

Chez Scheme Version 7.0 Release Notes

Copyright © 2005 Cadence Research Systems

All Rights Reserved

August 2005

1. Overview

This document outlines the changes made to *Chez Scheme* for Version 7.0 since Version 6.0.

Version 7.0 is available for the following platforms:

- Intel 80x86 Linux
- Intel 80x86 Linux, threaded
- Intel 80x86 Windows 95/98/ME/NT/2000/XP
- Intel 80x86 Windows 95/98/ME/NT/2000/XP, threaded
- Apple PowerPC Mac OS X
- Sun Sparc Solaris, 32-bit
- Sun Sparc Solaris, 32-bit threaded
- Sun Sparc Solaris, 64-bit
- Sun Sparc Solaris, 64-bit threaded

This document contains three sections describing significant (1) functionality changes, (2) bugs fixed, and (3) performance enhancements. A version number listed in parentheses in the header for a change indicates the first minor release or internal prerelease to support the change.

More information on *Chez Scheme* and *Petite Chez Scheme* can be found at <http://www.scheme.com>, and extensive documentation is available in *The Scheme Programming Language, 3rd edition* and the *Chez Scheme Version 7 User's Guide*.

2. Functionality Changes

2.1. Left-to-right evaluation of internal definitions (6.9d)

Internal variable definitions are now evaluated from left to right by default, using `letrec*` semantics.

The parameter `internal-defines-as-letrec*` may be set to `#f` to override this change.

2.2. Change in import semantics (6.9d)

An identifier *x* imported from a module *m* is treated in the importing context as if the corresponding export identifier had been present in the import form along with *m*.

With either the old or new semantics, the following `from` module works as expected.

```
(define-syntax from
  (syntax-rules ()
    [(- m v) (let () (import m) v)]))
```

```
(module a (x) (define x 'x-of-a))
(from a x) ⇒ x-of-a
```

With the new semantics, the following `x-from-a` macro also works.

```
(define-syntax x-from-a
  (syntax-rules ()
    [(-) (let () (import a) x)]))
(x-from-a) ⇒ x-of-a
```

With the old semantics, `(x-from-a)` would result in an unbound error for variable `x`.

The new semantics properly prevents imports from introduced module names from capturing named variables.

```
(let ([x 'local-x])
  (define-syntax alpha
    (syntax-rules ()
      [(- var) (let () (import a) (list x var))]))
  (alpha x)) ⇒ (x-of-a local-x)
```

It similarly prevents imports from named modules from capturing references to introduced variables.

```
(let ([x 'local-x])
  (define-syntax beta
    (syntax-rules ()
      [(- m var) (let () (import m) (list x var))]))
  (beta a x)) ⇒ (local-x x-of-a)
```

This new semantics extends to prefixed, renamed, and aliased bindings created by the new import syntax added in Version 6.9c.

2.3. Pretty-printer improvements (6.9d)

The pretty printer, used implicitly by the read-eval-print loop and explicitly via `pretty-print`, has been improved in several ways.

First, the printing of long lists, list-structured forms, and vectors now uses a more compact “fill mode” in which the elements of the list or vector are printed with as many elements as will fit in a line.

Second, a new mechanism (described in the *Chez Scheme Version 7 User's Guide*) has been added for advising the pretty printer how to print specific syntactic forms and procedure calls.

Third, the new mechanism has been used to improve how the built-in forms are printed, including those with multiple syntactic variants.

2.4. Macro helpers (6.9d)

Two new syntactic forms and a new procedure have been added to simplify the definition of syntactic abstractions. The new syntax

```
(datum template)
```

is equivalent to the more verbose

```
(syntax-object->datum (syntax template))
```

The new syntax

```
(with-implicit (id0 id1 ...) e1 e2 ...)
```

is equivalent to

```
(with-syntax ([id1 (datum->syntax-object #'id0 'id1)] ...) e1 e2 ...)
```

The new procedure `syntax->list` takes a syntax object representing a list-structured form and returns a list of syntax-objects, each representing the corresponding subform of the input form. Programmers are encouraged to use this procedure even when the current *Chez Scheme* implementation of `syntax-case` guarantees that the output of a `syntax` form is a list, since future versions of *Chez Scheme* may remove these guarantees in the interest of maintaining better source information.

2.5. `eqv?` and `equal?` change (6.9d)

The procedures `eqv?` and `equal?` now return `#t` for two inexact numbers if and only if they have precisely the same internal representation. In particular, the IEEE floating-point values `+0.0` and `-0.0`, which are considered numerically equal, are not considered equivalent by `eqv?` and `equal?`, since the two numbers behave differently in many circumstances. For example, `(/ 1.0 0.0)` evaluates to `+inf.0`, whereas `(/ 1.0 -0.0)` evaluates to `-inf.0`.

2.6. “Pretty” gensym printing option (6.9d)

If the parameter `print-gensym` is set to `pretty` rather than `#f` or `#t`, the gensym’s “pretty” name is printed with the old `#:` prefix. The reader recognizes this syntax as well, and constructs a new gensym with the given pretty name.

2.7. `expand/optimize` procedure (6.9d)

The new procedure `expand/optimize` accepts an object representing a Scheme form and an optional environment, like `expand`, and returns an object representing the result of expanding the form and passing it through the source optimizer used by both the compiler and interpreter. With this procedure, a programmer can get a rough idea what the compiler will do with a program. Unfortunately, however, the output does not reflect the optimizations performed by later passes of the compiler.

2.8. New logical operations (6.9d)

The new procedures `logbit?` and `fxlogbit?` test whether a specified bit is set in an integer. Both accept a nonnegative zero-based bit position and an exact integer. `fxlogbit?` is limited to fixnum arguments, while `logbit?` accepts both fixnum and bignum arguments.

The new procedures `logtest` and `fxlogtest` test whether any bits are set in common between two exact integer (fixnum or bignum) arguments. `fxlogtest` is limited to fixnum arguments, while `logtest` accepts both fixnum and bignum arguments.

The new procedures `logbit0` and `logbit1` each take two arguments: a nonnegative zero-based bit position *i* and an exact integer *n*. `logbit0` returns a new number formed by clearing bit *i* of *n*, and `logbit1` returns a new number formed by setting bit *i* of *n*.

The new procedures `logior` and `fxlogior` are equivalent to the old procedures `logor` and `fxlogor`. The old names are still defined as well.

2.9. Extended logical operations (6.9d)

The procedures `logand`, `logor`, `logxor`, `fxlogand`, `fxlogor`, and `fxlogxor` have been extended to accept zero or more arguments.

2.10. RSA example (6.9d)

A “new” example program, `rsa.ss`, now appears in the examples directory. It is a updated version of a example program that was removed several years back due to concerns over a patent that has now expired.

2.11. `apropos` (6.9d)

The recently added (see below) `apropos` and `apropos-list` procedures have been extended to take an optional environment argument. The environment defaults to the interaction environment if not specified. The related procedure `environment-symbols` has been improved to weed out some “false positives,” i.e., symbols representing variables that have been referenced in some code somewhere (perhaps by mistake) but not yet given values.

2.12. C-side thread activation/deactivation (6.9d)

Three new C entry points have been added to the C library interface for threaded versions of *Chez Scheme*: `Sactivate_thread`, `Sdeactivate_thread`, and `Sdestroy_thread`. `Sactivate_thread` must be called from a thread created by some other mechanism than Scheme’s `fork-thread` before the thread touches any Scheme data or runs any Scheme code. `Sdestroy_thread` should be called before a thread created by some other mechanism than `fork-thread` before it is destroyed. `Sdeactivate_thread` may be called to deactivate a thread when the thread might not be running Scheme code for an extended period, e.g., when the thread may block waiting for input from the operating system. Once deactivated, the thread should not touch any Scheme data or run any Scheme code until reactivated by calling `Sactivate_thread`.

2.13. `list-head` (6.9d)

A new procedure, `list-head`, has been added. It takes a list `ls` and an index `n` and returns a new list containing the first `n` elements of `ls`, in order.

2.14. Interactive inspector and multiple values (6.9c)

The interactive inspector `eval` and `=>` commands now now handle expressions that produce multiple values.

2.15. `trace-define-syntax` and `environment` macros (6.9c)

`trace-define-syntax` now handles macros that return a procedure to obtain the current compile-time environment, i.e., it now now displays the final output, not the returned procedure.

2.16. `trace-case-lambda` (6.9c)

A new `trace-case-lambda` form is now available for tracing anonymous variable-arity procedures. It is used by replacing the usual `case-lambda` keyword in a `case-lambda` expression with `trace-case-lambda` followed by a name. It is similar to `trace-lambda` for `lambda`.

