary

23.1.3.15 open/text

23.1.4 Permission bits

23.1.4.1 perm/irusr

23.1.4.2 perm/iwusr

23.1.4.3 perm/ixusr

23.1.4.4 perm/irgrp

23.1.4.5 perm/iwgrp

23.1.4.6 perm/ixgrp

23.1.4.7 perm/iroth

23.1.4.8 perm/iwoth

23.1.4.9 perm/ixoth

23.1.4.10 perm/irwxu

23.1.4.11 perm/irwxg

23.1.4.12 perm/irwxo

23.1.4.13 perm/isvtx

23.1.4.14 perm/isuid

23.1.4.15 perm/isgid

23.2 Directories

23.2.1 change-directory

[procedure] (change-directory NAME)

Changes the current working directory to NAME.

23.2.2 current-directory

[procedure] (current-directory [DIR])

Returns the name of the current working directory. If the optional argument DIR is given, then (current-directory DIR) is equivalent to (change-directory DIR).

23.2.3 create-directory

[procedure] (create-directory NAME)

Creates a directory with the pathname NAME.

23.2.4 delete-directory

[procedure] (delete-directory NAME)

Deletes the directory with the pathname NAME. The directory has to be empty.

23.2.5 directory

[procedure] (directory [PATHNAME [SHOW-DOTFILES?]])

Returns a list with all files that are contained in the directory with the name PATHNAME (which defaults to the value of (current-directory)). Files beginning with . are included only if SHOW-DOTFILES? is given and not #f.

23.2.6 directory?

[procedure] (directory? NAME)

Returns #t if there exists a file with the name NAME and if that file is a directory, or #f otherwise.

23.2.7 glob

[procedure] (glob PATTERN1 ...)

Returns a list of the pathnames of all existing files matching PATTERN1 ..., which should be strings containing the usual file-patterns (with * matching zero or more characters and ? matching zero or one

character).

23.2.8 set-root-directory!

[procedure] (set-root-directory! STRING)

Sets the root directory for the current process to the path given in STRING (using the chroot function). If the current process has no root permissions, the operation will fail.

23.3 Pipes

23.3.1 call-with-input-pipe

23.3.2 call-with-output-pipe

[procedure] (call-with-input-pipe CMDLINE PROC [MODE])
[procedure] (call-with-output-pipe CMDLINE PROC [MODE])

Call PROC with a single argument: a input- or output port for a pipe connected to the subprocess named in CMDLINE. If PROC returns normally, the pipe is closed and any result values are returned.

23.3.3 close-input-pipe

23.3.4 close-output-pipe

[procedure] (close-input-pipe PORT)
[procedure] (close-output-pipe PORT)

Closes the pipe given in PORT and waits until the connected subprocess finishes. The exit-status code of the invoked process is returned.

23.3.5 create-pipe

[procedure] (create-pipe)

The fundamental pipe-creation operator. Calls the C function `pipe()` and returns 2 values: the file-descriptors of the input- and output-ends of the pipe.

23.3.6 open-input-pipe

`[procedure] (open-input-pipe CMDLINE [MODE])`

Spawns a subprocess with the command-line string `CMDLINE` and returns a port, from which the output of the process can be read. If `MODE` is specified, it should be the keyword `#:text` (the default) or `#:binary`.

23.3.7 open-output-pipe

`[procedure] (open-output-pipe CMDLINE [MODE])`

Spawns a subprocess with the command-line string `CMDLINE` and returns a port. Anything written to that port is treated as the input for the process. If `MODE` is specified, it should be the keyword `#:text` (the default) or `#:binary`.

23.3.8 pipe/buf

This variable contains the maximal number of bytes that can be written atomically into a pipe or FIFO.

23.3.9 with-input-from-pipe

23.3.10 with-output-to-pipe

`[procedure] (with-input-from-pipe CMDLINE THUNK [MODE])`
`[procedure] (with-output-to-pipe CMDLINE THUNK [MODE])`

Temporarily set the value of `current-input-port`/`current-output-port` to a port for a pipe connected to the subprocess named in `CMDLINE` and call the procedure `THUNK` with no arguments. After `THUNK` returns normally the pipe is closed and the standard input-/output port is restored to its previous value and any result values are returned.

```
(with-output-to-pipe
  "gs -dNOPAUSE -sDEVICE=jpeg -dBATCH -sOutputFile=signballs.jpg -g600x600 -q -
  (lambda ()
    (print #<<EOF
%!IOPSC-1993 %%Creator: HAYAKAWA Takashi<xxxxxxxx@xx.xxxxxx.xx.xx>
/C/neg/d/mul/R/rlineto/E/exp/H{{cvx def}repeat}def/T/dup/g/gt/r/roll/J/ifelse
H/A/copy(z&v4QX&93r9AxYQOZomQalxS2w!!O&vMYa43d6r93rMYvx2dca!D&cjSnjSnjjS3o!v&6
X&55SAxM1CD7AjYxTTd62rmxCnTdSST0g&12wECST!&!J0g&D1!&xM0!J0g!l&544dC2Ac96ra!m&3
```

```
F&&vGoGSnCT0g&wDmlvGoS8wpn6wpS2wTCpS1Sd7ov7Uk7o4Qkdw!&Mvlx1S7oZES3w!J!J!Q&7185
Z&lx1CS9d9nE4!k&X&MY7!&1!J!x&jdnjdS3odS!N&mmx1C2wEc!G&150Nx4!n&2o!j&43r!U&0777
]&2AY2A776ddT4oS3oSnMVC00VV0RRR45E42063rNz&v7UX&UOzF!F!J![&44ETCnVn!a&1CDN!Y&0
V1c&j2AYdjmMdjjd!o&1r!M){( )T 0 4 3 r put T(/)g{T(9)g{cvn}{cvi}J}{($)g[]J}J
cvx}forall/moveto/p/floor/w/div/S/add 29 H[{[{]setgray fill}for Y}for showpage
EOF
) ) )
```

23.4 Fifos

23.4.1 create-fifo

[procedure] (create-fifo FILENAME [MODE])

Creates a FIFO with the name FILENAME and the permission bits MODE, which defaults to

[procedure] (+ perm/irwxu perm/irwxg perm/irwxo)

23.4.2 fifo?

[procedure] (fifo? FILENAME)

Returns #t if the file with the name FILENAME names a FIFO.

23.5 File descriptors and low-level I/O

23.5.1 duplicate-fileno

[procedure] (duplicate-fileno OLD [NEW])

If NEW is given, then the file-descriptor NEW is opened to access the file with the file-descriptor OLD. Otherwise a fresh file-descriptor accessing the same file as OLD is returned.

23.5.2 file-close

[procedure] (file-close FILENO)

Closes the input/output file with the file-descriptor FILENO.

23.5.3 file-open

[procedure] (file-open FILENAME FLAGS [MODE])

Opens the file specified with the string `FILENAME` and open-flags `FLAGS` using the C function `open()`. On success a file-descriptor for the opened file is returned. `FLAGS` should be a bitmask containing one or more of the `open/...` values **or**ed together using `bitwise-ior` (or simply added together). The optional `MODE` should be a bitmask composed of one or more permission values like `perm/irusr` and is only relevant when a new file is created. The default mode is `perm/irwxu | perm/irgrp | perm/iroth`.

23.5.4 file-mkstemp

[procedure] (file-mkstemp TEMPLATE-FILENAME)

Create a file based on the given `TEMPLATE-FILENAME`, in which the six last characters must be *XXXXXX*. These will be replaced with a string that makes the filename unique. The file descriptor of the created file and the generated filename is returned. See the `mkstemp(3)` manual page for details on how this function works. The template string given is not modified.

Example usage:

```
(let-values (((fd temp-path) (file-mkstemp "/tmp/mytemporary.XXXXXX")))
  (let ((temp-port (open-output-file* fd)))
    (format temp-port "This file is ~A.~%" temp-path)
    (close-output-port temp-port)))
```

23.5.5 file-read

[procedure] (file-read FILENO SIZE [BUFFER])

Reads `SIZE` bytes from the file with the file-descriptor `FILENO`. If a string or bytevector is passed in the optional argument `BUFFER`, then this string will be destructively modified to contain the read data. This procedure returns a list with two values: the buffer containing the data and the number of bytes read.

23.5.6 file-select

[procedure] (file-select READFDLIST WRITEFDLIST [TIMEOUT])

Waits until any of the file-descriptors given in the lists `READFDLIST` and `WRITEFDLIST` is ready for input or output, respectively. If the optional argument `TIMEOUT` is given and not false, then it should specify the number of seconds after which the wait is to be aborted (the value may be a floating point number). This procedure returns two values: the lists of file-descriptors ready for input and output, respectively. `READFDLIST` and **WRITEFDLIST** may also by file-descriptors instead of lists. In this case the returned values are booleans indicating whether input/output is ready by `#t` or `#f` otherwise. You can also pass `#f` as

READFDLIST or WRITEFDLIST argument, which is equivalent to ().

23.5.7 file-write

[procedure] (file-write FILENO BUFFER [SIZE])

Writes the contents of the string or bytevector BUFFER into the file with the file-descriptor FILENO. If the optional argument SIZE is given, then only the specified number of bytes are written.

23.5.8 file-control

[procedure] (file-control FILENO COMMAND [ARGUMENT])

Performs the fcntl operation COMMAND with the given FILENO and optional ARGUMENT. The return value is meaningful depending on the COMMAND.

23.5.9 open-input-file*

23.5.10 open-output-file*

[procedure] (open-input-file* FILENO [OPENMODE])
[procedure] (open-output-file* FILENO [OPENMODE])

Opens file for the file-descriptor FILENO for input or output and returns a port. FILENO should be a positive exact integer. OPENMODE specifies an additional mode for opening the file (currently only the keyword #:append is supported, which opens an output-file for appending).

23.5.11 portfileno

[procedure] (port->fileno PORT)

If PORT is a file- or tcp-port, then a file-descriptor is returned for this port. Otherwise an error is signaled.

23.6 Retrieving file attributes

23.6.1 file-access-time

23.6.2 file-change-time

23.6.3 file-modification-time

```
[procedure] (file-access-time FILE)
[procedure] (file-change-time FILE)
[procedure] (file-modification-time FILE)
```

Returns time (in seconds) of the last access, modification or change of `FILE`. `FILE` may be a filename or a file-descriptor. If the file does not exist, an error is signaled.

23.6.4 file-stat

```
[procedure] (file-stat FILE [LINK])
```

Returns a 13-element vector with the following contents: inode-number, mode (as with `file-permissions`), number of hard links, uid of owner (as with `file-owner`), gid of owner, size (as with `file-size`) and access-, change- and modification-time (as with `file-access-time`, `file-change-time` and `file-modification-time`, device id, device type (for special file inode, blocksize and blocks allocated. On Windows systems the last 4 values are undefined. If the optional argument `LINK` is given and not `#f`, then the file-statistics vector will be resolved for symbolic links (otherwise symbolic links are not resolved). Note that for very large files, the `file-size` value may be an inexact integer.

23.6.5 file-position

```
[procedure] (file-position FILE)
```

Returns the current file position of `FILE`, which should be a port or a file-descriptor.

23.6.6 file-size

```
[procedure] (file-size FILENAME)
```

Returns the size of the file designated by `FILE`. `FILE` may be a filename or a file-descriptor. If the file does not exist, an error is signaled. Note that for very large files, `file-size` may return an inexact integer.

23.6.7 regular-file?

[procedure] (regular-file? FILENAME)

Returns true, if FILENAME names a regular file (not a directory or symbolic link).

23.6.8 file-owner

[procedure] (file-owner FILE)

Returns the user-id of FILE. FILE may be a filename or a file-descriptor.

23.6.9 file-permissions

[procedure] (file-permissions FILE)

Returns the permission bits for FILE. You can test this value by performing bitwise operations on the result and the perm/... values. FILE may be a filename or a file-descriptor.

23.6.10 file-read-access?

23.6.11 file-write-access?

23.6.12 file-execute-access?

[procedure] (file-read-access? FILENAME)
[procedure] (file-write-access? FILENAME)
[procedure] (file-execute-access? FILENAME)

These procedures return #t if the current user has read, write or execute permissions on the file named FILENAME.

23.7 Changing file attributes

23.7.1 file-truncate

[procedure] (file-truncate FILE OFFSET)

Truncates the file `FILE` to the length `OFFSET`, which should be an integer. If the file-size is smaller or equal to `OFFSET` then nothing is done. `FILE` should be a filename or a file-descriptor.

23.7.2 set-file-position!

[procedure] (set-file-position! FILE POSITION [WHENCE])
[procedure] (set! (file-position FILE) POSITION)

Sets the current read/write position of `FILE` to `POSITION`, which should be an exact integer. `FILE` should be a port or a file-descriptor. `WHENCE` specifies how the position is to interpreted and should be one of the values `seek/set, seek/cur` and `seek/end`. It defaults to `seek/set`.

Exceptions: `(exn bounds)`, `(exn i/o file)`

23.7.3 change-file-mode

[procedure] (change-file-mode FILENAME MODE)

Changes the current file mode of the file named `FILENAME` to `MODE` using the `chmod()` system call. The `perm/...` variables contain the various permission bits and can be combined with the `bitwise-ior` procedure.

23.7.4 change-file-owner

[procedure] (change-file-owner FILENAME UID GID)

Changes the owner information of the file named `FILENAME` to the user- and group-ids `UID` and `GID` (which should be exact integers) using the `chown()` system call.

23.8 Processes

23.8.1 current-process-id

[procedure] (current-process-id)

Returns the process ID of the current process.

23.8.2 parent-process-id

[procedure] (parent-process-id)

Returns the process ID of the parent of the current process.

23.8.3 process-group-id

[procedure] (process-group-id PID)

Returns the process group ID of the process specified by PID.

23.8.4 process-execute

[procedure] (process-execute PATHNAME [ARGUMENT-LIST [ENVIRONMENT-LIST]])

Creates a new child process and replaces the running process with it using the C library function
execvp(3). If the optional argument ARGUMENT-LIST is given, then it should contain a list of strings
which are passed as arguments to the subprocess. If the optional argument ENVIRONMENT-LIST is
supplied, then the library function execve(2) is used, and the environment passed in
ENVIRONMENT-LIST (which should be of the form ("<NAME>=<VALUE>" ...) is given to the
invoked process. Note that execvp(3) respects the current setting of the PATH environment variable while
execve(3) does not.

23.8.5 process-fork

[procedure] (process-fork [THUNK])

Creates a new child process with the UNIX system call fork(). Returns either the PID of the child process
or 0. If THUNK is given, then the child process calls it as a procedure with no arguments and terminates.

23.8.6 process-run

[procedure] (process-run COMMANDLINE])
[procedure] (process-run COMMAND ARGUMENT-LIST)

Creates a new child process. The PID of the new process is returned.

- The single parameter version passes the COMMANDLINE to the system shell, so usual argument
  expansion can take place.
- The multiple parameter version directly invokes the COMMAND with the ARGUMENT-LIST.

23.8.7 process-signal

[procedure] (process-signal PID [SIGNAL])

Sends `SIGNAL` to the process with the id `PID` using the UNIX system call `kill()`. `SIGNAL` defaults to the value of the variable `signal/term`.

23.8.8 process-wait

[procedure] (process-wait [PID [NOHANG]])

Suspends the current process until the child process with the id `PID` has terminated using the UNIX system call `waitpid()`. If `PID` is not given, then this procedure waits for any child process. If `NOHANG` is given and not `#f` then the current process is not suspended. This procedure returns three values:

- `PID` or 0, if `NOHANG` is true and the child process has not terminated yet.
- `#t` if the process exited normally or `#f` otherwise.
- either the exit status, if the process terminated normally or the signal number that terminated/stopped the process.

23.8.9 process

[procedure] (process COMMANDLINE)
[procedure] (process COMMAND ARGUMENT-LIST [ENVIRONMENT-LIST])

Creates a subprocess and returns three values: an input port from which data written by the sub-process can be read, an output port from which any data written to will be received as input in the sub-process and the process-id of the started sub-process. Blocking reads and writes to or from the ports returned by `process` only block the current thread, not other threads executing concurrently.

- The single parameter version passes the string `COMMANDLINE` to the host-system's shell that is invoked as a subprocess.
- The multiple parameter version directly invokes the `COMMAND` as a subprocess. The `ARGUMENT-LIST` is directly passed, as is `ENVIRONMENT-LIST`.

Not using the shell may be preferrable for security reasons.

23.8.10 process*

[procedure] (process* COMMANDLINE)
[procedure] (process* COMMAND ARGUMENT-LIST [ENVIRONMENT-LIST])

Like `process` but returns 4 values: an input port from which data written by the sub-process can be read, an output port from which any data written to will be received as input in the sub-process, the process-id of the

started sub-process, and an input port from which data written by the sub-process to `stderr` can be read.

23.8.11 sleep

[procedure] (sleep SECONDS)

Puts the process to sleep for `SECONDS`. Returns either 0 if the time has completely elapsed, or the number of remaining seconds, if a signal occurred.

23.8.12 create-session

[procedure] (create-session)

Creates a new session if the calling process is not a process group leader and returns the session ID.

23.9 Hard and symbolic links

23.9.1 symbolic-link?

[procedure] (symbolic-link? FILENAME)

Returns true, if `FILENAME` names a symbolic link.

23.9.2 create-symbolic-link

[procedure] (create-symbolic-link OLDNAME NEWNAME)

Creates a symbolic link with the filename `NEWNAME` that points to the file named `OLDNAME`.

23.9.3 read-symbolic-link

[procedure] (read-symbolic-link FILENAME)

Returns the filename to which the symbolic link `FILENAME` points.

23.9.4 file-link

[procedure] (file-link OLDNAME NEWNAME)

Creates a hard link from OLDNAME to NEWNAME (both strings).

23.10 Retrieving user & group information

23.10.1 current-user-id

[procedure] (current-user-id)
[setter] (set! (current-user-id) UID)

Get or set the real user-id of the current process.

23.10.2 current-effective-user-id

[procedure] (current-effective-user-id)
[setter] (set! (current-effective-user-id) UID)

Get or set the effective user-id of the current process.

23.10.3 user-information

[procedure] (user-information USER [AS-VECTOR])

If USER specifes a valid username (as a string) or user ID, then the user database is consulted and a list of 7 values are returned: the user-name, the encrypted password, the user ID, the group ID, a user-specific string, the home directory and the default shell. When AS-VECTOR is #t a vector of 7 elements is returned instead of a list. If no user with this name or id then #f is returned.

23.10.4 current-group-id

[procedure] (current-group-id)
[setter] (set! (current-group-id) GID)

Get or set the real group-id of the current process.

23.10.5 current-effective-group-id

```
[procedure] (current-effective-group-id)
[setter] (set! (current-effective-group-id) GID)
```

Get or set the effective group-id of the current process. ID can be found, then `#f` is returned.

23.10.6 group-information

```
[procedure] (group-information GROUP)
```

If `GROUP` specifies a valid group-name or group-id, then this procedure returns a list of four values: the group-name, the encrypted group password, the group ID and a list of the names of all group members. If no group with the given name or ID exists, then `#f` is returned.

23.10.7 get-groups

```
[procedure] (get-groups)
```

Returns a list with the supplementary group IDs of the current user.

23.11 Changing user & group information

23.11.1 set-groups!

```
[procedure] (set-groups! GIDLIST)
```

Sets the supplementrary group IDs of the current user to the IDs given in the list `GIDLIST`.

Only the superuser may invoke this procedure.

