CARNEGIE-MELLON UNIVERSITY

DEPARTMENT OF COMPUTER SCIENCE

SPICE PROJECT

Common Lisp Reference Manual

Guy L. Steele Jr.

29 July 1982

Colander Edition
Even More Holes Than Before — But They're Smaller!

Spice Document S061

Keywords and index categories: PE Lisp & DS External

Location of machine-readable file: clm.mss @ CMU

Supported by the Defense Advanced Research Projects Agency, Department of Defense, ARPA Order 3597, monitored by the Air Force Avionics Laboratory under contract F33615-78-C-1551. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

Table of Contents

List of Tables

Acknowledgements

The many people who have contributed to the design of COMMON LISP are hereby gratefully acknowledged:

Alan Bawden[2]
Rodney A. Brooks[3]
Richard L. Bryan[2]
Glenn S. Burke[3]
Howard I. Cannon[2]
George J. Carrette[3]
David Dill[1]
Scott E. Fahlman[1]
Richard J. Fateman[4]
Neal Feinberg[1]
John Foderaro[4]
Richard P. Gabriel[5,6]
Joseph Ginder[1]
Richard Greenblatt[7]
Martin L. Griss[8]
Charles L. Hedrick[9]
Earl A. Killian[6]
John L. Kulp[2]
Larry M. Masinter[10]
John McCarthy[5]
Don Morrison[8]
David A. Moon[2]
William L. Scherlis[1]
Richard M. Stallman[3]
Barbara K. Steele[1]
Guy L. Steele Jr.[1], *editor*
William vanMelle[10]
Walter van Roggen[1]
Allan C. Wechsler[2]
Daniel L. Weinreb[2]
Jon L White[10]
Richard Zippel[3]
Leonard Zubkoff[1]

1. Computer Science Department, Carnegie-Mellon University, Schenley Park, Pittsburgh, Pennsylvania 15213
2. Symbolics, Inc., Cambridge, Massachusetts 02139
3. Massachusetts Institute of Technology, 545 Technology Square, Cambridge, Massachusetts 02139
4. Computer Science Division, Department of EECS, University of California, Berkeley, California 94720
5. Computer Science Department, Stanford University, Stanford, California 94305
6. University of California, Lawrence Livermore National Laboratory, Livermore, California 94550
7. Lisp Machines Incorporated (LMI), Cambridge, Massachusetts 02139
8. Department of Computer Science, University of Utah, Salt Lake City, Utah 84112
9. Laboratory for Computer Science Research, Rutgers University, New Brunswick, New Jersey 08903
10. Xerox Palo Alto Research Center, Palo Alto, California 94306

As can be seen from the list of affiliations, COMMON LISP was designed by a diverse group of people representing many institutions.

The organization, typography, and content of this document were inspired in large part by the *MacLISP Reference Manual* by David A. Moon and others [6], and by the *LISP Machine Manual* by Daniel Weinreb and David Moon [11], which in turn acknowledges the efforts of Richard Stallman, Mike McMahon, Alan Bawden, Glenn Burke, and "many people too numerous to list".

This edition is still in draft form. Please send remarks, corrections, and criticisms to:

Guy L. Steele Jr.
Computer Science Department
Carnegie-Mellon University
Schenley Park
Pittsburgh, Pennsylvania 15213

Chapter 1

Introduction

This manual documents a dialect of LISP called "COMMON LISP", which is a successor to MACLISP [6], influenced strongly by Lisp Machine LISP [11] and also to some extent by SCHEME [9] and INTERLISP [10].

1.1. Purpose

COMMON LISP is intended to meet these goals:

Commonality. COMMON LISP originated in an attempt to focus the work of several implementation groups each of which was constructing successor implementations of MACLISP for different computers. These implementations had begun to diverge because of the differences in the implementation environments: microcoded personal computers (Lisp Machine LISP, SPICE LISP), commercial timeshared computers (NIL), and supercomputers (S-1 LISP). While the differences among the several implementation environments will of necessity force incompatibilities among the implementations, nevertheless COMMON LISP can serve as a common dialect of which each implementation can be an upward-compatible superset.

Portability. COMMON LISP intentionally excludes features that cannot easily be implemented on a broad class of machines. On the one hand, features that are difficult or expensive to implement on hardware without special microcode are avoided or provided in a more abstract and efficiently implementable form. (Examples of this are the forwarding (invisible) pointers and locatives of Lisp Machine LISP. Some of the problems that they solve are addressed in different ways in COMMON LISP.) On the other hand, features that are useful only on certain "ordinary" or "commercial" processors are avoided or made optional. (An example of this is the type declaration facility, which is useful in some implementations and completely ignored in others; type declarations are completely optional and for correct programs affect only efficiency, never semantics.) Moreover, attention has been paid to making it easy to write programs in such a way as to depend as little as possible on machine-specific characteristics such as word length, while allowing some variety of implementation techniques.

Consistency. Most LISP implementations are internally inconsistent in that by default the interpreter and compiler may assign *different* semantics to correct programs; this stems primarily from the fact that the interpreter assumes all variables to be dynamically scoped, while the compiler assumes all variables to be local unless forced to assume otherwise. This has been done for the sake of convenience and efficiency, but can lead to very subtle bugs. The definition of

COMMON LISP avoids such anomalies by explicitly requiring the interpreter and compiler to impose identical semantics on correct programs.

Power. COMMON LISP is a descendant of MACLISP, which has always placed emphasis on providing system-building tools. Such tools may in turn be used to build the user-level packages such as INTERLISP provides; these packages are not, however, part of the COMMON LISP core specification. It is expected such packages will be built on top of the COMMON LISP core.

Expressiveness. COMMON LISP culls not only from MACLISP but from INTERLISP, other LISP dialects, and other programming languages what we believe from experience to be the most useful and understandable constructs. Constructs that have proved to be awkward or less useful are being eliminated (an example is the `store` construct of MACLISP).

Compatibility. Unless there is a good reason to the contrary, COMMON LISP strives to be compatible with Lisp Machine LISP, MACLISP, and INTERLISP, roughly in that order.

Efficiency. COMMON LISP has a number of features designed to facilitate the production of high-quality compiled code in those implementations that care to invest effort in an optimizing compiler. One implementation of COMMON LISP (namely S-1 LISP) already has a compiler that produces code for numerical computations that is competitive in execution speed to that produced by a FORTRAN compiler [1]. (This extends the work done in MACLISP to produce extremely efficient numerical code [4].)

Stability. It is intended that COMMON LISP, once defined and agreed upon, will change only slowly and with due deliberation. The various dialects that are supersets of COMMON LISP may serve as laboratories within which to test language extensions, but such extensions will be added to COMMON LISP only after careful examination and experimentation.

The COMMON LISP documentation is divided into four parts, known for now as the white pages, the yellow pages, the red pages, and the blue pages. (This document is the white pages.)

- The *white pages* (this document) is a language specification rather than an implementation specification. It defines a set of standard language concepts and constructs that may be used for communication of data structures and algorithms in the COMMON LISP dialect. This is sometimes referred to as the "core COMMON LISP language", because it contains conceptually necessary or important features. It is not necessarily implementationally minimal. While some features could be defined in terms of others by writing LISP code (and indeed may be implemented that way), it was felt that these features should be conceptually primitive so that there might be agreement among all users as to their usage. (For example, bignums and rational numbers could be implemented as LISP code given operations on fixnums. However, it is important to the conceptual integrity of the language that they be regarded by the user as primitive, and they are useful enough to warrant a standard definition.)

- The *yellow pages* is a program library document, containing documentation for assorted and relatively independent packages of code. While the white pages are to be relatively stable, the yellow pages are extensible; new programs of sufficient usefulness and quality will routinely be added from time to time. The primary advantage of the division into white and yellow pages is this relative stability; a package written solely in the white-pages language should not break if

changes are made to the yellow-pages library.

- The *red pages* is implementation-dependent documentation; there will be one set for each implementation. Here are specified such implementation-dependent parameters as word size, maximum array size, sizes of floating-point exponents and fractions, and so on, as well as implementation-dependent functions such as input/output primitives.

- The *blue pages* constitutes an implementation guide in the spirit of the INTERLISP virtual machine specification [7]. It specifies a subset of the white pages that an implementor must construct, and indicates a quantity of LISP code written in that subset that implements the remainder of the white pages. In principle there could be more than one set of blue pages, each with a companion file of LISP code. (For example, one might assume `if` to be primitive and define `cond` as a macro in terms of `if`, while another might do it the other way around.)

1.2. Notational Conventions

In COMMON LISP, as in most LISP dialects, the symbol `nil` (page 45) is used to represent both the empty list and the "false" value for Boolean tests. An empty list may, of course, also be written "`()`"; this normally denotes the same object as "`nil`". (It is possible, by extremely perverse manipulation of the package system, to cause the sequence of letters "`nil`" to be recognized not as the symbol that represents the empty list but as another symbol with the same name. However, "`()`" *always* denotes the empty list. This obscure possibility will be ignored in this document.) These two notations may be used interchangeably as far as the LISP system is concerned. However, as a matter of style, this document will prefer the notation "`()`" when it is desirable to emphasize its use as an empty list, and will prefer the notation "`nil`" when it is desirable to emphasize its use as the Boolean "false" or as a symbol. Moreover, an explicit quote mark is used to emphasize its use as a symbol rather than as Boolean "false".

For example:

```
(append '() '()) => ()          ;Emphasize use of empty lists.
(not nil) => t                  ;Emphasize use as Boolean "false".
(get 'nil 'color)               ;Emphasize use as a symbol.
```

Any data object other than `nil` is construed to be Boolean "not false", that is, "true". The symbol `t` is conventionally used to mean "true" when no other value is more appropriate. When a function is said to "return *false*" or to "be *false*" in some circumstance, this means that it returns `nil`. However, when a function is said to "return *true*" or to "be *true*" in some circumstance, this means that it returns some value other than `nil`, but not necessarily `t`.

All numbers in this document are in decimal notation unless there is an explicit indication to the contrary.

Execution of code in LISP is called *evaluation*, because executing a piece of code normally results in a data object called the *value* produced by the code. The symbol "`=>`" will be used in examples to indicate evaluation. For example:

```
(+ 4 5) => 9
```

means "the result of evaluating the code `(+ 4 5)` is (or would be, or would have been) 9".

The symbol "==>" will be used in examples to indicate macro expansion. For example:

```
(push x v) ==> (setf v (cons x v))
```

means "the result of expanding the macro-call form `(push x v)` is `(setf v (cons x v))`". This implies that the two pieces of code do the same thing; the second piece of code is the definition of what the first does.

The symbol "<=>" will be used in examples to indicate code equivalence. For example:

```
(- x y) <=> (+ x (- y))
```

means "the value and effects of `(- x y)` is always the same as the value and effects of `(+ x (- y))` for any values of the variables `x` and `y`". This implies that the two pieces of code do the same thing; however, neither directly defines the other in the way macro-expansion does.

When this document specifies that it "is an error" for some situation to occur, this means that:

- No valid COMMON LISP program should cause this situation to occur.
- If this situation occurs, the effects and results are completely undefined as far as adherence to the COMMON LISP specification is concerned.
- No COMMON LISP implementation is required to detect such an error.

This is not to say that some particular implementation might not define the effects and results for such a situation; it is merely that no program conforming to the COMMON LISP specification may correctly depend on such effects or results.

On the other hand, if it is specified in this document that in some situation "an error is *signalled*", this means that:

- If this situation occurs, an error (see `error` (page ERROR-FUN)) will be signalled.
- Valid COMMON LISP programs may rely on the fact that an error will be signalled.
- Every COMMON LISP implementation is required to detect such an error.

Functions, variables, named constants, special forms, and macros are described using a distinctive typographical format. Table 1-1 illustrates the manner in which COMMON LISP functions are documented. The first line specifies the name of the function, the manner in which it accepts arguments, and the fact that it is a function. Following indented paragraphs explain the definition and uses of the function and often present examples or related functions.

In general, actual code (including actual names of functions) appears in this typeface: `(cons a b)`. Names that stand for pieces of code (meta-variables) are written in *italics*. In a function description, the names of the parameters appear in italics for expository purposes. The word "`&optional`" in the list of parameters indicates that all arguments past that point are optional; the default values for the parameters are described in the text. Parameter lists may also contain "`&rest`", indicating that an indefinite number of

---

`sample-function` *arg1 arg2* `&optional` *arg3 arg4* [*Function*]

The function `sample-function` adds together *arg1* and *arg2*, and then multiplies the result by *arg3*. If *arg3* is not provided or is `nil`, the multiplication isn't done. `sample-function` then returns a list whose first element is this result and whose second element is *arg4* (which defaults to the symbol `foo`).

For example:

```
(function-name 3 4) => (7 foo)
(function-name 1 2 2 'bar) => (6 bar)
```

As a rule, `(sample-function` *x y*`) <=> (list (+` *x y*`) 'foo)`.

Table 1-1: Sample Function Description

`sample-variable` [*Variable*]

The variable `sample-variable` specifies how many times the special form `sample-special-form` should iterate. The value should always be a non-negative integer or `nil` (which means iterate indefinitely many times). The initial value is `0`.

Table 1-2: Sample Variable Description

`sample-constant` [*Constant*]

The named constant `sample-constant` has as its value the height of the terminal screen in furlongs times the base-2 logarithm of the implementation's total disk capacity in bytes, as a floating-point number.

Table 1-3: Sample Constant Description

---

arguments may appear, or "`&key`", indicating that keyword arguments are accepted. (The `&optional/&rest/&key` syntax is actually used in COMMON LISP function definitions for these purposes.)

Table 1-2 illustrates the manner in which a global variable is documented. The first line specifies the name of the variable and the fact that it is a variable.

Table 1-3 illustrates the manner in which a named constant is documented. The first line specifies the name of the constant and the fact that it is a constant. (A constant is just like a global variable, except that it is an error ever to alter its value or to bind it to a new value.)

`sample-special-form` [*name*] ({*var*}*) {*form*}+ [*Special form*]

This evaluates each form in sequence as an implicit `progn`, and does this as many times as specified by the global variable `sample-variable`. Each variable *var* is bound and initialized to 43 before the first iteration, and unbound after the last iteration. The name *name*, if supplied, may be used in a `return-from` (page 72) form to exit from the loop prematurely. If the loop ends normally, `sample-special-form` returns `nil`.

For example:

```
(setq sample-variable 3)
(sample-special-form () form1 form2)
```

This evaluates *form1*, *form2*, *form1*, *form2*, *form1*, *form2* in that order.

Table 1-4: Sample Special Form Description

`sample-macro` *var* {*tag* | *statement*}* [*Macro*]

This evaluates the statements as a `prog` body, with the variable *var* bound to 43.

```
(sample-macro x (+ x x)) => 86
(sample-macro var . body) ==> (prog ((var 43)) . body)
```

Table 1-5: Sample Macro Description

Tables 1-4 and 1-5 illustrate the documentation of special forms and macros (which are closely related in purpose). These are very different from functions. Functions are called according to a single,specific, consistent syntax; the `&optional/&rest/&key` syntax specifies how the function uses its arguments internally, but does not affect the syntax of a call. In contrast, each special form or macro can have its own idiosyncratic syntax. It is by special forms and macros that the syntax of COMMON LISP is defined and extended.

In the description of a special form or macro, an italicized word names a corresponding part of the form that invokes the special form or macro. Parentheses ("(" and ")") stand for themselves, and should be written as such when invoking the special form or macro. Square brackets ("[" and "]") indicate that what they enclose is optional (may appear zero times or one time in that place); the square brackets should not be written in code. Curly braces ("{" and "}") simply parenthesize what they enclose, but may be followed by a star ("*") or a plus sign ("+"); a star indicates that what the braces enclose may appear any number of times (including zero, that is, not at all), while a plus sign indicates that what the braces enclose may appear any non-zero number of times (that is, must appear at least once). Within braces or brackets, vertical bars ("|") separate mutually exclusive choices.

In the last example in Table 1-5, notice the use of "dot notation". The "." appearing in the expression `(sample-macro` *var* `.` *body*) means that the name *body* stands for a list of forms, not just a single form, at the end of a list. This notation is often used in examples.

The term "LISP reader" refers not to you, the reader of this document, nor to some person reading LISP code, but specifically to a LISP program (the function `read` (page 237)) that reads characters from an input stream and interprets them by parsing as representations of LISP objects.

Certain characters are used in special ways in the syntax of COMMON LISP. The complete syntax is explained in detail in Chapter 21, but a quick summary here may be useful:

' An accent acute ("single quote") followed by an expression *form* is an abbreviation for `(quote` *form*`)`. Thus `'foo` means `(quote foo)` and `'(cons 'a 'b)` means `(quote (cons (quote a) (quote b)))`.

; Semicolon is the comment character. It and all characters up to the end of the line are discarded.

" Double quotes surround character strings: `"This is a thirty-nine character string."`.

\ Backslash is an escape character. As a rule, it causes the next character to be treated as a letter rather than for its usual syntactic purpose. For example, `A\(B` denotes a symbol whose name is "`A(B`", and `"\""` denotes a character string containing one character, a double-quote.

The number sign begins a more complex syntax. The next character designates the precise syntax to follow. For example, `#o105` means 105_8 (105 in octal notation); `#\L` denotes a character object for the character "`L`"; and `#(a b c)` denotes a vector of three elements `a`, `b`, and `c`. A particularly important case is that `#'`*fn* means `(function` *fn*`)`, in a manner analogous to `'`*form* meaning `(quote` *form*`)`.

| Vertical bars surround the name of a symbol that has special characters in it.

` Accent grave ("backquote") signals that the next expression is a template that may contain commas. The backquote syntax represents a program that will construct a data structure according to the template.

, Commas are used within the backquote syntax.

: Colon is used to indicate which package a symbol belongs to. For example, `chaos:reset` denotes the symbol named `reset` in the package named `chaos`. A leading colon indicates a *keyword*, a symbol that always evaluates to itself.

All code in this manual is written in lower case. COMMON LISP is generally insensitive to the case in which code is written. Internally, names of symbols are ordinarily converted to and stored in upper-case form. There are ways to force case conversion on output if desired. In this document, wherever an interactive exchange between a user and the LISP system is shown, the input is exhibited in lower case and the output in upper case.

Some symbols are written with the colon (`:`) character apparently in their names. In particular, all *keyword* symbols have names starting with a colon. The colon character is not actually part of the print name, but is a package prefix indicating that the symbol belongs to the keyword package. This is all explained in Chapter 11; until you read that, just make believe that the colons are part of the names of the symbols.

Chapter 2

Scope and Extent

In describing various features of the COMMON LISP language, the notions of *scope* and *extent* are frequently useful. These arise when some object or construct must be referred to from some distant part of a program. *Scope* refers to the spatial or textual region of the program within which references may occur. *Extent* refers to the interval of time within which references may occur.

As a simple example, consider this program:

```
(defun copycell (x) (cons (car x) (cdr x)))
```

The scope of the parameter named `x` is the body of the `defun` form. There is no way to refer to this parameter from any other place but within the body of the `defun`. Similarly, the extent of the parameter `x` (for any particular call to `copycell`) is the interval from the time the function is invoked to the time it is exited. (In the general case, the extent of a parameter may last beyond the time of function exit, but that cannot occur in this simple case.)

Within COMMON LISP, a referenceable entity is *established* by the execution of some language construct, and the scope and extent of the entity are described relative to the construct and the time (during execution of the construct) at which the entity is established. There are a few kinds of scope and extent that are particularly useful in describing COMMON LISP:

- *Lexical scope.* Here references to the established entity can occur only within certain program portions that are lexically (that is, textually) contained within the establishing construct. Typically the construct will have a part designated the *body*, and the scope of all entities established will be (or include) the body.

  Example: the names of parameters to a function normally are lexically scoped.

- *Local scope.* Here references to the established entity can occur only within certain program portions that are lexically (that is, textually) contained within the establishing construct, but moreover may *not* occur nested within certain other constructs, namely `function` (page 56), the definition portions of `flet` (page 67) and `labels` (page 67), and such function-defining constructs as `defun` (page 42), `deftype` (page 31), `defmacro` (page 91), and `defstruct` (page 199).

- *Indefinite scope.* References may occur anywhere, in any program.

- *Dynamic extent.* References may occur at any time in the interval between establishment of the entity and the explicit disestablishment of the entity. As a rule, the entity is disestablished when execution of the establishing construct completes or is otherwise terminated. Therefore entities with dynamic extent obey a stack-like discipline, paralleling the nested executions of their establishing constructs.

  Example: the `with-open-file` (page 267) creates opens a connection to a file and creates a stream object to represent the connection. The stream object has indefinite extent, but the connection to the open file has dynamic extent: when control exits the `with-open-file` construct, either normally or abnormally, the file is automatically closed.

  Example: the binding of a "special" variable has dynamic extent.

- *Indefinite extent.* The entity continues to exist so long as the possibility of reference remains. (An implementation is free to destory the entity if it can prove that reference to it is no longer possible.)

  Example: most COMMON LISP data objects have indefinite extent. (By contrast, the list produced for a `&rest` parameter in Lisp Machine LISP has dynamic extent [11].)

  Example: the names of lexically scoped parameters to a function have indefinite extent. (By contrast, in ALGOL the names of lexically scoped parameters to a procedure have dynamic extent.) This function definition:

```
(defun compose (f g)
  #'(lambda (x) (f (g x))))
```

  when given two arguments, immediately returns a function as its value. The parameter bindings for `f` and `g` do not disappear, because the returned function, when called, could still refer to those bindings. Therefore

```
(funcall (compose #'sqrt #'abs) -9.0)
```

  produces the value `3.0`. (An analogous procedure would not work correctly in typical ALGOL implementations.)

In addition, to the above terms, it is convenient to define *dynamic scope* to mean *indefinite scope and dynamic extent.* Thus we speak of "special" variables as having dynamic scope, or being dynamically scoped, because they have indefinite scope and dynamic extent: a special variable can be referred to anywhere as long as its binding is currently in effect.

The above definitions do not take into account the possibility of *shadowing.* Remote reference of entities is accomplished by using *names* of one kind or another. If two entities have the same name, then the second (say) may shadow the first, in which case an occurrence of the name will refer to the second and cannot refer to the first.

In the case of lexical or local scope, if two constructs that establish entities with the same name are textually nested, then references within the inner construct refer to the entity established by the inner one; the inner one shadows the outer one. Outside the inner one but inside the outer one, references refer to the entity established by the outer construct. For example:

```
(defun test (x z)
  (let ((z (* x 2))) (print z))
  z)
```

The binding of the variable z by the `let` (page 65) construct shadows the parameter binding for the function `test`. The reference to the variable z in the `print` form refers to the `let` binding. The reference to z at the end of the function refers to the parameter named z.

In the case of dynamic extent, if the time intervals of two entities with the same name overlap, then one interval will necessarily be nested within the other one (this is a property of the design of COMMON LISP). A reference will always refer to the entity that has been most recently established that has not yet been disestablished. For example:

```
(defun fun1 (x)
  (catch 'trap (+ 3 (fun2 x))))

(defun fun2 (y)
  (catch 'trap (* 5 (fun3 y))))

(defun fun3 (z)
  (throw 'trap z))
```

Consider the call `(fun1 7)`. The result will be `10`. At the time the `throw` (page 87) is executed, there are two outstanding catchers with the name `trap`: one established within procedure `fun1`, and the other within procedure `fun2`. The latter is the more recent, and so the value 7 is returned from the `catch` form in `fun2`. Viewed from within `fun3`, the `catch` in `fun2` shadows the one in `fun1`. (Had `fun2` been defined as

```
(defun fun2 (y)
  (catch 'snare (* 5 (fun3 y))))
```

then the two catchers would have different names, and therefore the one in `fun1` would not be shadowed. The result would then have been 7.)

As a rule this document will simply speak of the scope or extent of an entity; the possibility is shadowing will be left implicit.

A list of the important scope and extent rules in COMMON LISP:

- Variable bindings normally have lexical scope and indefinite extent.
- Variable bindings that are declared to be `special` have dynamic scope (indefinite scope and dynamic extent).
- A catcher established by a `catch` (page 85), `catch-all` (page 85), `unwind-all` (page 85), or `unwind-protect` (page 86) special form has dynamic scope.
- An exit point established by a `block` (page 71) construct has lexical scope and dynamic extent. (Such exit points are also established by `do` (page 73), `prog` (page 78), and other iteration constructs.)
- The tags established by a `prog` (page 78) and referenced by `go` (page 80) have lexical scope and

dynamic extent.

Constructs that use lexical scope effectively generate a new name for each established entity on each execution. Therefore dynamic shadowing cannot occur (though lexical shadowing may). This is of particular importance when dynamic extent is involved. For example:

```
(defun contorted-example (f g x)
  (if (= x 0)
      (funcall f)
      (block here
        (+ 5 (contorted-example g
                                #'(lambda () (return-from here 4))
                                (- x 1))))))
```

Consider the call `(contorted-example nil nil 2)`. This produces the result `4`. At the time the `funcall` is executed there are three `block` (page 71) exit points outstanding, each apparently named `here`. However, the `return-from` (page 72) form executed refers to the *outermost* of the outstanding exit points, not the innermost, as a consequence of the rules of lexical scoping: it refers to that exit point textually visible at the point the `function` (page 56) construct (here abbreviated by the `#'` syntax) was executed.

Chapter 3

Data Types

COMMON LISP provides a variety of types of data objects. It is important to note that in LISP it is data objects that are typed, not variables. Any variable can have any LISP object as its value. (It is possible to make an explicit declaration that a variable will in fact take on one of only a limited set of values. However, such a declaration may always be omitted, and the program will still run correctly. Such a declaration merely consititutes advice from the user that may be useful in gaining efficiency. See `declare` (page 95).)

In COMMON LISP, a data type is a (possibly infinite) set of LISP objects. Many LISP objects belong to more than one such set, and so it doesn't always make sense to ask what *the* type of an object is; instead, one usually asks only whether an object belongs to a given type. The predicate `typep` (page 46) may be used to ask either of these questions.

The data types defined in COMMON LISP are arranged into an almost-hierarchy (a hierarchy with shared subtrees) defined by the subset relationship. Certain sets of objects are interesting enough to deserve labels (such as the set of numbers or the set of strings). Symbols are used for most such labels (here, and throughout this document, the word *symbol* refers to atomic symbols, one kind of LISP object). See Chapter 4 for a complete description of type specifiers.

The root of the hierarchy, which is the set of all objects, is specified by the symbol `t`. The empty data type, which contains no objects, is denoted by `nil`.

COMMON LISP objects may be roughly divided into the following categories: numbers, characters, symbols, lists, arrays, structures, functions, and "random" objects. Some of these categories have many subdivisions. There are also standard types that are the union of two or more of these categories. The categories listed above, while they are data types, are neither more nor less "real" than other data types; they simply constitute a particularly useful slice across the type hierarchy for expository purposes.

Each of these categories is described briefly below. Then one section of this chapter is devoted to each, going into more detail, and briefly describing notations for objects of each type. Descriptions of LISP functions that operate on data objects are in later chapters.

- *Numbers* are provided in various forms and representations. COMMON LISP provides a true integer data type: any integer, positive or negative, has in principle a representation as a COMMON

LISP data object, subject only to total memory limitations (rather than machine word width). A true rational data type is provided: the quotient of two integers, if not an integer, is a ratio. Floating-point numbers of various ranges and precisions are also provided. Some implementations may choose to provide Cartesian complex numbers.

- *Characters* represent printed glyphs such as letters or text formatting operations. Strings are particular one-dimensional arrays of characters. COMMON LISP provides for a rich character set, including ways to represent characters of various type styles.

- *Symbols* (sometimes called *atomic symbols* for emphasis or clarity) are named data objects. LISP provides machinery for locating a symbol object, given its name (in the form of a string). Symbols have *property lists*, which in effect allow symbols to be treated as record structures with an extensible set of named components, each of which may be any LISP object.

- *Lists* are sequences represented in the form of linked cells called *conses*. There is a special object (the symbol `nil`) that is the empty list. All other lists are built recursively by adding a new element to the front of an existing list. This is done by creating a new *cons*, which is an object having two components called the *car* and the *cdr*. The *car* may hold anything, and the *cdr* is made to point to the previously existing list. (Conses may actually be used completely generally as two-element record structures, but their most important use is to represent lists.)

- *Arrays* are dimensioned collections of objects. An array can have any non-negative number of dimensions, and is indexed by a sequence of integers. General arrays can have any LISP object as a component; others are specialized for efficiency, and can hold only certain types of LISP objects. It is possible for two arrays, possibly with differing dimension information, to share the same set of elements (such that modifying one array modifies the other also).

- *Vectors* are a special class of arrays. They have exactly one dimension, and two vectors cannot have shared data. For critical applications in some implementations, vectors may be significantly more efficient than arrays. Two important special cases are *strings*, which are one-dimensional vectors of characters, and *bit-vectors*, which are vectors that can contain only the integers `0` and `1`.

- *Structures* are user-defined record structures, objects that have named components. The `defstruct` (page 199) facility is used to define new structure types. Some COMMON LISP implementations may choose to implement certain system-supplied data types as structures; these might include *bignums*, *readtables*, *streams*, *hashtables*, and *pathnames*.

- *Functions* are objects that can be invoked as procedures; these may take arguments, and return values. (All LISP procedures can be construed to return a value, and therefore treated as functions. Those that have nothing better to return usually return `nil`.) Such objects include *closures* (functions that have retained bindings from some environment) and *subrs* (compiled code objects). Some functions are represented as a list whose *car* is a particular symbol such as `lambda`. Symbols may also be used as functions.

- *Random* objects are those that do not fit into any other category. This is a catch-all data type that primarily covers implementation-dependent objects for internal use.

These categories are not always mutually exclusive. As noted above, an implementation may choose to implement certain kinds of objects (such as the more arcane numerical types) as structures. Every vector is an

array, though not every array is a vector.

3.1. Numbers

```
number
    rational
        integer
            fixnum
            bignum
        ratio
    float
        short-float
        single-float
        double-float
        long-float
    complex
```

Table 3-1: Hierarchy of Numeric Types

There are several kinds of numbers defined in COMMON LISP. Table 3-1 shows the hierarchy of number types.

3.1.1. Integers

The *integer* data type is intended to represent mathematical integers. Unlike most programming languages, COMMON LISP in principle imposes no limit on the magnitude of an integer; storage is automatically allocated as necessary to represent large integers.

In every COMMON LISP implementation there is a range of integers that are represented more efficiently than others; each such integer is called a *fixnum*, and an integer that is not a fixnum is called a *bignum*. The distinction between fixnums and bignums is visible to the user in only a few places where the efficiency of representation is important; in particular, it is guaranteed that the rank of an array, as well as any dimension of an array (and therefore any index into an array), can be represented as a fixnum. Exactly which integers are fixnums is implementation-dependent; typically they will be those integers in the range -2^n to 2^n-1, inclusive, for some n not less than 15. See `most-positive-fixnum` (page 142) and `most-negative-fixnum` (page 142).

Integers are ordinarily written in decimal notation, as a sequence of decimal digits, optionally preceded by a sign and optionally followed by a decimal point.

For example:

```
                0   ;Zero.
               -0   ;This always means the same as 0.
               +6   ;The first perfect number.
               28   ;The second perfect number.
             1024.  ;Two to the tenth power.
               -1   ;e^(πi)
155112100433309859840000000.  ;25 factorial (25!). Probably a bignum.
```

Compatibility note: MACLISP and Lisp Machine LISP normally assume that integers are written in *octal* (radix-8) notation unless a decimal point is present. INTERLISP assumes integers are written in decimal notation, and uses a trailing `Q` to indicate octal radix; however, a decimal point, even in trailing position, *always* indicates a floating-point number. This is of course consistent with FORTRAN; ADA does not permit trailing decimal points, but instead requires them to be embedded. In COMMON LISP, integers written as described above are always construed to be in decimal notation, whether or not the decimal point is present; allowing the decimal point to be present permits compatibility with MACLISP.

Integers may be notated in radices other than ten. The notation

#nnr ddddd or *#nnR ddddd*

means the integer in radix-*nn* notation denoted by the digits *ddddd*. More precisely, one may write "`#`", a non-empty sequence of decimal digits representing an unsigned decimal integer *n*, "`r`" (or "`R`"), an optional sign, and a sequence of radix-*n* digits, to indicate an integer written in radix *n* (which must be between 2 and 36, inclusive). Only legal digits for the specified radix may be used; for example, an octal number may contain only the digits 0 through 7. Letters of the alphabet of either case may be used in order for digits above 9. Binary, octal, and hexadecimal radices are useful enough to warrant the special abbreviations "`#b`" for "`#2r`", "`#o`" for "`#8r`", and "`#x`" for "`#16r`".

For example:

```
#2r11010101  ;Another way of writing 213 decimal.
 #b11010101  ;Ditto.
#b+11010101  ;Ditto.
      #o325  ;Ditto, in octal radix.
       #xD5  ;Ditto, in hexadecimal radix.
    #16r+D5  ;Ditto.
      #o-300 ;Decimal -192, written in base 8.
   #3r-12010 ;Same thing in base 3.
     #25R-7H ;Same thing in base 25.
```

3.1.2. Ratios

A *ratio* is a number representing the mathematical ratio of two integers. Integers and ratios are collectively called rationals. The canonical printed representation of a rational number is as an integer if its value is integral, and otherwise as the ratio of two integers, the *numerator* and *denominator*, whose greatest common divisor is one, and of which the denominator is positive (and in fact greater than 1, or else the value would be integral), written with "`/`" as a separator thus: "`3/5`". It is possible to notate ratios in non-canonical (unreduced) forms, such as "`4/6`", but the LISP function `prin1` (page 242) always prints the canonical form for a ratio.

Implementation note: While each implementation of COMMON LISP will probably choose to maintain all ratios in reduced form, there is no requirement for this as long as its effects are not visible to the user. Note that while it may at first glance appear to save computation for the reader and various arithmetic operations not to have to produce reduced forms, this savings is likely to be counteracted by the increased cost of operating on larger numerators and denominators.

Rational numbers may be written as the possibly signed quotient of decimal numerals: an optional sign followed by two non-empty sequences of digits separated by a "/". The second sequence may not consist entirely of zeros.

For example:

```
2/3                          ;This is in canonical form.
4/6                          ;A non-canonical form for the same number.
-17/23
-30517578125/32768           ;This is (-5/2)^15.
10/5                         ;The canonical form for this is 2.
```

To notate rational numbers in radices other than ten, one uses the same radix specifiers (one of #*nn*R, #O, #B, or #X) as for integers.

For example:

```
#o-101/75                    ;Octal notation for -65/61.
#3r120/21                    ;Ternary notation for 15/7.
#Xbc/ad                      ;Hexadecimal notation for 188/173.
```

3.1.3. Floating-point Numbers

Generally speaking, a floating-point number is a (mathematical) rational number of the form $(-1)^{s}*f*b^{e-p}$, where s is a bit (0 or 1), the *sign*; b is an integer greater than 1, the *base* or *radix* of the representation; p is a positive integer, the *precision* (in base-b digits) of the floating-point number; f is a positive integer between b^{p-1} and $b^{p}-1$ (inclusive), the *fraction* (properly speaking, the fraction is actually f/b^{p}); and e is an integer, the *exponent*. In addition, there is a floating-point zero. (Depending on the implementation, there may also be a "minus zero".) The value of p and the range of e depends on the implementation and on the type of floating-point number within that implementation.