2.17. Meta definitions (6.9c)

A new `meta` definition prefix is now supported by the expander. This prefix tells the expander that any variable definition resulting from the definition is to be an expand-time definition available only on the right-hand sides of other meta definitions and, most importantly, to transformer expressions. It is used to define expand-time helpers and other information for use by one or more `syntax-case` transformers.

```
(module M (helper1 a b)
  (meta define helper1
    (lambda (---)
      ---))
  (meta define helper2
    (lambda (---)
      --- (helper2 ---) ---))
  (define-syntax a
    (lambda (x)
      --- (helper1 ---) ---))
  (define-syntax b
    (lambda (x)
      --- (helper1 ---) ---
      --- (helper2 ---) ---)))
```

As with `define-syntax` rhs expressions, meta expressions can evaluate references only to identifiers whose values are (already) available in the compile-time environment, e.g., macros and meta variables. They can, however, like `define-syntax` rhs expressions, build syntax objects containing occurrences of any identifiers in their scope.

Meta definitions propagate through macro expansion, so one can write, for example:

```
(module (a)
  (meta define-record foo (x))
  (define-syntax a
    (let ([q (make-foo (syntax 'q))])
      (lambda (x)
        (foo-x q))))))
a ⇒ q
```

where `define-record` is a macro that expands into a set of defines.

It is also sometimes convenient to write

```
(meta begin defn ...)
```

or

```
(meta module {exports} defn ...)
```

to create groups of meta bindings.

2.18. Identifier aliases (6.9c)

The new definition form `alias` may be used to create aliases from one identifier to another.

```
(let ((x 3)) (alias y x) y) ⇒ 3
```

Aliases appearing at top level track the current binding for the aliased identifier.

```
(alias foo lambda)
(define lambda 3)
foo ⇒ 3
```

2.19. New import syntax (6.9c)

The syntax for `import` has been extended to allow for selection and renaming of imports. An `import` form now has the following syntax:

```
(import spec ...)
```

Where *spec* is defined by the following grammar:

```
spec  →  id
      |  (only spec id*)
      |  (except spec id*)
      |  (add-prefix spec id)
      |  (drop-prefix spec id)
      |  (rename spec (new-id old-id)*)
      |  (alias spec (new-id old-id)*)
```

`only` imports only the listed identifiers; `except` imports all but the listed identifiers; `add-prefix` adds the given prefix to the imports; `drop-prefix` drops the given prefix to the imports, which must all share the given prefix; `rename` provides new names for the listed imports; and `alias` provides additional names for the listed imports.

2.20. apropos (6.9c)

Two new procedures have been added for obtaining or displaying subsets of the set of “interesting” symbols known to the system. An interesting symbol is one that has some sort of top-level binding. The first, `apropos-list`, returns a list of interesting symbols. The second, `apropos`, displays the interesting symbols, one per line. Each takes a single argument, a symbol or string, and returns or displays only those interesting symbols whose names contain the argument symbol’s name or the string as a substring.

```
(apropos-list 'apropos) ⇒ (apropos apropos-list)
```

```
(apropos 'car) ⇒ prints car and set-car!
```

2.21. putenv (6.9c)

A new `putenv` primitive has been added. `putenv` takes two string arguments, a key and value, and stores the key, value pair in the environment of the process, where it is available to the current process and any spawned processes. The key and value are copied into storage allocated outside of the Scheme heap; this space is never reclaimed.

2.22. Hash tables (6.9c)

Support for hash tables has been added. Hash tables allow programs to associate values with arbitrary Scheme objects. They serve essentially the same purpose as association lists but are much faster when dealing with large numbers of objects. The new primitives supporting hash tables are listed below.

`(make-hash-table)` or `(make-hash-table weak?)` creates a hash table.

If *weak?* is provided and is non-false, the hash table is a weak hash table, which means that it does not protect keys from the garbage collector. Keys reclaimed by the garbage collector are removed from the table,

along with their associated values.

`(hash-table? x)` returns `#t` if x is a hash table, otherwise `#f`.

`(put-hash-table! ht k v)` associates the value v with the key k in ht .

`(get-hash-table ht k d)` returns the value associated with k in ht or d if no value is associated with k . Key comparisons are performed with `eq?`.

`(remove-hash-table! ht k)` drops any association for k from ht .

`(hash-table-map ht p)` applies p to each key, value association in ht and returns a list of the resulting values, in no particular order. p should accept two arguments, a key and a value.

`(hash-table-for-each ht p)` applies p to each key, value association in ht . Unlike `hash-table-map`, it does not create a list of the values; instead, it's value is unspecified. p should accept two arguments, a key and a value.

2.23. Critical dynamic-wind (6.9c)

`dynamic-wind` now accepts an optional first argument:

`(dynamic-wind in body out)`

`(dynamic-wind critical? in body out)`

When present and non-`false`, the *in* thunk is invoked in a critical section along with the code that records that the body has been entered, and the *out* thunk is invoked in a critical section section along with the code that records that the body has been exited. Extreme caution must be taken with this form of `dynamic-wind`, since an error or long-running computation can leave interrupts and automatic garbage collection disabled.

2.24. Heap and boot-file changes (6.9c)

Heap- and boot-file searching. *Chez Scheme* now looks for boot files if it cannot find a compatible heap file for heap and boot files that are not specified on the command line. In a standard installation, this process is transparent to the user, but users who wish to create custom applications can now use boot files in place of heap files without resorting to shell scripts.

When *Chez Scheme* is run, it looks for one or more boot files and/or one or more heap files to load. Boot files contain the compiled Scheme code that implements most of the Scheme system, including the interpreter, compiler, and most libraries. Heap files contain prebuilt heap images, i.e., saved heap images into which the boot code has already been loaded. Heap and boot files may be specified explicitly on the command line via `-b` and `-h` options or implicitly. In the simplest case, no `-b` and `-h` options are given and the necessary heap or boot files are loaded automatically based on the name of the executable. For example, if the executable name is `myapp`, the system looks for `myapp.heap` in a set of standard directories, then for `myapp.boot`. It also looks for and loads any subordinate heaps or boot files required by `myapp.heap` or `myapp.boot`. Subordinate heap and boot files are also loaded automatically for the first boot file and any heap files explicitly specified via the command line. When boot and heap files are specified via the command line, all heap files must come before all boot files, and each file must be listed before those that depend upon it.

The `-verbose` command-line option may be used to trace the heap and boot file searching process and must appear before any boot or heap arguments for which search tracing is desired.

Ordinarily, the search for heap and boot files is limited to a set of default installation directories, but this may be overridden by setting the environment variable `SCHEMEHEAPDIRS`. `SCHEMEHEAPDIRS` should be a colon-separated list of directories, listed in the order in which they should be searched. Within each directory, the two-character escape sequence `%v` is replaced by the current version, and the two-character escape sequence `%m` is replaced by the machine type. A percent followed by any other character is replaced by the second character; in particular, `%%` is replaced by `%`, and `%.:` is replaced by `..`.

If SCHEMEHEAPDIRS ends in a non-escaped colon, the default directories are searched after those in SCHEMEHEAPDIRS; otherwise, only those listed in SCHEMEHEAPDIRS are searched. Under Windows, semi-colons are used in place of colons.

Creating boot files. Boot files consist of ordinary compiled code and may be created by concatenating a heap header onto the compiled code for one or more source files, e.g.:

```
% cat myapp.header myapp1.so myapp2.so > myapp.boot
```

Boot file header files are created with the Scheme procedure `make-boot-header`, which takes an output file name and the names of one or more alternatives for the boot file upon which the new boot file immediately depends, e.g.:

```
% echo "(make-boot-header "myapp.header" 'scheme 'petite)" | scheme
```

In ordinary usage, one or both of “scheme“ and “petite“ should be specified, as shown above, when creating the base boot file for an application. If the application requires the use of the compiler, just “scheme“ should be specified.

Creating heap files. Heap files are created via the “-s” command-line option, which now behaves slightly differently from before.

If the “-s” option is present and *Chez Scheme* exits normally, it will save the heap on the specified file.

The *heap level* is determined by the *level* argument, if any, which follows “-s”, if any. A level zero heap contains the entire heap, a level one heap contains only those portions that differ from the corresponding level zero heap, and so on. A level one heap may be created in a session in which a level zero heap has been loaded and subsequently be loaded with the “-h” option.

```
scheme -h ${HOME}/.heap -s1 ${HOME}/.heap.1 $*
scheme -h ${HOME}/.heap -h ${HOME}/.heap.1 $*
```

If no level follows the “-s” option, the level defaults to the level of the highest level heap loaded, or zero if no heaps have been loaded. If the special level “+” follows the “-s” option, the level is one level higher than the highest loaded heap, or zero if no heaps have been loaded. Thus, the following three lines create level zero, one, one, and two heaps.

```
SCHEMEHEAPDIRS=".: "
echo | scheme -b petite.boot -s heap.0
echo | scheme -h heap.0 -s+ heap.1
echo | scheme -h heap.1 -s heap.1
echo | scheme -h heap.1 -s+ heap.2
```

2.25. Windows library change (6.9c)

Windows builds of *Chez Scheme* now link against `msvcrt.dll`.

2.26. Windows registry primitives (6.9c)

Three new primitives for accessing the Windows registry have been added.