23.11.2 initialize-groups

```
[procedure] (initialize-groups USERNAME BASEGID)
```

Sets the supplementrary group IDs of the current user to the IDs from the user with name `USERNAME` (a string), including `BASEGID`.

Only the superuser may invoke this procedure.

23.11.3 set-process-group-id!

```
[procedure] (set-process-group-id! PID PGID)
[setter] (set! (process-group-id PID) PGID)
```

Sets the process group ID of the process specifed by `PID` to `PGID`.

23.12 Record locking

23.12.1 file-lock

```
[procedure] (file-lock PORT [START [LEN]])
```

Locks the file associated with `PORT` for reading or writing (according to whether `PORT` is an input- or output-port). `START` specifies the starting position in the file to be locked and defaults to 0. `LEN` specifies the length of the portion to be locked and defaults to `#t`, which means the complete file. `file-lock` returns a *lock*-object.

23.12.2 file-lock/blocking

```
[procedure] (file-lock/blocking PORT [START [LEN]])
```

Similar to `file-lock`, but if a lock is held on the file, the current process blocks (including all threads) until the lock is released.

23.12.3 file-test-lock

```
[procedure] (file-test-lock PORT [START [LEN]])
```

Tests whether the file associated with `PORT` is locked for reading or writing (according to whether `PORT` is an input- or output-port) and returns either `#f` or the process-id of the locking process.

23.12.4 file-unlock

```
[procedure] (file-unlock LOCK)
```

Unlocks the previously locked portion of a file given in `LOCK`.

23.13 Signal handling

23.13.1 set-alarm!

[procedure] (set-alarm! SECONDS)

Sets an internal timer to raise the `signal/alrm` after `SECONDS` are elapsed. You can use the `set-signal-handler!` procedure to write a handler for this signal.

23.13.2 set-signal-handler!

[procedure] (set-signal-handler! SIGNUM PROC)

Establishes the procedure of one argument `PROC` as the handler for the signal with the code `SIGNUM`. `PROC` is called with the signal number as its sole argument. If the argument `PROC` is `#f` then any signal handler will be removed.

Note that is is unspecified in which thread of execution the signal handler will be invoked.

23.13.3 signal-handler

[procedure] (signal-handler SIGNUM)

Returns the signal handler for the code `SIGNUM` or `#f`.

23.13.4 set-signal-mask!

[procedure] (set-signal-mask! SIGLIST)

Sets the signal mask of the current process to block all signals given in the list `SIGLIST`. Signals masked in that way will not be delivered to the current process.

23.13.5 signal-mask

[procedure] (signal-mask)

Returns the signal mask of the current process.

23.13.6 signal-masked?

[procedure] (signal-masked? SIGNUM)

Returns whether the signal for the code SIGNUM is currently masked.

23.13.7 signal-mask!

[procedure] (signal-mask! SIGNUM)

Masks (blocks) the signal for the code SIGNUM.

23.13.8 signal-unmask!

[procedure] (signal-unmask! SIGNUM)

Unmasks (unblocks) the signal for the code SIGNUM.

23.13.9 signal/term

23.13.10 signal/kill

23.13.11 signal/int

23.13.12 signal/hup

23.13.13 signal/fpe

23.13.14 signal/ill

23.13.15 signal/segv

23.13.16 signal/abrt

23.13.17 signal/trap

23.13.18 signal/quit

23.13.19 signal/alrm

23.13.20 signal/vtalrm

23.13.21 signal/prof

23.13.22 signal/io

23.13.23 signal/urg

23.13.24 signal/chld

23.13.25 signal/cont

23.13.26 signal/stop

23.13.27 signal/tstp

23.13.28 signal/pipe

23.13.29 signal/xcpu

23.13.30 signal/xfsz

23.13.31 signal/usr1

23.13.32 signal/usr2

23.13.33 signal/winch

These variables contain signal codes for use with `process-signal`, `set-signal-handler!`, `signal-handler`, `signal-masked?`, `signal-mask!`, or `signal-unmask!`.

23.14 Environment access

23.14.1 current-environment

[procedure] (current-environment)

Returns a association list of the environment variables and their current values.

23.14.2 setenv

[procedure] (setenv VARIABLE VALUE)

Sets the environment variable named `VARIABLE` to `VALUE`. Both arguments should be strings. If the variable is not defined in the environment, a new definition is created.

23.14.3 unsetenv

[procedure] (unsetenv VARIABLE)

Removes the definition of the environment variable `VARIABLE` from the environment of the current process. If the variable is not defined, nothing happens.

23.15 Memory mapped I/O

23.15.1 memory-mapped-file?

[pocedure] (memory-mapped-file? X)

Returns `#t`, if `X` is an object representing a memory mapped file, or `#f` otherwise.

23.15.2 map-file-to-memory

[procedure] (map-file-to-memory ADDRESS LEN PROTECTION FLAG FILENO [OFFSET])

Maps a section of a file to memory using the C function `mmap()`. `ADDRESS` should be a foreign pointer object or `#f`; `LEN` specifies the size of the section to be mapped; `PROTECTION` should be one or more of the flags `prot/read, prot/write, prot/exec` or `prot/none` **bitwise-ior**ed together; `FLAG` should be one or more of the flags `map/fixed, map/shared, map/private, map/anonymous` or `map/file`; `FILENO` should be the file-descriptor of the mapped file. The optional argument `OFFSET` gives the offset of the section of the file to be mapped and defaults to 0. This procedure returns an object representing the mapped file section. The procedure `move-memory!` can be used to access the mapped memory.

23.15.3 memory-mapped-file-pointer

[procedure] (memory-mapped-file-pointer MMAP)

Returns a machine pointer to the start of the memory region to which the file is mapped.

23.15.4 unmap-file-from-memory

[procedure] (unmap-file-from-memory MMAP [LEN])

Unmaps the section of a file mapped to memory using the C function `munmap()`. MMAP should be a mapped file as returned by the procedure `map-file-to-memory`. The optional argument LEN specifies the length of the section to be unmapped and defaults to the complete length given when the file was mapped.

23.16 Date and time routines

23.16.1 secondslocal-time

[procedure] (seconds->local-time SECONDS)

Breaks down the time value represented in SECONDS into a 10 element vector of the form `#(seconds minutes hours mday month year wday yday dstflag timezone)`, in the following format:

seconds (0)
>   the number of seconds after the minute (0 - 59)

minutes (1)
>   the number of minutes after the hour (0 - 59)

hours (2)
>   the number of hours past midnight (0 - 23)

mday (3)
>   the day of the month (1 - 31)

month (4)
>   the number of months since january (0 - 11)

year (5)
>   the number of years since 1900

wday (6)
>   the number of days since Sunday (0 - 6)

yday (7)
>   the number of days since January 1 (0 - 365)

dstflag (8)
>   a flag that is true if Daylight Saving Time is in effect at the time described.

timezone (9)
>   the difference between UTC and the latest local standard time, in seconds west of UTC.

23.16.2 local-timeseconds

[procedure] (local-time->seconds VECTOR)

Converts the ten-element vector `VECTOR` representing the time value relative to the current timezone into the number of seconds since the first of January, 1970 UTC.

23.16.3 local-timezone-abbreviation

[procedure] (local-timezone-abbreviation)

Returns the abbreviation for the local timezone as a string.

23.16.4 secondsstring

[procedure] (seconds->string SECONDS)

Converts the local time represented in `SECONDS` into a string of the form `"Tue May 21 13:46:22 1991"`.

23.16.5 secondsutc-time

[procedure] (seconds->utc-time SECONDS)

Similar to `seconds->local-time`, but interpretes `SECONDS` as UTC time.

23.16.6 utc-timeseconds

[procedure] (utc-time->seconds VECTOR)

Converts the ten-element vector `VECTOR` representing the UTC time value into the number of seconds since the first of January, 1970 UTC.

23.16.7 timestring

[procedure] (time->string VECTOR)

Converts the broken down time represented in the 10 element vector `VECTOR` into a string of the form `"Tue May 21 13:46:22 1991"`.

23.17 Raw exit

23.17.1 _exit

```
[procedure] (_exit [CODE])
```

Exits the current process without flushing any buffered output (using the C function `_exit`). Note that the `exit-handler` is not called when this procedure is invoked. The optional return-code `CODE` defaults to `0`.

23.18 ERRNO values

23.18.1 errno/perm

23.18.2 errno/noent

23.18.3 errno/srch

23.18.4 errno/intr

23.18.5 errno/io

23.18.6 errno/noexec

23.18.7 errno/badf

23.18.8 errno/child

23.18.9 errno/nomem

23.18.10 errno/acces

23.18.11 errno/fault

23.18.12 errno/busy

23.18.13 errno/notdir

23.18.14 errno/isdir

23.18.15 errno/inval

23.18.16 errno/mfile

23.18.17 errno/nospc

23.18.18 errno/spipe

23.18.19 errno/pipe

23.18.20 errno/again

23.18.21 errno/rofs

23.18.22 errno/exist

23.18.23 errno/wouldblock

These variables contain error codes as returned by `errno`.

23.19 Finding files

23.19.1 find-files

[procedure] (find-files DIRECTORY PREDICATE [ACTION [IDENTITY [LIMIT]]])

Recursively traverses the contents of `DIRECTORY` (which should be a string) and invokes the procedure `ACTION` for all files in which the procedure `PREDICATE` is true. `PREDICATE` may me a procedure of one argument or a regular-expression string. `ACTION` should be a procedure of two arguments: the currently encountered file and the result of the previous invocation of `ACTION`, or, if this is the first invocation, the value of `IDENTITY`. `ACTION` defaults to `cons`, `IDENTITY` defaults to `()`. `LIMIT` should be a procedure of one argument that is called for each nested directory and which should return true, if that directory is to be traversed recursively. `LIMIT` may also be an exact integer that gives the maximum recursion depth. For example, a depth of `0` means that only files in the top-level, specified directory are to be traversed. In this case, all nested directories are ignored. `LIMIT` may also be `#f` (the default), which is equivalent to `(constantly #t)`.

Note that `ACTION` is called with the full pathname of each file, including the directory prefix.

23.20 Getting the hostname and system information

23.20.1 get-host-name

[procedure] (get-host-name)

Returns the hostname of the machine that this process is running on.

23.20.2 system-information

[procedure] (system-information)

Invokes the UNIX system call `uname()` and returns a list of 5 values: system-name, node-name, OS release, OS version and machine.

23.21 Setting the file buffering mode

23.21.1 set-buffering-mode!

[procedure] (set-buffering-mode! PORT MODE [BUFSIZE])

Sets the buffering-mode for the file associated with `PORT` to `MODE`, which should be one of the keywords `#:full`, `#:line` or `#:none`. If `BUFSIZE` is specified it determines the size of the buffer to be used (if any).

23.22 Terminal ports

23.22.1 terminal-name

[procedure] (terminal-name PORT)

Returns the name of the terminal that is connected to `PORT`.

23.22.2 terminal-port?

[procedure] (terminal-port? PORT)

Returns #t if PORT is connected to a terminal and #f otherwise.

23.23 How Scheme procedures relate to UNIX C functions

change-directory

chdir
change-file-mode

chmod
change-file-owner

chown
create-directory

mkdir
create-fifo

mkfifo
create-pipe

pipe
create-session

setsid
create-symbolic-link

link
current-directory

curdir
current-effective-groupd-id

getegid
current-effective-user-id

geteuid
current-group-id

getgid
current-parent-id

getppid
current-process-id

getpid
```
current-user-id
```

getuid
```
delete-directory
```

rmdir
```
duplicate-fileno
```

dup/dup2
```
_exit
```

_exit
```
file-close
```

close
```
file-access-time
```

stat
```
file-change-time
```

stat
```
file-modification-time
```

stat
```
file-execute-access?
```

access
```
file-open
```

open
```
file-lock
```

fcntl
```
file-position
```

ftell/lseek
```
file-read
```

read
```
file-read-access?
```

access
```
file-select
```

select
```
file-control
```

fcntl
```
file-stat
```

stat
```
file-test-lock
```

fcntl
```
file-truncate
```

truncate/ftruncate
```
file-unlock
```

fcntl
```
file-write
```

write
```
file-write-access?
```

access
```
get-groups
```

getgroups
```
get-host-name
```

gethostname
```
initialize-groups
```

initgroups
```
local-time->seconds
```

mktime
```
local-timezone-abbreviation
```

localtime
```
map-file-to-memory
```

mmap
```
open-input-file*
```

fdopen
```
open-output-file*
```

fdopen
```
open-input-pipe
```

popen
```
open-output-pipe
```

popen
```
port->fileno
```

fileno
```
process-execute
```

execvp
`process-fork`

fork
`process-group-id`

getpgid
`process-signal`

kill
`process-wait`

waitpid
`close-input-pipe`

pclose
`close-output-pipe`

pclose
`read-symbolic-link`

readlink
`seconds->local-time`

localtime
`seconds->string`

ctime
`seconds->utc-time`

gmtime
`set-alarm!`

alarm
`set-buffering-mode!`

setvbuf
`set-file-position!`

fseek/seek
`set-groups!`

setgroups
`set-signal-mask!`

sigprocmask
`set-group-id!`

setgid
`set-process-group-id!`

setpgid
```
set-user-id!
```

setuid
```
set-root-directory!
```

chroot
```
setenv
```

setenv/putenv
```
sleep
```

sleep
```
system-information
```

uname
```
terminal-name
```

ttyname
```
terminal-port?
```

isatty
```
time->string
```

asctime
```
unsetenv
```

putenv
```
unmap-file-from-memory
```

munmap
```
user-information
```

getpwnam/getpwuid
```
utc-time->seconds
```

timegm

23.24 Windows specific notes

Use of UTF8 encoded strings is for pathnames is not supported. Windows uses a 16-bit UNICODE encoding with special system calls for wide-character support. Only single-byte string encoding can be used.

23.24.1 Procedure Changes

Exceptions to the above procedure definitions.

[procedure] (create-pipe [MODE])

The optional parameter MODE, default open/binary | open/noinherit. This can be open/binary or open/text, optionally or'ed with open/noinherit.

[procedure] (process-wait [PID [NOHANG]])

process-wait always returns #t for a terminated process and only the exit status is available. (Windows does not provide signals as an interprocess communication method.)

[procedure] (process-execute PATHNAME [ARGUMENT-LIST [ENVIRONMENT-LIST [EXACT-F
[procedure] (process COMMAND ARGUMENT-LIST [ENVIRONMENT-LIST [EXACT-FLAG]])
[procedure] (process* COMMAND ARGUMENT-LIST [ENVIRONMENT-LIST [EXACT-FLAG]])

The optional parameter EXACT-FLAG, default #f. When #f any argument string with embedded whitespace will be wrapped in quotes. When #t no such wrapping occurs.

23.24.2 Unsupported Definitions

The following definitions are not supported for native Windows builds (compiled with the Microsoft tools or with MinGW):

```
open/noctty  open/nonblock  open/fsync  open/sync
perm/isvtx  perm/isuid  perm/isgid
file-select file-control
signal/... (except signal/term, signal/int, signal/fpe, signal/ill, signal/segv
set-signal-mask!  signal-mask  signal-masked?  signal-mask!  signal-unmask!
user-information  group-information  get-groups  set-groups!  initialize-groups
errno/wouldblock
change-file-owner
current-user-id  current-group-id  current-effective-user-id  current-effective
set-user-id!  set-group-id!
create-session
process-group-id  set-process-group-id!
create-symbolic-link  read-symbolic-link
file-truncate
file-lock  file-lock/blocking  file-unlock  file-test-lock
create-fifo  fifo?
prot/...
map/...
map-file-to-memory  unmap-file-from-memory  memory-mapped-file-pointer  memory-
set-alarm!
terminal-port?  terminal-name
process-fork  process-signal
parent-process-id
set-root-directory!
utc-time->seconds
```

23.24.3 Additional Definitions

Only available for Windows

- open/noinherit

This variable is a mode value for `create-pipe`. Useful when spawning a child process.

- spawn/overlay
- spawn/wait
- spawn/nowait
- spawn/nowaito
- spawn/detach

These variables contains special flags that specify the exact semantics of `process-spawn`: `spawn/overlay` replaces the current process with the new one. `spawn/wait` suspends execution of the current process until the spawned process returns. `spawn/nowait` does the opposite (`spawn/nowaito` is identical, according to the Microsoft documentation) and runs the process asynchronously. `spawn/detach` runs the new process in the background, without being attached to a console.

23.24.4 process-spawn

[procedure] (`process-spawn MODE COMMAND [ARGUMENT-LIST [ENVIRONMENT-LIST [EXACT`

Creates and runs a new process with the given `COMMAND` filename and the optional `ARGUMENT-LIST` and `ENVIRONMENT-LIST`. `MODE` specifies how exactly the process should be executed and must be one or more of the `spawn/...` flags defined above.

The `EXACT-FLAG`, default `#f`, controls quote-wrapping of argument strings. When `#t` quote-wrapping is not performed.

Returns:

- the exit status when synchronous
- the PID when asynchronous
- -1 when failure

24 Unit utils

This unit contains file/pathname oriented procedures, apropos, plus acts as a "grab bag" for procedures without a good home, and which don't have to be available by default (as compared to the <u>extras</u> unit).

This unit uses the `extras` and `regex` units.

24.1 Environment Query

24.1.1 apropos

[procedure] (apropos SYMBOL-PATTERN [ENVIRONMENT] [#:MACROS?])

Displays symbols & type matching `SYMBOL-PATTERN` in the `ENVIRONMENT` on the `(current-output-port)`.

SYMBOL-PATTERN
    A symbol, string, or regex. When symbol or string substring matching is performed.
ENVIRONMENT
    An environment. When missing the `(interaction-environment)` is assumed.
#:MACROS?
    Keyword argument. A boolean. Include macro symbols? When missing `#f` is assumed.

24.1.2 apropos-list

[procedure] (apropos-list SYMBOL-PATTERN [ENVIRONMENT] [#:MACROS?])

Like `apropos` but returns a list of matching symbols.

24.2 Pathname operations

24.2.1 absolute-pathname?

[procedure] (absolute-pathname? PATHNAME)

Returns `#t` if the string `PATHNAME` names an absolute pathname, and returns `#f` otherwise.

24.2.2 decompose-pathname

[procedure] (decompose-pathname PATHNAME)

Returns three values: the directory-, filename- and extension-components of the file named by the string PATHNAME. For any component that is not contained in PATHNAME, #f is returned.

24.2.3 make-pathname

24.2.4 make-absolute-pathname

[procedure] (make-pathname DIRECTORY FILENAME [EXTENSION [SEPARATOR]])
[procedure] (make-absolute-pathname DIRECTORY FILENAME [EXTENSION [SEPARATOR]])

Returns a string that names the file with the components DIRECTORY, FILENAME and (optionally) EXTENSION with SEPARATOR being the directory separation indicator (usually / on UNIX systems and \ on Windows, defaulting to whatever platform this is running on). DIRECTORY can be #f (meaning no directory component), a string or a list of strings. FILENAME and EXTENSION should be strings or #f. make-absolute-pathname returns always an absolute pathname.

24.2.5 pathname-directory

[procedure] (pathname-directory PATHNAME)

24.2.6 pathname-file

[procedure] (pathname-file PATHNAME)

24.2.7 pathname-extension

[procedure] (pathname-extension PATHNAME)

Accessors for the components of PATHNAME. If the pathname does not contain the accessed component, then #f is returned.

24.2.8 pathname-replace-directory

[procedure] (pathname-replace-directory PATHNAME DIRECTORY)

24.2.9 pathname-replace-file

[procedure] (pathname-replace-file PATHNAME FILENAME)

24.2.10 pathname-replace-extension

[procedure] (pathname-replace-extension PATHNAME EXTENSION)

Return a new pathname with the specified component of PATHNAME replaced by a new value.