> **Implementation note:** The form of the above description should not be construed to require the internal representation to be in sign-magnitude form. Two's-complement and other representations are also acceptable. Note that the radix of the internal representation may be other than 2, as on the IBM 360 and 370, which use radix 16; see `short-float-radix` (page 143) and friends.

Floating-point numbers may be provided in a variety of precisions and sizes, depending on the implementation. High-quality floating-point software tends to depend critically on the precise nature of the floating-point arithmetic, and so may not always be completely portable. To aid in writing programs that are moderately portable, however, certain definitions are made here:

- A *short* floating-point number is of the representation of smallest fixed precision provided by an implementation.
- A *long* floating-point number is of the representation of the largest fixed precision provided by an implementation.
- Intermediate between short and long formats are two others, arbitrarily called *single* and *double*.

The precise definition of these categories is implementation-dependent. However, the rough intent is that short floating-point numbers be precise at least to about five decimal places; single floating-point numbers, at least to about seven decimal places; and double floating-point numbers, at least to about fourteen decimal

places. Therefore the following minimum requirements are suggested for these formats: the precision (measured in "bits", computed as $p*\log_2 b$) and the exponent size (also measured in "bits", computed as the base-2 logarithm of one plus the maximum exponent value) must be at least as great as the values in Table 3-2.

Format	Minimum Precision	Minimum Exponent Size
Short	20 bits	7 bits
Single	24 bits	8 bits
Double	50 bits	8 bits

Table 3-2: Minimum Floating-Point Precision and Exponent Size Requirements

In any given implementation the categories may overlap or coincide. For example, short might mean the same as single, and long might mean the same as double.

> **Implementation note:** Where it is feasible, it is recommended that an implementation provide at least two types of floating-point number, and preferably three. Ideally, short-format floating-point numbers should have an "immediate" representation that does not require consing, single-format floating-point numbers should approximate IEEE proposed standard single-format floating-point numbers, and double-format floating-point numbers should approximate IEEE proposed standard double-format floating-point numbers [5, 2, 3].

Floating point numbers are written in either decimal fraction or "computerized scientific" notation: an optional sign, then a non-empty sequence of digits with an embedded decimal point, then an optional decimal exponent specification. The decimal point is required, and there must be digits either before or after it; moreover, digits are required after the decimal point if there is no exponent specifier. The exponent specifier consists of an exponent marker, an optional sign, and a non-empty sequence of digits. For preciseness, here is a modified-BNF decription of floating-point notation. The notation "{*x*}*" means zero or more occurrences of "x", the notation "{*x*}+" means one or more occurrences of "x", and the notation "[*x*]" means zero or one occurrences of "x".

floating-point-number ::= [*sign*] {*digit*}* . {*digit*}+ [*exponent*] | [*sign*] {*digit*}+ . {*digit*}* *exponent*
sign ::= + | -
digit ::= 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
exponent ::= *exponent-marker* [sign] {digit}+
exponent-marker ::= e | s | f | d | l | b | E | F | D | S | L | B

If no exponent specifier is present, or if the exponent marker "e" (or "E") is used, then the precise format to be used is not specified. When such a floating-point number representation is read and converted to an internal floating-point data object, the format specified by the variable `read-default-float-format` (page 237) is used; the initial value of this variable is `single`.

The letters "s", "f", "d", and "l" (or their respective upper-case equivalents) specify explicitly the use of *short*, *single*, *double*, and *long* format, respectively. The letters "b" and "B" are reserved for future definition. For example:

```
0.0                                ;Floating-point zero in default format.
-.0                                ;Also a floating-point zero.
0.                                 ;The integer zero, not a floating-point number!
0.0s0                              ;A floating-point zero in short format.
3.14159265358979323846d0           ;A double-format approximation to π.
6.02E+23                           ;Avogadro's number, in default format.
3.1010299957f-1                    ;log10 2, in single format.
-0.000000001s9                     ;e^(πi) in short format, the hard way.
```

3.1.4. Complex Numbers

Complex numbers may or may not be supported by a COMMON LISP implementation. They are represented in Cartesian form, with a real part and an imaginary part each of which is a non-complex number (integer, floating-point number, or ratio). It should be emphasized that the parts of a complex number are not necessarily floating-point numbers; in this COMMON LISP is like PL/I and differs from FORTRAN. In general, these identities hold:

```
(eql (realpart (complex x y)) x)
(eql (imagpart (complex x y)) y)
```

Complex numbers may be notated by writing the characters "`#C`" followed by a list of the real and imaginary parts. (Indeed, "`#C(`*a b*`)`" is equivalent to "`#,(complex` *a b*`)`"; see the description of the function `complex` (page 134).)

For example:

```
#C(3.0s1 2.0s-1)
#C(5 -3)                           ;A Gaussian integer.
#C(5/3 7.0)
#C(0 1)                            ;The imaginary unit.
```

??? Query: This notation is truly bletcherous. What would people think of adopting the notation suggested for APL, namely to write the real and imaginary parts separated by "`J`" (or "`j`")? The above examples would then be written as "`3.0s1j2.0s-1`", "`5j-3`", "`5/3J7.0`", and "`0J1`". Note particularly that the latter is a concise (three-character) notation for the imaginary unit *i*, much easier to type than "`#C(0 1)`".

Some implementations furthermore provide specialized representations of complex numbers for efficiency. In such representations the real part and imaginary part are of the same specialized numeric type. The "`#C`" construct will produce the most specialized representation that will correctly represent the two notated parts. The type of a specialized complex number is indicated by a list of the word `complex` and the type of the components; for example, a specialized representation for complex numbers with short floating-point parts would be of type `(complex short-float)`. The type `complex` encompasses all complex representations; the particular representation that allows parts of any numeric type is referred to as type `(complex t)`.

3.2. Characters

Every character object has three attributes: *code*, *bits*, and *font*. The code attribute is intended to distinguish among the printed glyphs and formatting functions for characters. The bits attribute allows extra flags to be associated with a character. The font attribute permits a specification of the style of the glyphs

(such as italics). Each of these attributes may be understood to be a non-negative integer.

A character object can be notated by writing "#\" followed by the character itself. For example, "#\g" means the character object for a lower-case "g". This works well enough for "printing characters". Non-printing characters have names, and can be notated by writing "#\" and then the name; for example, "#\return" (or "#\RETURN" or "#\Return", for example) means the <return> character. The syntax for character names after "#\" is the same as that for symbols.

The font attribute may be notated in unsigned decimal notation between the "#" and the "\". For example, #3\A means the letter "A" in font 3. Note that not all COMMON LISP implementations provide for non-zero font attributes; see `char-font-limit` (page 145).

The bits attribute may be notated by preceding the name of the character by the names or initials of the bits, separated by hyphens. The character itself may be written instead of the name, preceded if necessary by "\". For example:

```
#\Control-Meta-Return
#\Hyper-Space
#\Control-A
#\Meta-\β
#\C-M-Return
```

Note that not all COMMON LISP implementations provide for non-zero bits attributes; see `char-font-limit` (page 145).

Any character whose bits and font attributes are zero may be contained in strings. All such characters together constitute a subtype of the characters; this subtype is called `string-char`.

3.3. Symbols

Symbols are LISP data objects that serve several purposes and have several interesting characteristics. Every symbol has a name, called its *print name*, or *pname*. Given a symbol, one can obtain its name in the form of a string. More interesting, given the name of a symbol as a string one can obtain the symbol itself. (More precisely, symbols are organized into *packages*, and all the symbols in a package are uniquely identified by name.)

Symbols have a component called the *property list*, or *plist*. By convention this is always a list whose even-numbered components (calling the initial one component zero) are symbols, here functioning as property names, and whose odd-numbered components are associated property values. Functions are provided for manipulating this property list; in effect, these allow a symbol to be treated as an extensible record structure.

Symbols are also used to represent certain kinds of variables in LISP programs, and there are functions for dealing with the values associated with symbols in this role.

A symbol can be notated simply by writing its name. If its name is not empty, and if the name consists only of upper-case alphabetic, numeric, or certain "pseudo-alphabetic" special characters (but not delimiter characters such as parentheses or space), and if the name of the symbol cannot be mistaken for a number, then the symbol can be notated by the sequence of characters in its name.

For example:

```
FROBBOZ                 ;The symbol whose name is "FROBBOZ".
frobboz                 ;Another way to notate the same symbol.
fRObBoz                 ;Yet another way to notate it.
unwind-protect          ;A symbol with a "-" in its name.
+$                      ;The symbol named "+$".
1+                      ;The symbol named "1+".
+1                      ;This is the integer 1, not a symbol.
pascal_style            ;This symbol has an underscore in its name.
b^2-4*a*c               ;This is a single symbol!
                        ;    It has several special characters in its name.
file.rel.43             ;This symbol has periods in its name.
/usr/games/zork         ;This symbol has slashes in its name.
```

Besides letters and numbers, the following characters are normally considered to be "alphabetic" for the purposes of notating symbols:

```
+ - * / ! @ $ % ^ & _ = < > ? ~ .
```

Some of these characters have conventional purposes for naming things; for example, symbols that name functions having extremely implementation-dependent semantics generally have names beginning with "%". The last character, ".", is considered alphabetic *provided* that it does not stand alone. By itself, it has a role in the notation of conses. (It also serves as the decimal point.)

A symbol may have upper-case letters, lower-case letters, or both in its print name. However, the LISP reader normally converts lower-case letters to the corresponding upper-case letters when reading symbols. The net effect is that most of the time case makes no difference when *notating* symbols. However, case *does* make a difference internally and when printing a symbol. Internally the symbols that name all standard COMMON LISP functions, variables, and keywords have upper-case names; their names appear in lower case in this document for readability. Typing such names in lower case works because the function `read` will convert them to upper case.

If a symbol cannot be notated simply by the characters of its name, because the (internal) name contains special characters or lower-case letters, then there are two "escape" conventions for notating them. Writing a "\" character before any character causes the character to be treated itself as an ordinary character for use in a symbol name. If any character in a notation is preceded by \, then that notation can never be interpreted as a number.

For example:

```
\(                      ;The symbol whose name is "(".
\+1                     ;The symbol whose name is "+1".
+\1                     ;Also the symbol whose name is "+1".
\frobboz                ;The symbol whose name is "fROBBOZ".
3.14159265\s0           ;The symbol whose name is "3.14159265s0".
3.14159265\S0           ;The symbol whose name is "3.14159265S0".
3.14159265s0            ;A short-format floating-point approximation to π.
APL\\360                ;The symbol whose name is "APL\360".
apl\\360                ;Also the symbol whose name is "APL\360".
\(b^2\)\ -\ 4*a*c       ;The name is "(B^2) - 4*A*C".
                        ;   It has parentheses and two spaces in it.
```

It may be tedious to insert a "\" before *every* delimiter character in the name of a symbol if there are many of them. An alternative convention is to surround the name of a symbol with vertical bars; these cause every character between them to be taken as part of the symbol's name, as if "\" had been written before each one, excepting only | itself and \, which must nevertheless be preceded by \.

For example:

```
|"|                    ;The same as writing \".
|(b^2) - 4*a*c|        ;The name is "(b^2) - 4*a*c".
|frobboz|              ;The name is "frobboz", not "FROBBOZ".
|APL\360|              ;The name is "APL360", because
                       ;   the "\" quotes the "3".
|APL\\360|             ;The name is "APL\360".
|apl\\360|             ;The name is "apl\360".
|\|\||                 ;Same as \|\|: the name is "||".
```

3.4. Lists and Conses

A *cons* is a little record structure containing two components, called the *car* and the *cdr*. Conses are used primarily to represent lists.

A *list* is recursively defined to be either the empty list (which is represented by the symbol `nil`, but can also be written as "()") or a cons whose *cdr* component is a list. A list is therefore a chain of conses linked by their *cdr* components and terminated by `nil`. The *car* components of the conses are called the *elements* of the list. For each element of the list there is a cons. The empty list has no elements at all.

A list is notated by writing the elements of the list in order, separated by blank space (space, tab, or return characters) and surrounded by parentheses.

For example:

```
(a b c)                ;A list of three symbols.
(2.0s0 (a 1) #\*)      ;A list of three things: a short floating-point number,
                       ;   another list, and a character object.
```

This is why the empty list can be written as "()"; it is a list with no elements.

A *dotted list* is one whose last cons does not have `nil` for its *cdr*, but some other data object (which is also not a cons, or the first-mentioned cons would not be the last cons of the list). Such a list is called "dotted"

because of the special notation used for it: the elements of the list are written between parentheses as before, but after the last element and before the right parenthesis are written a dot (surrounded by blank space) and then the *cdr* of the last cons. As a special case, a single cons is notated by writing the car and the cdr between parentheses and separated by a space-surrounded dot.

For example:

```
(a . 4)          ;A cons whose car is a symbol
                 ;   and whose cdr is an integer.
(a b c . d)      ;A list with three elements whose last cons
                 ;   has the symbol d in its cdr.
```

Compatibility note: In MACLISP, the dot in dotted-list notation needed not be surrounded by white space or other delimiters. The dot is required to be delimited in Lisp Machine LISP.

It is legitimate to write something like `(a b . (c d))`; this means the same as `(a b c d)`. The standard LISP output routines will never print a list in the first form, however; they will avoid dot notation wherever possible.

Often the term *list* is used to refer either to true lists or to dotted lists. The term "true list" will be used to refer to a list terminated by `nil`, when the distinction is important. Most functions advertised to operate on lists will work on dotted lists and ignore the non-`nil` *cdr* at the end.

Sometimes the term *tree* is used to refer to some cons and all the other conses transitively accessible to it through *car* and *cdr* links until non-conses are reached; these non-conses are called the *leaves* of the tree.

Lists, dotted lists, and trees are not mutually exclusive data types; they are simply useful points of view about structures of conses. There are yet other terms, such as *association list*. None of these are true LISP data types. Conses are a data type, and `nil` is the sole object of type `null`. The LISP data type `list` is taken to mean the union of the `cons` and `null` data types, and therefore encompasses both true lists and dotted lists.

3.5. Arrays

An *array* is an object with components arranged according to a rectilinear coordinate system. In general, these components may be any LISP data objects.

The number of dimensions of an array is called its *rank* (this terminology is borrowed from APL). This is a non-negative integer; for convenience, it is in fact required to be a fixnum (an integer of limited magnitude). Likewise, each dimension has a length that is a non-negative fixnum. The total number of elements in the array is the product of all the dimensions.

It is permissible for a dimension to be zero. In this case, the array has no elements, and any attempt to access an element in in error. However, other properties of the array (such as the dimensions themselves) may be used. If the rank is zero, then there are no dimensions, and the product of the dimensions is then by definition 1. A zero-rank array therefore has a single element.

An array element is specified by a sequence of indices. The length of the sequence must equal the rank of the array. Each index must be a non-negative integer strictly less than the corresponding array dimension. Array indexing is therefore zero-origin, not one-origin as in (the default case of) FORTRAN.

As an example, suppose that the variable `foo` names a 3-by-5 array. Then the first index may be 0, 1, or 2, and then second index may be 0, 1, 2, 3, or 4. One may refer to array elements using the function `aref` (page 185):

```
(aref foo 2 1)
```

refers to element (2, 1) of the array. Note that `aref` takes a variable number of arguments: an array, and as many indices as the array has dimensions. A zero-rank array has no dimensions, and therefore `aref` would take such an array and no indices, and return the sole element of the array.

One-dimensional arrays and lists are collectively considered to be *sequences*. They differ in that any component of a one-dimensional array can be accessed in constant time, while the average component access time for a list is linear in the length of the list; on the other hand, adding a new element to the front of a list takes constant time, while the same operation on an array takes time linear in the length of the array.

In general, arrays can be multi-dimensional, can have *fill pointers*, can share their contents with other array objects, and can have their size altered dynamically after creation.

Multidimensional arrays store their components in row-major order; that is, internally a multidimensional array is stored as a one-dimensional array, with the multidimensional index sets ordered lexicographically, last index varying fastest. This is important in two situations: (1) when arrays with different dimensions share their contents, and (2) when accessing very large arrays in virtual-memory implementation. (The first situation is semantic, the second pragmatic.)

If for some purpose an array is needed that is one-dimensional, unshared with any other array, and is not to have its size increased later, one may request that a *vector* be created. A vector is a limited kind of array. Some implementations can handle vectors in an especially efficient manner. Any operation that works for an array works on a vector, but certain operations such as `vref` (page 187) operate only on vectors and may therefore be made more efficient. Moreover, vectors may have a more compact representation than typical arrays.

A general vector (a one-dimensional array of S-expressions with no additional paraphernalia) can be notated by notating the components in order, separated by whitespace and surrounded by "`#(`" and "`)`". For example:

```
#(a b c)                          ;A vector of length 3.
#(2 3 5 7 11 13 17 19 23 29 31 37 41 43 47)
                                  ;A vector containing the primes below 50.
#()                               ;An empty vector.
```

Rationale: Numerous people have suggested that square brackets be used to notate vectors: "`[a b c]`" instead of "`#(a b c)`". This would be shorter, perhaps more readable, and certainly in accord with cultural conventions in other parts of computer science and mathematics. However, to preserve the usefulness of the user-definable macro-character feature of

the function `read` (page 237), it is necessary to leave some characters to the user for this purpose. Experience in MACLISP has shown that users, especially implementors of AI languages, often want to define special kinds of brackets. Therefore COMMON LISP avoids using these characters in its syntax so that the user may freely redefine their syntax: "`[]{}!?`".

Implementations may provide certain specialized representations of arrays for efficiency in the case where all the components are of the same specialized (typically numeric) type. All implementations provide specialized arrays for the cases when the components are characters or when the components are always 0 or 1; the one-dimensional instances of these specializations are respectively called *strings* and *bit-vectors*. Special notations are provided for the further restriction of these types to the vector case. A string vector can be written as the sequence of characters contained in the string, preceded and followed by a "`"`" (double-quote) character. Any "`"`" or "`\`" character in the sequence must additionally have a "`\`" character before it.

For example:

```
"Foo"                          ;A string with three characters in it.
""                             ;An empty string.
"\"APL\\360?\" he cried."      ;A string with twenty characters.
"|x| = |-x|"                   ;A ten-character string.
```

Notice that any vertical bar "`|`" in a string need not be preceded by a "`\`". Similarly, any double-quote in the name of a symbol written using vertical-bar notation need not be preceded by a "`\`". The double-quote and vertical-bar notations are similar but distinct: double-quotes indicate a character string containing the sequence of characters, while vertical bars indicate a symbol whose name is the contained sequence of characters.

A bit-vector is written much like a string, using double-quotes; however, a "`#`" is written before it, and the elements of the bit vector must be 0 or 1.

For example:

```
#"10110"                       ;A bit vector with five bits. Bit 0 is 1.
#""                            ;A null bit vector.
#"110101000101000101"          ;Bit n of this bit vector is 1 iff n+2 is prime.
```

3.6. Structures

Different structures may print out in different ways; the definition of a structure type may specify a print procedure to use for objects of that type (see the `:printer` (page DEFSTRUCT-PRINTER-KWD) option to `defstruct` (page 199)). The default notation for structures is:

#S(*structure-name*
    slot-name-1 *slot-value-1*
    slot-name-2 *slot-value-2*
    ...)

where "`#S`" indicates structure syntax, *structure-name* is the name (a symbol) of the structure type, each *slot-name* is the name (also a symbol) of a component, and each corresponding *slot-value* is the representation of the LISP object in that slot.

3.7. Functions

A *function* is anything that may be correctly given to the `funcall` (page 64) or `apply` (page 63) function, to be executed as code when arguments are supplied.

A *subr* (pronounced "subber") is a compiled code object. A *closure* is an object that represents an inner function together with environmental information about variable bindings of indefinite extent to which the function may refer.

A list whose *car* is `lambda` or `select` may serve as a function; see Chapter 5.

A symbol may serve as a function; an attempt to invoke a symbol as a function causes the contents of the symbol's function cell to be used. See `fsymeval` (page 57).

3.8. Randoms

Objects of type `random` tend to have implementation-dependent semantics, and so may print in implementation-dependent ways. As a rule, such objects cannot reliably be reconstructed from a printed representation, and so they are printed usually in a format informative to the user but not acceptable to the `read` function:

#<useful information>

A hypothetical example might be:

```
#<stack-pointer si:rename-within-new-definition-maybe 311037552>
```

The LISP reader will signal an error on encountering "#<".

It is not necessarily the case that all objects that are printed in the form "#<...>" are of type `random`; however, any object of type `random` will be printed in that form.

Chapter 4

Type Specifiers

In COMMON LISP, types are named by LISP objects, specifically symbols and lists, called *type specifiers*. Symbols name predefined classes of objects, while lists usually indicate combinations or specializations of simpler types. Symbols or lists may also be abbreviations for types that could be specified in other ways.

4.1. Type Specifier Symbols

The type symbols defined by the system include those shown in Table 4-1. In addition, when a structure type is defined using `defstruct` (page 199), the name of the structure type becomes a valid type symbol.

If a type specifier is a list, the *car* of the list is a symbol, and the rest of the list is subsidiary type information. As a general convention, any subsidiary item may be replaced by *, or simply omitted if it is the last item of the list; in any of these cases the item is said to be unspecified.

??? **Query:** Formerly ? was used to indicate an unspecified item, but that conflicted with the convention that the characters "`!?[]{}`" should be reserved to the user for possible use as macro characters. Is this change satisfactory?

4.2. Type Specifiers That Combine

The following type specifier lists define a data type in terms of other types or objects.

`null`	`cons`	`list`	`symbol`
`vector`	`string`	`bit-string`	`array`
`function`	`sequence`	`random`	`character`
`number`	`stream`	`float`	`string-char`
`integer`	`fixnum`	`bignum`	`bit`
`short-float`	`single-float`	`double-float`	`long-float`
`complex`	`ratio`	`readtable`	`package`
`subr`	`closure`		

Table 4-1: Standard Type Specifier Symbols

`(oneof` *object1 object2* `...)`

This denotes the set containing precisely those objects named. An object is of this type if and only if it is `eql` (page 49) to one of the specified objects.

> **Compatibility note:** This is approximately equivalent to what the INTERLISP DECL package calls `memq`. What INTERLISP calls `oneof`, COMMON LISP calls `or` (see below).

`(not` *type*`)` This denotes the set of all those objects that are *not* of the specified type.

`(or` *type1 type2* `...)`
This denotes the union of the specified types. For example, the type `list` by definition is the same as `(or null cons)`. Also, the value returned by the function `position` (page 163) is always of type `(or null (integer 0 *))` (either `nil` or a non-negative integer).

> **Compatibility note:** This is equivalent to what the INTERLISP DECL package calls `oneof`.

`(and` *type1 type2* `...)`
This denotes the intersection of the specified types.

> **Compatibility note:** This is equivalent to what the INTERLISP DECL package calls `allof`.

4.3. Type Specifiers That Specialize

Some type specifier lists denote *specializations* of data types named by symbols. These specializations may be reflected by more efficient representations in the underlying implementation. As an example, consider the type `(array short-float)`. Implementation A may choose to provide a specialized representation for arrays of short floating-point numbers, and implementation B may choose not to.

If you should want to create a array for the express purpose of holding only short-float objects, you may optionally specify to `make-array` (page 183) the element type `short-float`. This does not *require* `make-array` to create an object of type `(array short-float)`; it merely *permits* it. The request is construed to mean "Produce the most specialized array representation capable of holding short-floats that the implementation can provide." Implementation A will then produce a specialized short-float array (of type `(array short-float)`), and implementation B will produce an ordinary array (one of type `(array t)`).

If one were then to ask whether the array were actually of type `(array short-float)`, implementation A would say "yes", but implementation B would say "no". This is a property of `make-array` and similar functions: what you ask for is not necessarily what you get.

Types can therefore be used for two different purposes: *declaration* and *discrimination*. Declaring to `make-array` that elements will always be of type `short-float` permits optimization. Similarly, declaring that a variable takes on values of type `(array short-float)` amounts to saying that the variable will take on values that might be produced by specifying element type `short-float` to `make-array`. On the other hand, if the predicate `typep` is used to test whether an object is of type `(array short-float)`, only objects actually of that specialized type can satisfy the test; in implementation B no object can pass that test.

The valid list-format names for data types are:

`(array` *type dimensions*`)`

This denotes the set of specialized arrays whose elements are all members of the type *type* and whose dimensions match *dimensions*. For declaration purposes, this type encompasses those arrays that can result by specifying *type* as the element type to the function `make-array` (page 183); this may be different from what the type means for discrimination purposes. *type* must be a valid type specifier or unspecified. *dimensions* may be a non-negative integer, which is the number of dimensions, or it may be a list of non-negative integers representing the length of each dimension (any dimension may be unspecified instead), or it may be unspecified.

For example:

```
(array integer 3)          ;Three-dimensional arrays of integers.
(array integer (* * *))    ;Three-dimensional arrays of integers.
(array * (4 5 6))          ;4-by-5-by-6 arrays.
(array character (3 *))    ;Two-dimensional arrays of characters
                           ;    that have exactly three rows.
(array short-float ())     ;Zero-rank arrays of short floating-point numb
```

`(vector` *type size*`)`

This denotes the set of specialized vectors whose elements are all members of the type *type* and whose lengths match *size*. For declaration purposes, this type encompasses those vectors that can result by specifying *type* as the element type to the function `make-vector` (page 185); this may be different from what the type means for discrimination purposes. *type* must be a valid type specifier or unspecified. *size* may be a non-negative integer or unspecified.

For example:

```
(vector double-float)      ;Vectors of double-format floating-point numb
(vector * 5)               ;Vectors of length 5.
(vector t 5)               ;General vectors of length 5.
(vector (mod 32) *)        ;Vectors of integers between 0 and 31.
```

Note that `(vector t 5)` is a subset of `(vector * 5)`.

The specialized types `(vector string-char)` and `(vector bit)` are so useful that they have the special names `string` and `bit-string`; every COMMON LISP implementation must provide distinct representations for these as distinct specialized data types.

Rationale: NIL had been using the name `bits` for a bit vector. This tended to lead to awkward prose: one had to speak of "a `bits`". The singular noun `bit-vector` is easier to discuss.

`(complex` *rtype itype*`)`

Every element of this type is a complex number whose real part is of type *rtype* and whose imaginary part is of type *itype*. For declaration purposes, this type encompasses those complex numbers that can result by giving numbers of the specified type to the function `complex` (page 134); this may be different from what the type means for discrimination purposes.

In a break with the usual convention on omitted items, if *itype* is omitted (but not if it is explicitly unspecified) then it is taken to be the same as *rtype*. As examples, Gaussian

integers might be described as (`complex integer`), and the result of the complex logarithm function might be described as being of type (`complex float (float #.(- pi) #.pi)`).

(`function` (*arg1-type arg2-type ...*) *value1-type value2-type ...*)
This type may be used only for declaration and not for discrimination; `typep` (page 46) will signal an error if it encounters a specifier of this form. Every element of this type is a function that accepts arguments at *least* of the types specified by the *argj-type* forms, and returns values that are members of the types specified by the *valuej-type* forms. The `&optional`, `&rest`, and `&key` keywords may appear in *either* list of types; in the list of values, they indicate the parameter list of another function that, when given to `mvcall` (page 82) along with the values, woiuld be suitable for receiving those values. As an example, the function `cons` (page 168) is of type (`function (t t) cons`), because it can accept any two arguments and always returns a cons. It is also of type (`function (float string) list`), because it can certainly accept a floating-point number and a string (among other things), and its result is always of type `list` (in fact a `cons` and never `null`, but that does not matter for this type declaration).

4.4. Type Specifiers That Abbreviate

The following type specifiers are, for the most part, abbreviations for other type specifiers that would be far too verbose to write out explicitly (using, for example, `oneof`).

(`integer` *low high*)
This denotes the integers between *low* and *high*. The limits *low* and *high* must each be an integer, a list of an integer, or unspecified. An integer is an inclusive limit, a list of an integer is an exclusive limit, and * means that a limit does not exist and so effectively denotes minus or plus infinity, respectively. The type `fixnum` is simply a name for (`integer` *smallest largest*) for implementation-dependent values of *smallest* and *largest*. The type (`integer 0 1`) is so useful that it has the special name `bit`.

(`mod` *n*) The set of non-negative integers less than *n*. This is equivalent to (`integer 0` $n-1$) or to (`integer 0 (`*n*`)`).

(`signed-byte` *s*)
The set of integers that can be represented in two's-complement form in a byte of *s* bits. This is equivalent to (`integer` -2^{s-1} $2^{s-1}-1$).

(`unsigned-byte` *s*)
The set of non-negative integers that can be represented in a byte of *s* bits. This is equivalent to (`mod` 2^s), that is, (`integer 0` 2^s-1).

(`rational` *low high*)
This denotes the rationals between *low* and *high*. The limits *low* and *high* must each be a rational, a list of a rational, or unspecified. A rational is an inclusive limit, a list of a rational is an exclusive limit, and * means that a limit does not exist and so effectively denotes minus or plus infinity, respectively.

`(float` *low high*`)`

The set of floating-point numbers between *low* and *high*. The limits *low* and *high* must each be a floating-point number, a list of a floating-point number, or unspecified; a floating-point number is an inclusive limit, a list of a floating-point number is an exclusive limit, and * means that a limit does not exist and so effectively denotes minus or plus infinity, respectively.

In a similar manner one may use:

```
(short-float low high)
(single-float low high)
(double-float low high)
(long-float low high)
```

In this case, if a limit is a floating-point number (or a list of one), it must be one of the appropriate format.

`(string` *size*`)` This means the same as `(vector string-char` *size*`)`: the set of strings of the indicated size.

`(bit-vector` *size*`)`

This means the same as `(vector bit` *size*`)`: the set of bit-vectors of the indicated size.

4.5. Defining New Type Specifiers

New type specifiers can come into existence in two ways. First, defining a new structure type with `defstruct` (page 199) automatically causes the name of the structure to be a new type specifier symbol. Second, the `deftype` special form can be used to declare new abbreviations.

`deftype` *name varlist* {*form*}* [*Special form*]

This is very similar to a `defmacro` (page 91) form: *name* is the symbol that identifies the type specifier being defined, *varlist* is similar in form to a lambda-list (and may contain `&optional` and `&rest` tokens), and *body* is the body of the expander function. If we view a type specifier list as a list containing the type specifier name and some argument forms, the argument forms (unevaluated) are bound to the corresponding parameters in *varlist*. Then the body forms are evaluated as an implicit `progn`, and the value of the last form is interpreted as a new type specifier for which the original specifier was an abbreviation.

`deftype` differs from `defmacro` in that if no *initform* is specified for an `&optional` parameter, the default value is *, not `nil`.

For example:

```
(deftype mod (n) (integer 0 (,n)))
(deftype list () '(or null cons))
(deftype square-matrix (&optional type size)
  (array ,type (,size ,size)))

(square-matrix short-float 7)  means  (array short-float (7 7))
(square-matrix bit)  means  (array bit (* *))
```

If the type name defined by `deftype` is used simply as a type specifier symbol, it is interpreted as a type specifier list with no argument forms. Thus, in the example above, `square-matrix` would mean `(array * (* *))`, the set of two-dimensional arrays. This would unfortunately fail to convey the constraint that the two dimensions be the same; `(square-matrix bit)` has the same problem. This is an inherent limitation of the type definition system in COMMON LISP.

??? **Query:** Can this be fixed without too much hair? Should we have the INTERLISP `satisfies` clause?

Chapter 5

Program Structure

In the previous chapter the syntax was sketched for notating data objects in COMMON LISP. The same syntax is used for notating programs, because all COMMON LISP programs have a representation as COMMON LISP data objects.

5.1. Forms

The standard unit of interaction with a COMMON LISP implementation is the *form*, which is simply an S-expression meant to be *evaluated* as a program to produce one or more *values* (which are also data objects). One may request evaluation of *any* data object, but only certain ones (such as symbols and lists) are meaningful forms, while others (such as most arrays) are not. Examples of meaningful forms are `3`, whose value is `3`, and `(+ 3 4)`, whose value is `7`. We write "`3 => 3`" and "`(+ 3 4) => 7`" to indicate these facts ("`=>`" means "evaluates to").

Meaningful forms may be divided into three categories: self-evaluating forms, such as numbers; symbols, which stand for variables; and lists. The lists in turn may be divided into three categories: special forms, macro calls, and function calls.

5.1.1. Self-Evaluating Forms

All numbers, strings, and bit-vectors are *self-evaluating* forms. When such an object is evaluated form, that object itself (or possibly a copy in the case of numbers) is returned as the value of the form. The empty list `()`, which is also the false value `nil`, is also a self-evaluating form: the value of `nil` is `nil`. Keywords (symbols written with a leading colon) also evaluate to themselves: the value of `:start` is `:start`.

5.1.2. Variables

Symbols are used as names of variables in COMMON LISP programs. When a symbol is evaluated as a form, the value of the variable it names is produced. For example, after doing `(setq items 3)`, which assigns the value 3 to the variable named `items`, then `items => 3`. Variables can be *assigned* to (as by `setq` (page 58)) or *bound*. Any program construct that binds a variable effectively saves the old value of the variable and causes it to have a new value, and on exit from the construct the old value is reinstated.

There are actually two kinds of variables in COMMON LISP, called *lexical* (or *static*) variables and *special* (or *dynamic*) variables. At any given time either or both kinds of variable with the same name may have a current value. Which of the two kinds of variable is referred to when a symbol is evaluated depends on the context of the evaluation. The general rule is that if the symbol occurs textually within a program construct that creates a *binding* for a variable of the same name, then the reference is to the kind of variable specified by the binding; if no such program construct textually contains the reference, then it is taken to refer to the special variable of that name.

The distinction between the two kinds of variable is one of scope and access. A lexically bound variable can be referred to *only* by forms occurring at any *place* textually within the program construct that binds the variable. A dynamically bound (special) variable can be referred to at any *time* from the time the binding is made until the time evaluation of the construct that binds the variable terminates. Therefore lexical binding imposes spatial limitations on occurrences of references, whereas dynamic binding imposes temporal limitations.

The value a special variable has when there are currently no bindings of that variable is called the *global* value of the variable. A global value can be given to a variable only by assignment, because a value given by binding by definition is not global.

The symbols `t` and `nil` are reserved. One may not assign a value to `t` or `nil`, and one may not bind `t` or `nil`. The global value of `t` is always `t`, and the global value of `nil` is always `nil`. Constant symbols defined by `defconst` (page 44) also become reserved and may not be further assigned to or bound.

Rationale: It would seem appropriate for the compiler to be justified in issuing a warning if one does a `setq` on a constant defined by `defconst`. If one cannot assign, one should not be able to bind, either.

5.1.3. Special Forms

If a list is to be evaluated as a form, the first step is to examine the first element of the list. If the first element is one of the symbols appearing in Table 5-1, then the list is called a *special form*. (This use of the word "special" is unrelated to its use in the phrase "special variable".)