`(get-registry s)` returns a string containing the registry value of *s*, which must be a string. If no registry value for *s* exists, `get-registry` returns `#f`.

`(put-registry! s v)` sets the registry value of the string *s* to the string *v*. It signals an error if the value cannot be set, which may happen if the user has insufficient access.

(`remove-registry!` *s*) removes the registry key or value named by the string *s*. It returns an unspecified value. It signals an error if the value cannot be removed. Reasons for failure include the key not being present, the user having insufficient access, or *s* being a key with subkeys.

2.27. New configure script (6.9c)

Installation via the `custom` subdirectory of the distribution directory has been simplified by the addition of a `configure` script that sets the machine type and allows various installation parameters to be set.

2.28. `internal-defines-as-letrec*` (6.9c)

The `internal-defines-as-letrec*` parameter (see below) is now respected for definitions appearing within top-level modules.

2.29. Expanded format functionality (6.9b)

`format` and related procedures, including `printf`, `fprintf`, and the default warning and error handlers, have been enhanced to accept all of the Common Lisp format directives, except for those directly related to pretty printing.

Here are some scintillating examples of this over-engineered utility for your entertainment.

- floating-point printing—3 digits to right of decimal point

```
(format "~,3f" 3.14159) ⇒ "3.142"
```

- scientific notation—field of width 12 padded with asterisks, 3 digits to right of decimal point, 3 digits for exponent

```
(format "~12,3,3,,,*e" 1e23) ⇒ "***1.000e+023"
```

- monetary notation—default two digits to right of decimal point.

```
(format "$~$" (* 39.95 1.06)) ⇒ "$42.35"
```

- plurals

```
(format "~s bear~:p in ~s den~:p" 10 1) ⇒ "10 bears in 1 den"
```

- printed numbers

```
(format "~r" 2599) ⇒ "two thousand five hundred ninety-nine"  
(format "~:r" 99) ⇒ "ninety-ninth"
```

- roman numerals

```
(format "~@r" 2599) ⇒ "MMDXCIX"
```

- case conversion

```
(format "~@(~r~)" 2599) ⇒ "Two thousand five hundred ninety-nine"  
(format "~@:(~a~)" "Ouch!") ⇒ "ouch!"
```

There's much, much more, including conditionals, iteration, indirection, and justification.

In addition, `format` can now be used instead of `printf` and `fprintf` to print to the current output port or to another port, ala Common Lisp `format`. Here are the different modes supported by `format`.

```
(format string arg ...)    evaluates to formatted string (as always)
(format #f string arg ...)  evaluates to formatted string
(format #t string arg ...)  prints to current-output port
(format port string arg ...) prints to port
```

2.30. New environment procedures (6.9b)

Support for creating new environments has been added, in the form of a new `copy-environment` procedure. `copy-environment` takes from one to three arguments.

```
(copy-environment env)
(copy-environment env mutable?)
(copy-environment env mutable? syms)
```

`copy-environment` returns a copy of `env`, i.e., a new environment that contains the same bindings as `env`.

The environment is mutable if `mutable?` is omitted or true; if `mutable?` is false, the environment is immutable. A mutable environment can be extended with new bindings, its existing bindings can be modified, and its variables can be assigned. An immutable environment cannot be modified in any of these ways.

The set of bindings copied from `env` to the new environment is determined by `syms`, which defaults to the value of `(environment-symbols env)`. The binding, if any, for each element of `syms` is copied to the new environment, and no other bindings are present in the new environment.

```
(define e (copy-environment (scheme-environment)))
(eval '(define cons +) e)
(eval '(cons 3 4) e)           ⇒ 7
(eval '(cons 3 4) (scheme-environment)) ⇒ (3 . 4)
```

The procedure `environment-symbols` takes an environment and returns a list of the symbols representing variables bound in the environment. It is primarily useful in building the list of symbols to be copied from one environment to another.

```
(define listless-environment
  (copy-environment
    (scheme-environment)
    #t
    (remq 'list (environment-symbols (scheme-environment)))))
(eval '(let ((x (cons 3 4))) x) listless-environment) ⇒ (3 . 4)
(eval '(list 3 4) listless-environment) ⇒ Error: variable list is not bound.
```

2.31. interaction-environment is now a parameter (6.9b)

The procedure `interaction-environment` has been converted into a full parameter which can be altered to change the environment used by default by various evaluation and expansion procedures. This parameter is consulted by `compile`, `interpret`, `compile-file`, `sc-expand`, `top-level-value`, `top-level-bound?`, `define-top-level-value`, and `set-top-level-value!`. Because `eval` and `load` use the current evaluator, i.e., the value of `current-eval`, and the default evaluator for Chez Scheme is `compile` and for Petite Chez Scheme is `interpret`, `eval` and `load` are also affected by the setting of `interaction-environment`. Similarly, `expand` uses the current expander, i.e., the value of `current-expand`, which defaults to `sc-expand`, so `expand` is also affected by `interaction-environment`.

```

> (define e (copy-environment (interaction-environment)))
> (parameterize ((interaction-environment e)) (new-cafe))
>> (define cons +)
>> (cons 3 4)
7
>> (exit)
> (cons 3 4)
(3 . 4)

```

2.32. Extended top-level-value procedures (6.9b)

An optional *environment* argument has been added to the procedures `top-level-bound?`, `top-level-value`, `define-top-level-value`, and `set-top-level-value!`. The environment argument is always last, if present.

```

(define cons +)
(cons 3 4)                ⇒ 7
(top-level-value 'cons)  ⇒ #<procedure +>
(top-level-value 'cons (scheme-environment)) ⇒ #<procedure cons>

```

2.33. New trace-do syntactic form (6.9b)

Loops written using `do` may now be traced via `trace-do`.

```

> (trace-do ([i 4 (- i 1)]) ((= i 0) 'done) (printf "~s\n" i))
|(do 4)
4
|(do 3)
3
|(do 2)
2
|(do 1)
1
|(do 0)
|done
done

```

2.34. No more undefined variable warnings (6.9a)

The compiler no longer generates a warning whenever an attempt to reference or assign an undefined locally defined or `letrec`-bound might occur, since these warnings are sometimes false warnings and quickly get old. Actual attempts to reference undefined variables are still detected and caught at run time. (These warnings were produced only by Version 6.9.)

2.35. New string-port operations (6.9a)

Two new operations on string-ports have been added:

(`with-input-from-string` *string thunk*) binds the current-input port to a string-input port created from *string* during the invocation of `thunk` and returns the value or values returned by *thunk*.

(`with-output-to-string` *thunk*) binds the current-output port to a new string-output port during the invocation of `thunk` and returns the string extracted from the string-output port when *thunk* returns.

2.36. `mkdir` and `chmod` (6.9a)

`mkdir` and `chmod` primitives have been added. (`mkdir pathname mode`) creates a directory with the given `pathname`. `mode` is optional and is used to set the permissions on the directory. (`chmod pathname mode`) changes the permissions on the file or directory at the given `pathname` according to `mode`. Under Windows, the `mode` argument to `mkdir` is ignored, and `chmod` does not do anything.

2.37. Top-level `import-only` (6.9a)

`import-only` is now valid at top level, where it works like `import`.

2.38. Better warnings and error messages (6.9a)

The compiler and run-time system now generate better warning and error messages for argument-count errors.

2.39. New `compile-script` procedure (6.9a)

A new primitive procedure, `compile-script`, has been added. It is like `compile-file` but differs in two important ways: (1) it copies the leading `#!` line from the source-file script into the object file, and (2) it does not compress the resulting file. This permits compiled script files to be created from source script files for improved efficiency.

2.40. Environments (6.9a)

The values returned by `scheme-report-environment`, `null-environment`, and `interaction-environment` are now of type `environment`, a unique type distinct from all others. The predicate `environment?` may be used to identify environments.

2.41. Source annotations in compiled files (6.9a)

Source annotations are no longer kept within syntax objects written to compiled files when generation of inspector information is disabled. This makes the compiled files smaller and eliminates the possibility of confusing source file references in error messages produced by production code.

2.42. Bignum shift counts (6.9a)

`ash` now permits the use of very large (bignum) shift counts, although such large shift counts are of dubious value.

2.43. Incompatible records and compiled files (6.9a)

An error is now signaled when an incompatible record is loaded from a compiled file, just as it is when an incompatible record is created by source code.

2.44. Mac OS X Version (6.9)

Full support for *Chez Scheme* on PowerPC based Apple Computers running Mac OS X has been added.

2.45. Better feedback from the reader (6.9)

The reader now provides better feedback when certain lexical errors are detected during the loading or compiling of source files. A major improvement is that end-of-file errors now report what kind of datum was being read and the file position at which the datum appeared, rather than the (useless) position of the actual end of file.

2.46. Records and generativity (6.9)

Each record type is associated with a unique generated name, or gensym. This gensym is used to identify the record type in the printed and compiled representations of each record of the type. When the name argument to `make-record-type` is a gensym, this gensym is used; otherwise a new gensym is created. If `make-record-type` is invoked more than once with the same gensym and the record fields match, the same record type descriptor is returned. If the record fields do not match, an error is signaled. This permits the programmer to control the generativity of record types; full nongenerativity is achieved by passing a gensym, and various levels of generativity are achieved by passing in a non-gensym at expansion time or run time.