24.2.11 pathname-strip-directory

[procedure] (pathname-strip-directory PATHNAME)

24.2.12 pathname-strip-extension

[procedure] (pathname-strip-extension PATHNAME)

Return a new pathname with the specified component of PATHNAME stripped.

24.2.13 directory-null?

[procedure] (directory-null? DIRECTORY)

Does the DIRECTORY consist only of path separators and the period?

DIRECTORY may be a string or a list of strings.

24.3 Temporary files

24.3.1 create-temporary-file

[procedure] (create-temporary-file [EXTENSION])

Creates an empty temporary file and returns its pathname. If `EXTENSION` is not given, then `.tmp` is used. If the environment variable `TMPDIR, TEMP` or `TMP` is set, then the pathname names a file in that directory.

24.4 Deleting a file without signalling an error

24.4.1 delete-file*

[procedure] (delete-file* FILENAME)

If the file `FILENAME` exists, it is deleted and `#t` is returned. If the file does not exist, nothing happens and `#f` is returned.

24.5 Iterating over input lines and files

24.5.1 for-each-line

[procedure] (for-each-line PROCEDURE [PORT])

Calls `PROCEDURE` for each line read from `PORT` (which defaults to the value of `(current-input-port)`. The argument passed to `PROCEDURE` is a string with the contents of the line, excluding any line-terminators. When all input has been read from the port, `for-each-line` returns some unspecified value.

24.5.2 for-each-argv-line

[procedure] (for-each-argv-line PROCEDURE)

Opens each file listed on the command line in order, passing one line at a time into `PROCEDURE`. The filename - is interpreted as `(current-input-port)`. If no arguments are given on the command line it again uses the value of `(current-input-port)`. During execution of `PROCEDURE`, the current input port will be correctly bound to the current input source.

This code will act as a simple Unix cat(1) command:

```
(for-each-argv-line print)
```

24.5.3 port-for-each

`[procedure] (port-for-each FN THUNK)`

Apply `FN` to successive results of calling the zero argument procedure `THUNK` until it returns `#!eof`, discarding the results.

24.5.4 port-map

`[procedure] (port-map FN THUNK)`

Apply `FN` to successive results of calling the zero argument procedure `THUNK` until it returns `#!eof`, returning a list of the collected results.

24.5.5 port-fold

`[procedure] (port-map FN ACC THUNK)`

Apply `FN` to successive results of calling the zero argument procedure `THUNK`, passing the `ACC` value as the second argument. The `FN` result becomes the new `ACC` value. When `THUNK` returns `#!eof`, the last `FN` result is returned.

24.6 Executing shell commands with formatstring and error checking

24.6.1 system*

`[procedure] (system* FORMATSTRING ARGUMENT1 ...)`

Similar to `(system (sprintf FORMATSTRING ARGUMENT1 ...))`, but signals an error if the invoked program should return a nonzero exit status.

24.7 Reading a file's contents

24.7.1 read-all

[procedure] (read-all [FILE-OR-PORT])

If `FILE-OR-PORT` is a string, then this procedure returns the contents of the file as a string. If `FILE-OR-PORT` is a port, all remaining input is read and returned as a string. The port is not closed. If no argument is provided, input will be read from the port that is the current value of `(current-input-port)`.

24.8 Funky ports

24.8.1 make-broadcast-port

[procedure] (make-broadcast-port PORT ...)

Returns a custom output port that emits everything written into it to the ports given as `PORT ...`. Closing the broadcast port does not close any of the argument ports.

24.8.2 make-concatenated-port

[procedure] (make-concatenated-port PORT1 PORT2 ...)

Returns a custom input port that reads its input from `PORT1`, until it is empty, then from `PORT2` and so on. Closing the concatenated port does not close any of the argument ports.

24.9 Miscellaneous handy things

24.9.1 shift! DEPRECATED

[procedure] (shift! LIST [DEFAULT])

Returns the car of `LIST` (or `DEFAULT` if `LIST` is empty) and replaces the car of `LIST` with it's cadr and the cdr with the cddr. If `DEFAULT` is not given, and the list is empty, `#f` is returned. An example might be clearer, here:

```
(define lst '(1 2 3))
(shift! lst)              ==> 1, lst is now (2 3)
```

The list must contain at least 2 elements.

24.9.2 unshift! DEPRECATED

[procedure] (unshift! X PAIR)

Sets the car of PAIR to X and the cdr to its cddr. Returns PAIR:

```
(define lst '(2))
(unshift! 99 lst)        ; lst is now (99 2)
```

Previous: Unit posix

Next: Unit tcp

25 Unit tcp

This unit provides basic facilities for communicating over TCP sockets. The socket interface should be mostly compatible to the one found in PLT Scheme.

This unit uses the `extras` unit.

All errors related to failing network operations will raise a condition of kind `(exn i/o network)`.

25.1 tcp-listen

[procedure] (tcp-listen TCPPORT [BACKLOG [HOST]])

Creates and returns a TCP listener object that listens for connections on `TCPPORT`, which should be an exact integer. `BACKLOG` specifies the number of maximally pending connections (and defaults to 4). If the optional argument `HOST` is given and not `#f`, then only incoming connections for the given host (or IP) are accepted.

25.2 tcp-listener?

[procedure] (tcp-listener? X)

Returns `#t` if `X` is a TCP listener object, or `#f` otherwise.

25.3 tcp-close

[procedure] (tcp-close LISTENER)

Reclaims any resources associated with `LISTENER`.

25.4 tcp-accept

[procedure] (tcp-accept LISTENER)

Waits until a connection is established on the port on which `LISTENER` is listening and returns two values: an input- and output-port that can be used to communicate with the remote process. The current value of `tcp-accept-timeout` is used to determine the maximal number of milliseconds (if any) to wait until a connection is established. When a client connects any read- and write-operations on the returned ports will use the current values (at the time of the connection) of `tcp-read-timeout` and `tcp-write-timeout`, respectively, to determine the maximal number of milliseconds to wait for input/output before a timeout error is signalled.

Note: this operation and any I/O on the ports returned will not block other running threads.

25.5 tcp-accept-ready?

`[procedure] (tcp-accept-ready? LISTENER)`

Returns `#t` if there are any connections pending on `LISTENER`, or `#f` otherwise.

25.6 tcp-listener-port

`[procedure] (tcp-listener-port LISTENER)`

Returns the port number assigned to `LISTENER` (If you pass `0` to `tcp-listen`, then the system will choose a port-number for you).

25.7 tcp-listener-fileno

`[procedure] (tcp-listener-fileno LISTENER)`

Returns the file-descriptor associated with `LISTENER`.

25.8 tcp-connect

`[procedure] (tcp-connect HOSTNAME [TCPPORT])`

Establishes a client-side TCP connection to the machine with the name `HOSTNAME` (a string) at `TCPPORT` (an exact integer) and returns two values: an input- and output-port for communicating with the remote process. The current value of `tcp-connect-timeout` is used to determine the maximal number of milliseconds (if any) to wait until the connection is established. When the connection takes place any read- and write-operations on the returned ports will use the current values (at the time of the call to `tcp-connect`) of `tcp-read-timeout` and `tcp-write-timeout`, respectively, to determine the maximal number of milliseconds to wait for input/output before a timeout error is signalled.

If the `TCPPORT` is omitted, the port is parsed from the `HOSTNAME` string. The format expected is `HOSTNAME:PORT`. The `PORT` can either be a string representation of an integer or a service name which is translated to an integer using the POSIX function `getservbyname`.

Note: any I/O on the ports returned will not block other running threads.

25.9 tcp-addresses

`[procedure] (tcp-addresses PORT)`

Returns two values for the input- or output-port `PORT` (which should be a port returned by either `tcp-accept` or `tcp-connect`): the IP address of the local and the remote machine that are connected over the socket associated with `PORT`. The returned addresses are strings in `XXX.XXX.XXX.XXX` notation.

25.10 tcp-port-numbers

`[procedure] (tcp-port-numbers PORT)`

Returns two values for the input- or output-port `PORT` (which should be a port returned by either `tcp-accept` or `tcp-connect`): the TCP port numbers of the local and the remote machine that are connected over the socket associated with `PORT`.

25.11 tcp-abandon-port

`[procedure] (tcp-abandon-port PORT)`

Marks the socket port `PORT` as abandoned. This is mainly useful to close down a port without breaking the connection.

25.12 tcp-buffer-size

`[parameter] tcp-buffer-size`

Sets the size of the output buffer. By default no output-buffering for TCP output is done, but to improve performance by minimizing the number of TCP packets, buffering may be turned on by setting this parameter to an exact integer greater zero. A buffer size of zero or `#f` turns buffering off. The setting of this parameter takes effect at the time when the I/O ports for a particular socket are created, i.e. when `tcp-connect` or `tcp-accept` is called.

Note that since output is not immediately written to the associated socket, you may need to call `flush-output`, once you want the output to be transmitted. Closing the output port will flush automatically.

25.13 tcp-read-timeout

`[parameter] tcp-read-timeout`

Determines the timeout for TCP read operations in milliseconds. A timeout of `#f` disables timeout checking. The default read timeout is 60000, i.e. 1 minute.

25.14 tcp-write-timeout

`[parameter] tcp-write-timeout`

Determines the timeout for TCP write operations in milliseconds. A timeout of `#f` disables timeout checking. The default write timeout is 60000, i.e. 1 minute.

25.15 tcp-connect-timeout

`[parameter] tcp-connect-timeout`

Determines the timeout for `tcp-connect` operations in milliseconds. A timeout of `#f` disables timeout checking and is the default.

25.16 tcp-accept-timeout

`[parameter] tcp-accept-timeout`

Determines the timeout for `tcp-accept` operations in milliseconds. A timeout of `#f` disables timeout checking and is the default.

25.17 Example

A very simple example follows. Say we have the two files `client.scm` and `server.scm`:

```
; client.scm
(declare (uses tcp))
(define-values (i o) (tcp-connect "localhost" 4242))
(write-line "Good Bye!" o)
(print (read-line i))

; server.scm
(declare (uses tcp))
(define l (tcp-listen 4242))
(define-values (i o) (tcp-accept l))
(write-line "Hello!" o)
(print (read-line i))
(close-input-port i)
(close-output-port o)

% csc server.scm
% csc client.scm
% ./server &
```

```
% ./client
Good Bye!
Hello!
```

Previous: <u>Unit utils</u>

Next: <u>Unit lolevel</u>

26 Unit lolevel

This unit provides a number of handy low-level operations. **Use at your own risk.**

This unit uses the `srfi-4` and `extras` units.

26.1 Foreign pointers

26.1.1 addresspointer

[procedure] (address->pointer ADDRESS)

Creates a new foreign pointer object initialized to point to the address given in the integer ADDRESS.

26.1.2 allocate

[procedure] (allocate BYTES)

Returns a pointer to a freshly allocated region of static memory. This procedure could be defined as follows:

```
(define allocate (foreign-lambda c-pointer "malloc" integer))
```

26.1.3 free

[procedure] (free POINTER)

Frees the memory pointed to by POINTER. This procedure could be defined as follows:

```
(define free (foreign-lambda c-pointer "free" integer))
```

26.1.4 null-pointer

[procedure] (null-pointer)

Another way to say (address->pointer 0).

26.1.5 null-pointer?

`[procedure] (null-pointer? PTR)`

Returns #t if PTR contains a NULL pointer, or #f otherwise.

26.1.6 objectpointer

`[procedure] (object->pointer X)`

Returns a pointer pointing to the Scheme object X, which should be a non-immediate object. Note that data in the garbage collected heap moves during garbage collection.

26.1.7 pointer?

`[procedure] (pointer? X)`

Returns #t if X is a foreign pointer object, and #f otherwise.

26.1.8 pointer=?

`[procedure] (pointer=? PTR1 PTR2)`

Returns #t if the pointer-like objects PTR1 and PTR2 point to the same address.

26.1.9 pointeraddress

`[procedure] (pointer->address PTR)`

Returns the address, to which the pointer PTR points.

26.1.10 pointerobject

`[procedure] (pointer->object PTR)`

Returns the Scheme object pointed to by the pointer PTR.

26.1.11 pointer-offset

```
[procedure] (pointer-offset PTR N)
```

Returns a new pointer representing the pointer PTR increased by N.

26.1.12 pointer-u8-ref

```
[procedure] (pointer-u8-ref PTR)
```

Returns the unsigned byte at the address designated by PTR.

26.1.13 pointer-s8-ref

```
[procedure] (pointer-s8-ref PTR)
```

Returns the signed byte at the address designated by PTR.

26.1.14 pointer-u16-ref

```
[procedure] (pointer-u16-ref PTR)
```

Returns the unsigned 16-bit integer at the address designated by PTR.

26.1.15 pointer-s16-ref

```
[procedure] (pointer-s16-ref PTR)
```

Returns the signed 16-bit integer at the address designated by PTR.

26.1.16 pointer-u32-ref

```
[procedure] (pointer-u32-ref PTR)
```

Returns the unsigned 32-bit integer at the address designated by PTR.

26.1.17 pointer-s32-ref

```
[procedure] (pointer-s32-ref PTR)
```

Returns the signed 32-bit integer at the address designated by PTR.

26.1.18 pointer-f32-ref

```
[procedure] (pointer-f32-ref PTR)
```

Returns the 32-bit float at the address designated by PTR.

26.1.19 pointer-f64-ref

```
[procedure] (pointer-f64-ref PTR)
```

Returns the 64-bit double at the address designated by PTR.

26.1.20 pointer-u8-set!

```
[procedure] (pointer-u8-set! PTR N)
[procedure] (set! (pointer-u8-ref PTR) N)
```

Stores the unsigned byte N at the address designated by PTR.

26.1.21 pointer-s8-set!

```
[procedure] (pointer-s8-set! PTR N)
[procedure] (set! (pointer-s8-ref PTR) N)
```

Stores the signed byte N at the address designated by PTR.

26.1.22 pointer-u16-set!

```
[procedure] (pointer-u16-set! PTR N)
[procedure] (set! (pointer-u16-ref PTR) N)
```

Stores the unsigned 16-bit integer N at the address designated by PTR.

26.1.23 pointer-s16-set!

```
[procedure] (pointer-s16-set! PTR N)
[procedure] (set! (pointer-s16-ref PTR) N)
```

Stores the signed 16-bit integer N at the address designated by PTR.

26.1.24 pointer-u32-set!

```
[procedure] (pointer-u32-set! PTR N)
[procedure] (set! (pointer-u32-ref PTR) N)
```

Stores the unsigned 32-bit integer N at the address designated by PTR.

26.1.25 pointer-s32-set!

```
[procedure] (pointer-s32-set! PTR N)
[procedure] (set! (pointer-s32-ref PTR) N)
```

Stores the 32-bit integer N at the address designated by PTR.

26.1.26 pointer-f32-set!

```
[procedure] (pointer-f32-set! PTR N)
[procedure] (set! (pointer-f32-ref PTR) N)
```

Stores the 32-bit floating-point number N at the address designated by PTR.

26.1.27 pointer-f64-set!

```
[procedure] (pointer-f64-set! PTR N)
[procedure] (set! (pointer-f64-ref PTR) N)
```

Stores the 64-bit floating-point number N at the address designated by PTR.

26.1.28 align-to-word

```
[procedure] (align-to-word PTR-OR-INT)
```

Accepts either a machine pointer or an integer as argument and returns a new pointer or integer aligned to the native word size of the host platform.

26.2 Tagged pointers

Tagged pointers are foreign pointer objects with an extra tag object.

26.2.1 tag-pointer

[procedure] (tag-pointer PTR TAG)

Creates a new tagged pointer object from the foreign pointer PTR with the tag TAG, which may an arbitrary Scheme object.

26.2.2 tagged-pointer?

[procedure] (tagged-pointer? X TAG)

Returns #t, if X is a tagged pointer object with the tag TAG (using an eq? comparison), or #f otherwise.

26.2.3 pointer-tag

[procedure] (pointer-tag PTR)

If PTR is a tagged pointer object, its tag is returned. If PTR is a normal, untagged foreign pointer object #f is returned. Otherwise an error is signalled.

26.3 Extending procedures with data

26.3.1 extend-procedure

[procedure] (extend-procedure PROCEDURE X)

Returns a copy of the procedure PROCEDURE which contains an additional data slot initialized to X. If PROCEDURE is already an extended procedure, then its data slot is changed to contain X and the same procedure is returned.

26.3.2 extended-procedure?

[procedure] (extended-procedure? PROCEDURE)

Returns #t if PROCEDURE is an extended procedure, or #f otherwise.

26.3.3 procedure-data

[procedure] (procedure-data PROCEDURE)

Returns the data object contained in the extended procedure PROCEDURE, or #f if it is not an extended procedure.

26.3.4 set-procedure-data!

[procedure] (set-procedure-data! PROCEDURE X)

Changes the data object contained in the extended procedure PROCEDURE to X.

```
(define foo
  (letrec ((f (lambda () (procedure-data x)))
           (x #f) )
    (set! x (extend-procedure f 123))
    x) )
(foo)                                       ==> 123
(set-procedure-data! foo 'hello)
(foo)                                       ==> hello
```

26.4 Data in unmanaged memory

26.4.1 object-evict

[procedure] (object-evict X [ALLOCATOR])

Copies the object X recursively into the memory pointed to by the foreign pointer object returned by ALLOCATOR, which should be a procedure of a single argument (the number of bytes to allocate). The freshly copied object is returned. This facility allows moving arbitrary objects into static memory, but care should be taken when mutating evicted data: setting slots in evicted vector-like objects to non-evicted data is not allowed. It **is** possible to set characters/bytes in evicted strings or byte-vectors, though. It is advisable **not** to evict ports, because they might be mutated by certain file-operations. object-evict is able to handle circular and shared structures, but evicted symbols are no longer unique: a fresh copy of the symbol is created, so

```
(define x 'foo)
```

```
(define y (object-evict 'foo))
y                               ==> foo
(eq? x y)                       ==> #f
(define z (object-evict '(bar bar)))
(eq? (car z) (cadr z))          ==> #t
```

The `ALLOCATOR` defaults to `allocate`.

26.4.2 object-evict-to-location

[procedure] (object-evict-to-location X PTR [LIMIT])

As `object-evict` but moves the object at the address pointed to by the machine pointer `PTR`. If the number of copied bytes exceeds the optional `LIMIT` then an error is signalled (specifically a composite condition of types `exn` and `evict`. The latter provides a `limit` property which holds the exceeded limit. Two values are returned: the evicted object and a new pointer pointing to the first free address after the evicted object.

26.4.3 object-evicted?

[procedure] (object-evicted? X)

Returns `#t` if `X` is a non-immediate evicted data object, or `#f` otherwise.

26.4.4 object-size

[procedure] (object-size X)

Returns the number of bytes that would be needed to evict the data object `X`.

26.4.5 object-release

[procedure] (object-release X [RELEASER])

Frees memory occupied by the evicted object `X` recursively. `RELEASER` should be a procedure of a single argument (a foreign pointer object to the static memory to be freed) and defaults to `free`.

26.4.6 object-unevict

[procedure] (object-unevict X [FULL])

Copies the object `X` and nested objects back into the normal Scheme heap. Symbols are re-interned into the symbol table. Strings and byte-vectors are **not** copied, unless `FULL` is given and not `#f`.

26.5 Locatives

A *locative* is an object that points to an element of a containing object, much like a *pointer* in low-level, imperative programming languages like *C*. The element can be accessed and changed indirectly, by performing access or change operations on the locative. The container object can be computed by calling the `location->object` procedure.