Special forms are generally environment and control constructs. Every special form has its own idiosyncratic syntax. An example is the `if` special form: "`(if p (+ x 4) 5)`" in COMMON LISP means what "**if** *p* **then** $x+4$ **else** 5" would mean in ALGOL.

The evaluation of a special form normally produces a value (but it may instead call for a non-local exit (see `throw` (page 87)) or produce no values or more than one value (see `values` (page 82))).

The set of special forms is fixed in COMMON LISP; no way is provided for the user to define more. The user can create new syntactic constructs, however, by defining macros.

An implementation is free to implement as a macro any construct described herein as being a special form. Conversely, an implementation is free to implement as a special form any construct described herein as being

defun (page 42)	go (page 80)
defvar (page 43)	return (page 72)
defconst (page 44)	return-from (page 72)
and (page 52)	multiple-value-list (page 82)
or (page 52)	
quote (page 56)	mvcall (page 82)
function (page 56)	mvprog1 (page 82)
setq (page 58)	
psetq (page 58)	
progn (page 64)	multiple-value-bind (page 82)
prog1 (page 65)	multiple-value (page 83)
prog2 (page 65)	catch (page 85)
let* (page 66)	catch-all (page 85)
progv (page 67)	unwind-all (page 85)
cond (page 68)	unwind-protect (page 86)
if (page 69)	throw (page 87)
when (page 69)	declare (page 95)
unless (page 70)	locally (page 96)
case (page 70)	the (page 99)
typecase (page 70)	
do (page 73)	do-symbols (page 116)
do* (page 75)	do-external-symbols (page 116)
dolist (page 76)	do-internal-symbols (page 116)
dotimes (page 76)	do-all-symbols (page 116)
prog (page 78)	with-open-file (page 267)
prog* (page 80)	condition-bind (page 272)

(The page numbers indicate where the definitions of these special forms appear.)

Table 5-1: Names of All COMMON LISP Special Forms

a macro, provided that an equivalent macro definition is also provided.

5.1.4. Macros

If a form is a list and the first element is not the name of a special form, it may be the name of a *macro*; if so, the form is said to be a *macro call.* A macro is essentially a function from forms to forms that will, given a call to that macro, compute a new form to be evaluated in place of the macro call. (This computation is sometimes referred to as *macro expansion.*) For example, the macro named `push` (page 172) will take a form such as `(push x stack)` and from that form compute a new form `(setf stack (cons x stack))`. We say that the old form *expands* into the new form. The new form is then evaluated in place of the original form; the value of the new form is returned as the value of the original form.

There are a number of standard macros in COMMON LISP, and the user can define more by using `defmacro` (page 91).

Macros provided by a COMMON LISP implementation as described herein may expand into code that is not portable among differing implementations. That is, a macro call may be implementation-independent by

virtue of being so defined in this document, but the expansion need not be.

5.1.5. Function Calls

If a list is to be evaluated as a form and the first element is not a symbol that names a special form or macro, then the list is assumed to be a *function call.* The first element of the list is taken to name a function. Any and all remaining elements of the list are forms to be evaluated; one value is obtained from each form, and these values become the *arguments* to the function. The function is then *applied* to the arguments. The functional computation normally produces a value (but it may instead call for a non-local exit (see `throw` (page 87)) or produce no values or more than one value (see `values` (page 82))). If and when the function returns, whatever value(s) it returns becomes the value(s) of the function-call form.

For example, consider the evaluation of the form `(+ 3 (* 4 5))`. The symbol + names the addition function, not a special form or macro. Therefore the two forms `3` and `(* 4 5)` are evaluated to produce arguments. The form `3` evaluates to 3, and the form `(* 4 5)` is a function call (to the multiplication function). Therefore the forms `4` and `5` are evaluated, producing arguments `4` and `5` for the multiplication. The multiplication function calculates the number `20` and returns it. The values `3` and `20` are then given as arguments to the addition function, which calculates and returns the number `23`. Theerfore we say `(+ 3 (* 4 5)) => 23`.

5.2. Functions

There are two ways to indicate a function to be used in a function call form. One is to use a symbol that names the function. This use of symbols to name functions is completely independent of their use in naming special and lexical variables. The other way is to use a *lambda-expression,* which is a list whose first element is the symbol `lambda`. A *lambda-expression* is *not* a form; it cannot be meaningfully evaluated. Lambda-expressions and symbols as names of functions can appear only as the first element of a function-call form, or as the second element of the `function` (page 56) special form.

5.2.1. Named Functions

A name can be given to a function in one of two ways. A *global name* can be given to a function by using the `defun` (page 42) special form. A *local name* can be given to a function by using the `labels` (page 67) special form. If a symbol appears as the first element of a function-call form, then it refers to the definition established by the innermost `labels` construct that textually contains the reference, or if to the global definition (if any) if there is no such containing `labels` construct.

When a function is named, a lambda-expression is associated with that name (in effect). See `defun` (page 42) and `labels` (page 67) for an explanation of these lambda-expressions.

5.2.2. Lambda-Expressions

A *lambda-expression* is a list with the following syntax:

(`lambda` *lambda-list* . *body*)

The first element must be the symbol `lambda`. The second element must be a list. It is called the *lambda-list*, and specifies names for the *parameters* of the function. When the function denoted by the lambda-expression is applied to arguments, the arguments are matched with the parameters specified by the lambda-list. The *body* may then refer to the arguments by using the parameter names. The *body* consists of any number of forms (possibly zero). These forms are evaluated in sequence, and the value(s) of the *last* form only are returned as the value(s) of the application (the value `nil` is returned if there are zero forms in the body).

The complete syntax of a lambda-expression is:

```
(lambda ({var}*
         {[&optional {var | (var [initform [svar]])}*] [&rest var]
          | [&key {var | ({var | (keyword var)} [initform [svar]])}*]}
         [&aux {var | (var [initform])}*])
  {(declare {declaration}*)}*
  {form}*)
```

Each element of a lambda-list is either a *parameter specifier* or a *separator token*; separator tokens begin with "&". In all cases *var* must be a symbol, the name of a variable, and similarly for *svar* also; each *keyword* must be a keyword symbol. An *initform* may be any form.

A lambda-list has three parts, any or all of which may be empty:

- Specifiers for the *required* parameters. These are all the parameter specifiers up to the first separator token; if there is no such token, then all the specifiers are for required parameters.
- Either *optional* and *rest* parameters or *keyword* parameters (but not both).
  - If the token `&optional` is present, the *optional* parameter specifiers are those following the token `&optional` up to the next separator token or the end of the list. Following or instead of the `&optional` token and its following specifiers may be the token `&rest` followed by a single *rest* parameter specifier.
  - If the token `&key` is present, all specifiers up to the next separator token (which in this case must be `&aux`) or the end of the list are *keyword* parameter specifiers.
- If the token `&aux` is present, all specifiers after it are *auxiliary variable* specifiers.

Compatibility note: What is provided here is a subset of the functionality currently provided in Lisp Machine LISP. The principal restrictions here are:

- Keyword parameters may not be mixed with (positional) optional and rest parameters. The rationale for not mixing keyword parameters and positional optionals is that it would be very awkward to define a function in such a way that one could not specify any keyword parameters unless all positional optionals were specified. If the positional ones are to be non-trivially optional, then all the keyword parameters should also be optional, and as a matter of style it would be better for all the optional parameters to have keywords. (We know how to make interleaved required and optional positional parameters work, too, but as a matter of style we only allow optionals to follow required.) The rationale for not mixing keyword and rest parameters is less strong, and motivated primarily by a feeling of

awkwardness in letting more than one parameter receive the same argument. If we allow that, then why not (`&rest x a b &optional c d`)? There may be aliasing problems: can we guarantee, if a parameter is `setq`'d, that the corresponding part of a `&rest` list will or will not be correspondingly changed?

- No keyword argument may be provided for which there is no matching keyword parameter. This is a logical consequence of not mixing keyword and rest parameters, and also greatly improves program readability: the lambda-list enumerates all relevant keywords. Is non-trivial use made of `&allow-extra-keywords` in Lisp Machine LISP?

How do people feel about this? Lisp Machine LISP will run correct programs constructed according to the above specifications; it is a superset.

When the function represented by the lambda-expression is applied to arguments, the arguments and parameters are processed in order from left to right. In the simplest case, only required parameters are present in the lambda-list; each is specified simply by a name *var* for the parameter variable. When the function is applied, there must be exactly as many arguments as there are parameters, and each parameter is bound to one argument. Here, and in general, the parameter is bound as a lexical variable unless a declaration has been made that it should be a special binding (see `declare` (page 95)).

In the more general case, if there are *n* required parameters (*n* may be zero), there must be at least *n* arguments, and the required parameters are bound to the first *n* arguments. The other parameters are then processed using any remaining arguments.

If *optional* parameters are specified, then each one is processed as follows. If any unprocessed arguments remain, then the parameter variable *var* is bound to the next remaining argument, just as for a required parameter. If no arguments remain, however, the *initform* part of the parameter specifier is evaluated, and the parameter variable is bound to the resulting value (or to `nil` if no *initform* appears in the parameter specifier). If another variable name *svar* appears in the specifier, it is bound to *true* if an argument was available, and to *false* if no argument remained (and therefore *initform* had to be evaluated). The variable *svar* is called a *supplied-p* parameter; it is not bound to an argument, but to a value indicating whether or not an argument had been supplied for another parameter.

After all *optional* parameter specifiers have been processed, then there may or may not be a *rest* parameter. If there is none, then there should be no unprocessed arguments (it is an error if there are). If there is a *rest* parameter, it is bound to a list of all as-yet-unprocessed arguments. (If no unprocessed arguments remain, the *rest* parameter is bound to the empty list.)

Instead of *optional* and *rest* parameters, *keyword* parameters may be specified instead. In that case, after all required parameters (and an equal number of arguments) have been processed, there must remain an even number of arguments; these are processed in pairs, the first argument in each pair being interpreted as a keyword name and the second as the corresponding value. No two argument pairs should have the same keyword name.

In each keyword parameter specifier must be a name *var* for the parameter variable. If an explicit *keyword* is specified, that is the keyword name for the parameter. Otherwise the name *var* serves also as the keyword name, not of itself, but in that a keyword with the same name (in the `keyword` package) is used as the

keyword. Thus

```
(defun foo (&key radix (type 'integer)) ...)
```

means exactly the same as

```
(defun foo (&key ((:radix radix)) ((:type type) 'integer)) ...)
```

For each keyword parameter specifier, if there is an argument pair whose keyword name matches that specifier's keyword name, then the parameter variable for that specifier is bound to the second item (the value) of that argument pair. If no such argument pair exists, then the *initform* for that specifier is evaluated and the parameter variable is bound to that value (or to `nil` if no *initform* was specified). The variable *svar* is treated as for ordinary *optional* parameters: it is bound to *true* if there was a matching argument pair, and to *false* otherwise. It is an error if an argument pair has a keyword name not matched by any parameter specifier.

After all parameter specifiers have been processed, the auxiliary variable specifiers (those following the token `&aux`) are processed from left to right. For each one the *initform* is evaluated and the variable *var* bound to that value (or to `nil` if no *initform* was specified). (Nothing can be done with `&aux` variables that cannot be done with the special form `let` (page 65). Which to use is purely a matter of style.)

As a rule, whenever any *initform* is evaluated for any parameter specifier, that form may refer to any parameter variable to the left of the specifier in which the *initform* appears, including any supplied-p variables, and may rely on no other parameter variable having yet been bound (including its own parameter variable).

Compatibility note: At present, one cannot depend on this in Lisp Machine LISP for keyword parameters. It is the "obvious" generalization of the current state of affairs for optional parameters and aux variables. Opinions?

Once the lambda-list has been processed, the forms in the body of the lambda-expression are executed. These forms may refer to the arguments to the function by using the names of the parameters. On exit from the function, either by a normal return of the function's value(s) or by a non-local exit, the parameter bindings, whether lexical or special, are no longer in effect (but are not necessarily permanently discarded, for any such binding can later be reinstated only if a *closure* over that binding was created and saved before the exit occurred).

Examples of `&optional` and `&rest` parameters:

```
((lambda (a b) (+ a (* b 3))) 4 5) => 19
((lambda (a &optional (b 2)) (+ a (* b 3))) 4 5) => 19
((lambda (a &optional (b 2)) (+ a (* b 3))) 4) => 10
((lambda (&optional (a 2 b) (c 3 d) &rest x) (list a b c d x)))
   => (2 nil 3 nil nil)
((lambda (&optional (a 2 b) (c 3 d) &rest x) (list a b c d x)) 6)
   => (6 t 3 nil nil)
((lambda (&optional (a 2 b) (c 3 d) &rest x) (list a b c d x)) 6 3)
   => (6 t 3 t nil)
((lambda (&optional (a 2 b) (c 3 d) &rest x) (list a b c d x))
 6 3 8)
   => (6 t 3 t (8))
((lambda (&optional (a 2 b) (c 3 d) &rest x) (list a b c d x))
 6 3 8 9 10 11)
   => (6 t 3 t (8 9 10 11))
```

Examples of &key parameters:

```
((lambda (a b &key c d) (list a b c d)) 1 2) => (1 2 nil nil)
((lambda (a b &key c d) (list a b c d)) 1 2 :c 6) => (1 2 6 nil)
((lambda (a b &key c d) (list a b c d)) 1 2 :d 8) => (1 2 nil 8)
((lambda (a b &key c d) (list a b c d)) 1 2 :c 6 :d 8) => (1 2 6 8)
((lambda (a b &key c d) (list a b c d)) 1 2 :d 8 :c 6) => (1 2 6 8)
((lambda (a b &key c d) (list a b c d)) :a 1 :d 8 :c 6) => (:a 1 6 8)
((lambda (a b &key c d) (list a b c d)) :a :b :c :d)
   => (:a :b :d nil)
```

The &optional, &rest, and &key parameter specifiers are permitted, but not terribly useful, in lambda-expressions appearing explicitly as the first element of a function-call form. They are extremely useful, however, in functions given global names by **defun**.

5.2.3. Select-Expressions

A *select-expression* is a list with the following syntax:

(select {(*keys lambda-list* {(declare {*declaration*}*)}* {*form*}*)}*)

This is a function computationally equivalent to a lambda-expression containing a case (page 70) form (assuming the variables key and args to be names not used in any specified *lambda-list*, *declaration*, or *form*):

```
(lambda (key &rest args)
  (case key
    {(keys
       (apply #'(lambda lambda-list
                  {(declare {declaration}*)}*
                  {form}*)
                args))}*))
```

The function takes its first argument and dispatches on it to one of a set of sub-functions that can accept the remaining arguments.

Actually, there is another type of clause that may appear (that would have made the above description too complicated had it been included in the syntacical formula): if a select clause is simply (*keys symbol*),

then `symbol` is taken to be the name (that may be global, or lexically bound by a `labels` (page 67) or `flet` (page 67) construct) for a function to be called. In this case the named function is given *all* the arguments given to the select-function, not merely the arguments after the first one.

What makes `select` so useful is that the different sub-functions can accept the rest of the arguments in different ways, and that a good COMMON LISP compiler can easily produce better code for a `select`-defined function than indicated by the usage of `apply` above.

Compatibility note: This use of `select` as a `lambda`-like keyword does not conflict with its use in Lisp Machine LISP as the name of a special form.

`defselect` (page 42) is a convenient way of defining a globally named select-function.

Select-functions are handy for defining message-passing protocols. For example, here is an "actor" implementation of `cons` (page 168):

```
(defun qons (the-car the-cdr)
  #'(select (:car () the-car)
            (:cdr () the-cdr)
            (:rplaca (newvalue) (setq the-car newvalue))
            (:rplacd (newvalue) (setq the-cdr newvalue))
            (:consp () t)))
```

The result of the call `(qons 'a 5)` is a functional object; call it *x*. Then

```
(funcall x :cdr) => 5
(funcall x :rplacd "Hello") => "Hello"
(funcall x :cdr) => "Hello"
```

One could then define

```
(defun qar (x) (funcall x :car))
(defun qdr (x) (funcall x :cdr))
(defun rplaqa (x y) (funcall x :rplaca y))
(defun rplaqd (x y) (funcall x :rplacd y))
(defun qonsp (x) (funcall x :consp))
```

to complete the "actor" simulation of the properties of a cons cell.

5.3. Top-Level Forms

The standard way for the user to interact with a COMMON LISP implementation is via what is called a *read-eval-print loop*: the system repeatedly reads a form from some input source (such as a keyboard or a disk file), evaluates it, and then prints the value(s) to some output sink (such as a display screen or another disk file). As a rule any form (evaluable S-expression) is acceptable. However, certain special forms are specifically designed to be convenient for use as *top-level* forms, as opposed to form embedded within other forms, as `(+ 3 4)` is embedded within `(if p (+ 3 4) 6)`. These top-level special forms may be used to define globally named functions, to define macros, to make declarations, and to define global values for special variables.

5.3.1. Defining Named Functions

`defun` *name lambda-list* {(`declare` {*declaration*}*)}* [*doc-string*] {*form*}* [*Special form*]

Evaluating this special form causes the symbol *name* to be a global name for the function specified by the lambda-expression

(`lambda` *lambda-list* {(`declare` {*declaration*}*)}* {*form*}*)

defined in the lexical environment in which the `defun` form was executed (because `defun` forms normally appear at top level, this is normally the null lexical environment).

If the optional documentation string *doc-string* is present (it may be present only if at least one *form* is also specified, as it is otherwise taken to be a *form*), then it is put on the property list of the symbol *name* under the indicator `documentation` (see `putpr`). By convention, if the string contains multiple lines then the first line should be a complete summarizing sentence on which the remainder expands.

The body of the defined function is implicitly enclosed in a `block` (page 71) construct whose name is the same as the *name* of the function. Therefore `return` (page 72) and `return-from` (page 72) may be used to exit from the function.

Other implementation-dependent bookkeeping actions may be taken as well by `defun`. The *name* is returned as the value of the `defun` form.

For example:

```
(defun discriminant (a b c)
  (declare (number a b c))
  "Compute the discriminant for a quadratic equation.
   Given a, b, and c, the value b^2-4*a*c is calculated.
   The quadratic equation a*x^2+b*x+c=0 has real, multiple,
   or complex roots depending on whether this calculated
   value is positive, zero, or negative, respectively."
  (- (* b b) (* 4 a c)))
   => discriminant
   and now (discriminant 1 2/3 -2) => 76/9
```

It is permissible to redefine a function (for example, to install a corrected version of an incorrect definition!). It is not permissible to define as a function any symbol in use as the name of a special form or macro. To redefine a macro name as the name of a function, `fmakunbound` (page 59) must first be applied to the symbol.

??? **Query:** What do people think of this safety feature? The error handler could offer to do the `fmakunbound` for you and retry.

`defselect` *name* [*doc-string*] {(*keys lambda-list* {(`declare` {*declaration*}*)}* {*form*}*)}* [*Special*

Evaluating this special form causes the symbol *name* to be a global name for a function, as for `defun` (page 42). The function is defined in the lexical environment in which the `defselect` form was executed (because `defselect` forms normally appear at top level, this is normally the null lexical environment).

The function defined is the result of evaluating a `select` form

`(select {(`*keys lambda-list* `{(declare {`*declaration*`}*)}* {`*form*`}*)}*)`

See Section SELECT-FUNCTIONS.

Compatibility note: As defined here, this is incompatible with Lisp Machine LISP. The reason is the desire to define it in terms of `case` (page 70). This means that the default, fall-through case can always be specified by using `t` or `otherwise` as the key, and that one can associate several keys with one sub-function by using a list of keys. Also, I haven't allowed for an automatic `:which-operations` method. Finally, here a *doc-string* is allowed. Is this all right, or should we revert to the Lisp Machine LISP definition?

5.3.2. Defining Macros

Macros are usually defined by using the special form `defmacro` (page 91). This facility is fairly complicated, and is described in Chapter 8.

5.3.3. Declaring Global Variables and Named Constants

`defvar` *name* [*initial-value* [*documentation*]] [*Special form*]

`defvar` is the recommended way to declare the use of a special variable in a program. It is normally used only as a top-level form.

```
(defvar variable)
```

declares `variable` to be `special` (see `declare` (page 95)), and may perform other system-dependent bookkeeping actions. If a second "argument" is supplied:

`(defvar variable` *initial-value*`)`

then `variable` is initialized to the result of evaluating the form *initial-value* unless it already has a value. *initial-value* is not evaluated unless it is used; this is useful if it does something expensive like creating a large data structure. The initialization is performed by assignment, and so assigns the variable a global value unless there are currently special bindings of that variable.

`defvar` should be used only at top level, never in function definitions.

`defvar` also provides a good place to put a comment describing the meaning of the variable (whereas an ordinary `special` declaration offers the temptation to declare several variables at once and not have room to describe them all). This can be a simple LISP comment:

```
(defvar tv-height 768)  ;Height of TV screen in pixels.
```

or, better yet, a third "argument" to `defvar`, in which case various programs can access the documentation:

```
(defvar tv-height 768 "Height of TV screen in pixels")
```

The documentation should be a string.

`defconst` *name initial-value* [*documentation*] [*Special form*]

`defconst` is similar to `defvar`, but declares a global variable whose value is "constant". An initial value is always given to the variable. It is an error if there are currently any special bindings of the variable (but implementations may or may not check for this).

If the variable is already has a value, an error occurs unless the existing value is `equal` (page 50) to the specified *initial-value*.

Implementation note: Actually, a specific interaction should occur in which the user is asked whether it is permissible to alter the constant. Perhaps there should be some mechanism to discover who uses the constant.

Rationale: `defconst` declares a constant, whose value will "never" be changed. Other code may depend on this fact. On the other hand,`defvar` declares a global variable, whose value is initialized to something but will then be changed by the functions that use it to maintain some state.

Once a symbol has been declared by `defconst` to be constant, any further assignment to or binding of that variable is an error. This is the case for such system-supplied constants as `t` (page 45) and `most-positive-fixnum` (page 142).

Chapter 6

Predicates

A *predicate* is a function that tests for some condition involving its arguments and returns `nil` if the condition is false, or some non-`nil` value if the condition is true. One may think of a predicate as producing a Boolean value, where `nil` stands for *false* and anything else stands for *true*. Conditional control structures such as `cond` (page 68), `if` (page 69), `when` (page 69), and `unless` (page 70) test such Boolean values. We say that a predicate *is true* when it returns a non-`nil` value, and *is false* when it returns `nil`; that is, it is true or false according to whether the condition being tested is true or false.

By convention, the names of predicates usually end in the letter "`p`" (which stands for "predicate").

The control structures that test Boolean values only test for whether or not the value is `nil`, which is considered to be false. Any other value is considered to be true. A function that returns `nil` if it "fails" and some *useful* value when it "succeeds" is called a *pseudo-predicate*, because it can be used not only as a test but also for the useful value provided in case of success. An example of a pseudo-predicate is `member` (page 176).

If no better non-`nil` value is available for the purpose of indicating success, by convention the symbol `t` is used as the "standard" non-false value.

6.1. Logical Values

`nil` [*Constant*]

The value of `nil` is always `nil`. This object represents the logical *false* value and also the empty list. It can also be written "`()`".

`t` [*Constant*]

The value of `t` is always `t`.

6.2. Data Type Predicates

Perhaps the most important predicates in LISP are those that deal with data types; that is, given a data object one can determine whether or not it belongs to a given type, or one can compare two type specifiers.

6.2.1. General Type Predicate

`typep` *object* `&optional` *type* [*Function*]

(`typep` *object type*) is a predicate that is true if *object* is of type *type*, and is false otherwise. Note that an object can be "of" more than one type, since one type can include another. The *type* may be any of the type specifiers mentioned in Chapter 4 *except* that it may not be or contain a type specifier list whose first element is `function`.

(`typep` *object*) returns an implementation-dependent result: some *type* of which the *object* is a member. Implementations are encouraged to return the most specific type that can be conveniently computed and is likely to be useful to the user. It is required that if the argument is a named structure created by `defstruct` then `typep` will return the name of that structure and not the symbol `structure`. Because the result is implementation-dependent, it is usually better to use `typep` of one argument primarily for debugging purposes, and to use `typep` of two arguments or the `typecase` (page 70) special form in programs.

??? Query: One-argument `typep` remains as a hangover from MACLISP. Unfortunately, any use of it in COMMON LISP is unlikely to be portable because COMMON LISP has many more data types than MACLISP. Moreover, the results of one-argument `typep` must be somewhat implementation-dependent even among COMMON LISP implementations. Finally, it is not really a predicate. Perhaps the one-argument case should be split off and renamed to, say, `type-of` or `%data-type`?

`subtypep` *type1 type2* [*Function*]

The two type specifiers are compared; this predicate is true iff *type1* is a (not necessarily proper) subtype of *type2*. The arguments must be type specifiers that are acceptable to `typep` (page 46).

6.2.2. Specific Data Type Predicates

The following predicates are for testing for individual data types.

`null` *object* [*Function*]

`null` is true if its argument is `()`, and otherwise is false. This is the same operation performed by the function `not` (page 51); however, `not` is normally used to invert a Boolean value, while `null` is normally used to test for an empty list. The programmer can therefore express *intent* by the choice of function name.

```
(null x) <=> (typep x 'null) <=> (eq x '())
```

`symbolp` *object* [*Function*]

`symbolp` is true if its argument is a symbol, and otherwise is false.

```
(symbolp x) <=> (typep x 'symbol)
```

`atom` *object* [*Function*]

The predicate `atom` is true if its argument is not a cons, and otherwise is false. It is the inverse of `consp`. Note that `(atom '())` is true, because `()≡nil`.

```
(atom x) <=> (typep x 'atom) <=> (not (typep x 'cons))
```

`consp` *object* [*Function*]

The predicate `consp` is true if its argument is a cons, and otherwise is false. It is the inverse of `atom`. Note that `(consp '()) <=> (consp 'nil) => nil`.

```
(consp x) <=> (typep x 'cons) <=> (not (typep x 'atom))
```

Compatibility note: Some LISP implementations call this function `pairp` or `listp`. The name `pairp` was rejected for COMMON LISP because it emphasizes too strongly the dotted-pair notion rather than the usual usage of conses in lists. On the other hand, `listp` too strongly implies that the cons is in fact part of a list, which after all it might not be; moreover, `()` is a list, though not a cons. The name `consp` seems to be the appropriate compromise.

`listp` *object* [*Function*]

`listp` is true if its argument is a cons or the empty list `()`, and otherwise is false. It does not check for whether the list is a "true list" (one terminated by `nil`) or a "dotted list" (one terminated by a non-null atom).

```
(listp x) <=> (typep x 'list) <=> (typep x '(cons null))
```

Compatibility note: Lisp Machine LISP defines `listp` to mean the same as `pairp`, but this is under review. The definition given here is that adopted by NIL.

`numberp` *object* [*Function*]

`numberp` is true if its argument is any kind of number, and otherwise is false.

```
(numberp x) <=> (typep x 'number)
```

`integerp` *object* [*Function*]

`integerp` is true if its argument is an integer, and otherwise is false.

```
(integerp x) <=> (typep x 'integer)
```

Compatibility note: In MACLISP this is called `fixp`. Users have been confused as to whether this meant "`integerp`" or "`fixnump`", and so these names have been adopted here.

`rationalp` *object* [*Function*]

`rationalp` is true if its argument is a rational number (a ratio or an integer), and otherwise is false.

```
(rationalp x) <=> (typep x 'rational)
```

`floatp` *object* [*Function*]

`floatp` is true if its argument is a floating-point number, and otherwise is false.

```
(floatp x) <=> (typep x 'float)
```

`complexp` *object* [*Function*]

`complexp` is true if its argument is a complex number, and otherwise is false.

```
(complexp x) <=> (typep x 'complex)
```

`characterp` *object* [*Function*]

`characterp` is true if its argument is a character, and otherwise is false.

```
(characterp x) <=> (typep x 'character)
```

`stringp` *object* [*Function*]

`stringp` is true if its argument is a string, and otherwise is false.

```
(stringp x) <=> (typep x 'string)
```

`vectorp` *object* [*Function*]

`vectorp` is true if its argument is a vector, and otherwise is false.

```
(vectorp x) <=> (typep x 'vector)
```

`arrayp` *object* [*Function*]

`arrayp` is true if its argument is an array, and otherwise is false.

```
(arrayp x) <=> (typep x 'array)
```

`functionp` *object* [*Function*]

`functionp` is true if its argument is suitable for applying to arguments, using for example the `funcall` or `apply` function. Otherwise `functionp` is false.

`subrp` *object* [*Function*]

`subrp` is true if its argument is any compiled code object, and otherwise is false.

```
(subrp x) <=> (typep x 'subr)
```

`closurep` *object* [*Function*]

`closurep` is true if its argument is a closure, and otherwise is false.

6.3. Equality Predicates

COMMON LISP provides a spectrum of predicates for testing for equality of two objects: `eq` (the most specific), `eql`, `equal`, and `equalp` (the most general). `eq` and `equal` have the meanings traditional in LISP. `eql` was added because it is frequently needed, and `equalp` was added primarily to have a version of `equal` that would ignore type differences when comparing numbers and case differences when comparing characters. If two objects satisfy any one of these equality predicates, then they also satisfy all those that are more general.

`eq` *x y* [*Function*]

(`eq` *x y*) is true if and only if *x* and *y* are the same identical object. (Implementationally, *x* and *y* are usually `eq` if and only if they address the same identical memory location.)

It should be noted that things that print the same are not necessarily `eq` to each other. Symbols with the same print name usually are `eq` to each other, because of the use of the `intern` (page 112) function. However, numbers with the same value need not be `eq`, and two similar lists are usually not `eq`.

For example:

```
(eq 'a 'b) is false
(eq 'a 'a) is true
(eq 3 3) might be true or false, depending on the implementation
(eq 3 3.0) is false
(eq (cons 'a 'b) (cons 'a 'c)) is false
(eq (cons 'a 'b) (cons 'a 'b)) is false
(setq x '(a . b)) (eq x x) is true
(eq #\A #\A) might be true or false, depending on the implementation
(eq "Foo" "Foo") is false
(eq "FOO" "foo") is false
```

Implementation note: `eq` simply compares the two pointers given it, so any kind of object that is represented in an "immediate" fashion will indeed have like-valued instances satisfy `eq`. On the PERQ, for example, fixnums and characters happen to "work". However, no program should depend on this, as other implementations of COMMON LISP might not use an immediate representation for these data types.

`eql` *x y* [*Function*]

The `eql` predicate is true if its arguments are `eq`, or if they are numbers of the same type with the same value (that is, they are `=` (page 118)), or if they are character objects that represent the same character (that is, they are `char=` (page 148)).

For example:

```
(eql 'a 'b) is false
(eql 'a 'a) is true
(eql 3 3) is true
(eql 3 3.0) is false
(eql (cons 'a 'b) (cons 'a 'c)) is false
(eql (cons 'a 'b) (cons 'a 'b)) is false
(setq x '(a . b)) (eql x x) is true
(eql #\A #\A) is true
(eql "Foo" "Foo") is false
(eql "FOO" "foo") is false
```

`equal` *x y* [*Function*]

The `equal` predicate is true if its arguments are similar (isomorphic) objects. A rough rule of thumb is that two objects are `equal` if and only if their printed representations are the same.

Numbers and characters are compared as for `eql`. Symbols are compared as for `eq`. This can violate the rule of thumb about printed representations, but only in the case of two distinct symbols with the same print name, and this does not ordinarily occur.

Objects that have components are `equal` if they are of the same type and corresponding components are `equal`. This test is implemented in a recursive manner, and will fail to terminate for circular structures. For conses, `equal` is defined recursively as the two *car*'s being `equal` and the two *cdr*'s being `equal`.

Two arrays are `equal` if and only if they have the same number of dimensions, the dimensions match, the element types match, and the corresponding components are `equal`.

> **Compatibility note:** In Lisp Machine LISP, `equal` ignores the difference between upper and lower case in strings. This violates the rule of thumb about printed representations, however, which is very useful, especially to novices. It is also inconsistent with the treatment of single characters, which are represented as fixnums.

Two pathnames are `equal` iff corresponding components (host, device, and so on) are equivalent. Whether or not case is considered equivalent in strings depends on the file name conventions of the file system. The intent is that pathnames that are `equal` should be functionally equivalent.

For example:

```
(equal 'a 'b) is false
(equal 'a 'a) is true
(equal 3 3) is true
(equal 3 3.0) is false
(equal (cons 'a 'b) (cons 'a 'c)) is false
(equal (cons 'a 'b) (cons 'a 'b)) is true
(setq x '(a . b)) (equal x x) is true
(equal #\A #\A) is true
(equal "Foo" "Foo") is true
(equal "FOO" "foo") is false
```

To recursively compare only conses, and compare all atoms using `eq`, use `tree-equal` (page 168).

`equalp` *x y* `&optional` *fuzz* [*Function*]

Two objects are `equalp` if they are `eql`, if they are characters and differ only in alphabetic case (that is, they are `char-equal` (page 148)), if they are numbers and have the same numerical value, even if they are of different types, or if they have components that are all `equalp`. When comparing floating-point numbers, or comparing a floating-point number to any other kind of number, the optional argument *fuzz* is used; in effect the function `fuzzy=` (page 120) is used to perform such comparisons.

Objects that have components are `equalp` if they are of the same type and corresponding components are `equalp`. This test is implemented in a recursive manner, and will fail to terminate for circular structures. For conses, `equalp` is defined recursively as the two *car*'s being `equalp` and the two *cdr*'s being `equalp`.

Two arrays are `equalp` if and only if they have the same number of dimensions, the dimensions match, the element types match, and the corresponding components are `equalp`.

??? **Query:** How about eliminating the clause "the element types match" from the above specification? This would allow a string and a general array that happens to contain characters to be `equalp`, for example.

For example:

```
(equalp 'a 'b) is false
(equalp 'a 'a) is true
(equalp 3 3) is true
(equalp 3 3.0) is true
(equalp (cons 'a 'b) (cons 'a 'c)) is false
(equalp (cons 'a 'b) (cons 'a 'b)) is true
(setq x '(a . b)) (equalp x x) is true
(equalp #\A #\A) is true
(equalp "Foo" "Foo") is true
(equalp "FOO" "foo") is true
```

6.4. Logical Operators

COMMON LISP provides three operators on Boolean values: `and`, `or`, and `not`. Of these, `and` and `or` are also control structures, because their arguments are evaluated conditionally. `not` necessarily examines its single argument, and so is a simple function.

`not` *x* [*Function*]

`not` returns `t` if *x* is `nil`, and otherwise returns `nil`. It therefore inverts its argument, interpreted as a Boolean value.

`null` (page 46) is the same as `not`; both functions are included for the sake of clarity. As a matter of style, it is customary to use `null` to check whether something is the empty list, and to use `not` to invert the sense of a logical value.

`and` {*form*}* [*Special form*]

`(and` *form1 form2 ...* `)` evaluates each *form*, one at a time, from left to right. If any *form* evaluates to `nil`, `and` immediately is false without evaluating the remaining *forms*. If every *form* but the last evaluates to a non-`nil` value, `and` returns whatever the last *form* returns. Therefore in general `and` can be used both for logical operations, where `nil` stands for *false* and non-`nil` values stand for *true*, and as a conditional expression.