`define-record` is ordinarily expansion-time generative, meaning that each time an expression is expanded, e.g., when it is contained in a file loaded for source or compiled by `compile-file`, a new record type is constructed. If the name of the record is specified as a gensym, however, the record type is fully nongenerative. This allows the programmer to place two identical definitions of a record type in two files to be separately compiled or loaded. Only the pretty name of the gensym is used in forming the variables naming the record constructor, predicate, accessors, and setters. For example:

```
(define-record foo ((immutable a)))
```

and

```
(define-record #{foo |*5CNSgDOG8+wd\\%|} ((immutable a)))
```

each implicitly define `make-foo`, `foo?`, and `foo-a`, but the latter is nongenerative while the former is expansion-time generative.

2.47. `letrec` changes (6.9)

A new form of `letrec`, called `letrec*`, has been added. `letrec*` is like `letrec` except that the right-hand sides are evaluated in order from left to right. The scoping rules are the same: `letrec*` and `letrec` both support full mutual recursion among the bound variables.

A new parameter `internal-defines-as-letrec*`, which defaults to `#f` in this release, causes the internal definitions of a `lambda` or similar body to expand into a `letrec*` form rather than `letrec`, effectively sequentializing their evaluation from left to right. This parameter must be set at expansion time (usually evaluation, loading, or compile time) to have any effect.

The system now signals an error at run time for any attempt to use a `letrec` or `letrec*`-bound variable (including variables defined by internal definitions) before it is valid to do so, i.e., any attempt to evaluate a reference or assignment before entering the body of the `letrec` or `letrec*`. The compiler generates a warning whenever such an error might occur. (This is no longer true; see Section 2.34.)

2.48. New `let-values` syntactic form (6.9)

Support for a new `let-values` form has been added. `let-values` is a convenient way to receive multiple values and bind them to variables. It is structured like `let` but permits an arbitrary formals list (like `lambda`) on each right-hand side. For example:

```
> (let-values ([[a b] (values 1 2)] [c (values 1 2 3)]))
  (list a b c)
(1 2 (1 2 3))
```

2.49. open-input-output-file (6.9)

The new procedure `open-input-output-file` opens the file named by its first argument for reading and writing. It accepts an optional additional argument specifying the set of open options. Currently supported options include `exclusive`, `nonexclusive`, `buffered`, `unbuffered`, and `mode`. Input/output files are usually closed using `close-port` but may also be closed with either `close-input-port` or `close-output-port`.

2.50. New file open options (6.9)

The various file opening operations now accept an `exclusive` option, which establishes a lock on the opened file for as long as it remains open. On some systems the lock is advisory, i.e., it inhibits access by other processes only if they also attempt to open exclusively. A complementary `nonexclusive` option is supported as well, but is never needed since files are opened nonexclusively by default.

Operations that open output files (or input/output files) also support a new `mode` option, which determines the permission bits on Unix systems when the file is created by the operation. The following element in the options list must be an exact integer specifying the permissions in the manner of the Unix `open` function.

The following example demonstrates the use of both new options, along with the existing `truncate` option, to open a file exclusively for writing, truncating the file if it exists, and setting the permissions if the file must be created to allow read and write access by the owner, read access by the group, and no access by others.

```
(define p
  (open-output-file "/tmp/shared"
    '(exclusive truncate mode #o640)))
```

2.51. truncate-file (6.9)

The new procedure `truncate-file` truncates to zero length the file associated with its port argument and repositions the port to the beginning of the truncated file.

2.52. Error handling change (6.9)

It is no longer an error to return from the error handler. Instead, this is now the appropriate way to reset to the current cafe after an error is handled, if a reset is desired. Returning rather than explicitly resetting allows the system to clean up any aborted operations, e.g., to close any files opened by `load` or `compile-file` that are aborted as a result of the reset. The default error handler used to take care of this explicitly; this change allows new error handlers to be written independently of the default error handler that perform this same cleanup.

2.53. Logical operations on exact integers (6.9)

Version 7.0 implements the binary logical operations `logand`, `logor`, and `logxor`, along with unary `lognot`, for exact integers, including arbitrary-precision integers (bignums). Support for the fixnum versions of these operators is still present.

2.54. Generalized syntax-case patterns (6.9)

`syntax-case` patterns have been generalized to allow a fixed number of items after an ellipsis, allowing patterns like `(a b ... c d)` or `(a b ... c . d)`. In essence, the restriction that ellipses appear only at the end of a list or vector pattern has been replaced by a lesser restriction that only one ellipsis appear within a given list or vector pattern.

A pattern variable following an ellipsis and a dot, e.g., `d` in `(a b ... c . d)` matches only non-list items. For example, `(a b)` matched against `(1 2 3)` binds `a` to `(1 2 3)` and `b` to `()`.

Because `syntax-rules` is defined in terms of `syntax-case`, `syntax-rules` patterns have been generalized similarly.

2.55. “Defun” syntax for define-syntax (6.9)

`define-syntax` now supports the so-called “defun” syntax of `define`, e.g.,

```
(define-syntax (when x)
  (syntax-case x ()
    [(_ e0 e1 e2 ...) #'(if e0 (begin e1 e2 ...))]))
```

This syntax is not useful for when `syntax-rules` is used, since the argument to a `syntax-rules` transformer is implicit.

2.56. Heap search change (6.9)

Search for a heap does not terminate when a heap of the expected name is found unless the heap is compatible with the executable. This allows the user to have heaps for multiple builds of the system in the heap search path, without interference.

2.57. Machine-type check (6.8)

The `load` procedure now verifies that compiled object files loaded with `load` were produced for the current machine type and signals an error if there is a machine-type mismatch.

2.58. Support for scripting (6.8)

Several changes have been made to support the use of *Chez Scheme* and *Petite Chez Scheme* for shell scripts:

- A new “`--script scriptfile`” command-line argument tells *Chez Scheme* to run in script mode, with *scriptfile* as the script-file name. In script mode, the interactive *Chez Scheme* greeting and prompts are suppressed, i.e., `--quiet` is implied, and the procedure value of the parameter `scheme-script` is applied to the script-file name and the other command-line arguments not otherwise interpreted by *Chez Scheme*. These include earlier filename arguments and all remaining command-line arguments, as if `--` had been supplied after *scriptfile*.
- The default value of the parameter `scheme-script` sets the parameter *command-line-arguments* to the additional command-line arguments, then loads the script file.
- `load` ignores the first line of any file that begins with `#!` followed by a space or slash.

Taken together, these allow Scheme-coded shell scripts to be written, such as the following implementation of the traditional Unix `echo` command, and run in shells that support such shell scripts.

```

#! /usr/bin/scheme --script
(let ([args (command-line-arguments)])
  (unless (null? args)
    (let-values ([(newline? args)
                  (if (equal? (car args) "-n")
                      (values #f (cdr args))
                      (values #t args))])
      (do ([args args (cdr args)] [sep "" " "])
          ((null? args)
           (printf "~a~a" sep (car args)))
          (when newline? (newline))))))

```

The location of the `scheme` varies depending upon how it was installed, so the first line may have to be adjusted. `petite` will work as well as long as the compiler is not required to run the shell script.

2.59. More command-line option changes (6.8)

The set of command-line arguments has been expanded and refined. Here is the present set as presented by `scheme --help` or `petite --help`.

<code>-b path, --boot path</code>	load boot file
<code>-c, --compact</code>	toggle compaction flag
<code>-h path, --heap path</code>	load heap file
<code>-q, --quiet</code>	suppress greeting and prompt
<code>-s[n] path, --saveheap[n] path</code>	save heap file
<code>--script path</code>	run as shell script
<code>--version</code>	print version and exit
<code>--help</code>	print help and exit
<code>--</code>	pass through remaining args

2.60. Saved-heap changes (6.7)

Heaps explicitly named by the `-h` (`--heap`) option are now resolved using the `SCHEMEHEAPDIRS` path unless the full (absolute or relative) path is specified.

Each heap is now associated with a date stamp created when the corresponding executable image is created, and this date stamp is checked to insure heap compatibility.

2.61. String escapes (6.7)

New reader syntax for including special characters in strings has been added, along with character syntax for audible bell and vertical tabs. The new and existing character names and string escapes are summarized in the table below.

<i>string escape</i>	<i>character equivalent</i>	<i>description</i>
<code>\a</code>	<code>#\bel</code>	audible bell
<code>\b</code>	<code>#\backspace</code>	back space
<code>\f</code>	<code>#\page</code>	form feed
<code>\n</code>	<code>#\newline</code>	newline
<code>\r</code>	<code>#\return</code>	carriage return
<code>\t</code>	<code>#\tab</code>	tab
<code>\v</code>	<code>#\vt</code>	vertical tab
<code>\\</code>	<code>#\\</code>	back slash
<code>\"</code>	<code>#\"</code>	double quote
<code>\'</code>	<code>#\'</code>	single quote
<code>\nnn</code>	<code>#\nnn</code>	ascii code nnn (octal)
<code>\xnn</code>		ascii code nn (hexadecimal)

The reader signals an error for any escape sequence not listed above, e.g., for `"\c"`.