Locatives may be passed to foreign procedures that expect pointer arguments. The effect of creating locatives for evicted data (see `object-evict`) is undefined.

26.5.1 make-locative

[procedure] (make-locative EXP [INDEX])

Creates a locative that refers to the element of the non-immediate object `EXP` at position `INDEX`. `EXP` may be a vector, pair, string, blob, SRFI-4 number-vector, or record. `INDEX` should be a fixnum. `INDEX` defaults to 0.

26.5.2 make-weak-locative

[procedure] (make-weak-locative EXP [INDEX])

Creates a *weak* locative. Even though the locative refers to an element of a container object, the container object will still be reclaimed by garbage collection if no other references to it exist.

26.5.3 locative?

[procedure] (locative? X)

Returns `#t` if `X` is a locative, or `#f` otherwise.

26.5.4 locative-ref

[procedure] (locative-ref LOC)

Returns the element to which the locative `LOC` refers. If the containing object has been reclaimed by garbage collection, an error is signalled.

26.5.5 locative-set!

```
[procedure] (locative-set! LOC X)
[procedure] (set! (locative-ref LOC) X)
```

Changes the element to which the locative `LOC` refers to `X`. If the containing object has been reclaimed by garbage collection, an error is signalled.

26.5.6 locativeobject

```
[procedure] (locative->object LOC)
```

Returns the object that contains the element referred to by `LOC` or `#f` if the container has been reclaimed by garbage collection.

26.6 Accessing toplevel variables

26.6.1 global-bound?

```
[procedure] (global-bound? SYMBOL)
```

Returns `#t`, if the global (*toplevel*) variable with the name `SYMBOL` is bound to a value, or `#f` otherwise.

26.6.2 global-ref

```
[procedure] (global-ref SYMBOL)
```

Returns the value of the global variable `SYMBOL`. If no variable under that name is bound, an error is signalled.

Note that it is not possible to access a toplevel binding with `global-ref` or `global-set!` if it has been hidden in compiled code via `(declare (hide ...))`, or if the code has been compiled in `block` mode.

26.6.3 global-set!

```
[procedure] (global-set! SYMBOL X)
[procedure] (set! (global-ref SYMBOL) X)
```

Sets the global variable named `SYMBOL` to the value `X`.

26.7 Low-level data access

26.7.1 block-ref

```
[procedure] (block-ref BLOCK INDEX)
```

Returns the contents of the INDEXth slot of the object BLOCK. BLOCK may be a vector, record structure, pair or symbol.

26.7.2 block-set!

```
[procedure] (block-set! BLOCK INDEX X)
[procedure] (set! (block-ref BLOCK INDEX) X)
```

Sets the contents of the INDEXth slot of the object BLOCK to the value of X. BLOCK may be a vector, record structure, pair or symbol.

26.7.3 object-copy

```
[procedure] (object-copy X)
```

Copies X recursively and returns the fresh copy. Objects allocated in static memory are copied back into garbage collected storage.

26.7.4 make-record-instance

```
[procedure] (make-record-instance SYMBOL ARG1 ...)
```

Returns a new instance of the record type SYMBOL, with its slots initialized to ARG1 .... To illustrate:

```
(define-record point x y)
```

expands into something quite similar to:

```
(begin
  (define (make-point x y)
    (make-record-instance 'point x y) )
  (define (point? x)
    (and (record-instance? x)
         (eq? 'point (block-ref x 0)) ) )
  (define (point-x p) (block-ref p 1))
```

```scheme
(define (point-x-set! p x) (block-set! p 1 x))
(define (point-y p) (block-ref p 2))
(define (point-y-set! p y) (block-set! p 1 y)) )
```

26.7.5 move-memory!

[procedure] (move-memory! FROM TO [BYTES [FROM-OFFSET [TO-OFFSET]]])

Copies BYTES bytes of memory from FROM to TO. FROM and TO may be strings, primitive byte-vectors, SRFI-4 byte-vectors (see: @ref{Unit srfi-4}), memory mapped files, foreign pointers (as obtained from a call to foreign-lambda, for example) or locatives. if BYTES is not given and the size of the source or destination operand is known then the maximal number of bytes will be copied. Moving memory to the storage returned by locatives will cause havoc, if the locative refers to containers of non-immediate data, like vectors or pairs.

The additional fourth and fifth argument specify starting offsets (in bytes) for the source and destination arguments.

26.7.6 number-of-bytes

[procedure] (number-of-bytes BLOCK)

Returns the number of bytes that the object BLOCK contains. BLOCK may be any non-immediate value.

26.7.7 number-of-slots

[procedure] (number-of-slots BLOCK)

Returns the number of slots that the object BLOCK contains. BLOCK may be a vector, record structure, pair or symbol.

26.7.8 record-instance?

[procedure] (record-instance? X)

Returns #t if X is an instance of a record type. See also: make-record-instance.

26.7.9 recordvector

[procedure] (record->vector BLOCK)

Returns a new vector with the type and the elements of the record `BLOCK`.

26.8 Procedure-call- and variable reference hooks

26.8.1 set-invalid-procedure-call-handler!

[procedure] (set-invalid-procedure-call-handler! PROC)

Sets an internal hook that is invoked when a call to an object other than a procedure is executed at runtime. The procedure `PROC` will in that case be called with two arguments: the object being called and a list of the passed arguments.

```
;;; Access sequence-elements as in ARC:

(set-invalid-procedure-call-handler!
  (lambda (proc args)
    (cond [(string? proc) (apply string-ref proc args)]
          [(vector? proc) (apply vector-ref proc args)]
          [else (error "call of non-procedure" proc)] ) ) )

("hello" 4)     ==>  #\o
```

This facility does not work in code compiled with the *unsafe* setting.

26.8.2 unbound-variable-value

[procedure] (unbound-variable-value [X])

Defines the value that is returned for unbound variables. Normally an error is signalled, use this procedure to override the check and return `X` instead. To set the default behavior (of signalling an error), call `unbound-variable-value` with no arguments.

This facility does not work in code compiled with the *unsafe* setting.

26.9 Magic

26.9.1 object-become!

[procedure] (object-become! ALIST)

Changes the identity of the value of the car of each pair in `ALIST` to the value of the cdr. Both values may not be immediate (i.e. exact integers, characters, booleans or the empty list).

```
(define x "i used to be a string")
(define y '#(and now i am a vector))
(object-become! (list (cons x y)))
x                                      ==> #(and now i am a vector)
y                                      ==> #(and now i am a vector)
(eq? x y)                              ==> #t
```

Note: this operation invokes a major garbage collection.

The effect of using `object-become!` on evicted data (see `object-evict`) is undefined.

26.9.2 mutate-procedure

```
[procedure] (mutate-procedure OLD PROC)
```

Replaces the procedure `OLD` with the result of calling the one-argument procedure `PROC`. `PROC` will receive a copy of `OLD` that will be identical in behaviour to the result of `PROC`:

```
;;; Replace arbitrary procedure with tracing one:

(mutate-procedure my-proc
  (lambda (new)
    (lambda args
      (printf "~s called with arguments: ~s~%" new args)
      (apply new args) ) ) )
```

Previous: <u>Unit tcp</u>

Next: <u>Interface to external functions and variables</u>

27 Interface to external functions and variables

28 Accessing external objects

28.1 foreign-code

[syntax] (foreign-code STRING ...)

Executes the embedded C/C++ code `STRING ...`, which should be a sequence of C statements, which are executed and return an unspecified result.

```
(foreign-code "doSomeInitStuff();")    =>  #<unspecified>
```

Code wrapped inside `foreign-code` may not invoke callbacks into Scheme.

28.2 foreign-value

[syntax] (foreign-value STRING TYPE)

Evaluates the embedded C/C++ expression `STRING`, returning a value of type given in the foreign-type specifier `TYPE`.

```
(print (foreign-value "my_version_string" c-string))
```

28.3 foreign-declare

[syntax] (foreign-declare STRING ...)

Include given strings verbatim into header of generated file.

28.4 define-foreign-type

[syntax] (define-foreign-type NAME TYPE [ARGCONVERT [RETCONVERT]])

Defines an alias for `TYPE` with the name `NAME` (a symbol). `TYPE` may be a type-specifier or a string naming a C type. The namespace of foreign type specifiers is separate from the normal Scheme namespace. The optional arguments `ARGCONVERT` and `RETCONVERT` should evaluate to procedures that map argument- and result-values to a value that can be transformed to `TYPE`:

```
(define-foreign-type char-vector
  nonnull-c-string
  (compose list->string vector->list)
  (compose list->vector string->list) )

(define strlen
```

```
  (foreign-lambda int "strlen" char-vector) )

(strlen '#(#\a #\b #\c))                          ==> 3

(define memset
  (foreign-lambda char-vector "memset" char-vector char int) )

(memset '#(#_ #_ #_) #\X 3)                   ==> #(#\X #\X #\X)
```

Foreign type-definitions are only visible in the compilation-unit in which they are defined, so use `include` to use the same definitions in multiple files.

28.5 define-foreign-variable

`[syntax] (define-foreign-variable NAME TYPE [STRING])`

Defines a foreign variable of name `NAME` (a symbol). `STRING` should be the real name of a foreign variable or parameterless macro. If `STRING` is not given, then the variable name `NAME` will be converted to a string and used instead. All references and assignments (via `set!`) are modified to correctly convert values between Scheme and C representation. This foreign variable can only be accessed in the current compilation unit, but the name can be lexically shadowed. Note that `STRING` can name an arbitrary C expression. If no assignments are performed, then `STRING` doesn't even have to specify an lvalue.

```
#>
enum { abc=3, def, ghi };
<#

(define-macro (define-simple-foreign-enum . items)
  `(begin
     ,@(map (match-lambda
              [(name realname) `(define-foreign-variable ,name int ,realname)]
              [name `(define-foreign-variable ,name int)] )
        items) ) )

(define-simple-foreign-enum abc def ghi)

ghi                                    ==> 5
```

28.6 define-foreign-record

`[syntax] (define-foreign-record NAME [DECL ...] SLOT ...)`

Defines accessor procedures for a C structure definition. `NAME` should either be a symbol or a list of the form `(TYPENAME FOREIGNNAME)`. If `NAME` is a symbol, then a C declaration will be generated that defines a C struct named `struct NAME`. If `NAME` is a list, then no struct declaration will be generated and `FOREIGNNAME` should name an existing C record type. A foreign-type specifier named `NAME` (or `TYPENAME`) will be defined as a pointer to the given C structure. A `SLOT` definition should be a list of one of the following forms:

```
(TYPE SLOTNAME)
```

or

```
(TYPE SLOTNAME SIZE)
```

The latter form defines an array of `SIZE` elements of the type `TYPE` embedded in the structure. For every slot, the following accessor procedures will be generated:

28.6.1 TYPENAME-SLOTNAME

```
(TYPENAME-SLOTNAME FOREIGN-RECORD-POINTER [INDEX])
```

A procedure of one argument (a pointer to a C structure), that returns the slot value of the slot `SLOTNAME`. If a `SIZE` has been given in the slot definition, then an additional argument `INDEX` is required that specifies the index of an array-element.

28.6.2 TYPENAME-SLOTNAME-set!

```
(TYPENAME-SLOTNAME-set! FOREIGN-RECORD-POINTER [INXDEX] VALUE)
```

A procedure of two arguments (a pointer to a C structure) and a value, that sets the slot value of the slot `SLOTNAME` in the structure. If a `SIZE` has been given in the slot definition, then an additional argument `INDEX` is required for the array index.

If a slot type is of the form `(const ...)`, then no setter procedure will be generated. Slots of the types `(struct ...)` or `(union ...)` are accessed as pointers to the embedded struct (or union) and no setter will be generated.

Additionally, special record-declarations (`DECL ...`) may be given, where each declaration consists of a list of the form `(KEYWORD ARGUMENT ...)`. The available declarations are:

28.6.3 constructor

```
(constructor: NAME)
```

Generate a constructor-procedure with no arguments that has the name `NAME` (a symbol) that returns a pointer to a structure of this type. The storage will be allocated with `malloc(3)`.

28.6.4 destructor

```
(destructor: NAME)
```

Generate a destructor function with the name `NAME` that takes a pointer to a structure of this type as its single argument and releases the storage with `free(3)`. If the argument is `#f`, the destructor procedure does nothing.

28.6.5 rename

```
(rename: EXPRESSION)
```

Evaluates `EXPRESSION` at compile-/macro-expansion-time and applies the result, which should be a procedure, to the string-representation of the name of each accessor-procedure generated. Another (or the same) string should be returned, which in turn is taken as the actual name of the accessor.

An example:

```
(require-for-syntax 'srfi-13)

(define-foreign-record Some_Struct
  (rename: (compose string-downcase (cut string-translate <> "_" "-")))
  (constructor: make-some-struct)
  (destructor: free-some-struct)
  (int xCoord)
  (int yCoord) )
```

will generate the following procedures:

```
(make-some-struct)                  --> C-POINTER
(free-some-struct C-POINTER)

(some-struct-xcoord C-POINTER)  --> NUMBER
(some-struct-ycoord C-POINTER)  --> NUMBER

(some-struct-xcoord-set! C-POINTER NUMBER)
(some-struct-ycoord-set! C-POINTER NUMBER)
```

28.7 define-foreign-enum

```
[syntax] (define-foreign-enum TYPESPEC [USE-ALIASES] ENUMSPEC ...)
```

Defines a foreign type (as with `define-foreign-type`) that maps the elements of a C/C++ enum (or a enum-like list of constants) to and from a set of symbols.

`TYPESPEC` specifies a foreign type that converts a symbol argument from the set `ENUMSPEC ...` into the appropriate enum value when passed as an argument to a foreign function.

A list of symbols passed as an argument will be combined using `bitwise-ior`. An empty list will be passed as 0 (zero). Results of the enum type are automatically converted into a scheme value (note that combinations are not supported in this case).

TYPESPEC maybe a TYPENAME symbol or a list of the form `(SCHEMENAME REALTYPE [DEFAULT-SCHEME-VALUE])`, where REALTYPE designates the native type used. The default type specification is `(TYPENAME TYPENAME)`. The `DEFAULT-SCHEME-VALUE` overrides the default result of mapping from the native type; i.e. when no such mapping exists. When supplied the form is used unquoted, otherwise the result is `'()`.

ENUMSPEC is a TYPENAME symbol or a list of the form `(SCHEMENAME REALTYPE [SCHEME-VALUE])`, where REALTYPE designates the native type used. The default enum specification is `(TYPENAME TYPENAME)`. The `SCHEME-VALUE` overrides the result of mapping from the native type. When supplied the form is used unquoted, otherwise the SCHEMENAME symbol is returned.

USE-ALIASES is an optional boolean flag that determines whether an alias or the SCHEMENAME is used as the defined foreign variable name. The default is `#t`.

Additionally two procedures are defined named `SCHEMENAME->number` and `number->SCHEMENAME`. `SCHEMENAME->number` takes one argument and converts a symbol (or a list of symbols) into its numeric value. `number->SCHEMENAME` takes one argument and converts a numeric value into its scheme value.

Note that the specification of a scheme value override (`SCHEME-VALUE`) means the mapping may not be closed! `(number->SCHEMENAME (SCHEMENAME->number SCHEMENAME))` may not equal SCHEMENAME.

Here a heavily contrived example:

```
#>
enum foo { a_foo = 4, b_foo, c_foo };
enum foo bar(enum foo x) { printf("%d\n", x); return b_foo; }
<#

(define-foreign-enum (foo (enum "foo")) a_foo b_foo (c c_foo))

(define bar (foreign-lambda foo bar foo))

(pp (bar '()))
(pp (bar 'a_foo))
(pp (bar '(b_foo c)))
```

28.8 foreign-lambda

`[syntax] (foreign-lambda RETURNTYPE NAME ARGTYPE ...)`

Represents a binding to an external routine. This form can be used in the position of an ordinary `lambda` expression. NAME specifies the name of the external procedure and should be a string or a symbol.

28.9 foreign-lambda*

`[syntax] (foreign-lambda* RETURNTYPE ((ARGTYPE VARIABLE) ...) STRING ...)`

Similar to `foreign-lambda`, but instead of generating code to call an external function, the body of the C procedure is directly given in `STRING ...`:

```
(define my-strlen
  (foreign-lambda* int ((c-string str))
    "int n = 0; while(*(str++)) ++n; C_return(n);") )

(my-strlen "one two three")                ==> 13
```

For obscure technical reasons you should use the `C_return` macro instead of the normal `return` statement to return a result from the foreign lambda body as some cleanup code has to be run before execution commences in the calling code.

28.10 foreign-safe-lambda

[syntax] (foreign-safe-lambda RETURNTYPE NAME ARGTYPE ...)

This is similar to `foreign-lambda`, but also allows the called function to call Scheme functions and allocate Scheme data-objects. See Callbacks.

28.11 foreign-safe-lambda*

[syntax] (foreign-safe-lambda* RETURNTYPE ((ARGTYPE VARIABLE)...) STRING ...)

This is similar to `foreign-lambda*`, but also allows the called function to call Scheme functions and allocate Scheme data-objects. See Callbacks.

28.12 foreign-primitive

[syntax] (foreign-primitive [RETURNTYPE] ((ARGTYPE VARIABLE) ...) STRING ...)

This is also similar to `foreign-lambda*` but the code will be executed in a *primitive* CPS context, which means it will not actually return, but call it's continuation on exit. This means that code inside this form may allocate Scheme data on the C stack (the *nursery*) with `C_alloc` (see below). If the `RETURNTYPE` is omitted it defaults to `void`. You can return multiple values inside the body of the `foreign-primitive` form by calling this C function:

`C_values(N + 2, C_SCHEME_UNDEFINED, C_k, X1, ...)`

where `N` is the number of values to be returned, and `X1, ...` are the results, which should be Scheme data objects. When returning multiple values, the return-type should be omitted.

Previous: Interface to external functions and variables

Next: Foreign type specifiers

29 Foreign type specifiers

Here is a list of valid foreign type specifiers:

29.1 scheme-object

An arbitrary Scheme data object (immediate or non-immediate).

29.2 bool

As argument: any value (`#f` is false, anything else is true).

As result: anything different from 0 and the `NULL` pointer is `#t`.

29.3 byte unsigned-byte

A byte.

29.4 char unsigned-char

A character.

29.5 short unsigned-short

A short integer number.

29.6 int unsigned-int int32 unsigned-int32

An small integer number in fixnum range (at least 30 bit).

29.7 integer unsigned-integer integer32 unsigned-integer32 integer64

Either a fixnum or a flonum in the range of a (unsigned) machine *int* or with 32/64 bit width.

29.8 long unsigned-long

Either a fixnum or a flonum in the range of a (unsigned) machine *long* or with 32 bit width.

29.9 float double

A floating-point number. If an exact integer is passed as an argument, then it is automatically converted to a float.

29.10 number

A floating-point number. Similar to `double`, but when used as a result type, then either an exact integer or a floating-point number is returned, depending on whether the result fits into an exact integer or not.

29.11 symbol

A symbol, which will be passed to foreign code as a zero-terminated string.

When declared as the result of foreign code, the result should be a string and a symbol with the same name will be interned in the symbol table (and returned to the caller).

29.12 scheme-pointer

An untyped pointer to the contents of a non-immediate Scheme object (not allowed as return type). The value `#f` is also allowed and is passed as a `NULL` pointer.

Don't confuse this type with `(c-pointer ...)` which means something different (a machine-pointer object).

29.13 nonnull-scheme-pointer

As `scheme-pointer`, but guaranteed not to be `#f`.