For example:

```
(if (and (>= n 0)
         (lessp n (length a-vector))
         (eq (vref a-vector n) 'foo))
    (princ "Foo!"))
```

The above expression prints "`Foo!`" if element `n` of `a-vector` is the symbol `foo`, provided also that `n` is indeed a valid index for `a-vector`. Because `and` guarantees left-to-right testing of its parts, `vref` is not performed if `n` is out of range. (In this example writing

```
(and (>= n 0)
     (lessp n (length a-vector))
     (eq (vref a-vector n) 'foo)
     (princ "Foo!"))
```

would accomplish the same thing; the difference is purely stylistic.) Because of the guaranteed left-to-right ordering, `and` is like the **and then** operator in ADA, or what in some PASCAL-like languages is called **cand**, rather than the **and** operator.

See also `if` (page 69) and `when` (page 69), which are sometimes stylistically more appropriate than `and` for conditional purposes.

From the general definition, one can deduce that `(and` *x*`)` <=> *x*. Also, `(and)` is true, which is an identity for this operation.

`and` can be defined in terms of `cond` (page 68) as follows:

```
(and x y z ... w) <=>

(cond ((not x) nil)
      ((not y) nil)
      ((not z) nil)
      ...
      (t w))
```

`or` {*form*}* [*Special form*]

`(or` *form1 form2 ...* `)` evaluates each *form*, one at a time, from left to right. If any *form* evaluates to something other than `nil`, `or` immediately returns it without evaluating the remaining *forms*. If every *form* but the last evaluates to `nil`, `or` returns whatever evaluation of the last of the *forms* returns. Therefore in general `or` can be used both for logical operations, where `nil` stands for *false* and non-`nil` values stand for *true*, and as a conditional expression. Because of the guaranteed left-to-right ordering, `or` is like the **or else** operator in ADA, or what in some PASCAL-like languages is called **cor**, rather than the **or** operator.

See also `if` (page 69) and `unless` (page 70), which are sometimes stylistically more appropriate

than `or` for conditional purposes.

From the general definition, one can deduce that `(or` *x*`)` `<=>` *x*. Also, `(or)` is false, which is the identity for this operation.

`or` can be defined in terms of `cond` (page 68) as follows:

```
(or x y z ... w) <=> (cond (x) (y) (z) ... (t w))
```

Chapter 7

Control Structure

LISP provides a variety of special structures for organizing programs. Some have to do with flow of control (control structures), while others control access to variables (environment structures). Most of these features are implemented either as special forms or as macros (which typically expand into complex program fragments involving special forms).

Function application is the primary method for construction of LISP programs. Operations are written as the application of a function to its arguments. Usually, LISP programs are written as a large collection of small functions, each of which implements a simple operation. These functions operate by calling one another, and so larger operations are defined in terms of smaller ones. LISP functions may call upon themselves recursively, either directly or indirectly.

LISP, while more applicative in style than statement-oriented, nevertheless provides many operations which produce side-effects, and consequently requires constructs for controlling the sequencing of side-effects. The construct `progn` (page 64), which is roughly equivalent to an ALGOL **begin-end** block with all its semicolons, executes a number of forms sequentially, discarding the values of all but the last. Many LISP control constructs include sequencing implicitly, in which case they are said to provide an "implicit `progn`". Other sequencing constructs include `prog1` (page 65) and `prog2` (page 65).

For looping, COMMON LISP provides the general iteration facility `do` (page 73), as well as a variety of special-purpose iteration facilities for iterating or mapping over various data structures.

COMMON LISP provides the simple one-way conditionals `when` and `unless`, the simple two-way conditional `if`, and the more general multi-way conditionals such as `cond` and `selectq`. The choice of which form to use in any particular situation is a matter of taste and style.

Constructs for performing non-local exits with various scoping disciplines are provided: `block` (page 71), `return` (page 72), `catch` (page 85), and `throw` (page 87).

The multiple-value constructs provide an efficient way for a function to return more than one value; see `values` (page 82).

7.1. Constants and Variables

7.1.1. Reference

`quote` *object* [*Special form*]

`(quote x)` simply returns *x*. The argument is not evaluated, and may be any LISP object. This construct allows any LISP object to be written as a constant value in a program.

For example:

```
(setq a 43)
(list a (cons a 3)) => (43 (43 . 3))
(list (quote a) (quote (cons a 3)) => (a (cons a 3))
```

Since `quote` forms are so frequently useful but somewhat cumbersome to type, a standard abbreviation is defined for them: any form preceded by a single quote (') character is assumed to have "`(quote )`" wrapped around it.

For example:

```
(setq x '(the magic quote hack))
    is normally interpreted when read to mean
(setq x (quote (the magic quote hack)))
```

`function` *fn* [*Special form*]

The value of `function` is always the functional interpretation of *fn*; *fn* is interpreted as if it had appeared in the functional position of a function invocation. In particular, if *fn* is a symbol, the functional value of the variable whose name is that symbol is returned. If *fn* is a lambda expression or select expression, then a lexical closure is returned.

Since `function` forms are so frequently useful (for passing functions as arguments to other function) but somewhat cumbersome to type, a standard abbreviation is defined for them: any form preceded by a sharp sign and then a single quote (`#'`) is assumed to have "`(function )`" wrapped around it.

For example:

```
(remove-if #'numberp '(1 a b 3))
    is normally interpreted when read to mean
(remove-if (function numberp) '(1 a b 3))
```

`closure` *varlist function* [*Function*]

The function `closure` creates and returns a closure of the *function* over the special variables mentioned in the *varlist*.

The *varlist* must be a list of symbols. The *function* may be any functional object. The current bindings of the special (not lexical) variables named by the symbols are collected into a closure object along with the function. When the closure is invoked as a function, the saved bindings are re-established, and then *function* is invoked. The saved binding of a special variable is "shared"

with the current binding and with any other closures over the same variable binding; by "shared" it is meant that an assignment to one (via `setq` (page 58) or `set` (page 58)) is reflected in the others.

`symeval` *symbol* [*Function*]

`symeval` returns the current value of the dynamic (special) variable named by *symbol*. An error occurs if the symbol has no value; see `boundp` (page 57) and `makunbound` (page 59).

`symeval` cannot access the value of a local (lexically bound) variable.

This function is particularly useful for implementing interpreters for languages embedded in LISP. The corresponding assignment primitive is `set` (page 58).

`fsymeval` *symbol* [*Function*]

`fsymeval` returns the current global function definition named by *symbol*. An error occurs if the symbol has no function definition; see `fboundp` (page 57). Note that the definition may be a function, or may be an object representing a special form or macro. See `macro-p` (page 57) and `special-form-p` (page 57).

`fsymeval` cannot access the value of a local function name (lexically bound as by `flet` (page 67) or `labels` (page 67)).

This function is particularly useful for implementing interpreters for languages embedded in LISP. The corresponding assignment primitive is `fset` (page 59).

`boundp` *symbol* [*Function*]

`fboundp` *symbol* [*Function*]

`boundp` is true if the dynamic (special) variable named by *symbol* has a value; otherwise, it returns `nil`. `fboundp` is the analogous predicate for the global function definition named by *symbol*.

See also `set` (page 58), `fset` (page 59), `makunbound` (page 59), and `fmakunbound` (page 59).

`macro-p` *symbol* [*Function*]

`special-form-p` *symbol* [*Function*]

The function `macro-p` takes a symbol. If the symbol globally names a macro, then the expansion function (a function of one argument, a macro-call form) is returned; otherwise `nil` is returned.

The function `special-form-p` also takes a symbol. If the symbol globally names a special form (example: `quote` (page 56)), then a non-`nil` value is returned, typically a function of implementation-dependent nature that can be used to interpret a special form; otherwise `nil` is returned.

It is possible for *both* `macro-p` and `special-form-p` to be true of a symbol. This can arise because an implementation is free to implement any macro also as a special form for speed. On the

other hand, the macro definition must also be available for use by programs that understand only the standard special forms listed in Table 5-1.

7.1.2. Assignment

`setq` {*var form*}* [*Special form*]

The special form (`setq` *var1 form1 var2 form2 ...*) is the "simple variable assignment statement" of Lisp. First *form1* is evaluated and the result is assigned to *var1*, then *form2* is evaluated and the result is assigned to *var2*, and so forth. The variables are represented as symbols, of course, and are interpreted as referring to static or dynamic instances according to the usual rules. `setq` returns the last value assigned, that is, the result of the evaluation of its last argument. As a boundary case, the form (`setq`) is legal and returns `nil`. As a rule there must be an even number of argument forms.

For example:

```
(setq x (+ 3 2 1) y (cons x nil))
```

`x` is set to `6`, `y` is set to `(6)`, and the `setq` returns `(6)`. Note that the first assignment was performed before the second form was evaluated, allowing that form to use the new value of `x`.

See also the description of `setf` (page 60), which is the "general assignment statement", capable of assigning to variables, array elements, and other locations.

`psetq` {*var form*}* [*Special form*]

A `psetq` form is just like a `setq` form, except that the assignments happen in parallel; first all of the forms are evaluated, and then the variables are set to the resulting values. The value of the `psetq` form is `nil`.

For example:

```
(setq a 1)
(setq b 2)
(psetq a b  b a)
a => 2
b => 1
```

In this example, the values of `a` and `b` are exchanged by using parallel assignment. (Note that the `do` (page 73) iteration construct performs a very similar thing when stepping iteration variables.)

`set` *symbol value* [*Function*]

`set` allows alteration of the value of a dynamic (special) variable. `set` causes the dynamic variable named by *symbol* to take on *value* as its value. Only the value of the current dynamic binding is altered; if there are no bindings in effect, the most global value is altered.

For example:

```
(set (if (eq a b) 'c 'd) 'foo)
```

will either set `c` to `foo` or set `d` to `foo`, depending on the outcome of the test `(eq a b)`.

Both functions return *value* as the result value.

set cannot alter the value of a local (lexically bound) variable. The special form `setq` (page 58) is usually used for altering the values of variables (lexical or dynamic) in programs. `set` is particularly useful for implementing interpreters for languages embedded in LISP. See also `progv` (page 67), a construct which performs binding rather than assignment of dynamic variables.

`fset` *symbol value* [*Function*]

`fset` allows alteration of the global function definition named by *symbol* to be *value*. `fset` returns *value*.

fset cannot alter the value of a local (lexically bound) function definition, as made by `flet` (page 67) or `labels` (page 67). `fset` is particularly useful for implementing interpreters for languages embedded in LISP.

`makunbound` *symbol* [*Function*]

`fmakunbound` *symbol* [*Function*]

`makunbound` causes the dynamic (special) variable named by *symbol* to become unbound (have no value). `fmakunbound` does the analogous thing for the global function definition named by *symbol*.

For example:

```
(setq a 1)
a => 1
(makunbound 'a)
a => causes an error
(defun foo (x) (+ x 1))
(foo 4) => 5
(fmakunbound 'foo)
(foo 4) => causes an error
```

Both functions return *symbol* as the result value.

7.2. Generalized Variables

In LISP, a variable can remember one piece of data, a LISP object. The main operations on a variable are to recover that piece of data, and to alter the variable to remember a new object; these operations are often called *access* and *update* operations. The concept of variables named by symbols can be generalized to any storage location that can remember one piece of data, no matter how that location is named. Examples of such storage locations are the *car* and *cdr* of a cons, elements of an array, and components of a structure.

For each kind of generalized variable, there are typically two functions which implement the conceptual *access* and *update* operations. For a variable, merely mentioning the name of the variable accesses it, while the `setq` (page 58) special form can be used to update it. The function `car` (page 167) accesses the *car* of a

cons, and the function `rplaca` (page 174) updates it. The function `aref` (page 185) accesses an array element, and the function `aset` (page 185) updates it.

Rather than thinking about two distinct functions that respectively access and update a storage location somehow deduced from their arguments, we can instead simply think of a call to the access function with given arguments as a *name* for the storage location. Thus, just as `x` may be considered a name for a storage location (a variable), so `(car x)` is a name for the *car* of some cons (which is in turn named by `x`), and `(aref a 105)` is a name for element 105 of the array named `a`. Now, rather than having to remember two functions for each kind of generalized variable (having to remember, for example, that `aset` corresponds to `aref`), we adopt a uniform syntax for updating storage locations named in this way, using the `setf` special form. This is analogous to the way we use the `setq` special form to convert the name of a variable (which is also a form which accesses it) into a form which updates it. The uniformity of this approach may be seen from the following table:

Access function	Update function	Update using `setf`
`x`	`(setq x newvalue)`	`(setf x newvalue)`
`(car x)`	`(rplaca x newvalue)`	`(setf (car x) newvalue)`
`(aref a 105)`	`(aset newvalue a 105)`	`(setf (aref a 105) newvalue)`
`(nth n x)`	`(setnth n x newvalue)`	`(setf (nth n x) newvalue)`

`setf` is actually a macro that examines an access form and expands into the appropriate update function.

`setf` *place newvalue* [*Macro*]

`setf` takes a form *place* that when evaluated *accesses* a data object in some location, and "inverts" it to produce a corresponding form to *update* the location. A call to the `setf` macro therefore expands into an update form that stores the result of evaluating the form *newvalue* into the place referred to by the *access-form*.

For example:

```
(setf a 3) ==> (setq a 3)
(setf (plist 'a) '(foo bar)) ==> (setplist 'a '(foo bar))
(setf (aref q 2) 56) ==> (aset 56 q 2)
(setf (cadr w) x) ==> (rplaca (cdr w) x)
```

The form *place* may be any one of the following:

- The name of a variable (either lexical or dynamic).
- A function call form whose first element is the name of any one of the following functions:

`car` (page 167)	`caaaar` (page 167)	`cadddr` (page 167)
`cdr` (page 167)	`cdaaar` (page 167)	`cddddr` (page 167)
`caar` (page 167)	`cadaar` (page 167)	`elt` (page 157)
`cdar` (page 167)	`cddaar` (page 167)	`nth` (page 169)
`cadr` (page 167)	`caadar` (page 167)	`vref` (page 187)
`cddr` (page 167)	`cdadar` (page 167)	`aref` (page 185)
`caaar` (page 167)	`caddar` (page 167)	`symeval` (page 57)
`cdaar` (page 167)	`cdddar` (page 167)	`fsymeval` (page 57)

`cadar` (page 167)	`caaadr` (page 167)	`getpr` (page 102)
`cddar` (page 167)	`cdaadr` (page 167)	`gethash` (page 182)
`caadr` (page 167)	`cadadr` (page 167)	`plist` (page 103)
`cdadr` (page 167)	`cddadr` (page 167)	
`caddr` (page 167)	`caaddr` (page 167)	
`cdddr` (page 167)	`cdaddr` (page 167)	

- A function call form whose first element is the name of a selector function constructed by `defstruct` (page 199).

- A function call form whose first element is the name of any one of the following functions, provided that the new value is of the specified type so that it can be used to replace the specified "location" (which is in each of these cases not really a truly generalized variable):

Function name	Required type	Update function used
`char` (page 191)	`string-char`	`rplachar` (page 192)
`bit` (page 187)	`(mod 2)`	`rplacbit` (page 187)
`subseq` (page 157)	`sequence`	`replace` (page 160)

- A function call form whose first element is the name of any one of the following functions, provided that the specified argument to that function is in turn a *place* form; in this case the new *place* has stored back into it the result of applying the specified "update" function (which is in each of these cases not a true update function):

Function name	Argument that is a *place*	Update function used
`char-bit` (page 152)	First	`set-char-bit` (page 152)
`ldb` (page 139)	Second	`dpb` (page 140)
`mask-field` (page 140)	Second	`deposit-field` (page 140)

- A call on `getf` (page 103), in which case `(setf (getf x y) z)` expands into `(putf x y z)`.

- A `the` (page 99) type declaration form, in which case the declaration is transferred to the *newvalue* form, and the resulting `setf` form is analyzed. For example,

```
(setf (the integer (cadr x)) (+ y 3))
```

  is processed as if it were

```
(setf (cadr x) (the integer (+ y 3)))
```

- A macro call, in which case the macro call is expanded and `setf` then analyzes the resulting form.

`setf` carefully arranges to preserve the usual left-to-right order in which the various subforms are evaluated. For example,

```
(setf (aref (compute-an-array) 105) (compute-newvalue))
```

does not expand precisely into

```
(aset (compute-newvalue) (compute-an-array) 105)
```

lest side effects in the computations `(compute-an-array)` and `(compute-newvalue)` occur

in the wrong order. Instead this example will expand into something more like

```
(let ((G1 (compute-an-array))
      (G2 105)
      (G3 (compute-newvalue)))
  (aset G3 G1 G2))
```

The exact expansion for any particular form is not guaranteed and may even be implementation-dependent; all that is guaranteed is that the expansion of a `setf`-form will be an update form that works for that particular implementation, and that the left-to-right evaluation of subforms is preserved.

Compatibility note: Lisp Machine LISP, at least, officially does not preserve the order of evaluation, but also seems to regard this as a bug to be fixed. What shall COMMON LISP do?

The ultimate result of evaluating a `setf` form is the value of *newvalue*. (Therefore `(setf (car x) y)` does not expand into precisely `(rplaca x y)`, but into something more like

```
(let ((G1 x) (G2 y)) (rplaca x y) y)
```

the precise expansion being implementation-dependent.)

The user can define new `setf` expansions by using `defsetf` (page DEFSETF-FUN).

`swapf` *place newvalue* [*Macro*]

The datum in *place* is replaced by *newvalue*, and then the old value of *place* is returned. The form *place* may be any form acceptable as a generalized variable to `setf` (page 60).

For example:

```
(setq x '(a b c))
(swapf (cadr x) 'z) => b
   and now x => (a z c)
```

The effect of `(swapf` *place newvalue*`)` is roughly equivalent to

```
(prog1 place (setf place newvalue))
```

except that the latter would evaluate any subforms of *place* twice, while `swapf` takes care to evaluate them only once.

For example:

```
(setq n 0)
    (setq x '(a b c d))
    (swapf (nth (setq n (+ n 1)) x) 'z) => b
       and now x => (a z c d)
```

but

```
    (setq n 0)
    (setq x '(a b c d))
    (prog1 (nth (setq n (+ n 1)) x)
           (setf (nth (setq n (+ n 1)) x) 'z)) => b
  and now x => (a b z d)
```

Moreover, for certain *place* forms `swapf` may be significantly more efficient than the `prog1` version.

`exchf` *place1 place2* [*Macro*]

The data in *place1* and *place2* is exchanged, and then the old value of *place2* (which has become the new value of *place1*) is returned. The forms *place1* and *place2* may be any forms acceptable as generalized variables to `setf` (page 60). If *place1* and *place2* refer to the same generalized variable, then the effect is to leave it unchanged and return its value.

For example:

```
(setq x '(a b c))
(exchf (car x) (cadr x)) => b
   and now x => (b a c)
```

The effect of (`exchf` *place1 place2*) is roughly equivalent to

```
(setf place1 (prog1 place2 (setf place2 place1)))
```

except that the latter would evaluate any subforms of *place1* and *place2* twice, while `exchf` takes care to evaluate them only once. Moreover, for certain *place* forms `exchf` may be significantly more efficient than the `prog1` version.

Other macros that manipulate generalized variables include `getf` (page 103), `putf` (page 103), `remf` (page 104), `incf` (page 122), `decf` (page 122), `push` (page 172), and `pop` (page 173).

7.3. Function Invocation

The most primitive form for function invocation in LISP of course has no name; any list which which has no other interpretation as a macro call or special form is taken to be a function call. Other constructs are provided for less common but nevertheless frequently useful situations.

`apply` *function arglist* [*Function*]

This applies *function* to the list of arguments *arglist*. *arglist* should be a list; *function* can be a compiled-code object, or it may be a "lambda expression", that is, a list whose *car* is the symbol `lambda`, or it may a symbol, in which case the dynamic functional value of that symbol is used (but it is illegal in this case for that symbol to be the name of a macro or special form).

For example:

```
(setq f '+) (apply f '(1 2)) => 3
(setq f '-) (apply f '(1 2)) => -1
(apply 'cons '((+ 2 3) 4)) =>
        ((+ 2 3) . 4)   not (5 . 4)
```

Of course, *arglist* may be ( ) (in which case the function is given no arguments.) Note that if the function takes keyword arguments, the keywords as well as the corresponding values must appear in the *arglist*:

```
(apply #'(lambda (&key a b) (list a b)) '(:b 3)) => (nil 3)
```

Compatibility note: ???

See `eval` (page 209).

`funcall` *fn* `&rest` *arguments* [*Function*]

(`funcall` *fn a1 a2 ... an*) applies the function *fn* to the arguments *a1*, *a2*, ..., *an*. *fn* may not be a special form nor a macro; this would not be meaningful.

For example:

```
(cons 1 2) => (1 . 2)
(setq cons (fsymeval '+))
(funcall cons 1 2) => 3
```

The difference between `funcall` and an ordinary function call is that the function is obtained by ordinary LISP evaluation rather than by the special interpretation of the function position that normally occurs.

Compatibility note: This corresponds roughly to the INTERLISP primitive `apply*`.

`funcall*` *f* `&rest` *args* [*Function*]

`funcall*` is like a cross between `apply` and `funcall`. (`funcall*` *a1 a2 ... an list*) applies the function *f* to the arguments *a1* through *an* followed by the elements of *list*. Thus we have:

(`funcall` *f a1 ... an*) <=> (`funcall*` *f a1 ... an* '())
(`apply` *f list*) <=> (`funcall*` *f list*)

However, when `apply` or `funcall` fits the situation at hand, it may be stylistically clearer to use that than to use `funcall*`, whose use implies that something more complicated is going on.

```
(funcall* #'+ 1 1 1 '(1 1 1)) => 6

(defun report-error (&rest args)
       (funcall* (function format) error-output args))
```

Compatibility note: ???

7.4. Simple Sequencing

`progn` {*form*}* [*Special form*]

The `progn` construct takes a number of forms and evaluates them sequentially, in order, from left to right. The values of all the forms but the last are discarded; whatever the last form returns is returned by the `progn` form. One says that all the forms but the last are evaluated for *effect*, because their execution is useful only for the side effects caused, but the last form is executed for *value*.

`progn` is the primitive control structure construct for "compound statements"; it is analogous to **begin-end** blocks in ALGOL-like languages. Many LISP constructs are "implicit `progn`" forms, in that as part of their syntax each allows many forms to be written which are evaluated sequentially, the results of only the last of which are used for anything.

If the last form of the `progn` returns multiple values, then those multiple values are returned by the `progn` form. If there are no forms for the `progn`, then the result is `nil`. These rules generally

hold for implicit `progn` forms as well.

`prog1` *first* {*form*}* [*Special form*]

`prog1` is similar to `progn`, but it returns the value of its *first* form. All the argument forms are executed sequentially; the value the first form produces is saved while all the others are executed, and is then returned.

`prog1` is most commonly used to evaluate an expression with side effects, and return a value which must be computed *before* the side effects happen.

For example:

```
(prog1 (car x) (rplaca x 'foo))
```

alters the *car* of `x` to be `foo` and returns the old *car* of `x`.

`prog1` always returns a single value, even if the first form tries to return multiple values. A consequence of this is that (`prog1` *x*) and (`progn` *x*) may behave differently if *x* can produce multiple values. See `mvprog1` (page 82).

`prog2` *first second* {*form*}* [*Special form*]

`prog2` is similar to `prog1`, but it returns the value of its *second* form. All the argument forms are executed sequentially; the value of the second form is saved while all the other forms are executed, and is then returned.

`prog2` is provided mostly for historical compatibility.

```
(prog2 a b c ... z) <=> (progn a (prog1 b c ... z))
```

Occasionally it is desirable to perform one side effect, then a value-producing operation, then another side effect; in such a peculiar case `prog2` is fairly perspicuous.

For example:

```
(prog2 (open-a-file) (compute-on-file) (close-the-file))
  ;value is that of compute-on-file
```

`prog2`, like `prog1`, always returns a single value, even if the second form tries to return multiple values. A consequence of this is that (`prog2` *x y*) and (`progn` *x y*) may behave differently if *y* can produce multiple values.

7.5. Environment Manipulation

`let` ({*var* | (*var value*)}*) {*form*}* [*Macro*]

A `let` form can be used to execute a series of forms with specified variables bound to specified values.

For example:

```
(let ((var1 value1)
      (var2 value2)
      ...
      (varm valuem))
  body1
  body2
  ...
  bodyn)
```

first evaluates the expressions *value1*, *value2*, and so on, in that order, saving the resulting values. Then all of the variables *varj* are bound to the corresponding values in parallel; each binding will be a local binding unless there is a `:special` (page DECLARE-SPECIAL-KWD) declaration to the contrary. The expressions *bodyj* are then evaluated in order; the values of all but the last are discarded (that is, the body of a `let` form is an implicit `progn`). The `let` form returns what evaluating *bodyn* produces (if the body is empty, which is fairly useless, `let` returns `nil` as its value). The bindings of the variables disappear when the `let` form is exited.

Instead of a list (*varj valuej*) one may write simply *varj*. In this case *varj* is initialized to `nil`. As a matter of style, it is recommended that *varj* be written only when that variable will be stored into (such as by `setq` (page 58)) before its first use. If it is important that the initial value is `nil` rather than some undefined value, then it is clearer to write out (*varj* `nil`) (if the initial value is intended to mean "false") or (*varj* `'()`) (if the initial value is intended to be an empty list).

Declarations may appear at the beginning of the body of a `let`; they apply to the code in the body *and* to the bindings made by `let`, but not to the code which produces values for the bindings.

The `let` form shown above is entirely equivalent to:

```
((lambda (var1 var2 ... varm)
         body1 body2 ... bodyn)
 value1 value2 ... valuem)
```

but `let` allows each variable to be textually close to the expression which produces the corresponding value, thereby improving program readability.

`let*` ({*var* | (*var value*)}*) {*form*}* . [*Special form*]

`let*` is similar to `let` (page 65), but the bindings of variables are performed sequentially rather than in parallel. This allows the expression for the value of a variable to refer to variables previously bound in the `let*` form.

More precisely, the form:

```
(let* ((var1 value1)
       (var2 value2)
       ...
       (varm valuem))
  body1
  body2
  ...
  bodyn)
```

first evaluates the expression *value1*, then binds the variable *var1* to that value; then its evaluates

value2 and binds *var2*; and so on. The expressions *bodyj* are then evaluated in order; the values of all but the last are discarded (that is, the body of a `let*` form is an implicit `progn`). The `let*` form returns the results of evaluating *bodyn* (if the body is empty, which is fairly useless, `let*` returns `nil` as its value). The bindings of the variables disappear when the `let*` form is exited.

Instead of a list (*varj valuej*) one may write simply *varj*. In this case *varj* is initialized to `nil`. As a matter of style, it is recommended that *varj* be written only when that variable will be stored into (such as by `setq` (page 58)) before its first use. If it is important that the initial value is `nil` rather than some undefined value, then it is clearer to write out (*varj* `nil`) (if the initial value is intended to mean "false") or (*varj* `'()`) (if the initial value is intended to be an empty list).

Declarations may appear at the beginning of the body of a `let`; they apply to the code in the body *and* to the bindings made by `let`, but not to the code which produces values for the bindings.

`progv` *symbols values* {*form*}* [*Special form*]

`progv` is a special form which allows binding one or more dynamic variables whose names may be determined at run time. The sequence of forms (an implicit `progn`) is evaluated with the dynamic variables whose names are in the list *symbols* bound to corresponding values from the list *values*. (If too few values are supplied, the remaining symbols are bound to `nil`. If too many values are supplied, the excess values are ignored.) The results of the `progv` form are those of the last *form*. The bindings of the dynamic variables are undone on exit from the `progv` form. The lists of symbols and values are computed quantities; this is what makes `progv` different from, for example, `let` (page 65), where the variable names are stated explicitly in the program text.

`progv` is particularly useful for writing interpreters for languages embedded in LISP; it provides a handle on the mechanism for binding dynamic variables.

`flet` ({(*name lambda-list* {*declare-form*}* [*doc-string*] {*form*}*)}*) {*form*}* [*Special form*]
`labels` ({(*name lambda-list* {*declare-form*}* [*doc-string*] {*form*}*)}*) {*form*}* [*Special form*]
`macrolet` ({(*name varlist* {*form*}*)}*) {*form*}* [*Special form*]

`flet` may be used to define locally named functions. Within the body of the `flet` form, function names matching those declared by the `flet` refer to the locally defined functions rather than to the global function definitions of the same name.

Any number of functions may be simultaneously declared. Each declaration is similar in format to a `defun` (page 42) form: first a name, then a parameter list (which may contain `&optional`, `&rest`, or `&key` parameters), then optional declarations and documentation string, and finally a body.

The `labels` construct is identical in form to the `flet` construct. It differs in that the scope of the declared function names for `flet` encompasses only the body, while for `labels` it encompasses the function definitions themselves. That is, `labels` can be used to define mutually recursive functions, but `flet` cannot. This distinction is useful. Using `flet` one can locally redefine a global function name, and the new definition can refer to the global definition; the same

construction using `labels` would not have that effect.

```
(defun integer-power (n k)      ;A highly "bummed" integer
  (declare (integer n))         ; exponentiation routine.
  (declare (type (integer 0 *) k))
  (labels ((expt0 (x k a)
             (declare (integer x a) (type (integer 0 *) k))
             (cond ((zerop k) a)
                   ((evenp k) (expt1 (* x x) (floor k 2) a))
                   (t (expt0 (* x x) (floor k 2) (* x a)))))
           (expt1 (x k a)
             (declare (integer x a) (type (integer 0 *) k))
             (cond ((evenp k) (expt1 (* x x) (floor k 2) a))
                   (t (expt0 (* x x) (floor k 2) (* x a))))))
    (expt0 n k 1)))
```

`macrolet` is similar in form to `flet`, but defines local macros, using the same format used by `defmacro` (page 91).

7.6. Conditionals

`cond` {(*test* {*form*}*)}* [*Special form*]

The `cond` special form takes a number (possibly zero) of *clauses*, which are lists of forms. Each clause consists of a *test* followed by zero or more *consequents*.

For example:

```
(cond (test-1 consequent-1-1 consequent-1-2 ...)
      (test-2)
      (test-3 consequent-3-1 ...)
      ... )
```

The first clause whose *test* evaluates to non-`nil` is selected; all other clauses are ignored, and the consequents of the selected clause are evaluated in order (as an implicit `progn`).

More specifically, `cond` processes its clauses in order from left to right. For each clause, the *test* is evaluated. If the result is `nil`, `cond` advances to the next clause. Otherwise, the *cdr* of the clause is treated as a list of forms, or consequents, which are evaluated in order from left to right, as an implicit `progn`. After evaluating the consequents, `cond` returns without inspecting any remaining clauses. The `cond` special form returns the results of evaluating the last of the selected consequents; if there were no consequents in the selected clause, then the single (and necessarily non-null) value of the *test* is returned. If `cond` runs out of clauses (every test produced `nil`, and therefore no clause was selected), the value of the `cond` form is `nil`.

If it is desired to select the last clause unconditionally if all others fail, the standard convention is to use `t` for the *test*. As a matter of style, it is desirable to write a last clause "`(t nil)`" if the value of the *cond* form is to be used for something. Similarly, it is in questionable taste to let the last clause of a `cond` be a "singleton clause"; an explicit `t` should be provided. (Note moreover that `(cond ... (`*x*`))` may behave differently from `(cond ... (t` *x*`))` if *x* might produce multiple values; the former always returns a single value, while the latter returns whatever values *x*

returns.)

For example:

```
(setq z (cond (a 'foo) (b 'bar)))          ;Possibly confusing.
(setq z (cond (a 'foo) (b 'bar) (t nil))) ;Better.
(cond (a b) (c d) (e))                     ;Possibly confusing.
(cond (a b) (c d) (t e))                   ;Better.
(cond (a b) (c d) (t (values e)))          ;Better (if one value needed).
(cond (a b) (c))                           ;Possibly confusing.
(cond (a b) (t c))                         ;Better.
(if a b c)                                 ;Also better.
```

A LISP `cond` form may be compared to a continued **if-then-elseif** as found in many algebraic programming languages:

```
(cond (p ...)                          if p then ...
      (q ...)          roughly         else if q then ...
      (r ...)         corresponds      else if r then ...
      ...                to            ...
      (t ...))                         else ...
```

`if` *pred then* [*else*] [*Special form*]

The `if` special form corresponds to the **if-then-else** construct found in most algebraic programming languages. First the form *pred* is evaluated. If the result is not `nil`, then the form *then* is selected; otherwise the form *else* is selected. Whichever form is selected is then evaluated, and `if` returns whatever evaluation of the selected form returns.

(`if` *pred then else*) <=> (`cond` (*pred then*) (`t` *else*))

but `if` is considered more readable in some situations.

The *else* form may be omitted, in which case if the value of *pred* is `nil` then nothing is done and the value of the `if` form is `nil`. If the value of the `if` form is important in this situation, then the `and` (page 52) construct may be stylistically preferable, depending on the context. If the value is not important, but only the effect, then the `when` (page 69) construct may be stylistically preferable.

`when` *pred* {*form*}* [*Special form*]

(`when` *pred form1 form2* ...) first evaluates *pred*. If the result is `nil`, then no *form* is evaluated, and `nil` is returned. Otherwise the *forms* constitute an implicit `progn`, and so are evaluated sequentially from left to right, and the value of the last one is returned.

```
(when p a b c) <=> (and p (progn a b c))
(when p a b c) <=> (cond (p a b c))
(when p a b c) <=> (if p (progn a b c) 'nil)
(when p a b c) <=> (unless (not p) a b c)
```

As a matter of style, `when` is normally used to conditionally produce some side effects, and the value of the `when`-form is normally not used. If the value is relevant, then `and` (page 52) or `if` (page 69) may be stylistically more appropriate.

`unless` *pred* {*form*}* [*Special form*]

(`unless` *pred form1 form2 ...*) first evaluates *pred*. If the result is *not* `nil`, then the *forms* are not evaluated, and `nil` is returned. Otherwise the *forms* constitute an implicit `progn`, and so are evaluated sequentially from left to right, and the value of the last one is returned.

```
(unless p a b c) <=> (cond ((not p) a b c))
(unless p a b c) <=> (if p nil (progn a b c))
(unless p a b c) <=> (when (not p) a b c)
```

As a matter of style, `unless` is normally used to conditionally produce some side effects, and the value of the `unless`-form is normally not used. If the value is relevant, then `or` (page 52) or `if` (page 69) may be stylistically more appropriate.

`case` *keyform* {(({*key*}*) {*form*}*)}* [*Special form*]

`case` is a conditional that chooses one of its clauses to execute by comparing a value to various constants, which are typically keyword symbols, integers, or characters (but may be any objects). Its form is as follows:

```
(case keyform
  (keylist-1 consequent-1-1 consequent-1-2 ...)
  (keylist-2 consequent-2-1 ...)
  (keylist-3 consequent-3-1 ...)
  ...)
```

Structurally `case` is much like `cond` (page 68), and it behaves like `cond` in selecting one clause and then executing all consequents of that clause. It differs in the mechanism of clause selection.