The string escape for single quote is not useful in Scheme; it is included to simplify interoperability with ANSI C.

2.62. New `revisit` procedure (6.7)

The new `revisit` procedure loads and executes run-time code only from a compiled file, i.e., that portion of the code not loaded by `visit`.

Some of the code in a compiled object file can be classified as compile-time or run-time code. Compile-time code establishes compile-time identifier bindings or performs other tasks that are required for subsequent compilation. This includes keyword bindings for syntactic abstractions and module information, and may also include any code explicitly specified as “visit-time” code using `eval-when` (see below).

Run-time code is code that is not required for subsequent compilation; this includes the code to establish the run-time values of variables and other computations that are not specified as “visit-time” only.

The standard procedure `load` loads (and runs) all of the code in an object file, i.e., both compile-time code and run-time code. The *Chez Scheme* procedure `visit`, introduced in Version 6.0, loads only the compile-time portion of the code. The new *Chez Scheme* procedure `revisit`, loads only the run-time portion of the code.

Loading a file by visiting then revisiting the file is almost, but not quite, equivalent to loading it with `load`. The difference is the order in which the items in the file are processed. In the latter case, the loading of compile-time and run-time code is interleaved as it appeared in the original source file. This difference is typically unimportant.

2.63. New `eval-when` situations (6.7)

Two new `eval-when` situations, `visit` and `revisit`, are now supported. Like the existing situations `compile` and `load`, these new situations are relevant only when compiling a file (using `compile-file`). As implied by their names, these situations direct the compiler to generate code to be evaluated when the resulting object file is visited or revisited using the `visit` or `revisit` procedures, or using `load`, which both visits and revisits.

Refer to the *Chez Scheme Version 7 User's Guide* for details.

2.64. `eval-when` and definitions (6.7)

`eval-when` may now appear anywhere that definitions may appear and is considered to be a definition syntactically in those contexts.

An `eval-when` in an internal definition context is treated as a splicing construct, like `begin`, if the list of situations includes `eval`. That is, the sequence of forms within the `eval-when` form are treated as if they appeared in place of the `eval-when` form. If the situation list does not include `eval`, the `eval-when` form is ignored, i.e., the forms in the body are dropped.

2.65. record-writer and type-descriptor (6.6)

A print procedure may be established for record type via the `record-writer` procedure, which accepts a record-type descriptor and a procedure. The procedure is passed an object (an instance of the record type), an output port, and a write procedure. The write procedure must be used in situations where a call to `write` would ordinarily be used by the record writer. The print procedure may be called multiple times to print a single object to support cycle detection and graph printing. Thus, it should have no side effects other than writing to the given port argument.

When called with a record-type descriptor only, `record-writer` returns the record writer currently registered for the record type, which is the default record writer if no record writer has been registered for the record type.

This facility replaces the `print-method` option in `define-record`, which was eliminated when record types produced by `define-record` were made nongenerative. (See “Record inheritance and other changes” below.)

A new syntactic form, `type-descriptor`, has been introduced. When given the name of a record defined by `define-record`, `type-descriptor` returns the record-type descriptor for that record.

```
(define-record point (x y))
(type-descriptor point) ⇒ #<record type point>
```

In earlier versions, it was necessary to create an instance of the record type and extract the descriptor with `record-type-descriptor`. Using `type-descriptor` is both cleaner and more efficient.

2.66. File position on string ports (6.6)

`file-position` now works for string ports.

2.67. Record accessor changes (6.6)

`define-record` no longer defines accessors/mutators for inherited fields. It signals an error if duplicate field names are defined. Child field names can, however, overlap with parent field names, and duplicate field names can arise via the low-level `make-record-type` interface.

`record-field-accessor` and `record-field-mutator` now accept field ordinals as well as field names. A field’s ordinal is its zero-based position in the list of fields returned by `record-type-field-names`. Field ordinals are used by the inspector and printer to avoid problems with duplicate field names.

2.68. Changes to the debug and break handlers (6.6)

The debug handler, invoked as a result of an explicit call to `debug`, supports several new or altered options: `s` to show the current error message, `e` to exit and retain the error continuation, and `q` to quit and discard the error continuation. Discarding the error continuation can reduce storage retention, although this is often not a problem while debugging. When running multithreaded, the debug handler supports inspection of error continuations from any thread.

The break handler, invoked as the result of a keyboard interrupt or explicit call to `break`, now displays the prompt `break>` rather than `debug>`, which is reserved for the now separate debug handler. The break handler is otherwise unchanged.

2.69. Expander change (6.6)

An incompatible change to the `syntax-case` expander has been made to reduce the loss of source information when one macro generates another macro definition. List and vector structure in the subexpression of a syntax form is no longer guaranteed to be list and vector structure in the output form except where pattern variables are contained within that structure. For example, `#'(a ...)`, where `a` is a pattern variable, is guaranteed to evaluate to a list, but the constant structure `#'(a b c d)`, where none of `a`, `b`, `c`, and `d` are pattern variables, may not. The practical consequence of this change is that constant structures must be deconstructed using `syntax-case` or `syntax-rules` rather than `car`, `cdr`, and other list-processing operations.

2.70. More generated symbol changes (6.6)

(See also “Generated symbol changes” below.) The “pretty” name of a generated symbol (gensyms) is now determined when the name is printed or obtained via `symbol->string`, rather than when the gensym is created. Changes to the parameters `gensym-prefix` and `gensym-count` are effective only when the gensym’s name is determined. To get the old behavior, it is necessary to force the name to be determined when the gensym is created, using `symbol->string`.

2.71. Multithreading (6.5)

Chez Scheme now supports multithreading, including threads running on multiple processors. The thread support is implemented in terms of Posix Threads (pthreads). Information on the thread system may be found in the *Chez Scheme Version 7 User’s Guide*.

2.72. Unbuffered Ports (6.5)

Chez Scheme buffers file I/O operations for efficiency, but buffered I/O is not thread safe. Two threads that write to or read from the same port can corrupt the port, resulting in buffer overruns and, ultimately, invalid memory references. When running the threaded version of *Chez Scheme*, the console output port is unbuffered by default. This allows multiple threads to print error and/or debugging messages to the console. The output may be interleaved, even within the same line, but the port will not become corrupted.

Files created by `open-input-file`, `open-output-file`, `call-with-input-file`, `call-with-output-file`, `with-input-from-file`, and `with-output-to-file` are buffered by default: Threads that wish to allow multiple threads to write to or read from the same file via a single port should use the `unbuffered` flag when opening the file, e.g.:

```
(define p (open-output-file "prog.out" '(unbuffered)))
```

Refer to the *Chez Scheme Version 7 User’s Guide* for details.

2.73. 64-bit Sparc V9 architecture port of *Chez Scheme* (6.4)

Support for the Sparc 64-bit (V9) architecture has been added. Programs running in the 64-bit version (machine type “sparcv9”) can create heaps larger than the four-gigabyte limit imposed by the 32-bit architecture. Because data structures are larger and code sequences needed to perform various operations are longer in the 64-bit version, the 32-bit version (machine type “sps2”) is generally faster and should be used when heaps smaller than four gigabytes suffice. Because the data structures are incompatible, it is not possible to run both 32- and 64-bit code in the same process.

Foreign C code intended to be linked with the 64-bit version should be compiled using the C compiler’s `-xarch=v9` option.

Under the 64-bit version, two new foreign/record data types are available: `integer-64` and `unsigned-64`. The `iptr` and `uptr` data types (see below) map to `integer-64` and `unsigned-64` respectively.

2.74. New foreign-interface datatypes (6.4)

Because Scheme objects may no longer be the size of an `int` in C, but in fact generally map to the size of a C pointer, the definition of “`ptr`” in the header file `scheme.h` has changed from `int` to `void *`. Two new datatypes, `iptr` and `uptr`, are defined by the header file as well; these are defined to be signed and unsigned integers of the same size as a `ptr`.

The foreign interface (`foreign-procedure` and `foreign-callable`) also recognizes `iptr` and `uptr` as argument and result type specifiers, mapping them to the appropriate size of signed or unsigned integer.

To further facilitate consistency between Scheme and C code, the foreign interface supports an additional new type, `(void *)`, that maps to an unsigned integer of the appropriate size.

The argument-type specifiers supported by `foreign-procedure` and `foreign-callable` are summarized below.