Don't confuse this type with (`nonnull-c-pointer ...`) which means something different (a machine-pointer object).

29.14 c-pointer

An untyped operating-system pointer or a locative. The value `#f` is also allowed and is passed as a `NULL` pointer. If uses as the type of a return value, a `NULL` pointer will be returned as `#f`.

29.15 nonnull-c-pointer

As `c-pointer`, but guaranteed not to be `#f/NULL`.

29.16 blob

A blob object, passed as a pointer to its contents. Arguments of type `blob` may optionally be `#f`, which is passed as a NULL pointer.

This is not allowed as a return type.

29.17 nonnull-blob

As `blob`, but guaranteed not to be `#f`.

29.18 u8vector u16vector u32vector s8vector s16vector s32vector f32vector f64vector

A SRFI-4 number-vector object, passed as a pointer to its contents.

These type specifiers are not allowed as return types.

29.19 nonnull-u8vector nonnull-u16vector nonnull-u32vector nonnull-s8vector nonnull-s16vector nonnull-s32vector nonnull-f32vector nonnull-f64vector

As `u8vector ...`, but guaranteed not to be `#f`.

29.20 c-string

A C string (zero-terminated). The value `#f` is also allowed and is passed as a `NULL` pointer. If uses as the type of a return value, a `NULL` pointer will be returned as `#f`. Note that the string is copied (with a zero-byte appended) when passed as an argument to a foreign function. Also a return value of this type is copied into garbage collected memory.

29.21 nonnull-c-string

As `c-string`, but guaranteed not to be `#f/NULL`.

29.22 [nonnull-] c-string*

Similar to `[nonnull-] c-string`, but if used as a result-type, the pointer returned by the foreign code will be freed (using the C-libraries `free(1)`) after copying. This type specifier is not valid as a result type for callbacks defined with `define-external`.

29.23 [nonnull-] unsigned-c-string[*]

Same as `c-string`, but maps to the `unsigned char *` C type.

29.24 c-string-list

Expects a pointer to a list of C strings teminated by a `NULL` pointer and returns a list of strings.

Only valid as a result type of non-callback functions.

29.25 c-string-list*

Similar to `c-string-list` but releases the storage of each string and the pointer array using `free(1)`.

29.26 void

Specifies an undefined return value.

Not allowed as argument type.

29.27 (const TYPE)

The foreign type TYPE with an additional `const` specifier.

29.28 (enum NAME)

An enumeration type. Handled internally as an `integer`.

29.29 (c-pointer TYPE)

An operating-system pointer or a locative to an object of TYPE.

29.30 (nonnull-c-pointer TYPE)

As `(c-pointer TYPE)`, but guaranteed not to be `#f/NULL`.

29.31 (ref TYPE)

A C++ reference type. Reference types are handled the same way as pointers inside Scheme code.

29.32 (struct NAME)

A struct of the name `NAME`, which should be a string.

Structs cannot be directly passed as arguments to foreign function, neither can they be result values. Pointers to structs are allowed, though.

29.33 (template TYPE ARGTYPE ...)

A C++ template type. For example `vector<int>` would be specified as `(template "vector" int)`.

Template types cannot be directly passed as arguments or returned as results.

29.34 (union NAME)

A union of the name `NAME`, which should be a string.

Unions cannot be directly passed as arguments to foreign function, neither can they be result values. Pointers to unions are allowed, though.

29.35 (instance CNAME SCHEMECLASS)

A pointer to a C++ class instance. `CNAME` should designate the name of the C++ class, and `SCHEMECLASS` should be the class that wraps the instance pointer. Normally `SCHEMECLASS` should be a subclass of `<c++-object>`.

29.36 (instance-ref CNAME SCHEMECLASS)

A reference to a C++ class instance.

29.37 (function RESULTTYPE (ARGUMENTTYPE1 ... [...]) [CALLCONV])

A function pointer. `CALLCONV` specifies an optional calling convention and should be a string. The meaning of this string is entirely platform dependent. The value `#f` is also allowed and is passed as a `NULL` pointer.

29.38 Mappings

Foreign types are mapped to C types in the following manner:

bool	int
[unsigned-]char	[unsigned] char
[unsigned-]short	[unsigned] short
[unsigned-]int	[unsigned] int
[unsigned-]integer	[unsigned] int
[unsigned-]long	[unsigned] long
float	float
double	double
number	double
[nonnull-]c-pointer	void *
[nonnull-]blob	unsigned char *
[nonnull-]u8vector	unsigned char *
[nonnull-]s8vector	char *
[nonnull-]u16vector	unsigned short *
[nonnull-]s16vector	short *
[nonnull-]u32vector	uint32_t *
[nonnull-]s32vector	int32_t *
[nonnull-]f32vector	float *
[nonnull-]f64vector	double *
[nonnull-]c-string	char *
[nonnull-]unsigned-c-string	unsigned char *
c-string-list	char **
symbol	char *
void	void
([nonnull-]c-pointer TYPE)	TYPE *
(enum NAME)	enum NAME
(struct NAME)	struct NAME
(ref TYPE)	TYPE &
(template T1 T2 ...)	T1<T2, ...>
(union NAME)	union NAME
(function RTYPE (ATYPE ...) [CALLCONV])	[CALLCONV] RTYPE (*)(ATYPE, ...)
(instance CNAME SNAME)	CNAME *
(instance-ref CNAME SNAME)	CNAME &

Previous: <u>Accessing external objects</u>

Next: <u>Embedding</u>

30 Embedding

Compiled Scheme files can be linked with C code, provided the Scheme code was compiled in *embedded* mode by passing `-DC_EMBEDDED` to the C compiler (this will disable generation of a `main()` function). `csc` will do this, when given the `-embedded` option. Alternatively pass `-embedded` to `csc`.

The following C API is available:

30.1 CHICKEN_parse_command_line

[C function] void CHICKEN_parse_command_line (int argc, char *argv[], int *heap

Parse the programs command-line contained in `argc` and `argv` and return the heap-, stack- and symbol table limits given by runtime options of the form `-:...`, or choose default limits. The library procedure `argv` can access the command-line only if this function has been called by the containing application.

30.2 CHICKEN_initialize

[C function] int CHICKEN_initialize (int heap, int stack, int symbols, void *to

Initializes the Scheme execution context and memory. `heap` holds the number of bytes that are to be allocated for the secondary heap. `stack` holds the number of bytes for the primary heap. `symbols` contains the size of the symbol table. Passing `0` to one or more of these parameters will select a default size. `toplevel` should be a pointer to the toplevel entry point procedure. You should pass `C_toplevel` here. In any subsequent call to `CHICKEN_run` you can simply pass `NULL`. Calling this function more than once has no effect. If enough memory is available and initialization was successful, then `1` is returned, otherwise this function returns `0`.

30.3 CHICKEN_run

[C function] C_word CHICKEN_run (void *toplevel)

Starts the Scheme program. Call this function once to execute all toplevel expressions in your compiled Scheme program. If the runtime system was not initialized before, then `CHICKEN_initialize` is called with default sizes. `toplevel` is the toplevel entry-point procedure, you usually pass `C_toplevel` here. The result value is the continuation that can be used to re-invoke the Scheme code from the point after it called `return-to-host` (see below).

If you just need a Scheme interpreter, you can also pass `CHICKEN_default_toplevel` as the toplevel procedure, which just uses the default library units.

Once `CHICKEN_run` has been called, Scheme code is executing until all toplevel expressions have been evaluated or until `return-to-host` is called inside the Scheme program.

30.4 return-to-host

[procedure] (return-to-host)

Exits the Scheme code and returns to the invoking context that called `CHICKEN_run` or `CHICKEN_continue`.

After `return-to-host` has been executed and once `CHICKEN_run` returns, you can invoke callbacks which have been defined with `define-external`. The `eval` library unit also provides *boilerplate* callbacks, that simplify invoking Scheme code embedded in a C or C++ application a lot.

30.5 CHICKEN_eval

[C macro] int CHICKEN_eval (C_word exp, C_word *result)

Evaluates the Scheme object passed in `exp`, writing the result value to `result`. The return value is 1 if the operation succeeded, or 0 if an error occurred. Call `CHICKEN_get_error_message` to obtain a description of the error.

30.6 CHICKEN_eval_string

[C macro] int CHICKEN_eval_string (char *str, C_word *result)

Evaluates the Scheme expression passed in the string `str`, writing the result value to `result`.

30.7 CHICKEN_eval_to_string

[C macro] int CHICKEN_eval_to_string (C_word exp, char *result, int size)

Evaluates the Scheme expression passed in `exp`, writing a textual representation of the result into `result`. `size` should specify the maximal size of the result string.

30.8 CHICKEN_eval_string_to_string

[C macro] int CHICKEN_eval_string_to_string (char *str, char *result, int size)

Evaluates the Scheme expression passed in the string `str`, writing a textual representation of the result into `result`. `size` should specify the maximal size of the result string.

30.9 CHICKEN_apply

[C macro] int CHICKEN_apply (C_word func, C_word args, C_word *result)

Applies the procedure passed in `func` to the list of arguments `args`, writing the result value to `result`.

30.10 CHICKEN_apply_to_string

[C macro] int CHICKEN_apply_to_string (C_word func, C_word args, char *result,

Applies the procedure passed in `func` to the list of arguments `args`, writing a textual representation of the result into `result`.

30.11 CHICKEN_read

[C macro] int CHICKEN_read (char *str, C_word *result)

Reads a Scheme object from the string `str`, writing the result value to `result`.

30.12 CHICKEN_load

[C macro] int CHICKEN_load (char *filename)

Loads the Scheme file `filename` (either in source form or compiled).

30.13 CHICKEN_get_error_message

[C macro] void CHICKEN_get_error_message (char *result, int size)

Returns a textual description of the most recent error that occurred in executing embedded Scheme code.

30.14 CHICKEN_yield

[C macro] int CHICKEN_yield (int *status)

If threads have been spawned during earlier invocations of embedded Scheme code, then this function will run the next scheduled thread for one complete time-slice. This is useful, for example, inside an *idle* handler in a GUI application with background Scheme threads. Note that the `srfi-18` library unit has to be linked in for this.

An example:

```
% cat x.scm
;;; x.scm

(define (bar x) (gc) (* x x))

(define-external (baz (int i)) double
  (sqrt i))
(return-to-host)

% cat y.c
/* y.c */

#include <chicken.h>
#include <assert.h>

extern double baz(int);

int main() {
  char buffer[ 256 ];
  int status;
  C_word val = C_SCHEME_UNDEFINED;
  C_word *data[ 1 ];

  data[ 0 ] = &val;

  CHICKEN_run(C_toplevel);

  status = CHICKEN_read("(bar 99)", &val);
  assert(status);

  C_gc_protect(data, 1);

  printf("data: %08x\n", val);

  status = CHICKEN_eval_string_to_string("(bar)", buffer, 255);
  assert(!status);

  CHICKEN_get_error_message(buffer, 255);
  printf("ouch: %s\n", buffer);

  status = CHICKEN_eval_string_to_string("(bar 23)", buffer, 255);
  assert(status);

  printf("-> %s\n", buffer);
  printf("data: %08x\n", val);

  status = CHICKEN_eval_to_string(val, buffer, 255);
  assert(status);
  printf("-> %s\n", buffer);

  printf("->` %g\n", baz(22));

  return 0;
```

```
}

% csc x.scm y.c -embedded
```

It is also possible to re-enter the computation following the call to `return-to-host` by calling `CHICKEN_continue`:

30.15 CHICKEN_continue

```
[C function] C_word CHICKEN_continue (C_word k)
```

Re-enters Scheme execution. `k` is the continuation received from the previous invocation of `CHICKEN_run` or `CHICKEN_continue`. When `return-to-host` is called again, this function returns another continuation that can be used to restart again.

If you invoke callbacks prior to calling `CHICKEN_continue`, make sure that the continuation is not reclaimed by garbage collection. This can be avoided by using `C_gc_protect` or gc-roots.

Another example:

```
% cat x.scm
(require-extension srfi-18)

(define m (make-mutex))

(define (t)
  (mutex-lock! m)
  (thread-sleep! 1)
  (print (thread-name (current-thread)))
  (mutex-unlock! m)
  (t) )

(thread-start! (make-thread t 'PING!))
(thread-start! (make-thread t 'PONG!))

(let loop ()
  (return-to-host)
  (thread-yield!)
  (loop) )

% cat y.c
#include <chicken.h>

int main()
{
  C_word k = CHICKEN_run(C_toplevel);

  for(;;)
    k = CHICKEN_continue(k);

  return 0;
}
```

```
% csc x.scm y.c -embedded
```

It is advisable not to mix repeated uses of `CHICKEN_continue`/`return-to-host` (as in the example above) with callbacks. Once `return-to-host` is invoked, the runtime system and any Scheme code executed prior to the invocation is initialized and can be conveniently used via callbacks.

A simpler interface For handling GC-safe references to Scheme data are the so called *gc-roots*:

30.16 CHICKEN_new_gc_root

`[C function] void* CHICKEN_new_gc_root ()`

Returns a pointer to a *GC root*, which is an object that holds a reference to a Scheme value that will always be valid, even after a garbage collection. The content of the gc root is initialized to an unspecified value.

30.17 CHICKEN_delete_gc_root

`[C function] void CHICKEN_delete_gc_root (void *root)`

Deletes the gc root.

30.18 CHICKEN_gc_root_ref

`[C macro] C_word CHICKEN_gc_root_ref (void *root)`

Returns the value stored in the gc root.

30.19 CHICKEN_gc_root_set

`[C macro] void CHICKEN_gc_root_set (void *root, C_word value)`

Sets the content of the GC root to a new value.

Sometimes it is handy to access global variables from C code:

30.20 CHICKEN_global_lookup

[C function] void* CHICKEN_global_lookup (char *name)

Returns a GC root that holds the global variable with the name name. If no such variable exists, NULL is returned.

30.21 CHICKEN_global_ref

[C function] C_word CHICKEN_global_ref (void *global)

Returns the value of the global variable referenced by the GC root global.

30.22 CHICKEN_global_set

[C function] void CHICKEN_global_set (void *global, C_word value)

Sets the value of the global variable referenced by the GC root global to value.

Previous: Foreign type specifiers

Next: Callbacks

31 Callbacks

To enable an external C function to call back to Scheme, the form `foreign-safe-lambda` (or `foreign-safe-lambda*`) has to be used. This generates special code to save and restore important state information during execution of C code. There are two ways of calling Scheme procedures from C: the first is to invoke the runtime function `C_callback` with the closure to be called and the number of arguments. The second is to define an externally visible wrapper function around a Scheme procedure with the `define-external` form.

Note: the names of all functions, variables and macros exported by the CHICKEN runtime system start with `C_`. It is advisable to use a different naming scheme for your own code to avoid name clashes. Callbacks (defined by `define-external`) do not capture the lexical environment.

Non-local exits leaving the scope of the invocation of a callback from Scheme into C will not remove the C call-frame from the stack (and will result in a memory leak).

31.1 define-external

```
[syntax] (define-external [QUALIFIERS] (NAME (ARGUMENTTYPE1 VARIABLE1) ...) RET
[syntax] (define-external NAME TYPE [INIT])
```

The first form defines an externally callable Scheme procedure. `NAME` should be a symbol, which, when converted to a string, represents a legal C identifier. `ARGUMENTTYPE1 ...` and `RETURNTYPE` are foreign type specifiers for the argument variables `VAR1 ...` and the result, respectively. `QUALIFIERS` is an optional qualifier for the foreign procedure definition, like `__stdcall`.

```
(define-external (foo (c-string x)) int (string-length x))
```

The second form of `define-external` can be used to define variables that are accessible from foreign code. It declares a global variable named by the symbol `NAME` that has the type `TYPE`. `INIT` can be an arbitrary expression that is used to initialize the variable. `NAME` is accessible from Scheme just like any other foreign variable defined by `define-foreign-variable`.

```
(define-external foo int 42)
((foreign-lambda* int ()
  "C_return(foo);"))          ==> 42
```

Note: don't be tempted to assign strings or bytevectors to external variables. Garbage collection moves those objects around, so it is very bad idea to assign pointers to heap-data. If you have to do so, then copy the data object into statically allocated memory (for example by using `object-evict`).

Results of type `scheme-object` returned by `define-external` are always allocated in the secondary heap, that is, not in the stack.

31.2 C_callback

[C function] C_word C_callback (C_word closure, int argc)

This function can be used to invoke the Scheme procedure `closure`. `argc` should contain the number of arguments that are passed to the procedure on the temporary stack. Values are put onto the temporary stack with the `C_save` macro.

31.3 C_callback_adjust_stack

[C function] void C_callback_adjust_stack (C_word *ptr, int size)

The runtime-system uses the stack as a special allocation area and internally holds pointers to estimated limits to distinguish between Scheme data objects inside the stack from objects outside of it. If you invoke callbacks at wildly differing stack-levels, these limits may shift from invocation to invocation. Callbacks defined with `define-external` will perform appropriate adjustments automatically, but if you invoke `C_callback` manually, you should perform a `C_callback_adjust_stack` to make sure the internal limits are set properly. `ptr` should point to some data object on the stack and `size` is the number of words contained in the data object (or some estimate). The call will make sure the limits are adjusted so that the value pointed to by `ptr` is located in the stack.

Previous: Embedding

Next: Locations

32 Locations

It is also possible to define variables containing unboxed C data, so called *locations*. It should be noted that locations may only contain simple data, that is: everything that fits into a machine word, and double-precision floating point values.

32.1 define-location

[syntax] (define-location NAME TYPE [INIT])

Identical to (define-external NAME TYPE [INIT]), but the variable is not accessible from outside of the current compilation unit (it is declared static).

32.2 let-location

[syntax] (let-location ((NAME TYPE [INIT]) ...) BODY ...)

Defines a lexically bound location.

32.3 location

[syntax] (location NAME)
[syntax] (location X)

This form returns a pointer object that contains the address of the variable NAME. If the argument to location is not a location defined by define-location, define-external or let-location, then

(location X)

is essentially equivalent to

(make-locative X)

(See the manual chapter or locatives for more information about locatives.

Note that (location X) may be abbreviated as #$X.

```
(define-external foo int)
((foreign-lambda* void (((c-pointer int) ip)) "*ip = 123;")
  (location foo))
foo
```

This facility is especially useful in situations, where a C function returns more than one result value:

#>

```
#include <math.h>
<#

(define modf
  (foreign-lambda double "modf" double (c-pointer double)) )

(let-location ([i double])
  (let ([f (modf 1.99 (location i))])
    (print "i=" i ", f=" f) ) )
```

See <u>location and c-string*</u> for a tip on returning a c-string* type.

location returns a value of type c-pointer, when given the name of a callback-procedure defined with define-external.

Previous: <u>Callbacks</u>

Next: <u>Other support procedures</u>

33 Other support procedures

33.1 argc+argv

```
[procedure] (argc+argv)
```

Returns two values: an integer and a foreign-pointer object representing the `argc` and `argv` arguments passed to the current process.

Previous: Locations

Next: C interface

34 C interface

The following functions and macros are available for C code that invokes Scheme or foreign procedures that are called by Scheme:

34.1 C_save

[C macro] void C_save (C_word x) :

Saves the Scheme data object x on the temporary stack.

34.2 C_restore

[C macro] void C_restore

Pops and returns the topmost value from the temporary stack.