The first thing `case` does is to evaluate the form *keyform* to produce an object called the *key object*. Then `case` considers each of the clauses in turn. If *key* is in the *keylist* (that is, is `eql` to any item in the *keylist*) of a clause, the consequents of that clause are evaluated as an implicit `progn`, and `case` returns what was returned by the last consequent (or `nil` if there are no consequents in that clause). If no clause is satisfied, `case` returns `nil`.

It is an error for the same key to appear in more than one clause.

Instead of a *keylist*, one may write one of the symbols `t` and `otherwise`. A clause with such a symbol always succeeds, and must be the last clause.

> **Compatibility note:** Lisp Machine LISP uses `eq` for the comparison. In Lisp Machine LISP `case` therefore works for fixnums but not bignums. In the interest of hiding the fixnum-bignum distinction, `case` uses `eql` in COMMON LISP.

If there is only one key for a clause, then that key may be written in place of a list of that key, provided that no ambiguity results (the key should not be a cons or one of `nil` (which is confusable with `()`, a list of no keys), `t`, or `otherwise`).

`typecase` *keyform* {(*type* {*form*}*)}* [*Special form*]

`typecase` is a conditional which chooses one of its clauses by examining the type of an object. Its form is as follows:

```
(typecase keyform
  (type-1 consequent-1-1 consequent-1-2 ...)
  (type-2 consequent-2-1 ...)
  (type-3 consequent-3-1 ...)
  ...)
```

Structurally `typecase` is much like `cond` (page 68) or `case` (page 70), and it behaves like them in selecting one clause and then executing all consequents of that clause. It differs in the mechanism of clause selection.

The first thing `typecase` does is to evaluate the form *keyform* to produce an object called the key object. Then `typecase` considers each of the clauses in turn. The first clause for which the key is of that clause's specified *type* is selected, the consequents of this clause are evaluated as an implicit `progn`, and `typecaseq` returns what was returned by the last consequent (or `nil` if there are no consequents in that clause). If no clause is satisfied, `typecase` returns `nil`.

As for `case`, the symbol `t` or `otherwise` may be written for *type* to indicate that the clause should always be selected.

It is permissible for more than one clause to specify a given type, particularly if one is a subtype of another; the earliest applicable clause is chosen.

For example:

```
(typecase an-object
   (string ...)                    ;This clause handles strings.
   ((array t) ...)                 ;This clause handles general arrays.
   ((array bit) ...)               ;This clause handles bit arrays.
   (array ...)                     ;This handles all other arrays.
   ((or list number) ...)          ;This handles lists and numbers.
   (t ...))                        ;This handles all other objects.
```

A COMMON LISP compiler may choose to issue a warning if a clause cannot be selected because it is completely shadowed by earlier clauses.

7.7. Blocks and Exits

`block` *name* {*form*}* [*Special form*]

The `block` construct executes each *form* from left to right, returning whatever is returned by the last *form*. If, however, a `return` or `return-from` form is executed during the execution of some *form*, then the results specified by the `return` or `return-from` are immediately returned as the value of the `block` construct, and execution proceeds as if the `block` had terminated normally. In this `block` differs from `progn` (page 64); the latter has nothing to do with `return`.

The scope of the *name* is lexical; only a `return` or `return-from` textually contained in some *form* can exit from the block. The extent of the name is dynamic. Therefore it is only possible to exit from a given run-time incarnation of a block once, either normally or by explicit return.

The `defun` (page 42) form implicitly puts a `block` around the body of the function defined; the `block` has the same name as the function. Therefore one may use `return` or `return-from` to

return prematurely from a function defined by `defun`.

`return` *result* [*Special form*]

`return` is used to return from a `block`, `prog`, `do`, or similar construct. Whatever the evaluation of *result* produces is returned by

the construct being exited by `return`.

```
(defun member (item list)
       (do ((x list (cdr x)))
           ((null x) nil)
         (when (equal item (car x))
               (return x))))
```

`return` is, like `go`, a special form that does not return a value. Instead, it causes a containing construct to return a value. If the evaluation of *result* produces multiple values, those multiple values are returned by the construct exited.

`return` always exits from the innermost applicable construct that textually contains it. However, if the symbol `t` is used as the name of a block, then that block will be made "invisible" to `return` forms; any `return` inside that block will return to the next outermost level whose name is not `t`. `(return-from t ...)` will return from a block named `t`. This feature is not intended to be used by user-written code; it is for macros to expand into.

`return-from` *blockname result* [*Special form*]

This is just like `return`, except that before the *result* form is written a symbol (not evaluated), which is the name of the construct from which to return.

??? Query: Fahlman suggests a `restart` form that specifies a block and send control to the top of that block. Some kinds of loop can be done this way, especially error retries. Opinions?

7.8. Iteration

COMMON LISP provides a number of iteration constructs. The `do` (page 73) and `do*` (page 75) constructs provides a general iteration facility. For simple iterations over lists or *n* consecutive integers, `dolist` (page 76) and related constructs are provided. The `prog` (page 78) construct is the most general, permitting arbitrary `go` (page 80) statements within it. All of the iteration constructs permit statically defined non-local exits in the form of the `return` (page 72) statement and its variants.

7.8.1. General iteration

`do` ({(*var* [*init* [*step*]])}*) (*end-test* {*form*}*) {*tag* | *statement*}* [*Special form*]

The `do` special form provides a generalized iteration facility, with an arbitrary number of "index variables". These variables are bound within the iteration and stepped in parallel in specified ways. They may be used both to generate successive values of interest (such as successive integers) or to accumulate results. When an end condition is met, the iteration terminates with a specified value.

In general, a `do` loop looks like this:

```
(do ((var1 init1 step1)
     (var2 init2 step2)
     ...
     (varn initn stepn))
    (end-test . result)
    . progbody)
```

The first item in the form is a list of zero or more index-variable specifiers. Each index-variable specifier is a list of the name of a variable *var*, an initial value *init* (which defaults to `nil` if it is omitted) and a stepping form *step*. If *step* is omitted, the *var* is not changed by the `do` construct between repetitions (though code within the `do` is free to alter the value of the variable by using `setq` (page 58)).

An index-variable specifier can also be just the name of a variable. In this case, the variable has an initial value of `nil`, and is not changed between repetitions.

Before the first iteration, all the *init* forms are evaluated, and then each *var* is bound to the value of its respective *init*. This is a binding, not an assignment; when the loop terminates the old values of those variables will be restored. Note that *all* of the *init* forms are evaluated *before* any *var* is bound; hence *init* forms may refer to old values of the variables.

The second element of the `do`-form is a list of an end-testing predicate form *end-test*, and zero or more forms, called the *result* forms. This resembles a `cond` clause. At the beginning of each iteration, after processing the variables, the *end-test* is evaluated. If the result is `nil`, execution proceeds with the body of the `do`. If the result is not `nil`, the *result* forms are evaluated in order as an implicit `progn` (page 64), and then `do` returns. `do` returns the results of evaluating the last *result* form. If there are no *result* forms, the value of `do` is `nil`; note that this is *not* quite analogous to the treatment of clauses in a `cond` (page 68) special form.

At the beginning of each iteration other than the first, the index variables are updated as follows. First every *step* form is evaluated, from left to right. Then the resulting values are assigned (as with `psetq` (page 58)) to the respective index variables. Any variable which has no associated *step* form is not affected. Because *all* of the *step* forms are evaluated before *any* of the variables are altered, when a step form is evaluated it always has access to the *old* values of the index variables, even if other step forms precede it. After this process, the end-test is evaluated as described above.

If the end-test of a `do` form is `nil`, the test will never succeed. Therefore this provides an idiom for "do forever". The *body* of the `do` is executed repeatedly, stepping variables as usual, of course. The infinite loop can be terminated by the use of `return` (page 72), `return-from` (page 72), `go` (page 80) to an outer level, or `throw` (page 87).

For example:

```
(do ((j 0 (+ j 1)))
    (nil)                          ;Do forever.
  (format t "~%Input ~D:" j)
  (let ((item (read)))
    (if (null item) (return)       ;Process items until nil seen.
        (format t "~&Output ~D: ~S" j (process item)))))
```

The remainder of the do form constitutes a prog body. The function return (page 72) and its variants may be used within a do form to terminate it immediately, returning a specified result. Tags may appear within the body of a do loop for use by go (page 80) statements. When the end of a do body is reached, the next iteration cycle (beginning with the evaluation of *step* forms) occurs.

declare (page 95) forms may appear at the beginning of a do body. They apply to code in the do body, to the bindings of the do variables, to the *step* forms (but *not* the *init* forms), to the *end-test*, and to the *result* forms.

> **Compatibility note:** "Old-style" MACLISP do loops, of the form (do *var init step end-test . body*), are not supported. They are obsolete, and are easily converted to a new-style do with the insertion of three pairs of parentheses. In practice the compiler can catch nearly all instances of old-style do loops because they will not have a legal format anyway.

For example:

```
(do ((i 0 (+ i 1))          ;Sets every element of an-array to empty
     (n (array-length an-array)))
    ((= i n))
  (aset 'empty an-array i))
```

The construction

```
(do ((x e (cdr x))
     (oldx x x))
    ((null x))
  body)
```

exploits parallel assignment to index variables. On the first iteration, the value of oldx is whatever value x had before the do was entered. On succeeding iterations, oldx contains the value that x had on the previous iteration.

Very often an iterative algorithm can be most clearly expressed entirely in the *step* forms of a do, and the *body* is empty.

For example:

```
(do ((x foo (cdr x))
     (y bar (cdr y))
     (z '() (cons (f (car x) (car y)) z)))
    ((or (null x) (null y))
     (nreverse z)))
```

does the same thing as (mapcar #'f foo bar). Note that the *step* computation for z exploits the fact that variables are stepped in parallel. Also, the body of the loop is empty. Finally, the use of nreverse (page 158) to put an accumulated do loop result into the correct order is a standard idiom.

Other examples:

```
(defun length (list)
       (do ((x list (cdr x))
            (j 0 (+ j 1)))
           ((atom x) j)))

(defun reverse (list)
       (do ((x list (cdr x))
            (y '() (cons (car x) y)))
           ((atom x) y)))
```

Note the use of `atom` rather than `null` to test for the end of a list in the above two examples. This results in more robust code; it will not attempt to cdr the end of a dotted list.

As an example of nested loops, suppose that `env` holds a list of conses. The *car* of each cons is a list of symbols, and the *cdr* of each cons is a list of equal length containing corresponding values. Such a data structure is similar to an association list, but is divided into "frames"; the overall structure resembles a rib-cage. A lookup function on such a data structure might be:

```
(defun ribcage-lookup (sym ribcage)
       (do ((r ribcage (cdr r)))
           ((null r) nil)
         (do ((s (caar r) (cdr s))
              (v (cdar r) (cdr v)))
             ((null s))
           (when (eq (car s) sym)
             (return-from ribcage-lookup (car v))))))
```

(Notice the use of indentation in the above example to set off the bodies of the `do` loops.)

`do*` *bindspecs endtest* {*form*}* [*Special form*]

`do*` is exactly like `do` except that the bindings and steppings of the variables are performed sequentially rather than in parallel. At the beginning each variable is bound to the value of its *init* form before the *init* form for the next variable is evaluated. Similarly, between iterations each variable is given the new value computed by its *step* form before the *step* form of the next variable is evaluated.

7.8.2. Simple Iteration Constructs

The constructs `dolist` and `dotimes` perform a body of statements repeatedly. On each iteration a specified variable is bound to an element of interest which the body may examine. `dolist` examines successive elements of a list, and `dotimes` examines integers from *0* to *n-1* for some specified positive integer *n*.

The value of any of these constructs may be specified by an optional result form, which if omitted defaults to the value `nil`.

The `return` (page 72) statement may be used to return immediately from a `dolist` or `dotimes` form, discarding any following iterations which might have been performed; in effect, a `block` with an inaccessible

name surrounds the construct. The body of the loop is in fact a `prog` (page 78) body; it may contain tags to serve as the targets of `go` (page 80) statements, and may have `declare` (page 95) forms at the beginning.

`dolist` (*var listform* [*resultform*]) {*tag* | *statement*}* [*Special form*]

`dolist` provides straightforward iteration over the elements of a list. The expression (`dolist` (*var listform resultform*) . *progbody*) evaluates the form *listform*, which should produce a list. It then performs *progbody* once for each element in the list, in order, with the variable *var* bound to the element. Then *resultform* (a single form, *not* an implicit `progn`) is evaluated, and the result is the value of the `dolist` form. If *resultform* is omitted, the result is `nil`.

For example:

```
(dolist (x '(a b c d)) (prin1 x) (princ " ")) => nil
   after printing "a b c d "
```

An explicit `return` statement may be used to terminate the loop and return a specified value.

Compatibility note: The *resultform* part of a `dolist` is not currently supported in Lisp Machine LISP. It seems to improve the utility of the construct markedly.

`dotimes` (*var countform* [*resultform*]) {*tag* | *statement*}* [*Special form*]

`dotimes` provides straightforward iteration over a sequence of integers. The expression (`dotimes` (*var countform resultform*) *progbody*) evaluates the form *countform*, which should produce an integer. It then performs *progbody* once for each integer from zero (inclusive) to *count* (exclusive), in order, with the variable *var* bound to the integer; if the integer is zero or negative, then the *progbody* is performed zero times. Finally, *resultform* (a single form, *not* an implicit `progn`) is evaluated, and the result is the value of the `dotimes` form. If *resultform* is omitted, the result is `nil`.

Altering the value of *var* in the body of the loop (by using `setq` (page 58), for example) will have unpredictable, possibly implementation-dependent results. A COMMON LISP compiler may choose to issue a warning if such a variable appears in a `setq`.

For example:

```
(defun string-posq (char string &optional
                    (start 0)
                    (end (string-length string)))
   (dotimes (k (- end start) nil)
            (when (char= char (char string (+ start k)))
                  (return k))))
```

An explicit `return` statement may be used to terminate the loop and return a specified value.

See also `do-symbols` (page 116) and related constructs.

7.8.3. Mapping

Mapping is a type of iteration in which a function is successively applied to pieces of one or more sequences. The result of the iteration is a sequence containing the respective results of the function applications. There are several options for the way in which the pieces of the list are chosen and for what is done with the results returned by the applications of the function.

The function `map` (page 159) may be used to map over any kind of sequence. The following functions operate only on lists.

`mapcar` *function list* `&rest` *more-lists* [*Function*]
`maplist` *function list* `&rest` *more-lists* [*Function*]
`mapc` *function list* `&rest` *more-lists* [*Function*]
`mapl` *function list* `&rest` *more-lists* [*Function*]
`mapcan` *function list* `&rest` *more-lists* [*Function*]
`mapcon` *function list* `&rest` *more-lists* [*Function*]

For each these mapping functions, the first argument is a function and the rest must be lists. The function must take as many arguments as there are lists.

`mapcar` operates on successive elements of the lists. First the function is applied to the *car* of each list, then to the *cadr* of each list, and so on. (Ideally all the lists are the same length; if not, the iteration terminates when the shortest list runs out, and excess elements in other lists are ignored.) The value returned by `mapcar` is a list of the results of the successive calls to the function.

For example:

```
(mapcar #'abs '(3 -4 2 -5 -6)) => (3 4 2 5 6)
(mapcar #'cons '(a b c) '(1 2 3)) => ((a . 1) (b . 2) (c . 3))
```

`maplist` is like `mapcar` except that the function is applied to the list and successive cdr's of that list rather than to successive elements of the list.

For example:

```
(maplist #'(lambda (x) (cons 'foo x))
         '(a b c d))
   => ((foo a b c d) (foo b c d) (foo c d) (foo d))
(maplist #'(lambda (x) (if (member (car x) (cdr x)) 0 1)))
         '(a b a c d b c))
   => (0 0 1 0 1 1 1)
   ;An entry is 1 iff the corresponding element of the input
   ;    list was the last instance of that element in the input list.
```

`mapl` and `mapc` are like `maplist` and `mapcar` respectively, except that they do not accumulate the results of calling the function.

> **Compatibility note:** In all LISP systems since LISP 1.5, `mapl` has been called `map`. In the chapter on sequences it is explained why this was a bad choice. Here the name `map` is used for the far more useful generic sequence mapper, in closer accordance to the computer science literature, especially the growing body of papers on functional programming.

These functions are used when the function is being called merely for its side-effects, rather than its returned values. The value returned by `mapl` or `mapc` is the second argument, that is, the first

sequence argument.

`mapcan` and `mapcon` are like `mapcar` and `maplist` respectively, except that they combine the results of the function using `nconc` (page 171) instead of `list`. That is,

```
(mapcon f x1 ... xn)
   <=> (apply #'nconc (maplist f x1 ... xn))
```

and similarly for the relationship between `mapcan` and `mapcar`. Conceptually, these functions allow the mapped function to return a variable number of items to be put into the output list. This is particularly useful for effectively returning zero or one item:

```
(mapcan #'(lambda (x) (and (numberp x) (list x)))
        '(a 1 b c 3 4 d 5))
   => (1 3 4 5)
```

In this case the function serves as a filter; this is a standard LISP idiom using `mapcan`. (The function `remove-if-not` (page 160) might have been useful in this particular context, however.) Remember that `nconc` is a destructive operation, and therefore so are `mapcan` and `mapcon`; the lists returned by the *function* are altered in order to concatenate them.

Sometimes a `do` or a straightforward recursion is preferable to a mapping operation; however, the mapping functions should be used wherever they naturally apply because this increases the clarity of the code.

The functional argument to a mapping function must be acceptable to `apply`; it cannot be a macro or the name of a special form. Of course, there is nothing wrong with using functions which have `&optional` and `&rest` parameters.

7.8.4. The Program Feature

LISP implementations since LISP 1.5 have had what was originally called "the program feature", as if it were impossible to write programs without it! The `prog` construct allows one to write in an ALGOL-like or FORTRAN-like statement-oriented style, using `go` statements, which can refer to tags in the body of the `prog`. Modern LISP programming style tends to use `prog` rather infrequently. The various iteration constructs, such as `do` (page 73), have bodies with the characteristics of a `prog`.

`prog` ({*var* | (*var* [*init*])}*) {*tag* | *statement*}* [*Special form*]

`prog` is a special form that provides bound temporary variables, sequential evaluation of forms, and a "goto/return" facility. It is this latter characteristic that distinguishes `prog` from other LISP constructs; `lambda` and `let` (page 65) also provide local variable bindings, and `progn` (page 64) also evaluates forms sequentially.

A typical `prog` looks like:

```
(prog (var1 var2 (var3 init3) var4 (var5 init5))
      statement1
 tag1
      statement2
      statement3
      statement4
 tag2
      statement5
      ...
      )
```

The list after the keyword `prog` is a set of specifications for binding *var1*, *var2*, etc., which are temporary variables, bound locally to the `prog`. This list is processed exactly as the list in a `let` (page 65) statement: first all the *init* forms are evaluated from left to right (where `nil` is used for any omitted *init* form), and then the variables are all bound in parallel to the respective results. (`prog*` (page 80) is the same as `prog` except that this initialization is sequential rather than parallel.)

The part of a `prog` after the variable list is called the *body*. An item in the body may be a symbol or a number, in which case it is called a *tag*, or any other COMMON LISP form, in which case it is called a *statement*.

After `prog` binds the temporary variables, it processes each form in its body sequentially. *tags* are ignored; *statements* are evaluated, and their returned values discarded. If the end of the body is reached, the `prog` returns `nil`. However, two special forms may be used in `prog` bodies to alter the flow of control. If (`return` *x*) is evaluated, `prog` stops processing its body, evaluates *x*, and returns the result. If (`go` *tag*) is evaluated, `prog` jumps to the part of the body labelled with the *tag* (that is, with an atom `eql` (page 49) to *tag*). *tag* is not evaluated.

> **Compatibility note:** The "computed go" feature of MACLISP is not supported. The syntax of a computed `go` is idiosyncratic, and the feature is not supported by Lisp Machine LISP, NIL, or INTERLISP.

`go` and `return` forms must be *lexically* within the scope of the `prog`; it is not possible for one function to `return` to a `prog` that is in progress in its caller. Thus, a program that contains a `go` not contained within the body of a `prog` (or other constructs such as `do`, which have `prog` bodies) is in error. A dynamically scoped non-local exit mechanism is provided by `catch` (page 85) and `throw` (page 87) and other related operations.

Here is a fine example of what can be done with `prog`:

```
(defun king-of-confusion (w)
       (prog (x y z)     ;Initialize x, y, z to nil
             (setq y (car w) z (cdr w))
        loop
             (cond ((null y) (return x))
                   ((null z) (go err)))
        rejoin
             (setq x (cons (cons (car y) (car z)) x))
             (setq y (cdr y) z (cdr z))
             (go loop)
        err
             (error "Mismatch - gleep!")
             (setq z y)
             (go rejoin))
```

which is accomplished somewhat more perspicuously by:

```
(defun prince-of-clarity (w)
       (do ((y (car w) (cdr y))
            (z (cdr w) (cdr z))
            (x '() (cons (cons (car y) (car z)) x)))
           ((null y) x)
         (when (null z)
               (error "Mismatch - gleep!")
               (setq z y))))
```

Declarations may appear at the beginning of a `prog` body; see `declare` (page 95).

`prog*` [*Special form*]

The `prog*` special form is almost the same as `prog`. The only difference is that the binding and initialization of the temporary variables is done *sequentially*, so that the *init* form for each one can use the values of previous ones. Therefore `prog*` is to `prog` as `let*` (page 66) is to `let` (page 65).

For example:

```
(prog* ((y z) (x (car y)))
       (return x))
```

returns the car of the value of `z`.

`go` *tag* [*Special form*]

The (`go` *tag*) special form is used to do a "go to" within a a `prog` body. The *tag* must be a symbol or a number; *tag* is not evaluated. `go` transfers control to the point in the body labelled by a tag `equal` to the one given. If there is no such tag in the body, the bodies of lexically containing `prog` bodies (if any) are examined as well. It is an error if there is no matching tag.

The `go` form does not ever return a value. A `go` form may not appear as an argument to an ordinary function, but only at the top level of a `prog` body or within certain special forms such as conditionals which are within a `prog` body.

For example:

```
(prog ((n (string-length a-string)) (j 0))
 loop (when (= j n) (return a-string))
      (when (char= #\Space (char j a-string))
            (return (substring a-string 0 j)))
      (increment j)
      (go loop))
```

returns the first "word" in `a-string`, where words are separated by spaces. This could of course have been expressed more succinctly as:

```
(dotimes (j (string-length a-string) a-string)
  (when (char= #\Space (char j a-string))
        (return (substring a-string 0 j))))
```

As a matter of style, it is recommended that the user think twice before using a `go`. Most purposes of `go` can be accomplished with one of the iteration primitives, nested conditional forms, or `return-from` (page 72). If the use of `go` seems to be unavoidable, perhaps the control structure implemented by `go` should be packaged up as a macro definition. (If the use of `go` is avoidable, and `return` also is not needed, then `prog` probably is not needed either; `let` can be used to bind variables and then execute some statements.)

7.9. Multiple Values

Ordinarily the result of calling a LISP function is a single LISP object. Sometimes, however, it is convenient for a function to compute several quantities and return them. COMMON LISP provides a mechanism for handling multiple values directly. This mechanism is cleaner and more efficient than the usual tricks involving returning a list of results or stashing results in global variables.

7.9.1. Constructs for Handling Multiple Values

Normally multiple values are not used. Special forms are required both to *produce* multiple values and to *receive* them. If the caller of a function does not request multiple values, but the called function produces multiple values, then the first value is given to the caller and all others are discarded (and if the called function produces zero values then the caller gets `nil` as a value).

The primary primitive for producing multiple values is `values` (page 82), which takes any number of arguments and returns that many values. If the last form in the body of a function is a `values` with three arguments, then a call to that function will return three values. Other special forms also produce multiple values, but they can be described in terms of `values`. Some built-in COMMON LISP functions (such as `floor` (page 131)) return multiple values; those which do are so documented.

The special forms for receiving multiple values are `multiple-value-bind` (page 82), `multiple-value` (page 83), and `multiple-value-list` (page 82). These specify a form to evaluate and an indication of where to put the values returned by that form.

`values &rest` *args* [*Function*]

Returns all of its arguments, in order, as values.

For example:

```
(defun polar (x y)
       (values (sqrt (+ (* x x) (* y y))) (atan y x)))
(multiple-value-let (r theta) (polar 3.0 4.0)
     (list r theta))
  => (5.0 0.9272952)
```

The expression `(values)` returns zero values.

`values-list` *list* [*Function*]

Returns as multiple values all the elements of *list.*

For example:

```
(values-list (list a b c)) <=> (values a b c)
```

`multiple-value-list` *form* [*Special form*]

`multiple-value-list` evaluates *form,* and returns a list of the multiple values it returned.

For example:

```
(multiple-value-list (floor -3 4)) => (-1 1)
```

`mvcall` *function* {*form*}* [*Special form*]

`mvcall` first evaluates *function* to obtain a function, and then evaluates all of the *forms.* All the the values of the *forms* are gathered together (not just one value from each), and given as arguments to the function. The result of `mvcall` is whatever is returned by the function.

For example:

```
(mvcall #'+ (floor 5 3) (floor 7 3)) <=> (+ 1 2 2 1) => 6
(multiple-value-list form) <=> (mvcall #'list form)
```

`mvprog1` *form* {*form*}* [*Special form*]

`mvprog1` evaluates the first *form* and saves all the values produced by that form. It then evaluates the other *forms* from left to right, discarding their values. The values produced by the first *form* are returned by `mvprog1`. See `prog1` (page 65), which always returns a single value.

`multiple-value-bind` ({*var*}*) *values-form* {*form*}* [*Special form*]

The *values-form* is evaluated, and each of the variables *var* is bound to the respective value returned by that form. If there are more variables than values returned, extra values of `nil` are given to the remaining variables. If there are more values than variables, the excess values are simply discarded. The variables are bound to the values over the execution of the forms, which make up an implicit `progn`.

Compatibility note: This is compatible with Lisp Machine LISP.

For example:

```
(multiple-value-bind (x) (floor 5 3) (list x)) => (1)
(multiple-value-bind (x y) (floor 5 3) (list x y)) => (1 2)
(multiple-value-bind (x y z) (floor 5 3) (list x y z))
   => (1 2 nil)
```

In general,

```
(multiple-value-bind (x y z ...) form . body)
   <=>
(mvlet (&optional (x nil) (y nil) (z nil) ... &rest dummy)
  (declare (ignored dummy))
  form . body)
```

`multiple-value` *variables form* [*Special form*]

The *variables* must be a list of variables. The *form* is evaluated, and the variables are *set* (not bound) to the values returned by that form. If there are more variables than values returned, extra values of `nil` are assigned to the remaining variables. If there are more values than variables, the excess values are simply discarded.

Compatibility note: This is compatible with Lisp Machine LISP.

`multiple-value` always returns a single value, which is the first value returned by *form*, or `nil` if *form* produces zero values.

7.9.2. Rules for Tail-Recursive Situations

It is often the case that the value of a special form is defined to be the value of one of its sub-forms. For example, the value of a `cond` is the value of the last form in the selected clause. In most such cases, if the sub-form produces multiple values, the original form will also produce all of those values. This *passing-back* of multiple values of course has no effect unless eventually one of the special forms for receiving multiple values is reached.

To be explicit, multiple values can result from a special form under precisely these circumstances:

- `eval` (page 209) returns multiple values if the form given it to evaluate produces multiple values.
- `apply` (page 63), `funcall` (page 64), `funcall*` (page 64), `mvcall` (page 82), `subrcall` (page SUBRCALL-FUN), and `subrcall*` (page SUBRCALL*-FUN) pass back multiple values from the function applied or called.
- When a `lambda` (page LAMBDA-FUN)-expression is invoked, the function passes back multiple values from the last form of the `lambda` body (which is an implicit `progn`).
- Indeed, `progn` (page 64) itself passes back multiple values from its last form, as does any construct some part of which is defined to be an "implicit `progn`"; these include `progv` (page 67), `let` (page 65), `let*` (page 66),`when` (page 69), `unless` (page 70), `case` (page 70), `typecase` (page 70), `multiple-value-bind` (page 82), `multiple-value` (page 83), `catch` (page 85), and `catch-all` (page 85).

- `mvprog1` (page 82) passes back multiple values from its first form. However, `prog1` (page 65) always returns a single value.

- `unwind-protect` (page 86) returns multiple values if the form it protects does.

- `catch` (page 85) returns multiple values if the result form in a `throw` (page 87) exiting from such a catch produces multiple values.

- `cond` (page 68) passes back multiple values from the last form of the implicit `progn` of the selected clause. If, however, the clause selected is a singleton clause, then only a single value (the non-`nil` predicate value) is returned. This is true even if the singleton clause is the last clause of the `cond`. It is *not* permitted to treat a final clause "`(x)`" as being the same as "`(t x)`" for this reason; the latter passes back multiple values from the form `x`.

- `if` (page 69) passes back multiple values from whichever form is selected (the *then* form or the *else* form).

- `and` (page 52) and `or` (page 52) pass back multiple values from the last form, but not from forms other than the last.

- `do` (page 73), `prog` (page 78), `prog*` (page 80), and other constructs from which `return` (page 72) can return, each pass back the multiple values of the form appearing in In addition, `do` passes back multiple values from the last form of the exit clause, exactly as if the exit clause were a `cond` clause.

Among special forms which *never* pass back multiple values are `setq` (page 58), `multiple-value` (page 83), and `prog1` (page 65). A good way to force only one value to be returned from a form `x` is to write `(values x)`.

The most important rule about multiple values, however, is:

No matter how many values a form produces,
if the form is an argument form in a function call,
then exactly ONE value (the first one) is used.

For example, if you write `(cons (foo x))`, then `cons` will receive *exactly* one argument (which is of course an error), even if `foo` returns two values. To pass both values from `foo` to `cons`, one must use a special form, such as `(mvcall #'cons (foo x))`. In an ordinary function call, each argument form produces exactly *one* argument; if such a form returns zero values, `nil` is used for the argument, and if more than one value, all but the first are discarded. Similarly, conditional constructs which test the value of a form will use exactly one value (the first) from that form and discard the rest, or use `nil` if zero values are returned.

7.10. Dynamic Non-local Exits

COMMON LISP provides a facility for exiting from a complex process in a non-local, dynamically scoped manner. There are two classes of special forms for this purpose, called *catch* forms and *throw* forms, or simply *catches* and *throws*. A catch form evaluates some subforms in such a way that, if a throw form is executed during such evaluation, the evaluation is aborted at that point and the catch form immediately returns a value specified by the throw. Unlike `block` (page 71) and `return` (page 72), which allow for so exiting a `block` form from any point lexically within the body of the `block`, the catch/throw mechanism works even if the throw form is not textually within the body of the catch form. The throw need only occur within the extent (time span) of the evaluation of the body of the catch. This is analogous to the distinction between dynamically bound (special) variables and lexically bound (local) variables.

7.10.1. Catch Forms

`catch` *tag* {*form*}* [*Special form*]

The `catch` special form is the simplest catcher. The *tag* is evaluated first to produce an object that names the catch; it may be any LISP object. The *forms* are evaluated as an implicit `progn`, and the results of the last form are returned, except that if during the evaluation of the *forms* a throw should be executed, such that the *tag* of the throw matches (is `eq` to) the *tag* of the `catch`, then the evaluation of the *forms* is aborted and the results specified by the throw are immediately returned from the `catch` expression.

The tag is used to match up throws with catches (using `eq`, not `eql`; therefore numbers should not be used as catch tags). (`catch 'foo` *form*) will catch a (`throw 'foo` *form*) but not a (`throw 'bar` *form*). It is an error if `throw` is done when there is no suitable `catch` (or one of its variants) ready to catch it.

Compatibility note: This syntax for `catch` is not compatible with MACLISP. Lisp Machine LISP defines `catch` to be compatible with that of MACLISP, but discourages its use. The definition here is compatible with NIL.

`catch-all` *catch-function* {*form*}* [*Special form*]

`unwind-all` *catch-function* {*form*}* [*Special form*]

`catchall` behaves roughly like `catch`, except that instead of a *tag*, a *catch-function* is provided. If no throw occurs during the evaluation of the *forms*, then this behaves just as for `catch`: the `catchall` form returns what is returned from evaluation of the last of the *forms*. `catch-all` will catch *any* throw not caught by some inner catcher, however; if such a throw occurs, then the function is called, and whatever it returns is returned by `catch-all`. The *catch-function* will get one or more arguments; the first argument is always the throw tag, and the other arguments are the thrown results (there may be more than one if the *result* form for the throw produces multiple values).

The `catch-all` is not in force during execution of the *catch-function*. If a throw occurs within the *catch-function*, it will throw to some catch exterior to the `catch-all`. This is useful because the *catch-function* can examine the tag, and if it is not of interest can relay the throw.

```
(catch-all #'(lambda (tag &rest results)
               (caseq tag                          ;Check tag.
                 (win (values-list results))       ;If win, return results.
                 (lose (cleanup)                   ;If lose, clean up
                       (ferror "Lose lose!"))      ; and signal an error.
                 (otherwise                        ;Otherwise relay throw.
                  (throw tag (values-list results)))))
             (determine-win-or-lose))
```

`unwind-all` is just like `catch-all` except that the *catch-function* is always called, even if no throw occurs; in that case the first argument (the "tag") to the *catch-function* is `nil`, and the other arguments are the results from the last of the *forms*. Often `unwind-protect` is more suitable for a given task than `unwind-all`, however; the choice should be weighed for any particular application.

??? **Query:** Oooooops, there's a problem with these. What tag is supplied if what is causing the exit is a `go` or `return` from within the body to some tag or `block` outside the catcher? In MACLISP, a `go` from within a `catchall` quietly breaks up the `catchall` frame without invoking the `catchall` function, which means that it catches all throws but not all exits!

`unwind-protect` *protected-form* {*cleanup-form*}* [*Special form*]

Sometimes it is necessary to evaluate a form and make sure that certain side-effects take place after the form is evaluated; a typical example is:

```
(progn (start-motor)
       (drill-hole)
       (stop-motor))
```

The non-local exit facility of Lisp creates a situation in which the above code won't work, however: if `drill-hole` should do a throw to a catch which is outside of the `progn` form (perhaps because the drill bit broke), then `(stop-motor)` will never be evaluated (and the motor will presumably be left running). This is particularly likely if `drill-hole` causes a LISP error and the user tells the error-handler to give up and abort the computation. (A possibly more practical example might be:

```
(prog2 (open-a-file)
       (process-file)
       (close-the-file))
```

where it is desired always to close the file when the computation is terminated for whatever reason.)

In order to allow the above program to work, it can be rewritten using `unwind-protect` as follows:

```
(unwind-protect
  (progn (start-motor)
         (drill-hole))
  (stop-motor))
```

If `drill-hole` does a throw which attempts to quit out of the `unwind-protect`, then `(stop-motor)` will be executed.