<code>integer-32</code>	32-bit signed integer
<code>integer-64</code>	64-bit signed integer (64-bit platforms only)
<code>unsigned-32</code>	32-bit unsigned integer
<code>unsigned-64</code>	64-bit unsigned integer (64-bit platforms only)
<code>iptr</code>	32- or 64-bit signed integer on 32- or 64-bit platform
<code>uptr</code>	32- or 64-bit unsigned integer on 32- or 64-bit platform
<code>fixnum</code>	integer in fixnum range (maps to <code>iptr</code>)
<code>single-float</code>	single-precision (32-bit) floating-point number
<code>double-float</code>	double-precision (64-bit) floating-point number
<code>boolean</code>	boolean value (maps to C <code>int</code>)
<code>char</code>	character value (maps to C <code>char</code>)
<code>string</code>	string (maps to C <code>char *</code>)
<code>scheme-object</code>	any Scheme object
<code>(void *)</code>	32- or 64-bit unsigned integer on 32- or 64-bit platform

The result-type specifiers supported by `foreign-procedure` include all of the above plus `void`. The result-type specifiers supported by `foreign-callable` include all of the above except `string`, plus `void`.

2.75. Foreign-interface datatype extensions (6.4)

The range of Scheme integers accepted as `foreign-procedure` arguments specified to be `integer-32` or `unsigned-32` has been extended. For both types, the accepted range is now -2^{31} through 2^{32-1} . In either case, all inputs are converted into the corresponding 32-bit two’s complement representation, so that for `integer-32`, values of 2^{31} and above map to negative values and for `unsigned-32`, values below zero map to values of 2^{31} or greater.

The range values accepted by `integer-64` and `unsigned-64` is -2^{63} through 2^{64-1} , with a similar treatment of signed and unsigned values.

2.76. New record field types (6.3)

The record system (`define-record` and `make-record-type`) now supports signed and unsigned integers of size 8, 16, and, on 64-bit platforms, 64 bits, in addition to the previously supported 32-bit integers. It also supports single (32-bit) as well double (64-bit) floating-point numbers. The field type specifiers supported by the record system are summarized below.

<code>integer-8</code>	8-bit signed integer
<code>integer-16</code>	16-bit signed integer
<code>integer-32</code>	32-bit signed integer
<code>integer-64</code>	64-bit signed integer (64-bit platforms only)
<code>unsigned-8</code>	8-bit unsigned integer
<code>unsigned-16</code>	16-bit unsigned integer
<code>unsigned-32</code>	32-bit unsigned integer
<code>unsigned-64</code>	64-bit unsigned integer (64-bit platforms only)
<code>single-float</code>	single-precision (32-bit) floating-point number
<code>double-float</code>	double-precision (64-bit) floating-point number
<code>scheme-object</code>	any Scheme object (the default)

2.77. Fenders in `syntax-rules` (6.3)

The `syntax-rules` form has been extended to support fenders, with the same syntax and semantics of the fenders supported by `syntax-case`. This extension to `syntax-rules` should not be used in Scheme code intended to be portable to other Scheme systems, which do not generally support fenders.

2.78. Enhanced customization support (6.3)

A set of new C library routines (`Sscheme_init`, `Sregister_boot_file`, `Sbuild_heap`, etc.) have been added to allow Scheme programs to be run from other programs or to allow the construction custom executable images. Refer to the *Chez Scheme Version 7 User's Guide* for details.

2.79. Support for anonymous foreign procedures (6.3)

Support for anonymous foreign procedures has been added. When *entry-exp* in

```
(foreign-procedure entry-exp (arg-spec ...) result-spec)
```

evaluates to an integer, it is assumed to be the address of the foreign entry point rather than its name.

2.80. Support for calling Scheme from C (6.3)

A new `foreign-callable` form has been added that permits C or other languages that can make C-compatible calls to invoke Scheme procedures. As with `foreign-procedure`, `foreign-callable` permits the specification of input and output types and automatically performs the necessary conversions (marshaling) to and from Scheme datatypes. Information on `foreign-callable` and how to use it may be found in the *Chez Scheme Version 7 User's Guide*. (At present, `foreign-callable` is supported only under Windows 9x/NT and Intel Linux. Support for other machine types will be added later.)

2.81. Support for file compression (6.3)

Support for transparent compression/decompression when writing to/reading from files has been added via inclusion of zlib compression code. Compression/decompression is performed whenever a file is opened using the optional `compressed` argument to one of the file open routines, including `open-input-file` and `open-output-file`. See the *Chez Scheme Version 7 User's Guide* for details.

The compression and decompression is performed with the use of the zlib compression library developed by Jean-loup Gailly and Mark Adler. It is therefore compatible with the `gzip` program, which means that `gzip` may be used to uncompress files produced by *Chez Scheme* and visa versa.

Object files produced by the compiler are compressed by default. Object-file compression may be disabled by setting the parameter `compile-compressed` to false.

2.82. Enhancement to load and include (6.3)

`load` and `include` are now sensitive to the `source-directories` parameter for non-absolute pathname arguments. `load` consults this parameter for both source and object files. This parameter had previously been used only by the inspector.

2.83. New print-record parameter (6.3)

A new parameter, `print-record`, has been added. When set to `#t` (the default), the printer prints records fully. When set to `#f`, the printer uses an abbreviated form that shows a record's name but not its contents.

2.84. Record inheritance and other changes (6.2)

The record system now supports inheritance. If the optional first argument to `make-record-type` is a record-type descriptor, it identifies the parent record type for the new record type. With `define-record`, a parent record type is identified by placing the name of the parent record type (which must have been established by an earlier use of `define-record`) after the new record type name in the `define-record` syntax.

```
(define-record name (field ...))
(define-record name parent-name (field ...))
```

The new record type inherits the fields of the parent. The parent type predicate returns true for instances of the child (and any of its children). The parent accessors and mutators work on child instances as well.

`define-record` is also now nongenerative, meaning that each time a given `define-record` form is evaluated, it produces the same record type. That is, the record type is created when the `define-record` form is compiled, not each time the resulting code is run.

The print method argument to `make-record-type` has been eliminated, as have both the `print-method` and `reader-name define-record` options.

2.85. Generated symbol changes (6.2)

Several changes to generated symbols (gensyms) have been made. In previous versions, symbols created by `gensym` were “uninterned,” meaning they were not entered into the system's internal hash table. This is still true for gensyms in Version 7.0 until they are printed. When a gensym is printed, it is given a globally unique name and interned under that name. This allows them to be written in such a way that they can be read back and reliably commonized on input. A gensym's unique name can also be obtained by calling the new procedure `gensym->unique-string`, which, like printing, forces the unique name to be created if it has not already been created. Delaying the creation of unique names and the interning of gensyms as long as possible allows the system to reclaim gensyms that are no longer accessible even if they have top-level values or nonempty property lists.

Each gensym is also given a “pretty” name, which is formed as in past systems. By default, the pretty name consists of the letter “g” followed by a sequence number, e.g., `g25`. The pretty name is returned by `symbol->string`.

When `print-gensym` is set to true (the default), gensyms print as shown below.

```
#{pretty-name unique-name}
```

When set to false, gensyms print simply as *pretty-name*. The syntax `#:name` is no longer recognized by the reader.

The procedure `string->uninterned-symbol` no longer exists. In its place, `gensym` now accepts an optional second argument, a string to use as the pretty name. Similarly, the predicate `uninterned-symbol` has been replaced by `gensym?`. Compatibility definitions for both can be found in the file `examples/compat.ss` in the *Chez Scheme* distribution.

```
> (define g1 (gensym))
> (symbol->string g1)
"g1411"
> (symbol->string (gensym))
"g1412"
> g1
#{g1411 |"V*_da}f)+)Bp\\"}
> (gensym->unique-string g1)
"\\"V*_da}f)+)Bp\\\\""
```

2.86. Automatic heap loading (6.2)

Support for automatic heap loading has been added. This eliminates the need for shell scripts to start *Chez Scheme*, *Petite Chez Scheme*, and Scheme-coded applications in most cases. If the system is invoked as `progrname arg ...` with no heap arguments, the system looks for `progrname.heap` in a set of default locations, which can be overridden by setting the environment variable `SCHEMEHEAPDIRS`. Similarly, if lower-level heaps are not specified, the system looks for them in the same set of locations. `SCHEMEHEAPDIRS` is a colon-separated list of directories. Within `SCHEMEHEAPDIRS`, the current version, e.g., `6.2`, is substituted for `%v`, the machine type, e.g., `sps2`, is substituted for `%m`, and `char` is substituted for `%char` for all other characters to allow directory names to include `%v`, `%m`, and colons. If `SCHEMEHEAPDIRS` ends in a colon, the default directories are searched after the specified directories; otherwise only the specified directories are searched. When creating an application named `progrname`, link `scheme` to `progrname` and put `progrname.heap` in one of the standard locations. See the *Chez Scheme Version 7 User's Guide* for details.

2.87. Command-line option changes (6.2)

The obsolete `-o` option and an undocumented `-t` (test) option are no longer recognized by *Chez Scheme*. The `-`, `-b`, `-c`, `-h`, and `-sn` options must now be followed by a null character; that is, they are no longer recognized as options if there are additional characters. So, for example, `-help` and `-s99` are passed along the scheme startup routine. Multiple `-b` (boot file) options are now accepted. (The `-b` option is used for base heap creation and is not of interest to most users.)