34.3 C_fix

[C macro] C_word C_fix (int integer)

34.4 C_make_character

[C macro] C_word C_make_character (int char_code)

34.5 C_SCHEME_END_OF_LIST

[C macro] C_SCHEME_END_OF_LIST

34.6 C_word C_SCHEME_END_OF_FILE

[C macro] C_SCHEME_END_OF_FILE

34.7 C_word C_SCHEME_FALSE

[C macro] C_SCHEME_FALSE

34.8 C_word C_SCHEME_TRUE

[C macro] C_SCHEME_TRUE

These macros return immediate Scheme data objects.

34.9 C_string

[C function] C_word C_string (C_word **ptr, int length, char *string)

34.10 C_string2

[C function] C_word C_string2 (C_word **ptr, char *zero_terminated_string)

34.11 C_intern2

[C function] C_word C_intern2 (C_word **ptr, char *zero_terminated_string)

34.12 C_intern3

[C function] C_word C_intern3 (C_word **ptr, char *zero_terminated_string, C_wo

34.13 C_pair

[C function] C_word C_pair (C_word **ptr, C_word car, C_word cdr)

34.14 C_flonum

[C function] C_word C_flonum (C_word **ptr, double number)

34.15 C_int_to_num

[C function] C_word C_int_to_num (C_word **ptr, int integer)

34.16 C_mpointer

[C function] C_word C_mpointer (C_word **ptr, void *pointer)

34.17 C_vector

[C function] C_word C_vector (C_word **ptr, int length, ...)

34.18 C_list

[C function] C_word C_list (C_word **ptr, int length, ...)

These functions allocate memory from `ptr` and initialize a fresh data object. The new data object is returned. `ptr` should be the **address** of an allocation pointer created with `C_alloc`.

34.19 C_alloc

[C macro] C_word* C_alloc (int words)

Allocates memory from the C stack (`C_alloc`) and returns a pointer to it. `words` should be the number of words needed for all data objects that are to be created in this function. Note that stack-allocated data objects have to be passed to Scheme callback functions, or they will not be seen by the garbage collector. This is really only usable for callback procedure invocations, make sure not to use it in normal code, because the allocated memory will be re-used after the foreign procedure returns. When invoking Scheme callback procedures a minor garbage collection is performed, so data allocated with `C_alloc` will already have moved to a safe place.

Note that `C_alloc` is really just a wrapper around `alloca`, and can also be simulated by declaring a stack-allocated array of `C_word`s:

34.20 C_SIZEOF_LIST

[C macro] int C_SIZEOF_LIST (int length)

34.21 C_SIZEOF_STRING

[C macro] int C_SIZEOF_STRING (int length)

34.22 C_SIZEOF_VECTOR

[C macro] int C_SIZEOF_VECTOR (int length)

34.23 C_SIZEOF_INTERNED_SYMBOL

[C macro] int C_SIZEOF_INTERNED_SYMBOL (int length)

34.24 C_SIZEOF_PAIR

[C macro] int C_SIZEOF_PAIR

34.25 C_SIZEOF_FLONUM

[C macro] int C_SIZEOF_FLONUM

34.26 C_SIZEOF_POINTER

[C macro] int C_SIZEOF_POINTER

34.27 C_SIZEOF_LOCATIVE

[C macro] int C_SIZEOF_LOCATIVE

34.28 C_SIZEOF_TAGGED_POINTER

[C macro] int C_SIZEOF_TAGGED_POINTER

These are macros that return the size in words needed for a data object of a given type.

34.29 C_character_code

[C macro] int C_character_code (C_word character)

34.30 C_unfix

[C macro] int C_unfix (C_word fixnum)

34.31 C_flonum_magnitude

[C macro] double C_flonum_magnitude (C_word flonum)

34.32 C_c_string

[C function] char* C_c_string (C_word string)

34.33 C_num_to_int

[C function] int C_num_to_int (C_word fixnum_or_flonum)

34.34 C_pointer_address

[C function] void* C_pointer_address (C_word pointer)

These macros and functions can be used to convert Scheme data objects back to C data. Note that C_c_string() returns a pointer to the character buffer of the actual Scheme object and is not zero-terminated.

34.35 C_header_size

[C macro] int C_header_size (C_word x)

34.36 C_header_bits

[C macro] int C_header_bits (C_word x)

Return the number of elements and the type-bits of the non-immediate Scheme data object x.

34.37 C_block_item

[C macro] C_word C_block_item (C_word x, int index)

This macro can be used to access slots of the non-immediate Scheme data object x. index specifies the index of the slot to be fetched, starting at 0. Pairs have 2 slots, one for the **car** and one for the **cdr**. Vectors have one slot for each element.

34.38 C_u_i_car

[C macro] C_word C_u_i_car (C_word x)

34.39 C_u_i_cdr

[C macro] C_word C_u_i_cdr (C_word x)

Aliases for C_block_item(x, 0) and C_block_item(x, 1), respectively.

34.40 C_data_pointer

[C macro] void* C_data_pointer (C_word x)

Returns a pointer to the data-section of a non-immediate Scheme object.

34.41 C_make_header

[C macro] C_word C_make_header (C_word bits, C_word size)

A macro to build a Scheme object header from its bits and size parts.

34.42 C_mutate

[C function] C_word C_mutate (C_word *slot, C_word val)

Assign the Scheme value `val` to the location specified by `slot`. If the value points to data inside the nursery (the first heap-generation), then the garbage collector will remember to handle the data appropriately. Assigning nursery-pointers directly will otherwise result in lost data. Note that no copying takes place at the moment when `C_mutate` is called, but later - at the next (minor) garbage collection.

34.43 C_symbol_value

[C macro] C_word C_symbol_value (C_word symbol)

Returns the global value of the variable with the name `symbol`. If the variable is unbound `C_SCHEME_UNBOUND` is returned. You can set a variable's value with `C_mutate(&C_symbol_value(SYMBOL), VALUE)`.

34.44 C_gc_protect

[C function] void C_gc_protect (C_word *ptrs[], int n)

Registers `n` variables at address `ptrs` to be garbage collection roots. The locations should not contain pointers to data allocated in the nursery, only immediate values or pointers to heap-data are valid. Any assignment of potential nursery data into a root-array should be done via `C_mutate()`. The variables have to be initialized to sensible values before the next garbage collection starts (when in doubt, set all locations in `ptrs` to `C_SCHEME_UNDEFINED`) `C_gc_protect` may not called before the runtime system has been initialized (either by `CHICKEN_initialize`, `CHICKEN_run` or `CHICKEN_invoke`.

For a slightly simpler interface to creating and using GC roots see `CHICKEN_new_gc_root`.

34.45 C_gc_unprotect

[C function] void C_gc_unprotect (int n)

Removes the last `n` registered variables from the set of root variables.

34.46 C_pre_gc_hook

`[C Variable] void (*C_pre_gc_hook)(int mode)`

If not NULL, the function pointed to by this variable will be called before each garbage collection with a flag indicating what kind of collection was performed (either 0 for a minor collection or 2 for a resizing collection). A "resizing" collection means a secondary collection that moves all live data into a enlarged (or shrinked) heap-space. Minor collections happen very frequently, so the hook function should not consume too much time. The hook function may not invoke Scheme callbacks.

Note that resizing collections may be nested in normal major collections.

34.47 C_post_gc_hook

`[C Variable] void (*C_post_gc_hook)(int mode, long ms)`

If not NULL, the function pointed to by this variable will be called after each garbage collection with a flag indicating what kind of collection was performed (either 0 for a minor collection, 1 for a major collection or 2 for a resizing collection). Minor collections happen very frequently, so the hook function should not consume too much time. The hook function may not invoke Scheme callbacks. The ms argument records the number of milliseconds required for the garbage collection, if the collection was a major one. For minor collections the value of the ms argument is undefined.

34.48 An example for simple calls to foreign code involving callbacks

```
% cat foo.scm
#>
extern int callout(int, int, int);
<#

(define callout (foreign-safe-lambda int "callout" int int int))

(define-external (callin (scheme-object xyz)) int
  (print "This is 'callin': " xyz)
  123)

(print (callout 1 2 3))

% cat bar.c
#include <stdio.h>
#include "chicken.h"

extern int callout(int, int, int);
extern int callin(C_word x);

int callout(int x, int y, int z)
{
```

```
  C_word *ptr = C_alloc(C_SIZEOF_LIST(3));
  C_word lst;

  printf("This is 'callout': %d, %d, %d\n", x, y, z);
  lst = C_list(&ptr, 3, C_fix(x), C_fix(y), C_fix(z));
  return callin(lst);  /* Note: `callin' will have GC'd the data in `ptr' */
}

% csc foo.scm bar.c -o foo
% foo
This is 'callout': 1, 2, 3
This is 'callin': (1 2 3)
123
```

34.49 Notes:

Scheme procedures can call C functions, and C functions can call Scheme procedures, but for every pending C stack frame, the available size of the first heap generation (the *nursery*) will be decreased, because the C stack is identical to the nursery. On systems with a small nursery this might result in thrashing, since the C code between the invocation of C from Scheme and the actual calling back to Scheme might build up several stack-frames or allocates large amounts of stack data. To prevent this it is advisable to increase the default nursery size, either when compiling the file (using the `-nursery` option) or when running the executable (using the `-:s` runtime option).

Calls to Scheme/C may be nested arbitrarily, and Scheme continuations can be invoked as usual, but keep in mind that C stack frames will not be recovered, when a Scheme procedure call from C does not return normally.

When multiple threads are running concurrently, and control switches from one thread to another, then the continuation of the current thread is captured and saved. Any pending C stack frame still active from a callback will remain on the stack until the threads is re-activated again. This means that in a multithreading situation, when C callbacks are involved, the available nursery space can be smaller than expected. So doing many nested SchemeCScheme calls can reduce the available memory up to the point of thrashing. It is advisable to have only a single thread with pending C stack-frames at any given time.

Pointers to Scheme data objects should not be stored in local or global variables while calling back to Scheme. Any Scheme object not passed back to Scheme will be reclaimed or moved by the garbage collector.

Calls from C to Scheme are never tail-recursive.

Continuations captured via `call-with-current-continuation` and passed to C code can be invoked like any other Scheme procedure.

35 chicken-setup

35.1 Extension libraries

Extension libraries (*eggs*) are extensions to the core functionality provided by the basic CHICKEN system, to be built and installed separately. The mechanism for loading compiled extensions is based on dynamically loadable code and as such is only available on systems on which loading compiled code at runtime is supported. Currently these are most UNIX-compatible platforms that provide the `libdl` functionality like Linux, Solaris, BSD, Mac OS X and Windows using Cygwin.

Note: Extension may also be normal applications or shell scripts, but are usually libraries.

`chicken-setup` will download the source code for extension automatically from the canonical server at http://www.call-with-current-continuation.org/eggs if the requested egg does not exist in the current directory. Various command-line options exist for customizing the process and/or retrieving the egg from other locations or in other formats.

35.2 Installing extensions

To install an extension library, run the `chicken-setup` program with the extension name as argument. The extension archive is downloaded, its contents extracted and the contained *setup* script is executed. This setup script is a normal Scheme source file, which will be interpreted by `chicken-setup`. The complete language supported by `csi` is available, and the library units `srfi-1 regex utils posix tcp` are loaded. Additional libraries can be loaded at run-time.

The setup script should perform all necessary steps to build the new library (or application). After a successful build, the extension can be installed by invoking one of the procedures `install-extension`, `install-program` or `install-script`. These procedures will copy a number of given files into the extension repository or in the path where the CHICKEN executables are located (in the case of executable programs or scripts). Additionally the list of installed files, and user-defined metadata is stored in the repository.

If no extension name is given on the command-line, and if none of the options `-list`, `-version`, `-repository` (without argument), `-program-path` (without argument), `-fetch`, `-fetch-tree` or `-docindex` is given, then all `.setup` scripts in the current directory are processed.

35.3 Creating extensions

Extensions can be created by creating an (optionally gzipped) `tar` archive named `EXTENSION.egg` containing all needed files plus a `.setup` script in the root directory. After `chicken-setup` has extracted the files, the setup script will be invoked. There are no additional constraints on the structure of the archive, but the setup script has to be in the root path of the archive.

35.4 Procedures and macros available in setup scripts

35.4.1 install-extension

```
(install-extension ID FILELIST [INFOLIST])
```

Installs the extension library with the name `ID`. All files given in the list of strings `FILELIST` will be copied to the extension repository. It should be noted here that the extension id has to be identical to the name of the file implementing the extension. The extension may load or include other files, or may load other extensions at runtime specified by the `require-at-runtime` property.

`FILELIST` may be a filename, a list of filenames, or a list of pairs of the form `(SOURCE DEST)` (if you want to copy into a particular sub-directory - the destination directory will be created as needed). If `DEST` is a relative pathname, < it will be copied into the extension repository.

The optional argument `INFOLIST` should be an association list that maps symbols to values, this list will be stored as `ID.setup-info` at the same location as the extension code. Currently the following properties are used:

35.4.1.1 syntax

```
[extension property] (syntax)
```

Marks the extension as syntax-only. No code is compiled, the extension is intended as a file containing macros to be loaded at compile/macro-expansion time.

35.4.1.2 require-at-runtime

```
[extension property] (require-at-runtime ID ...)
```

Specifies extensions that should be loaded (via `require`) at runtime. This is mostly useful for syntax extensions that need additional support code at runtime.

35.4.1.3 version

```
[extension property] (version STRING)
```

Specifies version string.

35.4.1.4 documentation

`[extension property] (documentation FILENAME)`

The filename of a HTML document containing extension-specific documentation. This file should be given in the file-list passed to `install-extension` and a link to it will be automatically included in the index page (accessible via `chicken-setup -docindex`).

35.4.1.5 examples

`[extension property] (examples FILENAME ...)`

Copies the given files into the examples directory, which is usually `$prefix/share/chicken/examples` or `(make-pathname (chicken-home) "examples"))`.

Note that the files listed in this property should not be listed in the normal list of files to install passed to `install-extension`. This is the only exception - other files that are installed in the repository must be given in the file list.

35.4.1.6 exports

`[extension property] (exports EXPORT ...)`

Add export-information to the generated extension-information. `EXPORT` may be a symbol naming an exported toplevel variable or a string designating a file with exported variables, as generated by the `-emit-exports` option or the `emit-exports` declaration specifier.

35.4.1.7 static

`[extension property] (static STRING)`

If the extension also provides a static library, then STRING should contain the name of that library. Used by `csc` when compiling with the `-static-extensions` option.

35.4.1.8 static-options

`[extension property] (static-options STRING)`

Additional options that should be passed to the linker when linking with the static version of an extension (see `static` above). Used by `csc` when compiling with the `-static-extensions` option.

All other properties are currently ignored. The `FILELIST` argument may also be a single string.

35.4.2 install-program

[procedure] (install-program ID FILELIST [INFOLIST])

Similar to `install-extension`, but installs an executable program in the executable path (usually `/usr/local/bin`).

35.4.3 install-script

[procedure] (install-script ID FILELIST [INFOLIST])

Similar to `install-program`, but additionally changes the file permissions of all files in `FILELIST` to executable (for installing shell-scripts).

35.4.4 run

[syntax] (run FORM ...)

Runs the shell command `FORM`, which is wrapped in an implicit `quasiquote`. `(run (csc ...))` is treated specially and passes `-v` (if `-verbose` has been given to `chicken-setup`) and `-feature compiling-extension` options to the compiler.

35.4.5 compile

[syntax] (compile FORM ...)

Equivalent to `(run (csc FORM ...))`.

35.4.6 make

[syntax] (make ((TARGET (DEPENDENT ...) COMMAND ...) ...) ARGUMENTS)

A *make* macro that executes the expressions `COMMAND ...`, when any of the dependents `DEPENDENT ...` have changed, to build `TARGET`. This is the same as the `make` extension, which is available separately. For more information, see make.

35.4.7 patch

[procedure] (patch WHICH REGEX SUBST)

Replaces all occurrences of the regular expression REGEX with the string SUBST, in the file given in WHICH. If WHICH is a string, the file will be patched and overwritten. If WHICH is a list of the form OLD NEW, then a different file named NEW will be generated.

35.4.8 copy-file

[procedure] (copy-file FROM TO)

Copies the file or directory (recursively) given in the string FROM to the destination file or directory TO.

35.4.9 move-file

[procedure] (move-file FROM TO)

Moves the file or directory (recursively) given in the string FROM to the destination file or directory TO.

35.4.10 remove-file*

[procedure] (remove-file* PATH)

Removes the file or directory given in the string PATH.

35.4.11 find-library

[procedure] (find-library NAME PROC)

Returns #t if the library named libNAME.[a|so] (unix) or NAME.lib (windows) could be found by compiling and linking a test program. PROC should be the name of a C function that must be provided by the library. If no such library was found or the function could not be resolved, #f is returned.

35.4.12 find-header

[procedure] (find-header NAME)

Returns #t if a C include-file with the given name is available, or #f otherwise.

35.4.13 try-compile

[procedure] (try-compile CODE #!key cc cflags ldflags compile-only c++)

Returns #t if the C code in CODE compiles and links successfully, or #f otherwise. The keyword parameters cc (compiler name, defaults to the C compiler used to build this system), cflags and ldflags accept additional compilation and linking options. If compile-only is true, then no linking step takes place. If the keyword argument c++ is given and true, then the code will be compiled in C++ mode.

35.4.14 create-directory

[procedure] (create-directory PATH)

Creates the directory given in the string PATH, with all parent directories as needed.

35.4.15 installation-prefix

[parameter] installation-prefix

Holds the prefix under which CHICKEN executables and libraries have been installed (either the value of the environment variable CHICKEN_PREFIX or whatever prefix was specified at the time the system was built.

35.4.16 program-path

[parameter] (program-path [PATH])

Holds the path where executables are installed and defaults to either $CHICKEN_PREFIX/bin, if the environment variable CHICKEN_PREFIX is set or the path where the CHICKEN binaries (chicken, csi, etc.) are installed.

35.4.17 setup-root-directory

[parameter] (setup-root-directory [PATH])

Contains the path of the directory where chicken-setup was invoked.

35.4.18 setup-build-directory

[parameter] (setup-build-directory [PATH])

Contains the path of the directory where the extension is built. This is not necessarily identical to `setup-root-directory`.

35.4.19 setup-verbose-flag

[parameter] (setup-verbose-flag [BOOL])

Reflects the setting of the `-verbose` option, i.e. is `#t`, if `-verbose` was given.

35.4.20 setup-install-flag

[parameter] (setup-install-flag [BOOL])

Reflects the setting of the `--no-install` option, i.e. is `#f`, if `-no-install` was given.

35.4.21 required-chicken-version

[procedure] (required-chicken-version VERSION)

Signals an error if the version of CHICKEN that this script runs under is lexicographically less than `VERSION` (the argument will be converted to a string, first).

35.4.22 required-extension-version

[procedure] (required-extension-version EXTENSION1 VERSION1 ...)

Checks whether the extensions `EXTENSION1 ...` are installed and at least of version `VERSION1 ...`. The test is made by lexicographically comparing the string-representations of the given version with the version of the installed extension. If one of the listed extensions is not installed, has no associated version information or is of a version older than the one specified.

35.4.23 cross-chicken

[procedure] (cross-chicken)

Returns `#t` if this system is configured for cross-compilation or `#f` otherwise.

35.4.24 host-extension

```
[parameter] host-extension
```

For a cross-compiling CHICKEN, when compiling an extension, then it should be built for the host environment (as opposed to the target environment). This parameter is controlled by the `-host-extension` command-line option. A setup script should perform the proper steps of compiling any code by passing `-host` when invoking `csc` or using the `compile` macro.