As a general rule, `unwind-protect` guarantees to execute all the *cleanup-forms* before exiting, whether it terminates normally or is aborted by a throw of some kind. `unwind-protect` returns

whatever results from evaluation of the *protected-form*, and discards all the results from the *cleanup-forms*.

7.10.2. Throw Forms

`throw` *tag result* [*Special form*]

The `throw` special form is the only explicit thrower in COMMON LISP. (However, errors may cause throws to occur also.) The *tag* is evaluated first to produce an object called the throw tag. The most recent outstanding catch whose tag matches the throw tag is exited. Some catches, such as a `catch-all`, will match any throw tag; a `catch` matches only if the catch tag is `eq` to the throw tag.

In the process dynamic variable bindings are undone back to the point of the catch, and any intervening `unwind-protect` cleanup code is executed. The *result* form is evaluated before the unwinding process commences, and whatever results it produces are returned from the catch (or given to the *catch-function*, if appropriate).

If there is no outstanding catch whose tag matches the throw tag, no unwinding of the stack is performed, and an error is signalled. When the error is signalled, the outstanding catches and the dynamic variable bindings are those in force at the point of the throw.

Implementation note: These requirements imply that throwing must be done by two passes over the control stack. In the first pass one simply searches for a matching catch. In this search every `catch`, `catch-all`, and `unwind-all` must be considered, but every `unwind-protect` should be ignored. On the second pass the stack is actually unwound, one frame at a time, undoing dynamic bindings and outstanding `unwind-protect` in reverse order of creation until the matching catch is reached.

Chapter 8

Macros

The COMMON LISP macro facility allows the user to define arbitrary functions that convert certain LISP forms into different forms before evaluating or compiling them. This is done at the S-expression level, not at the character-string level as in most other languages. Macros are important in the writing of good code: they make it possible to write code that is clear and elegant at the user level, but that is converted to a more complex or more efficient internal form for execution.

When `eval` (page 209) is given a list whose *car* is a symbol, it looks in the definition cell of that symbol. If the definition is itself an object that satisfies the pseudo-predicate `macro-p` (page 57), then the original list is said to be a *macro call.* The non-`nil` result of `macro-p` will be a function of one argument, called the *expansion function.* This function is called with the entire macro call as its one argument; it must return some new LISP form, called the *expansion* of the macro call. This expansion is then evaluated in place of the original form.

When a function is being compiled, any macros it contains are expanded at compilation time. This means that a macro definition must be seen by the compiler before the first use of the macro. Macros cannot be used as functional arguments to such things as `apply` (page 63), `funcall` (page 64), or `map` (page 159); in such situations, the list representing the "original macro call" does not exist, so the expansion function would not know what to work on.

8.1. Defining Macros

`macro` *name* (*var*) {*form*}* [*Macro*]

The primitive special form for defining a macro is `macro`. Note, however, that the use of `macro` is often very awkward, and it is preferable to use `defmacro` in almost all circumstances. A call to `macro` has the following form:

(`macro` *name* (*var*) . *body*)

This is very similar in form to a `defun` form: *name* is the symbol whose macro-definition we are creating, *var* is a single required parameter name that is bound to the *entire* calling form, and *body* is the body of the expansion function, which is executed as an implicit *progn.* The last form in body produces, as its value, the form that will be passed back to `eval` as the macro expansion; the

expansion is then evaluated in place of the macro call. (Note that the expansion could itself be a macro call, and the cycle would repeat.)

The `if` (page 69) construct could be defined in terms of `cond` (page 68) as a macro:

```
(macro if (call-form)
  '(cond (,(cadr call-form) ,(caddr call-form))
         (t ,(cadddr call-form))))
```

If the above form is executed by the interpreter, it will set the definition of the symbol `if` to an object such that `macro-p` of that object returns a one-argument function equivalent to:

```
(lambda (calling-form)
  (list 'cond
        (list (cadr calling-form) (caddr calling-form))
        (list 't (cadddr calling-form))))
```

(The lambda-expression is produced by the `macro` construct. The calls to `list` are the (hypothetical) result of the backquote (`) macro character and its associated commas.)

Now, if `eval` encounters

```
(if (null foo) bar (plus bar 3))
```

this will be expanded into

```
(cond ((null foo) bar)
      (t (plus bar 3)))
```

and `eval` tries again on this new form.

As you can see in the above example, the main disadvantage of using `macro` to define macros is that the user must decompose the argument into its constituents using `car` and `cdr`. In a complex macro, this process is confusing and error-prone. The use of `defmacro` (page 91) alleviates this problem. It should also be clear that the backquote facility (???) is very useful in writing macros, since the form to be returned is normally a complex list structure, mostly constant but with a few evaluated forms scattered through the structure.

Note that when `macro` is encountered by the compiler, the normal action is to add the definition to the compilation environment and also to place a compiled version of the expander-function into the load file, so that the macro will be defined at runtime as well as during the current compilation. If the macro is to be used only during the current compilation and not at runtime, this can be achieved by using the `eval-when` (page EVAL-WHEN-FUN) construct:

```
(eval-when (compile)
  (macro name (var)
    body))
```

`defmacro` *name varlist* {*form*}* [*Macro*]

`defmacro` is a macro-defining macro that, unlike `macro`, decomposes the calling form in a more elegant and useful way. A call to `defmacro` has the following form:

(`defmacro` *name varlist . body*)

This is very similar to a `defun` (page 42) form: *name* is the symbol whose macro-definition we are creating, *varlist* is similar in form to a lambda-list, and *body* is the body of the expander function. If we view the macro call as a list containing a function name and some argument forms, the argument forms (unevaluated) are bound to the corresponding parameters in *varlist*. Then the body forms are evaluated as an implicit `progn`, and the value of the last form is returned as the expansion of the macro call.

Like the lambda-list in a `defun`, a `defmacro` *varlist* may contain variable symbols and the `&optional`, `&rest`, and `&aux` tokens (but not `&key`).

??? **Query:** Should `&key` be allowed?

For `&optional` parameters, initialization forms and "supplied-p" parameters may be specified, just as for `defun`. Two additional tokens are allowed in *defmacro* variable lists only:

`&body` This is identical in function to `&rest`, but it informs certain pretty-printing and editing functions that the remainder of the form is a body rather than arguments, and should be indented accordingly. (Only one of `&body` or `&rest` may be used.)

`&whole` This is followed by a single variable that is bound to the entire macro call form; this is the same value that the single parameter in a `macro` definition form would receive.

Compatibility note: Some LISP implementations, notably Lisp Machine LISP, allow a "destructuring" pattern to be used instead of, or mixed with, the `defun`-like arglist specified here. Prior to the appearance of `&optional`, the pattern may contain not only top-level symbols, but an arbitrary list structure built from cons cells and symbols; this is matched against the macro call cell by cell, producing a binding wherever the `defmacro` pattern contains a symbol. This is not supported by COMMON LISP; it does not support destructuring in `defun`, and `defmacro` needs to parallel `defun` as closely as possible to minimize confusion in what is already a difficult area for new users. Some COMMON LISP implementations may choose to provide destructuring `defmacro` as an extension.

Using `defstruct`, the definition for three-argument `if` would look like this:

```
(defmacro if (pred result else-result)
  '(cond (,pred ,result)
        (t ,else-result)))
```

This would produce the same macro-definition for `if` as the definition using `macro` above. If `if` is to accept two or three arguments, with the `else-result` defaulting to `nil`, as in fact it does in COMMON LISP, the definition might look like this:

```
(defmacro if (pred result &optional (else-result 'nil))
  '(cond (,pred ,result)
        (t ,else-result)))
```

If the compiler encounters a `defmacro`, the normal effect is that same as for a `macro` form: the new macro is added to the compilation environment, and a compiled form of the expansion function is also added to the output file so that the new macro will be operative at runtime. If this

is not the desired effect, the `defmacro` form can be wrapped in an `eval-when`.

Several global variables affect the code that `defmacro` produces.

`defmacro-check-args` [*Variable*]

If `defmacro-check-args` is true (which it initially is) when a `defmacro` (page 91) form is executed or compiled, the resulting macro defined by `defmacro` will contain code that signals an error if it is called with the wrong number of "argument" forms, that is, if the number of items following the function name in the calling form is inconsistent with the number of required and optional arguments specified in the `defmacro` form.

`defmacro-maybe-displace` [*Variable*]

If `defmacro-maybe-displace` is true (which it initially is) when a `defmacro` (page 91) form is executed or compiled, the resulting macro defined by `defmacro` will contain code that checks the variable `macro-displacement-hook` at expansion time. If `macro-displacement-hook` is null, the expansion is used normally. Otherwise, the value of this variable must be a function of two arguments: the original macro-call form and the expansion. Whatever this function returns is passed back to eval as the macro expansion to use this time around.

`macro-expansion-hook` [*Variable*]

The value of this variable is initially `nil`. The purpose of this variable is described above under `defmacro-maybe-displace` (page 92). If the user wants to speed up interpreted code that makes heavy use of macros, this variable can be set to (the name of) the function `displace` (page 92):

```
(setq macro-expansion-hook 'displace)
```

This will destructively replace the macro call with its expansion. Alternatively, some more complex function may be used.

`displace` *macro-call expansion* [*Function*]

`displace` destructively replaces the *macro-call* with its *expansion*, returning the *expansion*. It is the simplest possible "memoizing" function, whose purpose is to speed up interpreted code by doing each macro expansion only once. Its disadvantages are that the original form of the macro is not available to printing and debugging packages, and that if the macro definition is altered, the displaced calls will retain their old expansions. (More complex memoizing packages, which eliminate these disadvantages, will be available in the COMMON LISP library. These are not included in the base language because their use must be coordinated with the use of particular printing, loading, and debugging facilities.)

`macroexpand` *form* [*Function*]

`macroexpand-1` *form* [*Function*]

If *form* is a macro call (with respect to global macro definitions, ignoring any established by `macrolet` (page 67)), then `macroexpand-1` will expand the macro call *once* and return the expansion. If *form* is not a macro call, it is simply returned. `macroexpand` is similar, but repeatedly expands *form* until it is no longer a macro call.

Chapter 9

Declarations

Declarations allow you to specify extra information about your program to the LISP system. All declarations are completely optional and do not affect the meaning of a correct program, with one exception: `special` declarations do affect the interpretation of variable bindings and references, and so must be specified where appropriate. All other declarations are of an advisory nature, and may be used by the LISP system to aid you by performing extra error checking or producing more efficient compiled code. Declarations are also a good way to add documentation to a program.

9.1. Declaration Syntax

`declare` {*declaration*}* [*Special form*]

This form may occur only at top level, or at the beginning of the bodies of certain special forms; that is, a `declare` form not at top level may occur only as a statement of such a form, and all statements preceding it (if any) must also be `declare` forms. If a declaration is found anywhere else an error will be signalled.

Each *declaration* form is a list whose *car* is a keyword specifying the kind of declaration it is. Declarations may be divided into two classes: those that concern the bindings of variables, and those that do not. Those which concern variable bindings apply only to the bindings made by the special form at the head of whose body they appear. For example, in

```
(defun foo (x)
  (declare (type float x)) ...
  (let ((x 'a)) ...)
  ...)
```

the type declaration applies only to the outer binding of `x`, and not to the binding made in the `let`.

Compatibility note: This is different from MACLISP, in which type declarations are pervasive.

If a declaration that applies only to variable bindings appears at top level, it applies to the dynamic value of the variable. For example, the top-level declaration

```
(declare (type float tolerance))
```

specifies that the dynamic value of `tolerance` should always be a floating-point number.

Declarations that do not concern themselves with variable bindings are pervasive, affecting all code

in the body of the special form (but not code in any initialization forms used to compute initial values for bound variables).

For example:

```
(defun foo (x y) (declare (notinline floor)) ...)
```

advises that everywhere within the body of `foo` the function `floor` should not be open-coded, but called as an out-of-line subroutine. Any pervasive declaration made at top level constitutes a universal declaration, always in force unless locally shadowed.

For example:

```
(declare (inline floor))
```

advises that `floor` should normally be open-coded in-line by the compiler (but within `foo` it will be compiled out-of-line anyway, because of the shadowing local declaration to that effect).

For example:

```
(defun nonsense (k x)
  (declare (type integer k))
  (let ((j (foo k x))
        (x (* k k)))
    (declare (inline foo) (special x))
    (foo x j)))
```

In this rather nonsensical example, `k` is declared to be of type `integer`. The `inline` declaration applies to the inner call to `foo`, but not to the one to whose value `j` is bound, because that is *code* in the binding part of the `let`. The `special` declaration of `x` causes the `let` form to make a special binding for `x`, and causes the reference to `x` in the body of the `let` to be a special reference. However, the reference to `x` in the first call to `foo` is a local reference, not a special one.

Compatibility note: In MACLISP, `declare` does nothing in interpreted code, and is defined to simply evaluate all the argument forms in the compilation environment. In COMMON LISP, `declare` does useful things for both interpreted code and compiled code, and therefore arbitrary forms are not permitted within it. The tricks played in MACLISP with `declare` are better done using `eval-when` (page EVAL-WHEN-FUN).

`locally` {*declare-form*}* {*form*}* [*Special form*]

This special form may be used to make local pervasive declarations where desired. It does not bind any variables, and so cannot be used meaningfully for declarations of variable bindings.

For example:

```
(locally (declare (inline floor))
         (declare (notinline car cdr))
  (floor (car x) (cdr y)))
```

9.2. Declaration Forms

Here is a list of valid declaration forms for use in `declare`. A construct is said to be "affected" by a declaration if it occurs within the scope of a declaration.

`special` `(special` *var1 var2* ...`)` declares that all of the variables named are to be considered *special*. This declaration affects variable bindings, but also pervasively affects references.

All variable bindings affected are made to be dynamic bindings, and affected variable references refer to the current dynamic binding rather than the current local binding. This declaration does *not* pervasively affect bindings unless it occurs at top level (this latter exception arising from convenience and compatibility with MACLISP). Inner bindings of a variable implicitly shadow a `special` declaration, and must be explicitly re-declared to be special.

For example:

```
(declare (special x))                ;x is always special.
(defun example (x y)
  (declare (special y))
  (let ((y 3))
    (print (+ y (locally (declare (special y)) y)))
    (let ((y 4)) (declare (special y)) (foo x))))
```

In the contorted code above, the outermost and innermost bindings of *y* are special, and therefore dynamically scoped, but the middle binding is lexically scoped. The two arguments to + are different, one being the value (which is 3) of the lexically bound variable *y*, and the other being the value of the special variable named `y` (a binding of which happens, coincidentally, to lexically surround it at an outer level).

`type` — `(type` *type var1 var2 ...*`)` affects only variable bindings, and declares that the specified variables will take on values only of the specified type.

type — (*type var1 var2 ...*) is an abbreviation for `(type` *type var1 var2 ...*`)` provided that *type* is one of the symbols appearing in Table 4-1 (page 27).

`ftype` — `(ftype` *type function-name-1 function-name-2 ...*`)` declares that the named functions will be of the functional type *type*.

For example:

```
(declare (ftype (function (integer list) t) nth)
         (ftype (function (number) float) sin cos))
```

`function` — `(function` *name arglist result-type1 result-type2 ...*`)` is entirely equivalent to

`(ftype (function` *name arglist result-type1 result-type2 ...*`)` *name*)

but may be more convenient for some purposes.

For example:

```
(declare (function nth (integer list) t)
         (function sin (number) float)
         (function cos (number) float))
```

The syntax mildly resembles that of `defun` (page 42): a function name, then an argument list, then a specification of results.

`inline` — `(inline` *function1 function2 ...*`)` declares that it is desirable for the compiler to open-code calls to the specified functions; that is, the code for a specified function should be integrated into the calling routine, appearing "in line", rather than a procedure call appearing there. This may achieve extra speed at the expense of debuggability (calls to functions compiled in-line cannot be traced, for example). This declaration is pervasive. Remember that a compiler is free to ignore this declaration.

`notinline` (`notinline` *function1 function2 ...*) declares that it is *undesirable* to compile the specified functions in-line. This declaration is pervasive. Remember that a compiler is free to ignore this declaration.

`ignore` (`ignore` *var1 var2 ... varn*) affects only variable bindings, and declares that the bindings of the specified variables are never used. It is desirable for a compiler to issue a warning if a variable so declared is ever referred to or is also declared special, or if a variable is lexical, never referred to, and not declared to be ignored.

> ??? **Query:** This is a new idea: what do people think? This is more mnemonic than writing `ignore` or `nil` for an ignored parameter because you can give a meaningful (and possibly conventional) name. It is more explicit and robust than simply mentioning the variable at the front of the lambda-body; the latter convention prevents the compiler from issuing a warning about a possibly malformed program.

`optimize` (`optimize` *quality1 quality2 ...*) advises the compiler that *quality1* should be given greatest attention in producing compiled code, then *quality2*, and so on. The qualities may include `speed` (of the compiled code), `space` (both code size and run-time space), and `safety` (run-time error checking); any qualities not mentioned are assumed to be of lower priority than those mentioned. The default situation is implementation-dependent, but implementors are encouraged to consider `(optimize safety speed space)` for the default. This declaration is pervasive.

For example:

```
(defun often-used-subroutine (x y)
  (error-check x y)
  (hairy-setup x)
  (locally
     ;; This inner loop really needs to burn.
     (declare (optimize speed))
     (do ((i 0 (+ i 1))
          (z x (cdr z)))
         ((null z))
       (declare (fixnum i)))))
```

> ??? **Query:** This is a new idea: what do people think? Actually, one needs finer control over this, such as whether type declarations should be assumed by the compiler or cause explicit checking code to be emitted.

An implementation is free to support other (implementation-dependent) declarations as well. On the other hand, a COMMON LISP compiler is free to ignore entire classes of declarations (for example, implementation-dependent declarations not supported by that compiler's implementation!). Compiler implementors are encouraged, however, to program the compiler by default to issue a warning if the compiler finds a declaration of a kind it never uses (as a hedge against spelling errors).

9.3. Type Declaration for Forms

Frequently it is useful to declare that the value produced by the evaluation of some form will be of a particular type. Using `declare` one can declare the type of the value held by a bound variable, but there is no easy way to declare the type of the value of an unnamed form. One could write something like

```
((lambda (x) (declare (type type x)) x) form)
```

but that would be rather clumsy. For this purpose the `the` special form is defined: `(the` *type form*`)` means essentially the same as the larger expression above.

`the` *type form* [*Special form*]

The *form* is evaluated; whatever it produces is returned by the `the` form. In addition, it is an error if what is produced by the *form* does not conform to the data type specified by *type* (which is not evaluated). (A given implementation may or may not actually check for this error. Implementations are encouraged to make an explicit error check when running interpretively.) In effect, this declares that the user undertakes to guarantee that the values of the form will always be of the specified type.

For example:

```
(the string (concat x y))          ;The result of concat will be a string.
(the integer (+ x 3))              ;The result of + will be an integer.
(+ (the integer x) 3)              ;The value of x will be an integer.
(the (complex rational) (* z 3))
(the (unsigned-byte 8) (logand x mask))
```

Compatibility note: This construct is borrowed from the INTERLISP DECL package; INTERLISP, however, allows an implicit `progn` after the type specifier rather than just a single form. The MACLISP `fixnum-identity` and `flonum-identity` constructs can be expressed as `(the fixnum` *x*`)` and `(the single-float` *x*`)`.

Chapter 10

Symbols

A LISP symbol is a data object which has three user-visible components:

- The *property list* is a list which effectively provides each symbol with many modifiable named components.

- The *print name* must be a string, which is the sequence of characters used to identify the symbol. Symbols are of great use because a symbol can be located given its name (typed, say, on a keyboard). It is ordinarily not permitted to alter a symbol's print name.

- The *package cell* must refer to a package object. A package is a data structure used to locate a symbol given its name. A symbol is uniquely identified by its name only when considered relative to a package. A symbol may be in many packages, but it can be *owned* by at most one package. The package cell points to the owner, if any.

A symbol may actually have other components as well for use by the implementation. One of the more important uses of symbols is as names for program variables; it is frequently desirable for the implementor to use certain components of a symbol to implement the semantics of variables. However, there are several possible implementation strategies, and so such possible components are not described here.

The three components named above and the functions related to them are described more individually and in more detail in the following sections.

10.1. The Property List

Since its inception, LISP has associated with each symbol a kind of tabular data structure called a *property list* (*plist* for short). A property list contains zero or more entries; each entry associates from a keyword symbol (called the *indicator*) to a Lisp object (called the *value* or, sometimes, the *property*). There are no duplications among the indicators; a property-list may only have one property at a time with a given name. In this way, given a symbol and an indicator (another symbol), an associated value can be retrieved.

A property list is very similar in purpose to an association list. The difference is that a property list is an object with a unique identity; the operations for adding and removing property-list entries are destructive operations that alter the property-list rather than making a new one. Association lists, on the other hand, are

normally augmented non-destructively (without side effects), by adding new entries to the front (see `acons` (page 179) and `pairlis` (page 179)).

A property list is implemented as a memory cell containing a list with an even number (possibly zero) of elements. (Usually this memory cell is the property-list cell of a symbol, but any memory cell acceptable to `setf` (page 60) can be used if certain special forms are used.) Each pair of elements in the list constitutes an entry; the first item is the indicator and the second is the value. Because property-list functions are given the symbol and not the list itself, modifications to the property list can be recorded by storing back into the property-list cell of the symbol.

When a symbol is created, its property list is initially empty. Properties are created by `putpr` (page 102) and related functions.

COMMON LISP does not use a symbol's property list as extensively as earlier LISP implementations did. Less-used data, such as compiler, debugging, and documentation information, is kept on property lists in COMMON LISP.

Compatibility note: In older Lisp implementations, the print name, value, and function definition of a symbol were kept on its property list. The value cell was introduced into MACLISP and INTERLISP to speed up access to variables; similarly for the print-name cell and function cell (MACLISP does not use a function cell). Recent LISP implementations such as SPICE LISP, Lisp Machine LISP, and NIL have introduced all of these cells plus the package cell. None of the MACLISP system property names (`expr`, `fexpr`, `macro`, `array`, `subr`, `lsubr`, `fsubr`, and in former times `value` and `pname`) exist in COMMON LISP.

Compatibility note: In COMMON LISP, the notion of "disembodied property list" introduced in MACLISP is eliminated. It tended to be used for rather kludgy things, and in Lisp Machine LISP is often associated with the use of locatives (to make it "off by one" for searching alternating keyword lists). In COMMON LISP special `setf`-like property list functions are introduced: `getf` (page 103), `putf` (page 103), and `remf` (page 104).

`getpr` *symbol indicator* `&optional` *default* [*Function*]

getpr searches the property list of *symbol* for an indicator `eq` to *indicator*. If one is found, then the corresponding value is returned; otherwise *default* is returned. If *default* is not specified, then `nil` is used for *default*. Note that there is no way to distinguish an absent property from one whose value is *default*.

```
(getpr x y) <=> (getf (plist x) y)
```

Suppose that the property list of `foo` is `(bar t baz 3 hunoz "Huh?")`. Then, for example:

```
(getpr 'foo 'baz) => 3
(getpr 'foo 'hunoz) => "Huh?"
(getpr 'foo 'zoo) => nil
```

`putpr` *symbol indicator newvalue* [*Function*]

This causes *symbol* to have a property whose indicator is *indicator* and whose value is *newvalue*. If the property list already already had a property with an indicator `eq` to *indicator*, then the value previously associated with that indicator is removed from the property list and replaced by *newvalue*. The property list is destructively altered by using side effects. After a `putpr` is done, (`getpr` *symbol indicator*) will return *value*. `putpr` returns the new *value*.

```
(putpr x y z) <=> (putf (plist x) y z)
```

For example:

```
(putpr 'Nixon 'crook 'no) => no
(getpr 'Nixon 'crook) => no
```

??? **Query:** Should there be an analogue for `defprop`, say `defpr`?

`rempr` *symbol indicator* [*Function*]

This removes from *symbol* the property with an indicator `eq` to *indicator*, by splicing it out of the property list. It returns `nil` if no such property was found, or non-`nil` if a property was found.

```
(rempr x y) <=> (remf (plist x) y)
```

For example:

If the property list of `foo` was
```
(color blue height 6.3 near-to bar)
```
then
```
(rempr 'foo 'height) => t
```
and `foo`'s property list would have been altered to be
```
(color blue near-to bar)
```

`plist` *symbol* [*Function*]

This returns the list which contains the property pairs of *symbol*; the contents of the property list cell are extracted and returned.

Note that using `get` on the result of `plist` does *not* work. One must give the symbol itself to `get`.

`getf` *place indicator* `&optional` *default* [*Function*]

getf searches the property list stored in *place* for an indicator `eq` to *indicator*. If one is found, then the corresponding value is returned; otherwise *default* is returned. If *default* is not specified, then `nil` is used for *default*. Note that there is no way to distinguish an absent property from one whose value is *default*. Normally *place* is computed from a generalized variable acceptable to `setf` (page 60). See `getpr` (page 102).

`putf` *place indicator newvalue* [*Macro*]

This causes the property list stored in *place* to have a property whose indicator is *indicator* and whose value is *newvalue*. If the property list already already had a property with an indicator `eq` to *indicator*, then the value previously associated with that indicator is removed from the property list and replaced by *newvalue*. The property list is destructively altered by using side effects. After a `putf` is done, (`getf` *place indicator*) will return *value*. `putf` returns the new *value*. The form *place* may be any generalized variable acceptable to `setf` (page 60). See `putpr` (page 102).

`remf` *place indicator* [*Macro*]

This removes from the property list stored in *place* the property with an indicator `eq` to *indicator*, by splicing it out of the property list. It returns `nil` if no such property was found, or `t` if a property was found. The form *place* may be any generalized variable acceptable to `setf` (page 60). See `rempr` (page 103).

`get-properties` *place indicator-list* [*Function*]

`get-properties` is like `getf` (page 103), except that the second argument is a list of indicators. `get-properties` searches the property list stored in *place* for any of the indicators in *indicator-list*, until it finds a property whose indicator is one of the elements of *indicator-list*. Normally *place* is computed from a generalized variable acceptable to `setf` (page 60).

`get-properties` returns three values. The third value is `t` if any property was found, in which case the first two values are the indicator and value for some property whose indicator was in *indicator-list*; if no property was found, all three values are `nil`.

When more than one of the indicators in *indicator-list* is present in the property list, which one `get-properties` returns depends on the implementation. All that is guaranteed is that if there are one or more properties whose indicators are in *indicator-list*, some one such property will be chosen and returned.

??? **Query:** Should there be a `do-properties` in addition to, or instead of, `map-properties`?

`map-properties` *function place* [*Function*]

The property list stored in *place* is accessed, and *function* is called once for each property in the property list. The *function* should accept two arguments: the indicator and the value for a property. `map-properties` returns `nil`; the *function* is useful only for its side effects.

The order in which properties are given to *function* is implementation-dependent. Also, if side effects modify the property list during the `map-properties` computation, the effects are unpredictable. All that is guaranteed is that if no side effects occur on the property list, then *function* is applied once to each property in the property list.

For example:

Assume array element `(aref a 105)` contains `nil`.

```
(putf (aref a 105) 'color 'yellow)
(putf (aref a 105) 'height 105)
(putf (aref a 105) 'shape 'pyramid)
(map-properties #'(lambda (i v) (format t "~S <-=-> ~S" i v))
                (aref a 105))
```

might print:	or it might print:
`color <-=-> yellow`	`height <-=-> 105`
`height <-=-> 105`	`shape <-=-> pyramid`
`shape <-=-> pyramid`	`color <-=-> yellow`

or it might print any of the other four possible permutations.

10.2. The Print Name

Every symbol has an associated string called the *print-name*, or *pname* for short. This string is used as the external representation of the symbol: if the characters in the string are typed in to `read` (with suitable escape conventions for certain characters), it is interpreted as a reference to that symbol (if it is interned); and if the symbol is printed, `print` types out the print-name. For more information, see the section on the *reader* (see page READER) and *printer* (see page PRINTER).

`get-pname` *sym* [*Function*]

This returns the print-name of the symbol *sym*.

For example:

```
(get-pname 'XYZ) => "XYZ"
```

It is an extremely bad idea to modify a string being used as the print name of a symbol. Such a modification may confuse the function `read` (page 237) and the package system tremendously.

`samepnamep` *sym1 sym2* [*Function*]

This predicate is true if the two symbols *sym1* and *sym2* have `equal` print-names; that is, if their printed representation is the same. Upper and lower case letters are considered to be different.

> **Compatibility note:** In Lisp Machine LISP, `samepnamep` normally considers upper and lower case to be the same. However, in MACLISP, which originated this function, the cases are distinguished; Lisp Machine LISP introduced the incompatibility. COMMON LISP is compatible with MACLISP here.

If either or both of the arguments is a string instead of a symbol, then that string is used in place of the print-name. `samepnamep` is useful for determining if two symbols would be the same except that they are not in the same package.

For example:

```
(samepnamep 'xyz (maknam '(x y z)) is true
(samepnamep 'xyz (maknam '(w x y)) is false
```

10.3. Creating Symbols

Symbols can be used in two rather different ways. An *interned* symbol is one which is indexed by its print-name in a catalog called a *package*. Every time anyone asks for a symbol with that print-name, he gets the same (`eq`) symbol. Every time input is read with the function `read` (page 237), and that print-name appears, it is read as the same symbol. This property of symbols makes them appropriate to use as names for things and as hooks on which to hang permanent data objects (using the property list, for example; it is no accident that symbols are both the only LISP objects which are cataloged and the only LISP objects which have property lists).

Interned symbols are normally created automatically; the first time someone (such as the function `read`) asks the package system for a symbol with a given print-name, that symbol is automatically created. The function to use to ask for an interned symbol is `intern` (page 112), or one of the functions related to

`intern`.

Although interned symbols are the most commonly used, they will not be discussed further here. For more information, turn to the chapter on packages.

An *uninterned* symbol is a symbol used simply as a data object, with no special cataloging (it belongs to no particular package). An uninterned symbol prints in the same way as an interned symbol with the same print-name, but cannot be read back in. The following are some functions for creating uninterned symbols.

`make-symbol` *pname* [*Function*]

(`make-symbol` *pname*) creates a new uninterned symbol, whose print-name is the string *pname*. The value and function bindings will be unbound and the property list will be empty.

The string actually installed in the symbol's print-name component may be the given string *pname* or may be a copy of it, at the implementation's discretion. The user should not assume that (`get-pname (make-symbol x)`) is `eq` to `x`, but also should not alter a string once it has been given as an argument to `make-symbol`.

Implementation note: An implementation might choose, for example, to copy the string to some read-only area, in the expectation that it will never be altered.

Compatibility note: Lisp Machine LISP uses the second argument for an odd flag related to areas. It is unclear what NIL does about this.

`copysymbol` *sym* `&optional` *copy-props* [*Function*]

This returns a new uninterned symbol with the same print-name as *sym*. If *copy-props* is non-`nil`, then the initial value and function-definition of the new symbol will be the same as those of *sym*, and the property list of the new symbol will be a copy of *sym*'s. If *copy-props* is `nil` (the default), then the new symbol will be unbound and undefined, and its property list will be empty.

`gensym` `&optional` *x* [*Function*]

`gensym` invents a print-name, and creates a new symbol with that print-name. It returns the new, uninterned symbol.

The invented print-name consists of a prefix (which defaults to "G"), followed by the decimal representation of a number. The number is increased by one every time `gensym` is called.

If the argument *x* is present and is an integer, then *x* must be non-negative, and the internal counter is set to `x` for future use; otherwise the internal counter is incremented. If *x* is a string, then that string is made the default prefix for this and future calls to `gensym`. After handling the argument, `gensym` creates a symbol as it would with no argument.

For example:

```
(gensym) => G7
(gensym "FOO-") => FOO-8
(gensym 32) => FOO-32
(gensym) => FOO-33
(gensym "GARBAGE-") => GARBAGE-34
```

`gensym` is usually used to create a symbol which should not normally be seen by the user, and whose print-name is unimportant, except to allow easy distinction by eye between two such symbols. The optional argument is rarely supplied. The name comes from "generate symbol", and the symbols produced by it are often called "gensyms".

If it is crucial that no two generated symbols have the same print name (rather than merely being distinct data structures), or if it is desirable for the generated symbols to be interned, then the function `gentemp` (page 107) may be more appropriate to use.

`gentemp` *prefix* `&optional` *package* [*Function*]

`gentemp`, like `gensym` (page 106), creates and returns a new symbol. `gentemp` differs from `gensym` in that it interns the symbol (see `intern` (page 112)) in the *package* (which defaults to the current package; see `package` (page 112)). `gentemp` guarantees the symbol will be a new one not already existing in the package; it does this by using a counter as `gensym` does, but if the generated symbol is not really new then the process is repeated until a new one is created. There is no provision for resetting the `gentemp` counter. Also, the prefix for `gentemp` is not remembered from one call to the next; if *prefix* is omitted, the default prefix `T` is used.

`symbol-package` *sym* [*Function*]

Given a symbol *sym*, `symbol-package` returns the contents of the package cell of that symbol. This will be a package object or `nil`.

Chapter 11

Packages

One problem with most LISP systems is the use of a single name space for all symbols. In large LISP systems, with modules written by many different programmers, accidental name collisions become a serious problem. In the past, this problem has been addressed by the use of a prefix on each symbol name in a module or by some sort of clumsy "obarray" switching to keep the names separated.

COMMON LISP addresses this problem through the *package system*, derived from an earlier package system developed for Lisp Machine LISP [11]. The COMMON LISP package system provides an *export* mechanism for easily dividing the symbols in a package into *external symbols*, which are part of the package's public interface to other packages, and *internal symbols*, which are for internal use only and are normally hidden from other packages.

A *package* is a data structure that establishes a mapping from print names (strings) to symbols. (The package thus replaces the "oblist" or "obarray" of earlier LISP systems.) A symbol may appear in many packages, but will always have the same name. On the other hand, the same name may refer to different symbols in different packages. No two symbols in the same package may have the same name.

Some of the symbols in a package may be marked as being *exported* by that package; these are the *external symbols*. Those symbols not exported are said to be *internal* to that package. Any symbol can be added to the set of external symbols by using the function `export` (page 113).

The value of the special variable `package` (page 112) must always be a package object or the name of a package object; this is referred to as the *current package*. Each package is named by a symbol.

When the LISP reader has, by parsing, determined a string of characters thought to name a symbol, that name looked up in the current package. If the name is found, the corresponding symbol is returned. If the name is not found there, a new symbol is created for it and is placed in the current package as an internal symbol; if the name is seen again while this same package is current, the same symbol will then be returned. When a new symbol is created, a pointer to the package in which it is initially placed is stored in the *package cell* of that symbol; the package is said to be the symbol's *home package*.

Often it is desirable, when typing an expression to be read by the LISP reader, to refer to a symbol in some package other than the current one. This is done through the use of a *qualified name*, which consists of the

package name, followed by a colon, followed by the print name of the symbol. This causes the symbol's name to be looked up in the specified package. For example, "`editor:buffer`" refers to the symbol named "`buffer`" in the package named "`editor`", regardless of whether there is a symbol named "`buffer`" in the current package. If "`buffer`" does not exist in package "`editor`", it is created there as a new internal symbol. (If, on the other hand, there is no package named "`editor`", an error is signalled.)