2.88. Compile-time checking of format strings (6.1)

Compile-time control-string and argument-count checking is now performed at optimize levels 2 and 3 for `format`, `printf`, `fprintf`, `error`, and `warning`.

2.89. Quasiquote extension (6.1)

`unquote` and `unquote-splicing` have been extended in an upwardly compatible way to allow multiple subforms in splicing (list or vector) contexts. (`unquote-splicing x ...`) now splices into the surrounding context the result of appending the values of `x ...`, and (`unquote x ...`) now inserts into the surrounding context the values of `x ...`. These changes support certain useful nested `quasiquote` idioms, such as `,@,@`, which has the effect of a doubly indirect splicing when used within a doubly nested and doubly evaluated `quasiquote` expression. For example,

```
(define x '(m n))
(define m '(b c))
(define n '(d e))
(eval '(a ,@,@x f) (interaction-environment))
```

evaluates to (a b c d e f).

2.90. Tracing syntax transformers (6.1)

Support for `trace-define-syntax` has been added. This form is like `define-syntax`, except that the transformer bound to the defined keyword is wrapped in tracing code that displays the input and output to the transformer in the same manner as `trace-lambda`. See the *Chez Scheme Version 7 User's Guide* for details.

2.91. Making ports from file descriptors (6.1)

`make-output-port`, `make-input-port`, and `make-input/output-port` now allow the first, *handler*, argument to be a nonnegative fixnum representing a low-level file descriptor, obtained via some operating-system dependent means. When *handler* is a file descriptor, the newly created port is given a handler that reads from and/or writes to the file descriptor. See the *Chez Scheme Version 7 User's Guide* for details.

2.92. Inspection of record-type descriptors (6.1)

The interactive and object inspectors now allow a record's record-type descriptor to be accessed and inspected. The command `rtd` is used while inspecting a record in the interactive inspector to inspect the record's record-type descriptor. The message `rtd` is used similarly in the object inspector.

2.93. Block Comments (6.0a)

The reader now allows a block of characters to be commented out within properly nested `#|`, `|#` pairs, without regard to line boundaries. For example, the expression

```
(+ #| This is a comment|# 3 4)
```

evaluates to 7, as does

```
(+ #| This is a comment
    with another #| nested
    inside, crossing several
    line|# boundaries|# 3 4)
```

2.94. Support for Windows "stdcall" Convention (6.0a)

Chez Scheme now supports the Windows `__stdcall` procedure calling convention on Intel 80x86 systems for foreign procedures created with `foreign-procedure`. The default is still the `__cdecl` convention, as with Microsoft's C compiler. `__stdcall` is specified by placing the keyword `__stdcall` before the foreign procedure name, e.g.:

```
(foreign-procedure __stdcall "GetVolumeInformationA"
  (string string unsigned-32 string string string string unsigned-32)
  boolean)
```


`--stdcall` is useful for calling Windows API functions, except for those with “varargs” interfaces, which use the default `--cdecl` convention.

3. Bug Fixes

3.1. Storage management bug (7.0)

A storage management bug that on rare occasions caused an unrecoverable invalid memory reference to occur has been fixed. [This bug dated back to Version 6.6.]

3.2. Three-plus argument subtraction and division bugs (7.0)

The various subtraction and division operators, like `/` and `f1-`, which were in some cases implemented by repeated divisions or subtractions, now consistently implement the documented semantics whereby, e.g., `(- x0 x1 x2 ...)` is treated as `(- x0 (+ x1 x2 ...))`. Difference can be reflected in round-off errors, subtractions involving `+0.0` and `-0.0`, and occurrence or nonoccurrence of arithmetic overflow and underflow. [These bugs dated back to Version 6.0 or before.]

3.3. Thread bugs (6.9d)

Several bugs in the support for threading, which manifested themselves in spurious lock-ups and invalid memory references, have been fixed. [These bugs dated back to Version 6.5.]

3.4. Rounding bug (6.9d)

A bug that caused the `round` procedure to round down some values that are slightly over half-way between two integers, e.g., `0.5000000000000001` has been fixed. [This bug dated back to Version 6.0 or before.]

3.5. Fasl writing/reading bugs (6.9c)

A bug in `fasl-write` which caused the error “incorrect number of arguments to #<procedure bld>” has been fixed. This bug was present in version 6.9b but not in version 6.9 or earlier. A separate bug that caused NaNs (`+nan.0`) to be read improperly on some systems by the `fasl` loader has also been fixed. A third bug dating back to Version 6.7 has also been fixed. This bug sometimes resulted in segmentation faults for `fasl` files containing very long strings. [The first bug dated back to Version 6.9c, the second to Version 6.0, and the third to Version 6.7.]

3.6. Format bug (6.9c)

A bug in `format` which resulted in an invalid memory reference when the `~d`, `x`, `~o`, `~b`, and `~nr` options were given non-numeric inputs has been fixed. [This bug dated back to Version 6.9b.]

3.7. string->number bugs (6.9c)

Several bugs in `string->number` have been fixed. Most had to do with division by zero signaling errors rather than returning a valid numeric or false output. Certain invalid inputs, like `3@-i`, for which `string->number` had incorrectly been returning numeric values, now properly result in a false value. Some of these fixes are also reflected in more appropriate behavior for some inputs on the part of the reader and printer. [These bugs dated back to Version 6.0.]

3.8. Printer bug (6.9c)

A bug which caused the printer to omit the backslash escape before a hash (#) character at the beginning of a symbol's name wasn't escaping # at beginning of a symbol's name in cases where bars are present elsewhere in the symbol's name has been fixed. For example, `'\#hash\|bar` now prints properly as `'\#hash\|bar`. [This bug dated back to Version 6.0.]

3.9. letrec* bug (6.9c)

An optimization bug in the handling of `letrec*` which could result in some out-of-order execution of global or assignable local variable references has been fixed. [This bug dated back to Version 6.9.]

3.10. eval-when and begin syntax errors (6.9c)

Syntax error messages for `eval-when` and `begin` when these forms are loaded or compiled from a file no longer display the internal syntax-object representation but instead display properly with line number and character position information. [This bug dated back to Version 6.0.]

3.11. Inspector invalid memory reference bug (6.9c)

The inspector now provides an appropriate error message instead of causing an invalid memory reference when the “file” command is given non-string, non-symbol argument. [This bug dated back to Version 6.0.]

3.12. expt call hanging the compiler (6.9c)

The `expt` procedure no longer causes the optimizer to hang or take a long time when it appears with large constant arguments or arguments the compiler can determine to be constant via static analysis. [This bug dated back to Version 6.0.]

3.13. Mapping procedures and mutated input lists (6.9c)

The mapping procedures `for-each`, `ormap`, and `andmap` now signal an error if their input lists are somehow truncated while they are running. `map` is no longer sensitive to whether its input lists are truncated. [This bug dated back to Version 6.0.]

3.14. R5RS environment and c...r primitives (6.9c)

The `c...r` procedures, like `caddar`, are now all properly classified as `r5rs` primitives and thus are now included in the `r5rs` module and environment. Most of these had been improperly classified. [This bug dated back to Version 6.0.]

3.15. Waiter bug mixing engines and continuations (6.9c)

An error in the waiter that could result in an invalid memory reference while running engine computations involving continuations has been fixed. (Mixing engines and continuations generally doesn't lead to satisfactory results, since engines are implemented using continuations, but mixing them should not cause invalid memory references.) [This bug dated back to Version 6.0.]

3.16. Heap saving bug (6.9c)

Saving the heap to the same file as it was loaded from no longer fails on systems that support memory mapping of heaps. [This bug dated back to Version 6.0.]

3.17. Undefined variable detection bug (6.9b)

A bug in the compiler's handling of `letrec` and `letrec*` that allowed some undefined-variable references to go undetected has been fixed. [This bug dated back to Version 6.9.]

3.18. Sparc 64-bit Dynamic linker bug (6.9a)

A bug in the dynamic linker that could cause an invalid memory reference or illegal instruction trap when very large heaps are used has been fixed. [This bug dated back to Version 6.4.]

3.19. Sparc 64-bit file-position bug (6.9a)

`file-position` no longer complains when handed positions outside of the 32-bit range when running the 64-bit Sparc version of the system. [This bug dated back to Version 6.4.]

3.20. `expt` bug (6.9a)

A bug which caused excessive round-off error in `expt` has been fixed. [This bug dated back to Version 6.0.]

3.21. `compile-profile` bug (6.9a)

`compile-file` now determines the value of `compile-profile` once for each expression so that portions of a file can be instrumented for profiling, e.g., via the insertion of

```
(eval-when (compile) (compile-profile #t))
```

into the file. [This bug dated back to Version 6.0.]

3.22. `foreign-callable` memory fault (6.9a)

A bug in `foreign-callable` that resulted in occasional memory faults has been fixed. [This bug dated back to Version 6.3.]