35.5 Examples for extensions

The simplest case is a single file that does not export any syntax. For example

```
;;;; hello.scm

(define (hello name)
  (print "Hello, " name " !") )
```

We need a `.setup` script to build and install our nifty extension:

```
;;;; hello.setup

;; compile the code into a dynamically loadable shared object
;; (will generate hello.so)
(compile -s hello.scm)

;; Install as extension library
(install-extension 'hello "hello.so")
```

After entering

```
$ chicken-setup hello
```

at the shell prompt (and in the same directory where the two files exist), the file `hello.scm` will be compiled into a dynamically loadable library. If the compilation succeeds, `hello.so` will be stored in the repository, together with a file named `hello.setup-info` containing an a-list with metadata. If no extension name is given to `chicken-setup`, it will simply execute the first file with the `.setup` extension it can find.

Use it like any other CHICKEN extension:

```
$ csi -q
#;1> (require-extension hello)
; loading /usr/local/lib/chicken/1/hello.so ...
#;2> (hello "me")
Hello, me!
#;3>
```

Here we create a simple application:

```
;;;; hello2.scm
```

```scheme
(print "Hello, ")
(for-each (lambda (x) (printf "~A " x)) (command-line-arguments))
(print "!")
```

We also need a setup script:

```scheme
;;;; hello2.setup

(compile hello2.scm)  ; compile `hello2'
(install-program 'hello2 "hello2") ; name of the extension and files to be inst
```

To use it, just run `chicken-setup` in the same directory:

```
$ chicken-setup
```

(Here we omit the extension name)

Now the program `hello2` will be installed in the same location as the other CHICKEN tools (like `chicken`, `csi`, etc.), which will normally be `/usr/local/bin`. Note that you need write-permissions for those locations and may have to run `chicken-setup` with administrative rights.

Uninstallation is just as easy:

```
$ chicken-setup -uninstall hello2
```

`chicken-setup` provides a `make` macro, so build operations can be of arbitrary complexity. When running `chicken-setup` with an argument `NAME`, for which no associated file `NAME.setup`, `NAME.egg` or `NAME.scm` exists will ask you to download the extension via HTTP from the default URL http://www.call-with-current-continuation.org/eggs. You can use the `-host` option to specify an alternative source location. Extensions that are required to compile and/or use the requested extension are downloaded and installed automatically.

If the given extension name contains a path prefix and the `-host` option is given, then `chicken-setup` can also download and install eggs from an arbitrary HTTP server. Alternatively you can pass a full URL (including the `http://` prefix. Note that no dependency checks are done when downloading eggs directly with the URL syntax.

Finally a somewhat more complex example: We want to package a syntax extension with additional support code that is to be loaded at run-time of any Scheme code that uses that extension. We create a *glass* lambda, a procedure with free variables that can be manipulated from outside:

```scheme
;;;; glass.scm

(define-macro (glass-lambda llist vars . body)
  ;; Low-level macros are fun!
  (let ([lvar (gensym)]
        [svar (gensym)]
        [x (gensym)]
        [y (gensym)]
        [yn (gensym)] )
    `(let ,(map (lambda (v) (list v #f)) vars)
       (define (,svar ,x . ,y)
  let* ([(yn (pair? ,y)]
     and ,yn (car ,yy)]) )
    case ,x(
```

```
              ,@(mbda (v)
               if ,yn  `([,v] (
  set! ,v ,y)                       (
                                     ,v) ) )
                    vars)
        else (error "variable not found" ,x)) ) ) )
         (define ,lvar (lambda ,llist ,@body))
         (extend-procedure ,lvar ,svar) ) ) )
```

Here some support code that needs to be loaded at runtime:

```
;;;; glass-support.scm

(require-extension lolevel)

(define glass-lambda-accessor procedure-data)
(define (glass-lambda-ref gl v) ((procedure-data gl) v))
(define (glass-lambda-set! gl v x) ((procedure-data gl) v x))
```

The setup script looks like this:

```
(compile -s glass-support.scm)

(install-extension
  'glass
  '("glass.scm" "glass-support.so")
  '((syntax) (require-at-runtime glass-support)) )
```

The invocation of `install-extension` provides the files that are to be copied into the extension repository, and a metadata list that specifies that the extension `glass` is a syntax extension and that, if it is declared to be used by other code (either with the `require-extension` or `require-for-syntax` form), then client code should perform an implicit (`require 'glass-support`) at startup.

This can be conveniently packaged as an *egg*:

```
$ tar cfz glass.egg glass.setup glass.scm glass-support.scm
```

And now we use it:

```
$ chicken-setup glass
$ csi -quiet
#;1> (require-extension glass)
; loading /usr/local/lib/chicken/1/glass.scm ...
; loading /usr/local/lib/chicken/1/glass-support.so ...
#;2> (define foo (glass-lambda (x) (y) (+ x y)))
#;3> (glass-lambda-set! foo 'y 99)
#;4> (foo 33)
132
```

35.6 chicken-setup reference

Available options:

`-h -help`
> Show usage information and exit.

`-V -version`
> Display version and exit.

`-R -repository [PATHNAME]`
> When used without an argument, the path of the extension repository is displayed on standard output. When given an argument, the repository pathname (and the `repository-path` parameter) will be set to `PATHNAME` for all subsequent operations. The default repository path is the installation library directory (usually `/usr/local/lib/chicken`), or (if set) the directory given in the environment variable `CHICKEN_REPOSITORY`. `PATHNAME` should be an absolute pathname.

`-P -program-path [PATHNAME]`
> When used without an argument, the path for executables is displayed on standard output. When given an argument, the program path for installing executables and scripts will be set to `PATHNAME` for all subsequent operations. `PATHNAME` should be an absolute pathname.

`-h -host HOSTNAME[:PORT]`
> Specifies alternative host for downloading extensions, optionally with a TCP port number (which defaults to 80).

`-u -uninstall EXTENSION`
> Removes all files that were installed for `EXTENSION` from the file-system, together with any metadata that has been stored.

`-l -list [NAME ...]`
> List all installed extensions or show extension information.

`-r -run FILENAME`
> Load and execute given file.

`-s -script FILENAME`
> Executes the given Scheme source file with all remaining arguments and exit. The *she-bang* shell script header is recognized, so you can write Scheme scripts that use `chicken-setup` just as with `csi`.

`-e -eval EXPRESSION`
> Evaluates the given expression(s)

`-v -verbose`
> Display additional debug information

`-k -keep`
> Keep temporary files and directories

`-c -csc-option OPTION`
> Passes `OPTION` as an extra argument to invocations of the compiler-driver (`csc`); this works only if `csc` is invoked as `(run (csc ...))`

`-d -dont-ask`
> Do not ask the user before trying to download required extensions

`-n -no-install`
> Do not install generated binaries and/or support files; any invocations of `install-program`, `install-extension` or `install-script` will be be no-ops

`-i -docindex`
> Displays the path to the index-page of any installed extension-documentation; if the index page does not exist, it is created

`-t -test EXTENSION ...`
> return success if all given extensions are installed

`-ls EXTENSION`
> List installed files for extension

`-fetch-tree`

> Download and print the repository catalog

`-create-tree DIRECTORY`

> Create a fresh, minimal repository catalog and writes it to stdout

`-t -test`

> If the extension sources contain a directory named `tests` and this directory includes a file named `run.scm` then this file is executed (with `tests` being the current working directory)

`-tree FILENAME`

> Download and show the repository catalog

`-svn URL`

> Fetch extension from <u>Subversion</u> repository

`-revision REV`

> Specifies SVN revision to check out

`-local PATHNAME`

> Fetch extension from local file

`-destdir PATHNAME`

> Specify alternative installation prefix (for packaging)

`-host-extension`

> Compile extension in "host" mode (sets the parameter `host-extension` to `#f`)

`--`

> Ignore all following arguments

Note that the options are processed exactly in the order in which they appear in the command-line.

35.7 Windows notes

`chicken-setup` works on Windows, when compiled with Visual C++, but depends on the `tar` and `gunzip` tools to extract the contents of an egg. The best way is to download an egg either manually (or with `chicken-setup -fetch`) and extract its contents with a separate program (like `winzip`). the `CHICKEN_REPOSITORY` environment variable has to be set to a directory where your compiled extensions should be located.

The `.setup` scripts will not always work under Windows, and the extensions may require libraries that are not provided for Windows or work differently. Under these circumstances it is recommended to perform the required steps to build an extension manually.

The required UNIX tools are also available as Windows binaries. Google or ask on the CHICKEN mailing list if you need help locating them.

35.8 Security

When extensions are downloaded and installed one is executing code from potentially compromised systems. This applies also when `chicken-setup` executes system tests for required extensions. As the code has been retrieved over the network effectively untrusted code is going to be evaluated. When `chicken-setup` is run as *root* the whole system is at the mercy of the build instructions (note that this is also the case every time you install software via `sudo make install`, so this is not specific to the CHICKEN extension mechanism).

Security-conscious users should never run `chicken-setup` as root. A simple remedy is to set the environment variable `CHICKEN_REPOSITORY`, which will transparently place the repository at an arbitrary

user-selected location. Alternatively obtain write/execute access to the default location of the repository (usually `/usr/local/lib/chicken`) to avoid running as root.

35.9 Other modes of installation

It is possible to install extensions directly from a <u>Subversion</u> repository or from a local checkout by using the `-svn` or `-local` options. By using either the `svn` client program (which must be installed) or file-system operations, all necessary files will be copied into the current directory (creating a subdirectory named `EXTENSIONNAME.egg-dir`), built and subsequently installed.

Dependency information, which is necessary to ensure required extensions are also installed, is downloaded automatically. If you have no internet connection or don't want to connect, you can also use a local file containing the necessary dependency information. The `-fetch-tree` option retrieves the canonical *repository file* at <u>http://www.call-with-current-continuation.org/eggs/repository</u>, writing it to stdout. Redirecting this output into a file and passing the file via the `-tree` option to `chicken-setup` allows you now to use the local repository file:

Retrieve complete extension repository (big):

```
% cd /opt
% svn co https://galinha.ucpel.tche.br/svn/chicken-eggs/release/3 eggs
```

Get your own copy of the repository file:

```
% chicken-setup -fetch-tree >~/my-repository-file
```

Now you can install eggs from your local checkout, with full dependency tracking and without being connected to the internet:

```
% cd ~/tmp
% chicken-setup -local /opt/eggs -tree ~/my-repository-file opengl
```

35.10 Linking extensions statically

The compiler and <u>chicken-setup</u> support statically linked eggs. The general approach is to generate an object file or static library (in addition to the usual shared library) in your `.setup` script and install it along with the dynamically loadable extension. The setup properties `static` should contain the name of the object file (or static library) to be linked, when `csc` gets passed the `-static-extensions` option:

```
(compile -s -O2 -d1 my-ext.scm)   ; dynamically loadable "normal" version
(compile -c -O2 -d1 my-ext -unit my-ext)  ; statically linkable version
(install-extension
  'my-ext
  '("my-ext.so" "my-ext.o")
  '((static "my-ext.o")) )
```

Note the use of the `-unit` option in the second compilation step: static linking must use static library units. `chicken-setup` will perform platform-dependent file-extension translation for the file list, but does currently not do that for the `static` extension property.

To actually link with the static version of `my-ext`, do:

```
% csc -static-extensions my-program.scm -uses my-ext
```

The compiler will try to do the right thing, but can not handle all extensions, since the ability to statically link eggs is relatively new. Eggs that support static linking are designated as being able to do so. If you require a statically linkable version of an egg that has not been converted yet, contact the extension author or the CHICKEN mailing list.

Previous: Interface to external functions and variables

Next: Data representation

36 Data representation

Note: In all cases below, bits are numbered starting at 1 and beginning with the lowest-order bit.

There exist two different kinds of data objects in the CHICKEN system: immediate and non-immediate objects.

36.1 Immediate objects

Immediate objects are represented by a single machine word, which is usually of 32 bits length, or 64 bits on 64-bit architectures. The immediate objects come in four different flavors:

fixnums, that is, small exact integers, where bit 1 is set to 1. This gives fixnums a range of 31 bits for the actual numeric value (63 bits on 64-bit architectures).

characters, where bits 1-4 are equal to `C_CHARACTER_BITS`. The Unicode code point of the character is encoded in bits 9 to 32.

booleans, where bits 1-4 are equal to `C_BOOLEAN_BITS`. Bit 5 is one for #t and zero for #f.

other values: the empty list, the value of unbound identifiers, the undefined value (void), and end-of-file. Bits 1-4 are equal to `C_SPECIAL_BITS`; bits 5 to 8 contain an identifying number for this type of object. The following constants are defined: `C_SCHEME_END_OF_LIST` `C_SCHEME_UNDEFINED` `C_SCHEME_UNBOUND` `C_SCHEME_END_OF_FILE`

Collectively, bits 1 and 2 are known as the *immediate mark bits*. When bit 1 is set, the object is a fixnum, as described above, and bit 2 is part of its value. When bit 1 is clear but bit 2 is set, it is an immediate object other than a fixnum. If neither bit 1 nor bit 2 is set, the object is non-immediate, as described below.

36.2 Non-immediate objects

Non-immediate objects are blocks of data represented by a pointer into the heap. The pointer's immediate mark bits (bits 1 and 2) must be zero to indicate the object is non-immediate; this guarantees the data block is aligned on a 4-byte boundary, at minimum. Alignment of data words is required on modern architectures anyway, so we get the ability to distinguish between immediate and non-immediate objects for free.

The first word of the data block contains a header, which gives information about the type of the object. The header has the size of a machine word, usually 32 bits (64 bits on 64 bit architectures).

Bits 1 to 24 contain the length of the data object, which is either the number of bytes in a string (or byte-vector) or the the number of elements for a vector or for a structure type.

Bits 25 to 28 contain the type code of the object.

Bits 29 to 32 contain miscellaneous flags used for garbage collection or internal data type dispatching. These flags are:

C_GC_FORWARDING_BIT

Flag used for forwarding garbage collected object pointers.

C_BYTEBLOCK_BIT

Flag that specifies whether this data object contains raw bytes (a string or byte-vector) or pointers to other data objects.

C_SPECIALBLOCK_BIT

Flag that specifies whether this object contains a *special* non-object pointer value in its first slot. An example for this kind of objects are closures, which are a vector-type object with the code-pointer as the first item.

C_8ALIGN_BIT

Flag that specifies whether the data area of this block should be aligned on an 8-byte boundary (floating-point numbers, for example).

The actual data follows immediately after the header. Note that block-addresses are always aligned to the native machine-word boundary. Scheme data objects map to blocks in the following manner:

pairs: vector-like object (type bits `C_PAIR_TYPE`), where the car and the cdr are contained in the first and second slots, respectively.

vectors: vector object (type bits `C_VECTOR_TYPE`).

strings: byte-vector object (type bits `C_STRING_TYPE`).

procedures: special vector object (type bits `C_CLOSURE_TYPE`). The first slot contains a pointer to a compiled C function. Any extra slots contain the free variables (since a flat closure representation is used).

flonums: a byte-vector object (type bits `C_FLONUM_BITS`). Slots one and two (or a single slot on 64 bit architectures) contain a 64-bit floating-point number, in the representation used by the host systems C compiler.

symbols: a vector object (type bits `C_SYMBOL_TYPE`). Slots one and two contain the toplevel variable value and the print-name (a string) of the symbol, respectively.

ports: a special vector object (type bits `C_PORT_TYPE`). The first slot contains a pointer to a file- stream, if this is a file-pointer, or NULL if not. The other slots contain housekeeping data used for this port.

structures: a vector object (type bits `C_STRUCTURE_TYPE`). The first slot contains a symbol that specifies the kind of structure this record is an instance of. The other slots contain the actual record items.

pointers: a special vector object (type bits `C_POINTER_TYPE`). The single slot contains a machine pointer.

tagged pointers: similar to a pointer (type bits `C_TAGGED_POINTER_TYPE`), but the object contains an additional slot with a tag (an arbitrary data object) that identifies the type of the pointer.

Data objects may be allocated outside of the garbage collected heap, as long as their layout follows the above mentioned scheme. But care has to be taken not to mutate these objects with heap-data (i.e. non-immediate objects), because this will confuse the garbage collector.

For more information see the header file `chicken.h`.

Previous: <u>chicken-setup</u>

Next: <u>Bugs and limitations</u>

37 Bugs and limitations

- Compiling large files takes too much time.
- If a known procedure has unused arguments, but is always called without those parameters, then the optimizer *repairs* the procedure in certain situations and removes the parameter from the lambda-list.
- `port-position` currently works only for input ports.
- Leaf routine optimization can theoretically result in code that thrashes, if tight loops perform excessively many mutations.

Previous: Data representation

Next: FAQ

38 FAQ

This is the list of Frequently Asked Questions about Chicken Scheme. If you have a question not answered here, feel free to post to the chicken-users mailing list; if you consider your question general enough, feel free to add it to this list.

38.1 General

38.1.1 Why yet another Scheme implementation?

Since Scheme is a relatively simple language, a large number of implementations exist and each has its specific advantages and disadvantages. Some are fast, some provide a rich programming environment. Some are free, others are tailored to specific domains, and so on. The reasons for the existence of CHICKEN are:

- CHICKEN is portable because it generates C code that runs on a large number of platforms.
- CHICKEN is extendable, since its code generation scheme and runtime system/garbage collector fits neatly into a C environment.
- CHICKEN is free and can be freely distributed, including its source code.
- CHICKEN offers better performance than nearly all interpreter based implementations, but still provides full Scheme semantics.
- As far as we know, CHICKEN is the first implementation of Scheme that uses Henry Baker's <u>Cheney on the M.T.A</u> concept.

38.1.2 What should I do if I find a bug?

Send e-mail to <u>felix@call-with-current-continuation.org</u> with some hints about the problem, like version/build of the compiler, platform, system configuration, code that causes the bug, etc.

38.1.3 Why are values defined with `define-foreign-variable` or `define-constant` or `define-inline` not seen outside of the containing source file?

Accesses to foreign variables are translated directly into C constructs that access the variable, so the Scheme name given to that variable does only exist during compile-time. The same goes for constant- and inline-definitions: The name is only there to tell the compiler that this reference is to be replaced with the actual value.

38.1.4 How does `cond-expand` know which features are registered in used units?

Each unit used via `(declare (uses ...))` is registered as a feature and so a symbol with the unit-name can be tested by `cond-expand` during macro-expansion-time. Features registered using the `register-feature!` procedure are only available during run-time of the compiled file. You can use the `eval-when` form to register features at compile time.

38.1.5 Why are constants defined by `define-constant` not honoured in `case` constructs?

`case` expands into a cascaded `if` expression, where the first item in each arm is treated as a quoted list. So the `case` macro can not infer wether a symbol is to be treated as a constant-name (defined via `define-constant`) or a literal symbol.

38.1.6 How can I enable case sensitive reading/writing in user code?

To enable the `read` procedure to read symbols and identifiers case sensitive, you can set the parameter `case-sensitivity` to `#t`.

38.1.7 How can I change `match-error-control` during compilation?

Use `eval-when`, like this:

```
(eval-when (compile)
(match-error-control #:unspecified) )
```

38.1.8 Why doesn't CHICKEN support the full numeric tower by default?

The short answer:

```
% chicken-setup numbers
% csi -q
#;1> (use numbers)
```

The long answer:

There are a number of reasons for this:

- For most applications of Scheme fixnums (exact word-sized integers) and flonums (64-bit floating-point numbers) are more than sufficient;

- Interfacing to C is simpler;

- Dispatching of arithmetic operations is more efficient.

There is an extension based on the GNU Multiprecision Package that implements most of the full numeric tower, see http://galinha.ucpel.tche.br/numbers.