The package name, if it is not itself qualified, is looked up in the special package named "`packages`", but this default may be overridden by recursive use of the colon convention. Thus the qualified name "`editor:display:buffer`" is deciphered by first finding the symbol "`editor:display`" in the usual way, then using this symbol (which must be the name of a package) as the package in which to look up the symbol named "`buffer`". (Because the *first* name is always looked up in the package `packages`, which is itself in the `packages` package, adding "`packages:`" to the front of an already qualified name does not change the meaning of the name. So, for example, "`editor:display:foo`" and "`packages:editor:display:foo`" both denote the same symbol.)

If a symbol names a package, then the package is stored on the property list of the symbol under the property name `:package`. Suppose the variable `x` has a symbol as its value; then `(get s :package)` will return the associated package. Given a package, the function `package-name` (page 112) will return the symbol that names the package.

Symbols from another package may be added to the current package in two ways. First, an individual symbol may be added by use of the `import` function. The form `(import 'editor:buffer)` takes the symbol `buffer` in the package `editor` (this symbol was located when the form was read by the LISP reader) and adds it to the current package as an internal symbol. The imported symbol is not automatically exported from the current package, but if it is already present and external, that is not changed. After the call to `import`, it is possible to refer to `buffer` in the current package without any qualifier. The status of the symbol `buffer` in the package named `editor` is unchanged, and `editor` remains the home package of this symbol. If the imported symbol already exists in the current package, the `import` operation effectively does nothing. If a distinct symbol with the name `buffer` already exists in the current package, a correctable error is signalled. The value returned from this error is the symbol that should remain in the package, the other being discarded.

The second mechanism, the `use` function, imports into the current package *all* of the *external* symbols of another package. These symbols can then be referred to from the current package without qualification. The *internal* symbols of the used package are not imported, and therefore cannot conflict with symbols in the current package. The status of the imported symbols in their original package is unchanged. Conflicts between symbols already in the current package and those imported by `use` are handled as in `import`: a correctable error occurs. However, `use` provides a mechanism for suppressing this error in case a few of the symbols are known in advance to be in conflict. The `use` function imports only those symbols that are exported by the used package at the time `use` is called; it is not a general inheritance mechanism and does not arrange for future changes in the used package to be imported.

Each symbol contains a package slot which is used to record the home package of the symbol. When the

symbol is printed, if it is present in the current package (either as an internal or an external symbol), it is printed without any qualification; otherwise, it is printed with the recorded package as the qualifier.

11.1. Built-in Packages

The following packages are built into the system and are treated as special in some way:

`lisp` The package named `lisp` contains the primitives of the COMMON LISP system. Its external symbols include all of the user-visible functions and global variables that are present in the basic LISP system, such as `car`, `cdr`, `package`, etc. Almost all other packages will want to "`use`" this one so that these symbols will be available without qualification.

`user` The `user` package is, by default, the current package at the time a COMMON LISP system starts up. It includes the external symbols from the LISP package at startup time.

`keyword` This package contains, as external symbols, all of the keywords used by built-in or user-defined LISP functions. It is not recommended that these keywords be loaded into other packages via the `use` function, as conflicts may result. Instead, a special syntax is provided to make it easy to access symbols in the `keyword` package: a null leading package name is treated as being identical to `keyword`. Thus `:foo` is the same as `keyword:foo`. By special arrangement, symbols in the `keyword` package always evaluate to themselves, so the user can type `:foo` instead of `':foo`.

`packages` This is the package that contains the symbols that name the other packages. If the LISP reader sees "`editor:buffer`", for example, it first looks up the name "`editor`" in the package named `packages`. This must produce a symbol that names a package, which is then used in looking up the name "`buffer`" to find the desired symbol.

`si` This package name is reserved to the implementation. (The name is an abbreviation for "system internals".)

11.2. Package System Functions and Variables

`make-package` *package-name* `&optional` *copy-from* [*Function*]

Creates and returns a new package with the specified package name. If the package name is a symbol, that symbol is used directly; a string is interned in the `packages` package to produce a symbol.

If a package of this name already exists, a correctable error is signalled. Copy-from may specify another package of which the new one will initially be a copy; if copy-from is `t`, the new package initially contains only the external symbols of the `lisp` package; if copy-from is `nil` (the default), the new package is empty.

`package` [*Variable*]

The value of this variable must be either a package or a symbol that names a package; this package is said to be the current package. The initial value of `package` is the `user` package.

`package` *package* [*Function*]

This converts its argument to be a package object. If the argument is already a package, it is a returned. If it is a symbol, the package it names is returned (it is an error if it does not name a package).

`package-name` *package* [*Function*]

This returns a symbol that names a package. If the argument is a package, its name is returned. If the argument is a symbol, it is returned if it names a package, but an error is signalled if it does not.

`begin-package` *package-name* [*Function*]
`end-package` *package-name* [*Function*]

The *package-name* must be the name of a package, in the form of a string or a symbol.

For `begin-package`, if no package currently exists with this name, one is created that imports all external symbols of the `lisp` package. `begin-package` rebinds the `package` variable to the specified package, saving the old value for restoration when the matching `end-package` is encountered. A call to `begin-package` is normally placed at the beginning of a file that is to be loaded into some package other than `user`.

For `end-package`, the package specified must be the current package or else an error is signalled. The package current before the matching call to `begin-package` was encountered is made current once again.

If a pair of `begin-package` and `end-package` are nested within another pair, there is no hierarchical relationship between the inner and outer pair. The inner pair merely temporarily shadows the outer pair.

Rationale: This is so that one package can be loaded during the loading of another one, as by a MACLISP-like autoload facility.

??? **Query:** Should `load` (page 270) arrange to bind things so that mismatched `begin-package` and `end-package` constructs don't screw things up outside the loaded file?

`intern` *string-or-symbol* `&optional` *package* [*Function*]

The *package* may be a package or a symbol that names a package, and defaults to the current package. It is searched for a symbol with the name specified by the first argument. If one is found, it is returned; note particularly that if the argument was symbol, and a *different* symbol with the same name is found in already in the package, the latter is returned and the argument is discarded.

If one is not found, then if the first argument is a string a symbol with that name is created; then the given or created symbol is installed in the package as an internal symbol and returned. Moreover, if the symbol has no home package, then *package* becomes its home package.

`remob` *string-or-symbol* `&optional` *package* [*Function*]

If the first argument is a string, the *package* is searched for a symbol of that name; if the first argument is a symbol, that symbol is used directly. If the symbol given or found is in fact in the package, it is removed from the package. Moreover, if *package* is the home package for the symbol, the symbol is made to have no home package. The *package* defaults to the current package. `remob` returns `t` if it actually removed a symbol, and `nil` otherwise.

??? Query: This name is traditional, but wouldn't `unintern` or `remove-symbol` be better?

`internedp` *string-or-symbol* `&optional` *package* [*Function*]

This is a predicate. If the first argument is a string, then `internedp` is true if the *package* contains a symbol whose name is the string. If the first argument is a symbol, then `internedp` is true if the *package* contains that very symbol. Otherwise `internedp` is false. The *package* may be a package or a symbol that names one, and defaults to the current package.

`externalp` *string-or-symbol* `&optional` *package* [*Function*]

This is a predicate. If the first argument is a string, then `externalp` is true if the *package* contains an *external* symbol whose name is the string. If the first argument is a symbol, then `externalp` is true if the *package* contains that very symbol as an external symbol. Otherwise `externalp` is false. The *package* may be a package or a symbol that names one, and defaults to the current package.

`export` *symbols* `&optional` *package* [*Function*]

The argument should be a list of symbols, or possibly a single symbol. The specified symbols become external symbols of the specified *package*. The *package* may be a package or a symbol that names one, and defaults to the current package. Any symbol not already in the package is first imported (see `import` (page 114)). If a specified symbol is already an external symbol of the package, it is unaffected. `export` returns `t`.

By convention, a call to `export` listing all exported symbols is placed near the start of a file, after a call to `begin-package` to advertise which of the symbols used in the file are intended to be external.

`unexport` *symbols* `&optional` *package* [*Function*]

The argument should be a list of symbols, or possibly a single symbol. The specified symbols are made to be no longer external symbols of the specified *package*. The *package* may be a package or a symbol that names one, and defaults to the current package. Any specified symbol that is an external symbol of the package is made an internal symbol of the package. Any specified symbol internal to the package or not already in the package not affected (see `import` (page 114)). `unexport` returns `t`.

`import` *symbols* `&optional` *to-package* [*Function*]

The argument should be a list of symbols, or possibly a single symbol. The specified symbols become internal symbols of the specified *to-package*. The *to-package* may be a package or a symbol that names one, and defaults to the current package. If, for some specified symbol, the package already contains another symbol of the same name, a correctable error is signalled. If a specified symbol is already in the package, it is unaffected, and in particular remains an external symbol of the package if it already was one. `import` returns `t`.

`shadow` *symbols* `&optional` *to-package* [*Function*]

The argument should be a list of symbols, or possibly a single symbol. For each specified symbol, if the package of that symbol is not the *to-package*, then a new symbol with the same name and no properties, value, or function definition is created and interned in the *to-package*. The net effect is that the *to-package* ends up with symbols of its own for all the specified names.

The *to-package* may be a package or a symbol that names one, and defaults to the current package. If, for some specified symbol not owned by the package, the package already contains another symbol of the same name, nothing happens; it is not an error.

The purpose of `shadow` is to provide a means for declaring that a particular symbol is to be used "locally" in the package, even though it might have been imported from some other package. For example, suppose one were writing an INTERLISP compatibility package for COMMON LISP. One difference between the two is the definition of the function `nth` (page 169). One might write:

```
(begin-package 'interlisp)
(provide 'interlisp)
(export '(masterscope helpsys dwimify ...))
(shadow '(nth ...))
(require 'odd-utilities)
  ...
(defun nth (x n)        ;InterLISP NTH function.
  ...)
```

`shadow` returns `t`.

`use` *from-package* `&optional` *to-package ignore-list force-list* [*Function*]

Each of the external symbols from the *from-package* is imported into the *to-package*, which defaults to the current package. The rules are the same as for `import` (page 114), except that if an imported symbol conflicts with one already present, there are three possible actions. If the imported symbol is on the *ignore-list*, it is not imported. If the imported symbol is on the *force-list*, it is added to the current package after removing the conflicting symbol from the package (see `remob` (page 113)). If the imported symbol is on neither list, a correctable error is signalled, as described for `import`. (If the symbol is on both the *ignore-list* and the *force-list*, the *ignore-list* takes precedence.) `use` returns `t`.

`provide` *package* [*Function*]

`require` *package* `&optional` *pathname* [*Function*]

Calling `provide` notes the fact that a program module associated with the named package has been loaded or otherwise instantiated. This is used in conjunction with `require`.

Calling `require` does nothing if the indicated package has already been "provided". If it has not, then the *pathname* is given to `load` (page 270) in an attempt to obtain the necessary module from the file system. After the loading process is done, if the package still has not been provided, then an error is signalled. Once the package has been provided, then `use` is applied to it to obtain its exported symbols for the current package. The *pathname* defaults in an implementation-dependent way that may depend on the name of the *package*. (Typically, the name of the package might be used as a file name to access a directory where the yellow-pages modules are stored.)

Here is an example of what a yellow-pages module might look like. The `timestamp` module exports three functions: `timestamp`, `moonprinc`, and `sunprinc`. (The purpose of the module is to print timestamps to a stream; a timestamp includes the time, date, day of week, phase of the moon, and position of the sun. This is a whimsical module.) The `timestamp` module requires two other modules for its operation, `moonphase` and `suncalc`; one is a standard library module, and the other is private. For reasons best ignored here, the `timestamp` module has its own function named `sqrt` that differs from the standard `sqrt` (page 124).

```
(begin-package 'timestamp)
(provide 'timestamp)
(export '(timestamp moonprinc sunprinc))
(require 'moonphase)
(require 'suncalc "/usr/gls/chutzpah/suncalc")
(shadow 'sqrt)

(defconst latitude 48.503)     ;Location of the University of
(defconst longitude 97.61)     ; Southern North Dakota at Hoople

(defun timestamp ...)

(defun moonprinc ...)

(defun sunprinc ...)

(defun stamp-utility ...)

(end-package 'timestamp)
```

It is important that the calls to `provide` and `export` precede the calls to `require`. For suppose that the `moonphase` module requires `timestamp`! When `timestamp` is loaded, if `moonphase` is loaded as a result, it had better find by that point that `timestamp` has already been provided (or will be very soon!), lest `timestamp` be recursively and redundantly loaded, causing an infinite loop. Similarly, by the time that the `moonphase` package tries to `use` the `timestamp` package, the exported symbols of the `timestamp` package must already have been declared, or else the `moonphase` package will not get them.

`package-use-conflicts` *from-package* `&optional` *to-package* [*Function*]

Returns a list of all external symbols in *from-package* that conflict with symbols in *to-package* (which defaults to the current package), or `nil` if there are none. Two symbols conflict if they are different but have the same print name.

`do-symbols` (*var* [*package*] [*result-form*]) {*tag* | *statement*}* [*Special form*]
`do-external-symbols` (*var* [*package*] [*result-form*]) {*tag* | *statement*}* [*Special form*]
`do-internal-symbols` (*var* [*package*] [*result-form*]) {*tag* | *statement*}* [*Special form*]

`dolist` provides straightforward iteration over the symbols of a package. The body is performed once for each symbol in the *package*, in no particular order, with the variable *var* bound to the symbol. Then *resultform* (a single form, *not* an implicit `progn`) is evaluated, and the result is the value of the `dolist` form. If *resultform* is omitted, the result is `nil`. If execution of the body affects which symbols are contained in the *package*, other than possibly to remove the symbol currently the value of *var*, the effects are unpredictable.

`do-external-symbols` is similar, but provides only the external symbols of the *package*. `do-internal-symbols` is similar, but provides only the internal symbols of the *package*.

`do-all-symbols` (*var* [*result-form*]) {*tag* | *statement*}* [*Special form*]

This executes the body once for every symbol contained in every package whose name is in the `packages` package. (This doesn't actually get all symbols whatsoever.) It is *not* in general the case that each symbol is processed only once, since a symbol may appear in many packages.

Chapter 12

Numbers

COMMON LISP provides several different representations for numbers. These representations may be divided into four categories: integers, ratios, floating-point numbers, and complex numbers. Many numeric functions will accept any kind of number; they are *generic*. Those functions which accept only certain kinds of numbers are so documented below.

A COMMON LISP implementation is permitted not to support complex numbers. If it does not, then all the functions defined here (such as `conjugate`) must be defined nevertheless, but whenever a function would have to construct and return a complex number, an error is signalled instead.

??? **Query:** Say! This is a glitch. Can everyone agree just to go ahead and support complex numbers? Or is that really too hard, even given sharing of LISP-level code?

In general, numbers in COMMON LISP are not true objects; `eq` cannot be counted upon to operate on them reliably. In particular, it is possible that the expression

```
(let ((x z) (y z)) (eq x y))
```

may be false rather than true, if the value of `z` is a number.

Rationale: This odd breakdown of `eq` in the case of numbers allows the implementor enough design freedom to produce exceptionally efficient numerical code on conventional architectures. MACLISP requires this freedom, for example, in order to produce compiled numerical code equal in speed to FORTRAN. If not for this freedom, then at least for the sake of compatibility, COMMON LISP makes this same restriction.

If two objects are to be compared for "identity", but either might be a number, then the predicate `eql` (page 49) is probably appropriate; if both objects are known to be numbers, then `=` (page 118) may be preferable.

As a rule, computations with floating-point numbers are only approximate. The *precision* of a floating-point number is not necessarily correlated at all with the *accuracy* of that number. The precision refers to the number of bits retained in the representation. When an operation combines a short floating-point number with a long one, the result will be a long floating-point number. This rule is made to ensure that as much accuracy as possible is preserved; however, it is by no means a guarantee. COMMON LISP numerical routines do assume, however, that the accuracy of an argument does not exceed its precision. Therefore when two small floating-point numbers are combined, the result will always be a small floating-point number. This assumption can be overridden by first explicitly converting a small floating-point number to a larger representation. (COMMON LISP never converts automatically from a larger size to a smaller one in an effort to save space.)

Rational computations cannot overflow in the usual sense (though of course there may not be enough storage to represent one), as integers and ratios may in principle be of any magnitude. Floating-point computations may get exponent overflow or underflow, in which case an error is signalled.

12.1. Predicates on Numbers

`zerop` *number* [*Function*]

This predicate is true if *number* is zero (either the integer zero, a floating-point zero, or a complex zero), and is false otherwise. It is an error if the argument *number* is not a number.

`plusp` *number* [*Function*]

This predicate is true if *number* is strictly greater than zero, and is false otherwise. It is an error if the argument *number* is not a non-complex number.

`minusp` *number* [*Function*]

This predicate is true if *number* is strictly less than zero; otherwise `nil` is returned. It is an error if the argument *number* is not a non-complex number.

`oddp` *integer* [*Function*]

This predicate is true if the argument *integer* is odd (not divisible by two), and otherwise is false. It is an error if the argument is not an integer.

`evenp` *integer* [*Function*]

This predicate is true if the argument *integer* is even (divisible by two), and otherwise is false. It is an error if the argument is not an integer.

See also the data-type predicates `integerp` (page 47), `rationalp` (page 47) `floatp` (page 48), `complexp` (page 48), and `numberp` (page 47).

12.2. Comparisons on Numbers

All of the functions in this section require that their arguments be numbers, and signal an error if given a non-number. They work on all types of numbers, automatically performing any required coercions.

`=` *number* `&rest` *more-numbers* [*Function*]
`/=` *number* `&rest` *more-numbers* [*Function*]
`<` *number* `&rest` *more-numbers* [*Function*]
`>` *number* `&rest` *more-numbers* [*Function*]
`<=` *number* `&rest` *more-numbers* [*Function*]

`>=` *number* `&rest` *more-numbers* [*Function*]

These functions each take one or more arguments. If the sequence of arguments satisfies a certain condition:

`=`	all the same
`/=`	all different
`<`	monotonically increasing
`>`	monotonically decreasing
`<=`	monotonically nondecreasing
`>=`	monotonically nonincreasing

then the predicate is true, and otherwise is false. Complex numbers may be compared using `=` and `/=`, but the others require non-complex arguments.

For example:

```
(= 3 3) is true                    (/= 3 3) is false
(= 3 5) is false                   (/= 3 5) is true
(= 3 3 3 3) is true                (/= 3 3 3 3) is false
(= 3 3 5 3) is false               (/= 3 3 5 3) is false
(= 3 6 5 2) is false               (/= 3 6 5 2) is true
(< 3 5) is true                    (<= 3 5) is true
(< 3 -5) is false                  (<= 3 -5) is false
(< 3 3) is false                   (<= 3 3) is true
(< 0 3 4 6 7) is true              (<= 0 3 4 6 7) is true
(< 0 3 4 4 6) is false             (<= 0 3 4 4 6) is true
(> 4 3) is true                    (>= 4 3) is true
(> 4 3 2 1 0) is true              (>= 4 3 2 1 0) is true
(> 4 3 3 2 0) is false             (>= 4 3 3 2 0) is true
(> 4 3 1 2 0) is false             (>= 4 3 1 2 0) is false
```

With two arguments, these functions perform the usual arithmetic comparison tests. With three or more arguments, they are useful for range checks.

For example:

```
(<= 0 x 9)                  ;true iff x is between 0 and 9, inclusive
(< 0.0 x 1.0)               ;true iff x is between 0.0 and 1.0, exclusive
(< -1 j (length s))         ;true iff j is a valid index for s
(<= 0 j k (- (length s) 1))       ;true iff j and k are each valid
                                  ;    indices for s and also j≤k
```

For two non-complex arguments *x* and *y*, the law of trichotomy holds. Exactly one of `(= x y)`, `(< x y)`, and `(> x y)` will be true. Also:

```
(/= x y) <=> (not (= x y)) <=> (or (< x y) (> x y))
(<= x y) <=> (not (> x y)) <=> (or (< x y) (= x y))
(>= x y) <=> (not (< x y)) <=> (or (< x y) (= x y))
```

These relationships do *not* generalize to more or fewer than two arguments.

Rationale: The "unequality" relation is called "`/=`" rather than "`<>`" (the name used in PASCAL) for two reasons. First, `/=` of more than two arguments is not the same as the `or` of `<` and `>` of those same arguments. Second, unequality is meaningful for complex numbers even though `<` and `>` are not. For both reasons it would be misleading to associate unequality with the names of `<` and `>`.

Compatibility note: In COMMON LISP, the comparison operations perform "mixed-mode" comparisons: `(= 3 3.0)` is true. In MACLISP, there must be exactly two arguments, and they must be either both fixnums or both

floating-point numbers. To compare two numbers for numerical equality and type equality, use `eql` (page 49).

`max` *number* `&rest` *more-numbers* [*Function*]

The arguments may be any non-complex numbers. `max` returns the argument which is greatest (closest to positive infinity).

For example:

```
(max 1 3 2 -7) => 3
(max -2 3 0 7) => 7
(max 3) => 3
(max 3.0 7 1) => 7 or 7.0
```

If the arguments are a mixture of integers and floating-point numbers, and the largest is a rational, then the implementation is free to produce either that rational or its floating-point approximation.

`min` *number* `&rest` *more-numbers* [*Function*]

The arguments may be any non-complex numbers. `min` returns the argument which is least (closest to negative infinity).

For example:

```
(max 1 3 2 -7) => -7
(max -2 3 0 7) => -2
(min 3) => 3
(min 3.0 7 1) => 1 or 1.0
```

If the arguments are a mixture of rationals and floating-point numbers, and the smallest is a rational, then the implementation is free to produce either that rational or its floating-point approximation.

`fuzzy=` *number1 number2* `&optional` *fuzz* [*Function*]

This predicate is true if *number1* and *number2* are "roughly equal". The optional argument *fuzz* allows nearly-equal numbers to be considered equal: two numbers x and y are considered to be equal if the absolute value of their difference is no greater than *fuzz* times the absolute value of the one with the larger absolute value, that is, if $abs(x-y) \leq fuzz * max(abs(x), abs(y))$. If no third argument is supplied, then *fuzz* defaults to

```
(max (fuzziness x) (fuzziness y))
```

For example:

`(fuzzy= 2/3 0.6666 0.001)` is true

`fuzziness` *number* [*Function*]

The accuracy of a number, in the absence of any context, is not really a mathematically well-defined concept, because it depends on how the number was calculated and on the accuracy of the givens. Nevertheless the following arbitrary definition of "fuzziness" is offered in its place, for use by `fuzzy=`.

The fuzziness of a rational number is zero. The fuzziness of a floating-point number is $2^{-2f/3}$ where f is the number of bits in the fraction of the floating-point number. The fuzziness of a complex number z is

```
(max (fuzziness (realpart z)) (fuzziness (imagpart z)))
```

12.3. Arithmetic Operations

All of the functions in this section require that their arguments be numbers, and signal an error if given a non-number. They work on all types of numbers, automatically performing any required coercions.

`+ &rest` *numbers* [*Function*]

Returns the sum of the arguments. If there are no arguments, the result is `0`, which is an identity for this operation.

Compatibility note: While + is compatible with its use in Lisp Machine LISP, it is incompatible with MACLISP, which uses + for fixnum-only addition.

`-` *number* `&rest` *more-numbers* [*Function*]

The function `-`, when given one argument, returns the negative of that argument.

The function `-`, when given more than one argument, subtracts from the first argument all the others, and returns the result.

Compatibility note: While - is compatible with its use in Lisp Machine LISP, it is incompatible with MACLISP, which uses - for fixnum-only subtraction. Also, - differs from `difference` as used in most LISP systems in the case of one argument.

`* &rest` *numbers* [*Function*]

Returns the product of the arguments. If there are no arguments, the result is `1`, which is an identity for this operation.

Compatibility note: While * is compatible with its use in Lisp Machine LISP, it is incompatible with MACLISP, which uses * for fixnum-only multiplication.

`/` *number* `&rest` *more-numbers* [*Function*]

The function `/`, when given more than one argument, divides the first argument by all the others, and returns the result.

With one argument, `/` reciprocates the argument.

`/` will produce a ratio if the mathematical quotient of two integers is not an exact integer.

For example:

```
(/ 12 4) => 3
(/ 13 4) => 13/4
(/ -8) => -1/8
```

To divide one integer by another producing an integer result, use one of the functions `floor`,

`ceil`, `trunc`, or `round` (page 131).

If any argument is a floating-point number, then the rules of floating-point contagion apply.

Compatibility note: What `/` does is totally unlike what the usual `//` or `quotient` operator does. In most LISP systems, `quotient` behaves like `/` except when dividing integers, in which case it behaves like `trunc` (page 131) of two arguments; this behavior is mathematically intractable, leading to such anomalies as

```
(quotient 1.0 2.0) => 0.5    but    (quotient 1 2) => 0
```

In practice `quotient` is used only when one is sure that both argument are integers, *or* when one is sure that at least one argument is a floating-point number. `/` is tractable for its purpose, and "works" for *any* numbers. For "integer division", `trunc` (page 131), `floor` (page 131), `ceil` (page 131), and `round` (page 131) are available in COMMON LISP.

`1+` *number* [*Function*]

`1-` *number* [*Function*]

`(1+ x)` is the same as `(+ x 1)`.

`(1- x)` is the same as `(- x 1)`. Note that the short name may be confusing: `(1- x)` does *not* mean $1-x$; rather, it means $x-1$.

Rationale: These are included primarily for compatibility with MACLISP and Lisp Machine LISP. Programmers may wish to avoid the possible confusion in new code.

Implementation note: Compiler writers are very strongly encouraged to ensure that `(1+ x)` and `(+ x 1)` compile into identical code, and similarly for `(1- x)` and `(- x 1)`, to avoid pressure on a LISP programmer to write possibly less clear code for the sake of efficiency. This can easily be done as a source-language transformation.

`incf` *place* [*delta*] [*Macro*]

`decf` *place* [*delta*] [*Macro*]

The number produced by the form *delta* is added to (`incf`) or subtracted from (`decf`) the number in the generalized variable named by *place* , and the sum is stored back into *place* and returned. The form *place* may be any form acceptable as a generalized variable to `setf` (page 60). If *delta* is not supplied, then the number in *place* is changed by 1.

For example:

```
(setq n 0)
(incf n) => 1        and now n => 1
(decf n 3) => -2     and now n => -2
(decf n -5) => 3     and now n => 3
(decf n) => 2        and now n => 2
```

The effect of (`incf` *place delta*) is roughly equivalent to

```
(setf place (+ place delta))
```

except that the latter would evaluate any subforms of *place* twice, while `incf` takes care to evaluate them only once. Moreover, for certain *place* forms `incf` may be significantly more efficient than the `setf` version.

`conjugate` *number* [*Function*]

This returns the complex conjugate of *number*. The conjugate of a non-complex number is itself. For a complex number z,

```
(conjugate z) <=> (complex (realpart z) (- (imagpart z)))
```

`gcd &rest` *rationals* [*Function*]

Returns the greatest common divisor of all the arguments, which must be rationals or complex rationals (complex numbers whose components are rational).

If the arguments are all integers, the result is always a non-negative integer.

??? **Query:** I, GLS, hereby recant all this complex rational nonsense. Shall we revert to `gcd` just supporting plain old integers?

If the arguments are all rationals, the result is always a non-negative rational.

If the arguments are all Gaussian integers (complex numbers with integer components), the result is always a first-quadrant Gaussian integer.

In the general case, the result is that complex rational of greatest possible magnitude that is in the first quadrant (including the positive real axis and zero, and excluding the positive imaginary axis) and that when divided into each argument produces a Gaussian integer.

If no arguments are given, `gcd` returns 0, which is an identity for this operation.

??? **Query:** Is `gcd` of more than two arguments ever really used? If not, is the overhead of the multiple-argument implementation worth the elegance? (Similarly for `lcm`.)

`lcm` *rational* `&rest` *more-rationals* [*Function*]

This returns the least common multiple of its arguments, which must be rationals or complex rationals. For two arguments,

```
(lcm a b) <=> (/ (* a b) (gcd a b))
```

For one argument, `lcm` returns that argument. For three or more arguments,

```
(lcm a b c ... z) <=> (lcm (lcm a b) c ... z)
```

For example:

```
(lcm 14 35) => 70
(lcm 3/4 2/5) => 6
```

12.4. Irrational and Transcendental Functions

COMMON LISP provides no data type that can accurately represent irrational values. The functions in this section are described as if the results were mathematically accurate, but they actually all produce floating-point approximations to the true mathematical result. In some places mathematical identities are set forth that are intended to elucidate the meanings of the functions; however, two mathematically identical expressions may be computationally different because of errors inherent in the floating-point approximation process.

12.4.1. Exponential and Logarithmic Functions

`exp` *number* [*Function*]

Returns *e* raised to the power `number`, where *e* is the base of the natural logarithms.

`expt` *base-number power-number* [*Function*]

Returns *base-number* raised to the power *power-number*. If the *base-number* is rational and the *power-number* is an integer, the calculation will be exact and the result will be rational; otherwise a floating-point approximation may result.

Implementation note: If the exponent is an integer a repeated-squaring algorithm may be used, while if the exponent is a floating-point number or complex the result may be calculated as:

```
(exp (* power-number (log base-number)))
```

or in any other reasonable manner.

`log` *number* `&optional` *base* [*Function*]

Returns the logarithm of *number* in the base *base*, which defaults to *e*, the base of the natural logarithms.

For example:

```
(log 8.0 2) => 3.0
(log 0.01 10) => -2.0
```

`sqrt` *number* [*Function*]

Returns the principal square root of *number*.

`isqrt` *integer* [*Function*]

Integer square-root: the argument must be a non-negative integer, and the result is the greatest integer less than or equal to the exact positive square root of the argument.

12.4.2. Trigonometric and Related Functions

`abs` *number* [*Function*]

Returns the absolute value of the argument. For a non-complex number,

```
(abs x) <=> (if (minusp x) (- x) x)
```

For a complex number *z*, the absolute value may be computed as

```
(sqrt (+ (expt (realpart z 2)) (expt (imagpart z 2))))
```

For non-complex numbers, `abs` is a rational function, but it may be irrational for complex arguments.

`phase` *number* [*Function*]

The phase of a number is the angle part of its polar representation as a complex number. That is,

```
(phase x) <=> (atan (realpart x) (imagpart x))
```

The result is in radians, in the range $-\pi$ (exclusive) to π (inclusive). The phase of zero is defined to be zero.

`signum` *number* [*Function*]

By definition,

```
(signum x) <=> (if (zerop x) x (/ x (abs x)))
```

For a rational number, `signum` will return one of `-1`, `0`, or `1` according to whether the number is negative, zero, or positive. For a floating-point number, the result will be a floating-point number of the same format with one of the mentioned three values. For a complex number *z*, (`signum` *z*) is a complex number of the same phase but with unit magnitude.

For non-complex numbers, `signum` is a rational function, but it may be irrational for complex arguments.

`sin` *radians* [*Function*]
`cos` *radians* [*Function*]
`tan` *radians* [*Function*]

`sin` returns the sine of the argument, `cos` the cosine, and `tan` the tangent. The argument is in radians. The argument may be complex.

`cis` *radians* [*Function*]

This computes $e^{i*radians}$. The name "`cis`" means "cos + *i* sin", because $e^{i\theta} = \cos\theta + i\sin\theta$. The argument is in radians, and may be any non-complex number. The result is a complex number whose real part is the cosine of the argument, and whose imaginary part is the sine. Put another way, the result is a complex number whose phase is the argument and whose magnitude is unity.

Implementation note: Often it is cheaper to calculate the sine and cosine of a single angle together than to perform two disjoint calculations.

`asin` *number* [*Function*]
`acos` *number* [*Function*]

`asin` returns the arcsine of the argument, and `cos` the arccosine. The result is in radians. The argument may be complex.

`atan` *y* `&optional` *x* [*Function*]

An arctangent is calculated and the result is returned in radians.

With two arguments *y* and *x*, neither argument may be complex. The result is the arctangent of the quantity *y*/*x*. The signs of *y* and *x* are used to derive quadrant information; moreover, *x* may be zero provided *y* is not zero. The value of `atan` is always between $-\pi$ (exclusive) and π (inclusive).

The following table details various special cases.

Condition		Cartesian locus	Range of result
$y = 0$	$x > 0$	Positive x-axis	0
$y > 0$	$x > 0$	Quadrant I	$0 < \text{result} < \pi/2$
$y > 0$	$x = 0$	Positive y-axis	$\pi/2$
$y > 0$	$x < 0$	Quadrant II	$\pi/2 < \text{result} < \pi$
$y = 0$	$x < 0$	Negative x-axis	π
$y < 0$	$x < 0$	Quadrant III	$-\pi < \text{result} < -\pi/2$
$y < 0$	$x = 0$	Negative y-axis	$-\pi/2$
$y < 0$	$x > 0$	Quadrant IV	$-\pi/2 < \text{result} < 0$
$y = 0$	$x = 0$	Origin	error

Actually, the < signs in the above table ought to be ≤ signs, because of rounding effects; if y is greater than zero but nevertheless very small, then the floating-point approximation to $\pi/2$ might be a more accurate result than any other floating-point number. (For that matter, when `y` = 0 the exact value $\pi/2$ cannot be produced anyway, but instead only an approximation.)

With only one argument y, the argument may be complex. The result is the arctangent of y. For non-complex arguments the result lies between $-\pi/2$ and $\pi/2$ (both exclusive).

Compatibility note: MACLISP has a function called `atan` which range from 0 to 2π. Every other language in the world (ANSI FORTRAN, IBM PL/I, InterLISP) has an arctangent function with range $-\pi$ to π. Lisp Machine LISP provides two functions, `atan` (compatible with MACLISP) and *atan2* (compatible with everyone else).

COMMON LISP makes `atan` the standard one with range $-\pi$ to π. Observe that this makes the one-argument and two-argument versions of `atan` compatible in the sense that the branch cuts do not fall in different places, which is probably why most languages use this definition. (An aside: the INTERLISP one-argument function `arctan` has a range from 0 to π, while every other language in the world provides the range $-\pi/2$ to $\pi/2$! Nevertheless, since INTERLISP uses the standard two-argument version, its branch cuts are inconsistent anyway.)

`pi` [*Variable*]

This global variable has as its value the best possible approximation to π in the largest floating-point format provided by the implementation.

For example:

```
(defun sind (x)                         ;The argument is in degrees.
  (sin (* x (/ (float pi x) 180))))
```

An approximation to π in some other precision can be obtained by writing `(float pi` *x*`)`, where *x* is a floating-point number of the desired precision; see `float` (page 130).