3.23. Laziness in local transformer evaluation (6.9)

To avoid unnecessary work at expansion time, the expander defers the evaluation of local transformers, i.e., those bound by `let-syntax`, `letrec-syntax` or internal `define-syntax` forms, until they are actually used. To avoid this laziness from being detectable, the expander now defers the evaluation only for transformer expressions that expand into `lambda` expressions. [This bug dated back to Version 6.0.]

3.24. Inspector output bug (6.9)

The inspector now properly prints the output of the `=>` and `eval` messages to the console output port rather than to the current output port, if different. [This bug dated back to Version 6.0.]

3.25. Reader bug (6.9)

A bug in the reading of records, which sometimes made the annotations created by the reader visible within record constants appearing in files loaded from source or compiled, has been fixed. The bug occurred when a record field was constructed using a graph reference to a marked portion of the structure outside of the record, e.g., `'#1=(#[foo #1#])` [This bug dated back to Version 6.0.]

3.26. Compiler bug (6.9)

A bug in the compiler that caused inappropriate “incorrect argument count” warnings and errors when invoking a locally bound procedure created with `case-lambda` has been fixed. [This bug dated back to Version 6.0.]

3.27. Expander bug (6.9)

A bug in the expander causing unexpected “out of context” errors for exports of identifiers with the same name as an internal unexported identifier has been fixed. [This bug dated back to Version 6.0.]

3.28. Interpreter evaluation-order bug (6.8)

An interpreter evaluation-order bug has been fixed, preventing the interpreter from signaling an error prematurely for attempts to invoke nonprocedures when the call is never actually made. For example, the interpreter used to signal an error for the following program, but now properly returns 1.

```
(call/cc (lambda (k) (0 (k 1))))
```

This fix affects only programs running under Petite Chez Scheme or programs running under Chez Scheme that explicitly use the interpreter. [This bug dated back to Version 6.0.]

3.29. Inlining bug (6.8)

A bug in the compiler’s inlining of certain indirect calls to procedures that are provided an incorrect number of arguments has been fixed. The bug previously resulted in an invalid memory reference during compilation; it now results in an appropriate run-time “incorrect number of arguments” error. [This bug dated back to Version 6.0.]

3.30. Generated record constructor fixes (6.8)

A bug in `record-constructor`’s algorithm for determining record sizes that resulted in more allocation than necessary has been fixed. A separate bug that could result in “invalid parameter list” errors while creating the record constructor has also been fixed. [These bugs dated back to Version 6.0.]

3.31. Locked object bugs (6.5/6.6)

Bugs in the handling of objects locked via `lock-object` have been fixed. These bugs sometimes resulted in invalid memory references or failure to properly lock objects. [These bugs dated back to Version 6.0.]

3.32. Block comment bug (6.4)

A bug in the reader that prevented certain block comments from being treated properly has been fixed. [This bug dated back to Version 6.0.]

3.33. Application of unbound variables (6.3)

A bug that resulted in an invalid memory reference rather than a sensible error message when applying the value of certain unbound variables has been fixed. [This bug dated back to Version 6.2.]

3.34. Assignments to unreferenced letrec-bound variables (6.3)

A bug that sometimes resulted in an invalid memory reference for programs that contain assignments to unreferenced `letrec`-bound variables has been fixed. [This bug dated back to Version 6.0.]

3.35. Setting primitive names to nonprocedures (6.2)

Setting a primitive name, e.g., `car` or `map`, to a nonprocedure values, then invoking it, resulted in an invalid memory reference. This bug has been fixed. [This bug dated back to Version 6.0.]

3.36. Trace error (6.2)

A bug in the trace package, which prevented further uses of `trace` and `untrace` if a traced procedure was redefined as a syntactic abstraction, has been fixed. [This bug dated back to Version 6.0.]

3.37. Module names introduced by macros (6.1)

The expander now properly renames module names introduced into the expansion of a macro at top level, as it already did for other defined identifiers. [This bug dated back to Version 6.0.]

3.38. Interpreter case-lambda bug (6.1)

A bug in the interpreter's treatment of `case-lambda`, which caused argument-count errors to be produced for certain correct applications of procedures created by `case-lambda`, has been fixed. [This bug dated back to Version 6.0.]

3.39. IRIX memory limitation (6.1)

A limitation on the maximum size of a heap under IRIX caused by inappropriate linker assumptions about the operating environment has been lifted. In theory, the size of a heap is now limited only by available (32-bit) virtual memory, and in practice we have been able to test heaps of up to about two gigabytes in size. [This bug dated back to Version 6.0.]

3.40. Storage-Management and Records (6.0a)

A bug in the garbage collector's handling of records containing both mutable pointer (`scheme-object`) and nonpointer (`integer-32`, `unsigned-32`, or `double-float`) fields has been fixed. The bug sometimes caused unrecoverable invalid memory references. [This bug dated back to Version 6.0.]

3.41. Source-Tracking Problems (6.0a)

A couple of problems with source tracking in the compiler have been fixed. These problems sometimes resulted in memory faults within the inspector or in misleading information being printed as part of syntax error message. An unrelated bug with finding source files with pathnames that start with a drive letter under Windows has also been fixed. [This bug dated back to Version 6.0.]

3.42. Overeager Argument-Count Checking (6.0a)

The compiler now signals a warning rather than an error when it detects a call to a locally bound procedure with an incorrect number of arguments. This prevents the compiler from rejecting programs such as the following, which might otherwise execute without errors:

```
(let ((f (lambda (x) x)))
  (if (read)
      (f 3)
      (f 3 4)))
```

It also prevents the compiler from rejecting programs with less obvious “errors,” uncovered during procedure inlining, such as the following:

```
(let ((f (lambda (g n)
           (if (= n 1)
               (g (lambda (x) x))
               (g (lambda (x y) x))))))
  (f (lambda (h) (h 3)) 1))
```

[This bug dated back to Version 6.0.]

4. Performance Enhancements

4.1. Locked-object handling (6.9d)

The garbage collector’s handling of locked objects has been improved, reducing the high penalty for having many objects locked at the time a collection occurs. The indirect penalty, which is that lots of memory may be tied up to hold just a few locked objects, still exists, so programs should lock objects only when absolutely necessary.

4.2. Top-level calls (6.9d)

Calls through variables bound at top level are now slightly faster, especially at optimize-level 0, due to an improvement in the representation of the top-level environment.

4.3. Source Optimization (6.9d)

The compiler performs various new source optimizations, including partial folding of numeric operations like `(+ 3 x 4)` to `(+ 7 x)` at optimization levels 2 and 3.

4.4. Variable-sized allocation requests (6.9c)

The code for variable-sized allocation requests, e.g., calls to `make-vector` with nonconstant size arguments, is now shorter and faster. This should have noticeable benefit only when allocating relatively small objects.

4.5. Improved mapping performance (6.9c)

The compiler now generates better inline or out-of-line code for the mapping procedures `map`, `for-each`, `andmap`, and `ormap` at optimization level 3. `for-each`, `andmap`, and `ormap` are also now “properly tail recursive” in that they tail call the input procedure on the last list element.

4.6. Eta reduction of primitive calls (6.9c)

The optimizer now eta reduces calls to primitives, e.g.:

```
((lambda (x1 ... xn)  
  (prim x1 ... xn))  
 e1 ... en)
```

is reduced to

```
(prim e1 ... en)
```

where *prim* is a recognized Scheme primitive. (Primitives are recognized by name at optimization level 2 and above, and they are also recognized within the scope of an import from the `scheme` or other built-in modules. Primitives may also be explicitly identified via the `##` prefix, e.g., `##car` or `##3%fx+`.)

4.7. Faster fixnum operations (6.9a)

Fixnum addition and subtraction operators, which had been open-coded only at optimize-level 3, are open-coded at optimize-level 2 as well.

4.8. Faster append and apply (6.9a)

The compiler now generates better code for `append` and `apply` at optimize-level 3.

4.9. Faster syntax-case expander (6.9a)

Improvements have been made in the implementation of `syntax-case` to make it significantly faster for some extreme expansion problems involving many levels of macro expansion.

4.10. Inexact numbers with large exponents (6.9a)

Reading of inexact numbers with very large positive or negative exponents is now much faster; the system checks for such large exponents and immediately returns the appropriate zero or infinity value.

4.11. Giving memory back to the O/S (6.6)

Version 7.0 returns memory to the operating system when possible following the collection of the oldest generation.

4.12. Faster bignum arithmetic (6.4)

Internal changes to the bignum representation and bignum arithmetic package have resulted in a modest speedup of bignum operations, typically by around 10-15%.

4.13. Vector handling in quasiquote (6.1)

Handling of certain vectors in `quasiquote` has been improved. For example, `'#(a 1 ,(+ 3 4) x y)` now expands into the following code.

```
(##2%vector 'a 1 (+ 3 4) 'x 'y)
```

Previously, the same expression expanded into code that generated a list and converted this list into a vector via `list->vector`.