38.1.9 How can I specialize a generic function method to match instances of every class?

Specializing a method on `<object>` doesn't work on primitive data objects like numbers, strings, etc. so for example

```
(define-method (foo (x <my-class>)) ...)
(define-method (foo (x <object>)) ...)
(foo 123)
```

will signal an error, because to applicable method can be found. To specialize a method for primitive objects, use `<top>`:

```
(define-method (foo (x <top>)) ...)
```

38.1.10 Does CHICKEN support native threads?

Currently native threads are not supported. The runtime system is not reentrant, and the garbage-collection algorithm would be made much more complicated, since the location of every object (whether it is allocated on the stack or on the heap or completely outside the GC-able data space) has to be checked - this would be rather complex and inefficient in a situation where multiple threads are involved.

38.1.11 Does CHICKEN support Unicode strings?

Yes, as an extension.

By default all string and character functions operate bytewise, so that characters with an iteger value greater than 255 don't make much sense and multibyte UTF-8 characters are seen and manipulated as separate bytes, analogous to what a C program would see.

You can enable UTF-8 support by placing the following two lines at the beginning of your source file (or in your ~/.csirc for interactive sessions) before any other code, including other use directives:

```
(use iset syntax-case utf8)
(import utf8)
```

This will replace all builtin string operators with UTF-8-aware versions, that will treat strings as sequences of multibyte UTF-8 characters, thus enabling you to represent and manipulate Unicode characters while remaining compatible with most C libraries and system interfaces.

Most eggs should work correctly in utf8 mode, including the regex extension, but you still have the option of working around incompatibilities of specific eggs by loading them before the (import utf8) directive. Keep in mind that some operations, such as string-length, are much more expensive in utf8 (multibyte) mode, and should be used with care. See the utf8 egg documentation for details.

38.2 Platform specific

38.2.1 How do I generate a DLL under MS Windows (tm) ?

Use `csc` in combination with the `-dll` option:

```
C:\> csc foo.scm -dll
```

38.2.2 How do I generate a GUI application under Windows(tm)?

Invoke `csc` with the `-windows` option. Or pass the `-DC_WINDOWS_GUI` option to the C compiler and link with the GUI version of the runtime system (that's `libchicken-gui[-static].lib`. The GUI runtime displays error messages in a message box and does some rudimentary command-line parsing.

38.2.3 Compiling very large files under Windows with the Microsoft C compiler fails with a message indicating insufficient heap space.

It seems that the Microsoft C compiler can only handle files up to a certain size, and it doesn't utilize virtual memory as well as the GNU C compiler, for example. Try closing running applications. If that fails, try to break up the Scheme code into several library units.

38.2.4 When I run `csi` inside an emacs buffer under Windows, nothing happens.

Invoke `csi` with the `-:c` runtime option. Under Windows the interpreter thinks it is not running under control of a terminal and doesn't print the prompt and does not flush the output stream properly.

38.2.5 I load compiled code dynamically in a Windows GUI application and it crashes.

Code compiled into a DLL to be loaded dynamically must be linked with the same runtime system as the loading application. That means that all dynamically loaded entities (including extensions built and installed with `chicken-setup`) must be compiled with the `-windows csc` option.

38.2.6 On Windows, `csc.exe` seems to be doing something wrong.

The Windows development tools include a C# compiler with the same name. Either invoke `csc.exe` with a full pathname, or put the directory where you installed CHICKEN in front of the MS development tool path in the `PATH` environment variable.

38.2.7 On Windows source and/or output filenames with embedded whitespace are not found.

There is no current workaround. Do not use filenames with embedded whitespace for code. However, command names with embedded whitespace will work correctly.

38.3 Customization

38.3.1 How do I run custom startup code before the runtime-system is invoked?

When you invoke the C compiler for your translated Scheme source program, add the C compiler option `-DC_EMBEDDED`, or pass `-embedded` to the `csc` driver program, so no entry-point function will be generated (`main()`). When your are finished with your startup processing, invoke:

```
CHICKEN_main(argc, argv, C_toplevel);
```

where `C_toplevel` is the entry-point into the compiled Scheme code. You should add the following declarations at the head of your code:

```
#include "chicken.h"
extern void C_toplevel(C_word,C_word,C_word) C_noret;
```

38.3.2 How can I add compiled user passes?

To add a compiled user pass instead of an interpreted one, create a library unit and recompile the main unit of the compiler (in the file `chicken.scm`) with an additional `uses` declaration. Then link all compiler modules and your (compiled) extension to create a new version of the compiler, like this (assuming all sources are in the current directory):

```
% cat userpass.scm
;;;; userpass.scm - My very own compiler pass

(declare (unit userpass))
```

```
;; Perhaps more user passes/extensions are added:
(let ([old (user-pass)])
(user-pass
(lambda (x)
(let ([x2 (do-something-with x)])
   if old (
              (old x2)
              x2) ) ) ) )
```

```
% csc -c -x userpass.scm
% csc chicken.scm -c -o chicken-extended.o -uses userpass
% gcc chicken-extended.o support.o easyffi.o compiler.o optimizer.o batch-drive
c-backend.o userpass.o `csc -ldflags -libs` -o chicken-extended
```

On platforms that support it (Linux ELF, Solaris, Windows + VC++), compiled code can be loaded via `-extend` just like source files (see `load` in the User's Manual).

38.4 Compiled macros

38.4.1 Why is `define-macro` complaining about unbound variables?

Macro bodies that are defined and used in a compiled source-file are evaluated during compilation and so have no access to anything created with `define`. Use `define-for-syntax` instead.

38.4.2 Why isn't `load` properly loading my library of macros?

During compile-time, macros are only available in the source file in which they are defined. Files included via `include` are considered part of the containing file.

38.4.3 Why is `include` unable to load my hygienic macros?

It is not sufficient for the included file to require the `syntax-case` extension. Call `(require-extension syntax-case)` *before* calling `include`.

38.4.4 Why are macros not visible outside of the compilation unit in which they are defined?

Macros are defined during compile time, so when a file has been compiled, the definitions are gone. An exception to this rule are macros defined with `define-macro`, which are also visible at run-time, i.e. in `eval`. To use macros defined in other files, use the `include` special form.

38.5 Warnings and errors

38.5.1 Why does my program crash when I use callback functions (from Scheme to C and back to Scheme again)?

There are two reasons why code involving callbacks can crash out of no apparent reason:

1. It is important to use `foreign-safe-lambda/foreign-safe-lambda*` for the C code that is to call back into Scheme. If this is not done than sooner or later the available stack space will be exhausted.
2. If the C code uses a large amount of stack storage, or if Scheme-to-C-to-Scheme calls are nested deeply, then the available nursery space on the stack will run low. To avoid this it might be advisable to run the compiled code with a larger nursery setting, i.e. run the code with `-:s...` and a larger value than the default (for example `-:s300k`), or use the `-nursery` compiler option. Note that this can decrease runtime performance on some platforms.

38.5.2 Why does the linker complain about a missing function `_C_..._toplevel`?

This message indicates that your program uses a library-unit, but that the object-file or library was not supplied to the linker. If you have the unit `foo`, which is contained in `foo.o` than you have to supply it to the linker like this (assuming a GCC environment):

```
% csc program.scm foo.o -o program
```

38.5.3 Why does the linker complain about a missing function `_C_toplevel`?

This means you have compiled a library unit as an application. When a unit-declaration (as in `(declare (unit ...)))`) is given, then this file has a specially named toplevel entry procedure. Just remove the declaration, or compile this file to an object-module and link it to your application code.

38.5.4 Why does my program crash when I compile a file with `-unsafe` or unsafe declarations?

The compiler option `-unsafe` or the declaration `(declare (unsafe))` disable certain safety-checks to improve performance, so code that would normally trigger an error will work unexpectedly or even crash the running application. It is advisable to develop and debug a program in safe mode (without unsafe declarations) and use this feature only if the application works properly.

38.5.5 Why do I get a warning when I define a global variable named `match`?

Even when the `match` unit is not used, the macros from that package are visible in the compiler. The reason for this is that macros can not be accessed from library units (only when explicitly evaluated in running code). To speed up macro-expansion time, the compiler and the interpreter both already provide the compiled `match-...` macro definitions. Macros shadowed lexically are no problem, but global definitions of variables named identically to (global) macros are useless - the macro definition shadows the global variable.

This problem can be solved using a different name or undefining the macro, like this:

```
(eval-when (compile eval) (undefine-macro! 'match))
```

38.5.6 Why don't toplevel-continuations captured in interpreted code work?

Consider the following piece of code:

```
(define k (call-with-current-continuation (lambda (k) k)))
(k k)
```

When compiled, this will loop endlessly. But when interpreted, `(k k)` will return to the read-eval-print loop! This happens because the continuation captured will eventually read the next toplevel expression from the standard-input (or an input-file if loading from a file). At the moment `k` was defined, the next expression was `(k k)`. But when `k` is invoked, the next expression will be whatever follows after `(k k)`. In other words, invoking a captured continuation will not rewind the file-position of the input source. A solution is to wrap the whole code into a `(begin ...)` expression, so all toplevel expressions will be loaded together.

38.5.7 Why does `define-reader-ctor` not work in my compiled program?

The following piece of code does not work as expected:

```
(eval-when (compile)
(define-reader-ctor 'integer->char integer->char) )
(print #,(integer->char 33))
```

The problem is that the compiler reads the complete source-file before doing any processing on it, so the sharp-comma form is encountered before the reader-ctor is defined. A possible solution is to include the file containing the sharp-comma form, like this:

```
(eval-when (compile)
(define-reader-ctor 'integer->char integer->char) )

(include "other-file")
```

```
;;; other-file.scm:
(print #,(integer->char 33))
```

38.5.8 Why do built-in units, such as srfi-1, srfi-18, and posix fail to load?

When you try to `use` a built-in unit such as `srfi-18`, you may get the following error:

```
#;1> (use srfi-18)
; loading library srfi-18 ...
Error: (load-library) unable to load library
srfi-18
"dlopen(libchicken.dylib, 9): image not found"                    ;; on a Mac
"libchicken.so: cannot open shared object file: No such file or directory"  ;;
```

Another symptom is that `(require 'srfi-18)` will silently fail.

This typically happens because the Chicken libraries have been installed in a non-standard location, such as your home directory. The workaround is to explicitly tell the dynamic linker where to look for your libraries:

```
export DYLD_LIBRARY_PATH=~/scheme/chicken/lib:$DYLD_LIBRARY_PATH ;; Mac
export LD_LIBRARY_PATH=~/scheme/chicken/lib:$LD_LIBRARY_PATH    ;; Linux
```

38.5.9 How can I increase the size of the trace shown when runtime errors are detected?

When a runtime error is detected, Chicken will print the last entries from the trace of functions called (unless your executable was compiled with the `-no-trace` option. By default, only 16 entries will be shown. To increase this number pass the `-:aN` parameter to your executable.

38.6 Optimizations

38.6.1 How can I obtain smaller executables?

If you don't need `eval` or the stuff in the `extras` library unit, you can just use the `library` unit:

```
        (declare (uses library))
        (display "Hello world!\n")
```

(Don't forget to compile with the `-explicit-use` option) Compiled with Visual C++ this generates an executable of around 240 kilobytes. It is theoretically possible to compile something without the library, but a program would have to implement quite a lot of support code on its own.

38.6.2 How can I obtain faster executables?

There are a number of declaration specifiers that should be used to speed up compiled files: declaring `(standard-bindings)` is mandatory, since this enables most optimizations. Even if some standard procedures should be redefined, you can list untouched bindings in the declaration. Declaring `(extended-bindings)` lets the compiler choose faster versions of certain internal library functions. This might give another speedup. You can also use the the `usual-integrations` declaration, which is identical to declaring `standard-bindings` and `extended-bindings` (note that `usual-integrations` is set by default). Declaring `(block)` tells the compiler that global procedures are not changed outside the current compilation unit, this gives the compiler some more opportunities for optimization. If no floating point arithmetic is required, then declaring `(number-type fixnum)` can give a big performance improvement, because the compiler can now inline most arithmetic operations. Declaring `(unsafe)` will switch off most safety checks. If threads are not used, you can declare `(disable-interrupts)`. You should always use maximum optimizations settings for your C compiler. Good GCC compiler options on Pentium (and compatible) hardware are: `-Os -fomit-frame-pointer -fno-strict-aliasing` Some programs are very sensitive to the setting of the nursery (the first heap-generation). You should experiment with different nursery settings (either by compiling with the `-nursery` option or by using the `-:s...` runtime option).

38.6.3 Which non-standard procedures are treated specially when the `extended-bindings` or `usual-integrations` declaration or compiler option is used?

The following standard bindings are handled specially, depending on optimization options and compiler settings:

```
+ * - / quotient eq? eqv? equal? apply c...r values call-with-values
list-ref null? length not char? string? symbol? vector? pair? procedure?
boolean? number? complex? rational? real? exact? inexact? list? eof-object?
string-ref string-set! vector-ref vector-set! char=? char<? char>? char<=? char
char-numeric? char-alphabetic? char-whitespace? char-upper-case?
char-lower-case? char-upcae char-downcase list-tail assv memv memq assoc
member set-car! set-cdr! abs exp sin cos tan log asin acos atan sqrt
zero? positive? negative? vector-length string-length char->integer
integer->char inexact->exact = > < >= <= for-each map substring
string-append gcd lcm list exact->inexact string->number number->string
even? odd? remainder floor ceiling truncate round cons vector string
string=? string-ci=? make-vector call-with-current-continuation
write-char read-string
```

The following extended bindings are handled specially:

```
bitwise-and bitwise-ior bitwise-xor bitwise-not bit-set? add1 sub1 fx+ fx- fx*
fx/ fxmod fx= fx> fx>= fixnum? fxneg fxmax fxmin fxand fxior fxxor fxnot fxshl
fxshr flonum? fp+ fp- fp* fp/ atom? fp= fp> fp>= fpneg fpmax fpmin
arithmetic-shift signum flush-output thread-specific thread-specific-set!
not-pair? null-list? print print* u8vector->blob/shared
s8vector->blob/shared u16vector->blob/shared s16vector->blob/shared
u32vector->blob/shared s32vector->blob/shared f32vector->blob/shared
```

```
f64vector->blob/shared block-ref blob-size u8vector-length s8vector-length
u16vector-length s16vector-length u32vector-length s32vector-length
f32vector-length f64vector-length u8vector-ref s8vector-ref u16vector-ref
s16vector-ref u32vector-ref s32vector-ref f32vector-ref f64vector-ref
u8vector-set! s8vector-set! u16vector-set! s16vector-set! u32vector-set!
s32vector-set! hash-table-ref block-set! number-of-slots first second third
fourth null-pointer? pointer->object make-record-instance locative-ref
locative-set! locative? locative->object identity cpu-time error call/cc any?
substring=? substring-ci=? substring-index substring-index-ci
```

38.6.4 Can I load compiled code at runtime?

Yes. You can load compiled at code at runtime with `load` just as well as you can load Scheme source code. Compiled code will, of course, run faster.

To do this, pass to `load` a path for a shared object. Use a form such as `(load "foo.so")` and run `csc -shared foo.scm` to produce `foo.so` from `foo.scm` (at which point `foo.scm` will no longer be required).

38.7 Garbage collection

38.7.1 Why does a loop that doesn't `cons` still trigger garbage collections?

Under CHICKENs implementation policy, tail recursion is achieved simply by avoiding to return from a function call. Since the programs are CPS converted, a continuous sequence of nested procedure calls is performed. At some stage the stack-space has to run out and the current procedure and its parameters (including the current continuation) are stored somewhere in the runtime system. Now a minor garbage collection occurs and rescues all live data from the stack (the first heap generation) and moves it into the the second heap generation. Then the stack is cleared (using a `longjmp`) and execution can continue from the saved state. With this method arbitrary recursion (in tail- or non-tail position) can happen, provided the application doesn't run out of heap-space. (The difference between a tail- and a non-tail call is that the tail-call has no live data after it invokes its continuation - and so the amount of heap-space needed stays constant)

38.7.2 Why do finalizers not seem to work in simple cases in the interpeter?

Consider the following interaction in CSI:

```
#;1> (define x '(1 2 3))
#;2> (define (yammer x) (print x " is dead"))
#;3> (set-finalizer! x yammer)
(1 2 3)
```

```
#;4> (gc #t)
157812
#;5> (define x #f)
#;6> (gc #t)
157812
#;7>
```

While you might expect objects to be reclaimed and "*(1 2 3) is dead*" printed, it won't happen: the literal list gets held in the interpreter history, because it is the result value of the set-finalizer! call. Running this in a normal program will work fine.

When testing finalizers from the interpreter, you might want to define a trivial macro such as

```
(define-macro (v x) `(begin (print ,x) (void)))
```

and wrap calls to `set-finalizer!` in it.

38.8 Interpreter

38.8.1 Does CSI support history and autocompletion?

CSI doesn't support it natively but it can be activated with the http://www.call-with-current-continuation.org/eggs/readline.html egg. After installing the egg, add the following to your `~/.csirc` or equivalent file:

```
(require-extension readline)
(current-input-port (make-gnu-readline-port))
(gnu-history-install-file-manager (string-append (or (getenv "HOME") ".") "/.cs
```

Users of *nix-like systems (including Cygwin), may also want to check out rlwrap. This program lets you "wrap" another process (e.g. `rlwrap csi`) with the readline library, giving you history, autocompletion, and the ability to set the keystroke set. Vi fans can get vi keystrokes by adding "set editing-mode vi" to their `.inputrc` file.

38.8.2 Does code loaded with `load` run compiled or interpreted?

If you compile a file with a call to `load`, the code will be loaded at runtime and, if the file loaded is a Scheme source code file (instead of a shared object), it will be interpreted (even if the caller program is compiled).

38.9 Extensions

38.9.1 How can I install Chicken eggs to a non-default location?

You can just set the CHICKEN_REPOSITORY environment variable. It should contain the path where you want eggs to be installed:

```
$ export CHICKEN_REPOSITORY=~/chicken/
$ chicken-setup extensionname
```

In order to make programs (including csi) see these eggs, you should set this variable when you run them. Alternatively, you can call the repository-path Scheme procedure before loading the eggs, as in:

```
(repository-path "/home/azul/chicken")
(use format-modular)
```

Note, however, that using repository-path as above hard-codes the location of your eggs in your source files. While this might not be an issue in your case, it might be safe to keep this configuration outside of the source code (that is, specifying it as an environment variable) to make it easier to maintain.

38.9.2 Can I install chicken eggs as a non-root user?

Yes, just install them in a directory you can write.

Previous: Bugs and limitations

Next: Acknowledgements

39 Acknowledgements

Eli Barzilay
> some performance tweaks used in TinyCLOS.

Mikael Djurfeldt
> topological sort used by compiler.

Marc Feeley
> pretty-printer.

Aubrey Jaffer
> implementation of `dynamic-wind`.

Richard O'Keefe
> sorting routines.

Olin Shivers
> implementation of `let-optionals[*]` and reference implementations of SRFI-1, SRFI-13 and SRFI-14.

Andrew Wilcox

queues.
Andrew Wright
        pattern matcher.
http://galinha.ucpel.tche.br/Alex Shinn
        `scheme-complete.el` emacs tab-completion

40 Bibliography

Henry Baker: *CONS Should Not CONS Its Arguments, Part II: Cheney on the M.T.A.*
http://home.pipeline.com/~hbaker1/CheneyMTA.html

Revised^5 Report on the Algorithmic Language Scheme
http://www.schemers.org/Documents/Standards/R5RS

Previous: Acknowledgements