`sinh` *number* [*Function*]

`cosh` *number* [*Function*]

`tanh` *number* [*Function*]

`asinh` *number* [*Function*]

`acosh` *number* [*Function*]

`atanh` *number* [*Function*]

These functions compute the hyperbolic sine, cosine, tangent, arcsine, arccosine, and arctangent functions, which are mathematically defined as follows:

Hyperbolic sine	$(e^x - e^{-x})/2$
Hyperbolic cosine	$(e^x + e^{-x})/2$
Hyperbolic tangent	$(e^x - e^{-x})/(e^x + e^{-x})$
Hyperbolic arcsine	$\log (x+\sqrt{1+x^2})$
Hyperbolic arccosine	$\log (x+(x+1)\sqrt{(x-1)/(x+1)})$
Hyperbolic arctangent	$\log ((1+x)\sqrt{1-1/x^2})$

Implementation note: These formulae are mathematically correct, assuming completely accurate computation. They may be terrible methods for floating-point computation! Implementors should consult a good text on numerical analysis. The formulas given above are not necessarily the simplest ones for real-valued computations, either; they are chosen to define the branch cuts in desirable ways for the complex case.

12.4.3. Branch Cuts, Principal Values, and Boundary Conditions in the Complex Plane

Many of the irrational and transcendental functions are multiply-defined in the complex domain; for example, there are in general an infinite number of complex values for the logarithm function. In each such case a principal value must be chosen for the function to return. In general, such values cannot be chosen so as to make the range continuous; lines of discontinuity called *branch cuts* must be defined.

COMMON LISP defines the branch cuts, principal values, and boundary conditions for the complex functions following a proposal for complex functions in APL [8]. The contents of this section are borrowed largely from that proposal.

`sqrt` The branch cut for square root lies along the negative real axis, continuous with quadrant II. The range consists of the right half-plane, including the non-negative imaginary axis and excluding the negative imaginary axis.

`phase` The branch cut for the phase function lies along the negative real axis, continuous with quadrant II. The range consists of that portion of the real axis between $-\pi$ (exclusive) and π (inclusive).

`log` The branch cut for the logarithm function of one argument (natural logarithm) lies along the negative real axis, continuous with quadrant II. The domain excludes the origin. For a complex number $z = x + y\,i$, log z is defined to be $(\log |z|) + i\,phase(z)$. Therefore the range of the one-argument logarithm function is that strip of the complex plane containing numbers with imaginary parts between $-\pi$ (exclusive) and π (inclusive).

The two-argument logarithm function is defined as $\log_b z = (\log z)/(\log b)$. This defines the principal values precisely. The range of the two-argument logarithm function is the entire complex plane. It is an error if z is zero. If z is nonzero and b is zero, the logarithm is taken to be zero.

`exp` The simple exponential function has no branch cut.

`expt` The two-argument exponential function is defined as $b^x = e^{x \log b}$. This defines the

principal values precisely. The range of the two-argument exponential function is the entire complex plane. Regarded as a function of x, with b fixed, there is no branch cut. Regarded as a function of b, with x fixed, there is, in general, a branch cut along the negative real axis, continuous with quadrant II, and the domain excludes the origin. By definition, $0^0=1$. If $b=0$ and the real part of x is strictly positive, then $b^x=0$. For all other values of x, 0^x is an error.

asin

The following definition for arcsine determines the range and branch cuts:

$$\arcsin z = -i \log (i z + \sqrt{1-z^2})$$

The branch cut for the arcsine function is in two pieces: one along the negative real axis to the left of -1 (inclusive), continuous with quadrant II, and one along the positive real axis to the right of 1 (inclusive), continuous with quadrant IV. The range is that strip of the complex plane containing numbers whose real part is between $-\pi/2$ and $\pi/2$. A number with real part equal to $-\pi/2$ is in the range iff its imaginary part is non-negative; a number with real part equal to $\pi/2$ is in the range iff its imaginary part is non-positive.

acos

The following definition for arccosine determines the range and branch cuts:

$$\arccos z = -i \log (z + i \sqrt{1-z^2})$$

or, which is equivalent,

$$\arccos z = (\pi/2) - \arcsin z$$

The branch cut for the arccosine function is in two pieces: one along the negative real axis to the left of -1 (inclusive), continuous with quadrant II, and one along the positive real axis to the right of 1 (inclusive), continuous with quadrant IV. This is the same branch cut as for arcsine. The range is that strip of the complex plane containing numbers whose real part is between 0 and π. A number with real part equal to 0 is in the range iff its imaginary part is non-negative; a number with real part equal to π is in the range iff its imaginary part is non-positive.

atan

The following definition for (one-argument) arctangent determines the range and branch cuts:

$$\arctan z = -i \log ((1+i z) \sqrt{1/(1+z^2)})$$

Beware of simplifying this formula; "obvious" simplifications are likely to alter the branch cuts or the values on the branch cuts incorrectly. The branch cut for the arctangent function is in two pieces: one along the positive imaginary axis above i (exclusive), continuous with quadrant II, and one along the negative imaginary axis below $-i$ (exclusive), continuous with quadrant IV. The points i and $-i$ are excluded from the domain. The range is that strip of the complex plane containing numbers whose real part is between $-\pi/2$ and $\pi/2$. A number with real part equal to $-\pi/2$ is in the range iff its imaginary part is strictly positive; a number with real part equal to $\pi/2$ is in the range iff its imaginary part is strictly negative. Thus the range of arctangent is identical to that of arcsine with the points $-\pi/2$ and $\pi/2$ excluded.

asinh The following definition for the inverse hyperbolic sine determines the range and branch cuts:

$$\text{arcsinh } z = \log\left(x + \sqrt{1+x^2}\right)$$

The branch cut for the inverse hyperbolic sine function is in two pieces: one along the positive imaginary axis above i (inclusive), continuous with quadrant I, and one along the negative imaginary axis below $-i$ (inclusive), continuous with quadrant III. The range is that strip of the complex plane containing numbers whose imaginary part is between $-\pi/2$ and $\pi/2$. A number with imaginary part equal to $-\pi/2$ is in the range iff its real part is non-positive; a number with imaginary part equal to $\pi/2$ is in the range iff its imaginary part is non-negative.

acosh The following definition for the inverse hyperbolic cosine determines the range and branch cuts:

$$\text{arccosh } z = \log\left(x + (x+1)\sqrt{(x-1)/(x+1)}\right)$$

The branch cut for the inverse hyperbolic cosine function lies along the real axis to the left of 1 (inclusive), extending indefinitely along the negative real axis, continuous with quadrant II and (between 0 and 1) with quadrant I. The range is that half-strip of the complex plane containing numbers whose real part is non-negative and whose imaginary part is between $-\pi$ (exclusive) and π (inclusive). A number with real part zero is in the range iff its imaginary part is between zero (inclusive) and π (inclusive).

atanh The following definition for the inverse hyperbolic tangent determines the range and branch cuts:

$$\text{arctanh } z = \log\left((1+x)\sqrt{1-1/x^2}\right)$$

Beware of simplifying this formula; "obvious" simplifications are likely to alter the branch cuts or the values on the branch cuts incorrectly. The branch cut for the inverse hyperbolic tangent function is in two pieces: one along the negative real axis to the left of -1 (inclusive), continuous with quadrant III, and one along the positive real axis to the right of 1 (inclusive), continuous with quadrant I. The range is that strip of the complex plane containing numbers whose imaginary part is between $-\pi/2$ and $\pi/2$. A number with imaginary part equal to $-\pi/2$ is in the range iff its real part is strictly negative; a number with imaginary part equal to $\pi/2$ is in the range iff its imaginary part is strictly positive. Thus the range of arctangent is identical to that of arcsine with the points $-\pi i/2$ and $\pi i/2$ excluded.

With these definitions, the following useful identities are obeyed throughout the applicable portion of the complex domain, even on the branch cuts:

$\sin iz = i\sinh z$	$\sinh iz = i\sin z$	$\arctan iz = i\,\text{arctanh } z$
$\cos iz = \cosh z$	$\cosh iz = \cos z$	$\text{arcsinh } iz = i\arcsin z$
$\tan iz = i\tanh z$	$\arcsin iz = i\,\text{arcsinh } z$	$\text{arctanh } iz = i\arctan z$

12.5. Type Conversions and Component Extractions on Numbers

While most arithmetic functions will operate on any kind of number, coercing types if necessary, the following functions are provided to allow specific conversions of data types to be forced, when desired.

`float` *number* `&optional` *other* [*Function*]

Converts any non-complex number to a floating-point number. With no second argument, then if a given format of floating-point number is sufficiently precise to represent the result, then the result may be of that format or of any larger format, depending on the implementation; but if no fixed format is sufficiently precise, then the format of greatest precision provided by the implementation is used.

If the argument *other* is provided, then it must be a floating-point number, and *number* is converted to the same format as *other*.

`rational` *number* [*Function*]

`rationalize` *number* `&optional` *tolerance* [*Function*]

Each of these functions converts any non-complex number to be a rational number. If the argument is already rational, that argument is returned. The two functions differ in their treatment of floating-point numbers.

`rational` assumes that the floating-point number is completely accurate, and returns a rational number mathematically equal to the precise value of the floating-point number. This is (probably) much faster than `rationalize`.

`rationalize` assumes that the floating-point number is accurate only to the precision of the floating-point representation, and may return any rational number for which the floating-point number is the best available approximation of its format; in doing this it attempts to keep both numerator and denominator small. It is always the case that

```
(eql (float (rationalize x) x) x)
```

That is, rationalizing a floating-point number and then converting it back to a floating-point number of the same format produces the original number.

The optional argument *tolerance* may be used to alter the assumption concerning precision. If *tolerance* is a positive integer, then *number* is assumed to be accurate to only that many bits. If it is a negative integer, then *number* is assumed to be accurate only to within that many bits of the low end of the fraction. If it is a positive floating-point number, then it is a relative tolerance; *number* is assumed to be precise only to an amount equal to *number* times *tolerance*.

??? **Query:** (1) Should *tolerance* be applied even if the argument is not a floating-point number? For example, `(rationalize 113/355 0.01)` might produce `22/7`.

(2) Should the third argument to `fuzzy=` be like the second argument to `rationalize`? Then perhaps we could make the claim that

```
(fuzzy= (float (rationalize x tol) x) x tol)
```

for all `x` and `tol`.

`numerator` *rational* [*Function*]
`denominator` *rational* [*Function*]

These functions take a rational number (an integer or ratio) and return as an integer the numerator or denominator of the canonical reduced form of the rational. The numerator of an integer is that integer, and the denominator of an integer is 1. Note that

```
(gcd (numerator x) (denominator x)) => 1
```

The denominator will always be a strictly positive integer; the numerator may be any integer.

For example:

```
(numerator (/ 8 -6)) => -4
(denominator (/ 8 -6)) => 3
```

There is no `fix` function in COMMON LISP, because there are several interesting ways to convert non-integral values to integers. These are provided by the functions below, which perform not only type-conversion but also some non-trivial calculations.

`floor` *number* `&optional` *divisor* [*Function*]
`ceil` *number* `&optional` *divisor* [*Function*]
`trunc` *number* `&optional` *divisor* [*Function*]
`round` *number* `&optional` *divisor* [*Function*]

??? **Query:** Should we rename `ceil` and `trunc` to be `ceiling` and `truncate`?

In the simple, one-argument case, each of these functions converts its argument *number* (which may not be complex) to be an integer. If the argument is already an integer, it is returned directly. If the argument is a ratio or floating-point number, the functions use different algorithms for the conversion.

`floor` converts its argument by truncating towards negative infinity; that is, the result is the largest integer which is not larger than the argument.

`ceil` converts its argument by truncating towards positive infinity; that is, the result is the smallest integer which is not smaller than the argument.

`trunc` converts its argument by truncating towards zero; that is, the result is the integer of the same sign as the argument and which has the greatest integral magnitude not greater than that of the argument.

`round` converts its argument by rounding to the nearest integer; if *number* is exactly halfway between two integers (that is, of the form *integer*+0.5) then it is rounded to the one which is even (divisible by two).

Here is a table showing what the four functions produce when given various arguments.

Argument	floor	ceiling	trunc	round
2.6	2	3	2	3
2.5	2	3	2	2
2.4	2	3	2	2
0.7	0	1	0	1
0.3	0	1	0	0
-0.3	-1	0	0	0
-0.7	-1	0	0	-1
-2.4	-3	-2	-2	-2
-2.5	-3	-2	-2	-2
-2.6	-3	-2	-2	-3

If a second argument *divisor* is supplied, then the result is the appropriate type of rounding or truncation applied to the result of dividing the *number* by the *divisor*. For example, `(floor 5 2) = (floor (/ 5 2))`, but is potentially more efficient. The *divisor* may be any non-complex number. The one-argument case is exactly like the two-argument case where the second argument is 1.

Each of the functions actually returns *two* values; the second result is the remainder, and may be obtained using `multiple-value-bind` (page 82) and related constructs. If any of these functions is given two arguments x and y and produces results q and r, then $q*y+r=x$. The remainder r is an integer if both arguments are integers, is rational if both arguments are rational, and is floating-point if either argument is floating-point. (In the one-argument case the remainder is a number of the same type as the argument.) The first result is always an integer.

Compatibility note: The names of the functions `floor`, `ceil`, `trunc`, and `round` are more accurate than names like `fix` which have heretofore been used in various LISP systems. The names used here are compatible with standard mathematical terminology (and with PL/I, as it happens). In FORTRAN `ifix` means `trunc`. ALGOL 68 provides `round`, and uses `entier` to mean `floor`. In MACLISP, `fix` and `ifix` both mean `floor` (one is generic, the other flonum-in/fixnum-out). In INTERLISP, `fix` means `trunc`. In Lisp Machine LISP, `fix` means `floor` and `fixr` means `round`. STANDARD LISP provides a `fix` function, but does not accurately specify what it does exactly. The existing usage of the name `fix` is so confused that it seems best to avoid it altogether.

The names and definitions given here have recently been adopted by Lisp Machine LISP, and MACLISP and NIL seem likely to follow suit.

`mod` *number* `&optional` *divisor tolerance* [*Function*]
`rem` *number* `&optional` *divisor tolerance* [*Function*]

If the optional argument *tolerance* is omitted, `mod` performs the operation `floor` (page 131) on its arguments, and returns the *second* result of `floor` as its only result. Similarly, `rem` performs the operation `trunc` (page 131) on its arguments, and returns the *second* result of `trunc` as its only result.

`mod` and `rem` are therefore the usual modulus and remainder functions when applied to two integer arguments. In general, however, the arguments may be integers or floating-point numbers.

With one argument, these functions perform the "mod 1" or "fractional part" operation, differing in the direction of rounding: the result of `mod` of one argument is always non-negative, while the result of `rem` of one argument always has the same sign as the argument.

```
(mod 13 4) => 1                    (rem 13 4) => 1
(mod -13 4) => 3                   (rem -13 4) => -1
(mod 13 -4) => -3                  (rem 13 -4) => 1
(mod -13 -4) => -1                 (rem -13 -4) => -1
(mod 13.4) => 0.4                  (rem 13.4) => 0.4
(mod -13.4) => 0.6                 (rem -13.4) => -0.4
```

If the optional argument *tolerance* is given, then it is handled in the following manner. Like the optional argument *tolerance* to `rationalize`, it may be a positive or negative integer or a positive floating-point number. For expository purposes define `(delta` *x*`)`, for a floating-point number *x*, to be one-half the value of the smallest floating-point number *y* of the same format as *x* such that

```
(> (+ y (float 1 x)) x)
```

Then define the function `compute-tolerance` as follows:

```
(defun compute-tolerance (x tol)
  (cond ((floatp tol) tol)
        ((minusp tol) (* (delta x) (expt 2 (- tol))))
        (t (expt 2 (- tol)))))
```

Now when `mod` or `rem` performs its computation, it is as if it called `floor` or `trunc` and returned the second result. Let the *first* result from `floor` or `trunc` be called *q*; this will be an integer. If `mod` or `rem` is given the optional argument *tolerance*, it will signal an error, rather than delivering a result, if *number* is a floating-point number and

```
(> (* q (delta number)) (compute-tolerance number tolerance))
```

The interpretation is that *tolerance* is a measure of the accuracy required of the computed remainder. If the quotient *q* is very large, then the original *number* must have been so large relative to the *divisor* that the remainder cannot be very accurate.

`ffloor` *number* `&optional` *divisor* [*Function*]
`fceil` *number* `&optional` *divisor* [*Function*]
`ftrunc` *number* `&optional` *divisor* [*Function*]
`fround` *number* `&optional` *divisor* [*Function*]

These functions are just like `floor`, `ceil`, `trunc`, and `round`, except that the result (the first result of two) is always a floating-point number rather than an integer. It is roughly as if `ffloor` gave its arguments to `floor`, and then applied `float` to the first result before passing them both back. In practice, however, `ffloor` may be implemented much more efficiently. Similar remarks apply to the other three functions. If the first argument is a floating-point number, and the second agrument is not a floating-point number of shorter format, then the first result will be a floating-point number of the same type as the first argument.

For example:

```
(ffloor -4.7) => -5.0 and 0.3
(ffloor 3.5d0) => 3.0d0 and 0.5d0
```

`float-fraction` *float* [*Function*]
`float-exponent` *float* [*Function*]
`scale-float` *float integer* [*Function*]

The function `float-fraction` takes a floating-point number and returns a new floating-point number of the same format. Let *b* be the radix for the floating-point representation (see `short-float-radix` (page 143) and friends); then float-fraction divides the argument by an integral power of *b* so as to bring its value between 1/*b* (inclusive) and 1 (inclusive), and returns the quotient.

The function `float-exponent` performs a similar operation, but then returns the integer exponent to which *b* must be raised to produce the appropriate power for the division.

The function `scale-float` takes a floating-point number *f* and an integer k, and returns `(* f (expt (float b f) k))`. (The use of `scale-float` may be much more efficient than using exponentiation and multiplication.)

Note that `(scale-float (float-fraction f) (float-exponent f))` <=> *f*.

Rationale: These functions allow the writing of machine-independent, or at least machine-parameterized, floating-point software of reasonable efficiency.

`complex` *realpart* `&optional` *imagpart* [*Function*]

The arguments must be non-complex numbers; a complex number is returned that has *realpart* as its real part and *imagpart* as its imaginary part. If *imagpart* is not specified then `(*` *realpart* `0)` is effectively used (this definition has the effect that in this case the two parts will be both rational or both floating-point numbers of the same format).

`realpart` *number* [*Function*]
`imagpart` *number* [*Function*]

These return the real and imaginary parts of a complex number. If *number* is a non-complex number, then `realpart` returns its argument *number* and `imagpart` returns `(*` *number* `0)` (this has the effect that the imaginary part of a rational is 0 and that of a floating-point number is a floating-point zero of the same format).

??? **Query:** What would be the pros and cons of *requiring* the two parts of a complex number to be either both rational or both floating-point numbers of the same format?

12.6. Logical Operations on Numbers

The logical operations in this section treat integers as if they were represented in two's-complement notation.

Implementation note: Internally, of course, an implementation of COMMON LISP may or may not use a two's-complement representation. All that is necessary is that the logical operations perform calculations so as to give this appearance to the user.

The logical operations provide a convenient way to represent an infinite vector of bits. Let such a conceptual vector be indexed by the non-negative integers. Then bit j is assigned a "weight" 2^j. Assume that only a finite number of bits are ones, or that only a finite number of bits are zeros. A vector with only a finite number of one-bits is represented as the sum of the weights of the one-bits, a positive integer. A vector with only a finite number of zero-bits is represented as `-1` minus the sum of the weights of the zero-bits, a negative integer.

This method of using integers to represent bit vectors can in turn be used to represent sets. Suppose that some (possibly countably infinite) universe of discourse for sets is mapped into the non-negative integers. Then a set can be represented as a bit vector; an element is in the set if the bit whose index corresponds to that element is a one-bit. In this way all finite sets can be represented (by positive integers), as well as all sets whose complements are finite (by negative integers). The functions `logior`, `logand`, and `logxor` defined below then compute the union, intersection, and symmetric difference operations on sets represented in this way.

`logior &rest` *integers* [*Function*]

Returns the bit-wise logical *inclusive or* of its arguments. If no argument is given, then the result is zero, which is an identity for this operation.

`logxor &rest` *integers* [*Function*]

Returns the bit-wise logical *exclusive or* of its arguments. If no argument is given, then the result is zero, which is an identity for this operation.

`logand &rest` *integers* [*Function*]

Returns the bit-wise logical *and* of its arguments. If no argument is given, then the result is `-1`, which is an identity for this operation.

`logeqv &rest` *integers* [*Function*]

Returns the bit-wise logical *equivalence* (also known as *exclusive nor*) of its arguments. If no argument is given, then the result is `-1`, which is an identity for this operation.

`lognand` *integer1 integer2* [*Function*]
`lognor` *integer1 integer2* [*Function*]
`logandc1` *integer1 integer2* [*Function*]
`logandc2` *integer1 integer2* [*Function*]
`logorc1` *integer1 integer2* [*Function*]
`logorc2` *integer1 integer2* [*Function*]

These are the other six non-trivial bit-wise logical operations on two arguments. Because they are not commutative or associative, they take exactly two arguments rather than any non-negative number of arguments.

```
(lognand n1 n2) <=> (lognot (logand n1 n2))
 (lognor n1 n2) <=> (lognot (logor n1 n2))
(logandc1 n1 n2) <=> (logand (lognot n1) n2)
(logandc2 n1 n2) <=> (logand n1 (lognot n2))
 (logorc1 n1 n2) <=> (logor (lognot n1) n2)
 (logorc2 n1 n2) <=> (logor n1 (lognot n2))
```

The ten bit-wise logical operations on two integers are summarized in this table:

Argument 1	0	0	1	1	
Argument 2	0	1	0	1	*Operation name*
`logand`	0	0	0	1	and
`logior`	0	1	1	1	inclusive or
`logxor`	0	1	1	0	exclusive or
`logeqv`	1	0	0	1	equivalence (exclusive nor)
`lognand`	1	1	1	0	not-and
`lognor`	1	0	0	0	not-or
`logandc1`	0	1	0	0	and complement of arg1 with arg2
`logandc2`	0	0	1	0	and arg1 with complement of arg2
`logorc1`	1	1	0	1	or complement of arg1 with arg2
`logorc2`	1	0	1	1	or arg1 with complement of arg2

`boole` *op integer1 integer2* [*Function*]

`boole-clr` [*Variable*]

`boole-set` [*Variable*]

`boole-1` [*Variable*]

`boole-2` [*Variable*]

`boole-c1` [*Variable*]

`boole-c2` [*Variable*]

`boole-and` [*Variable*]

`boole-ior` [*Variable*]

`boole-xor` [*Variable*]

`boole-eqv` [*Variable*]

`boole-nand` [*Variable*]

`boole-nor` [*Variable*]

`boole-andc1` [*Variable*]

`boole-andc2` [*Variable*]

`boole-orc1` [*Variable*]

`boole-orc2` [*Variable*]

The function `boole` takes an operation *op* and two integers, and returns an integer produced by performing the logical operation specified by *op* on the two integers. The precise values of the sixteen variables are implementation-dependent, but they are suitable for use as the first argument to `boole`:

integer1	0	0	1	1	
integer2	0	1	0	1	*Operation performed*
`boole-clr`	0	0	0	0	always 0
`boole-set`	1	1	1	1	always 1
`boole-1`	0	0	1	1	*integer1*
`boole-2`	0	1	0	1	*integer2*
`boole-c1`	1	1	0	0	complement of *integer1*
`boole-c2`	1	0	1	0	complement of *integer2*
`boole-and`	0	0	0	1	and
`boole-ior`	0	1	1	1	inclusive or
`boole-xor`	0	1	1	0	exclusive or
`boole-eqv`	1	0	0	1	equivalence (exclusive nor)
`boole-nand`	1	1	1	0	not-and
`boole-nor`	1	0	0	0	not-or
`boole-andc1`	0	1	0	0	and complement of *integer1* with *integer2*
`boole-andc2`	0	0	1	0	and *integer1* with complement of *integer2*
`boole-orc1`	1	1	0	1	or complement of *integer1* with *integer2*
`boole-orc2`	1	0	1	1	or *integer1* with complement of *integer2*

`boole` can therefore compute all sixteen logical functions on two arguments. In general,

```
(boole boole-and x y) <=> (logand x y)
```

and the latter is more perspicuous. However, `boole` is useful when it is necessary to parameterize a procedure so that it can use one of several logical operations.

`lognot` *integer* [*Function*]

Returns the bit-wise logical *not* of its argument. Every bit of the result is the complement of the corresponding bit in the argument.

```
(logbitp j (lognot x)) <=> (not (logbitp j x))
```

`logtest` *integer1 integer2* [*Function*]

`logtest` is a predicate which is true if any of the bits designated by the 1's in *integer1* are 1's in *integer2*.

```
(logtest x y) <=> (not (zerop (logand x y)))
```

`logbitp` *index integer* [*Function*]

logbitp is true if the bit in *integer* whose index is *index* (that is, its weight is 2^{index}) is a one-bit; otherwise it is false.

For example:

```
(logbitp 2 6) is true
(logbitp 0 6) is false
(logbitp k n) <=> (ldb-test (byte 1 k) n)
```

`ash` *integer count* [*Function*]

Shifts *integer* arithmetically left by *count* bit positions if *count* is positive, or right -*count* bit positions if *count* is negative. The sign of the result is always the same as the sign of *integer*.

Arithmetically, this operation performs the computation *floor*(*integer**2^{count}).

Logically, this moves all of the bits in *integer* to the left, adding zero-bits at the bottom, or moves them to the right, discarding bits. (In this context the question of what gets shifted in on the left is irrelevant; integers, viewed as strings of bits, are "half-infinite", that is, conceptually extend infinitely far to the left.)

For example:

```
(logbitp j (ash n k))
   <=> (and (>= j k) (logbitp (- j k) n))
```

`logcount` *integer* [*Function*]

The number of bits in *integer* is determined and returned. If *integer* is positive, then 1 bits in its binary representation are counted. If *integer* is negative, then the 0 bits in its two's-complement binary representation are counted. The result is always a non-negative integer.

For example:

```
(logcount 13) => 3          ; Binary representation is ...0001101
(logcount -13) => 2         ; Binary representation is ...1110011
(logcount 30) => 4          ; Binary representation is ...0011110
(logcount -30) => 4         ; Binary representation is ...1100010
```

As a rule,

```
(logcount x) <=> (logcount (- (+ x 1)))
```

`haulong` *integer* [*Function*]

This returns the number of significant bits in the absolute value of *integer*. The precise computation performed is *ceiling*(log2(*abs*(*integer*)+1)).

For example:

```
(haulong 0) => 0
(haulong 3) => 2
(haulong 4) => 3
(haulong -7) => 3
```

`haipart` *integer count* [*Function*]

Returns the high *count* bits of the binary representation of the absolute value of *integer*, or the low -*count* bits if *count* is negative. A possible definition of `haipart`:

```
(defun haipart (integer count)
       (let ((x (abs integer)))
            (if (minusp count)
                (ldb (byte (- count) 0) x)
                (ldb (byte count (max (- (haulong x) n) 0))
                     x))))
```

12.7. Byte Manipulation Functions

Several functions are provided for dealing with an arbitrary-width field of contiguous bits appearing anywhere in an integer. Such a contiguous set of bits is called a *byte*. Here the term *byte* does not imply some fixed number of bits (such as eight), but a field of arbitrary and user-specifiable width.

The byte-manipulation functions use objects called *byte specifiers* to designate a specific byte position within an integer. The representation of a byte specifier is implementation-dependent; it is sufficient to know that the function *byte* will construct one, and that the byte-manipulation functions will accept them. The function *byte* accepts two integers representing the *position* and *size* of the byte, and returns a byte specifier. Such a specifier designates a byte whose width is *size*, and whose right-hand bit has weight $2^{position}$, in the terminology of integers used as logical bit vectors.

`byte` *size position* [*Function*]

byte takes two integers representing the size and position of a byte, and returns a byte specifier suitable for use as an argument to byte-manipulation functions.

`byte-size` *bytespec* [*Function*]
`byte-position` *bytespec* [*Function*]

Given a byte specifier, `byte-size` returns the size specified as an integer; `byte-position` similarly returns the position.

For example:

```
(byte-size (byte j k)) <=> j
(byte-position (byte j k)) <=> k
```

`ldb` *bytespec integer* [*Function*]

bytespec specifies a byte of *integer* to be extracted. The result is returned as a positive integer.

For example:

```
(logbitp j (ldb (byte s p) n)
   <=> (and (< j s) (logbitp (+ j p) n))
```

The name of the function "`ldb`" means "load byte".

`ldb-test` *bytespec integer* [*Function*]

`ldb-test` is a predicate which is true if any of the bits designated by the byte specifier *bytespec* are 1's in *integer*; that is, it is true if the designated field is non-zero.

```
(ldb-test bytespec n) <=> (not (zerop (ldb bytespec n)))
```

`mask-field` *bytespec integer* [*Function*]

This is similar to `ldb`; however, the result contains the specified byte of *integer* in the position specified by *bytespec*, rather than in position 0 as with `ldb`. The result therefore agrees with *integer* in the byte specified, but has zero bits everywhere else.

For example:

```
(ldb bs (mask-field bs n)) <=> (ldb bs n)
(logbitp j (mask-field (byte s p) n))
   <=> (and (>= j p) (< j s) (logbitp j n))
(mask-field bs n) <=> (logand n (ldb bs -1))
```

`dpb` *newbyte bytespec integer* [*Function*]

Returns a number which is the same as *integer* except in the bits specified by *bytespec*. Let *s* be the size specified by *bytespec*; then the low *s* bits of *newbyte* appear in the result in the byte specified by *bytespec*. The integer *newbyte* is therefore interpreted as being right-justified, as if it were the result of `ldb`.

For example:

```
(logbitp j (dpb m (byte s p) n))
   <=> (if (and (>= j p) (< j (+ p s)))
           (logbitp (- j p) m)
           (logbitp j n))
```

The name of the function "`dpb`" means "deposit byte".

`deposit-field` *newbyte bytespec integer* [*Function*]

This function is to `mask-field` as `dpb` is to `ldb`. The result is an integer which contains the bits of *newbyte* within the byte specified by *bytespec*, and elsewhere contains the bits of *integer*.

For example:

```
(logbitp j (dpb m (byte s p) n))
   <=> (if (and (>= j p) (< j (+ p s)))
           (logbitp j m)
           (logbitp j n))
```

Implementation note: If the *bytespec* is a constant, one may of course construct, at compile time, an equivalent mask *m*, for example by computing `(deposit-field -1` *bytespec* `0)`. Given this mask *m*, one may then compute

`(deposit-field` *newbyte bytespec integer*`)]`

by computing

`(logor (logand` *newbyte m*`) (logand` *integer* `(lognot` *m*`)))`

where the result of `(lognot` *m*`)` can of course also be computed at compile time. However, the following expression (which I got indirectly from Knuth) may also be used, and may require fewer temporary registers in some situations:

`(logxor` *integer* `(logand` *m* `(logxor` *integer newbyte*`)))`

A related, though possibly less useful, trick is that

```
(let ((z (logand (logxor x y) m)))
  (setq x (logxor z x))
  (setq y (logxor z y)))
```

interchanges those bits of `x` and `y` for which the mask `m` is 1, and leaves alone those bits of `x` and `y` for which `m`

is 0.

12.8. Random Numbers

`random` *number1* `&optional` *number2* [*Function*]

(`random` *n*) accepts a positive number *n* and returns a number of the same kind between zero (inclusive) and *n* (exclusive). The number *n* may be an integer or a floating-point number. An approximately uniform choice distribution is used; If *n* is an integer, each of the possible results occurs with (approximate) probability $1/n$.

(`random` *low high*) is equivalent to (`+` *low* (`random` (`-` *high low*))); it provides a choice from the range *low* (inclusive) to *high* (exclusive).

Compatibility note: In INTERLISP, the range limits are both inclusive. Would this be more intuitive? It is easy to implement for integers, but much harder for floating-point numbers.

Compatibility note: `random` of zero arguments has been omitted because its value is too implementation-dependent (limited by fixnum range).

Implementation note: In general, it is not adequate to define (`random` *n*) for integral *n* to be simply (`mod` (`random`) *n*); this fails to be uniformly distributed if *n* is larger than the largest number produced by `random`, or even if *n* merely approaches this number. Assuming that the underlying mechanism produces "random bits" (possibly in chunks such as fixnums), the best approach is to produce enough random bits to construct an integer *k* some number *d* of bits larger than (`haulong` *n*) (see `haulong` (page 138)), and then compute (`mod` *k n*). The quantity *d* should be at least 7, and preferably 10 or more.

To produce random floating-point numbers in the range [*A*, *B*), accepted practice (as determined by a quick look through the *Collected Algorithms from the ACM*, particularly algorithms 133, 266, 294, and 370) is to compute $X*(B-A)+A$, where X is a floating-point number uniformly distributed over [0.0, 1.0) and computed by calculating a random integer N in the range [0, M) (typically by a multiplicative-congruential or linear-congruential method mod M) and then setting $X=N/M$. If one takes $M = 2^f$, where f is the length of the fraction of a floating-point number (and it is in fact common to choose M to be a power of two), then this method is equivalent to the following assembly-language-level procedure. Assume the representation has no hidden bit. Take a floating-point 0.5, and clobber its entire fraction with random bits. Normalize the result if necessary.

For example, on the PDP-10, assume that accumulator `T` is completely random (all 36 bits are random). Then the code sequence

```
LSH T,-9          ;Clear high 9 bits; low 27 are random.
FSC T,128.        ;Install exponent and normalize.
```

will produce in `T` a random floating-point number uniformly distributed over [0.0, 1.0). (Instead of the `LSH`, one could do "`TLZ T,777000`; but if the 36 random bits came from a congruential random-number generator, the high-order bits tend to be "more random" than the low-order ones, and so the `LSH` would be a bit better for uniform distribution. Ideally all the bits would be the result of high-quality randomness.)

With a hidden-bit representation, normalization is not a problem, but dealing with the hidden bit is. The method can be adapted as follows. Take a floating-point 1.0 and clobber the explicit fraction bits with random bits; this produces a random floating-point number in the range [1.0, 2.0). Then simply subtract 1.0. In effect, we let the hidden bit creep in and then subtract it away again.

For example, on the VAX, assume that register `T` is completely random (but a little less random than on the PDP-10, as it has only 32 random bits). Then the code sequence

```
INSV #^X81,#7,#9,T      ;Install correct sign bit and exponent.
SUBF #^F1.0,T           ;Subtract 1.0.
```

will produce in `T` a random floating-point number uniformly distributed over [0.0, 1.0). Again, if the low-order bits are not random enough, then "`ROTL #7,T`" should be performed first.

`random-state` [*Variable*]

This variable holds a data structure which encodes the internal state of the random-number generator used by `random`. The nature of this data structure is implementation-dependent. It may be printed out and successfully read back in, but may or may not function correctly as a random-number state object in another implementation. A call to `random` will perform a side effect on this data structure. Lambda-binding this variable to a different random-number state object will correctly save and restore the old state object, of course.

`random-state &optional` *state* [*Function*]

This function returns a new random-number state object, suitable for use as the value of the variable `random-state`. If *state* is `nil` or omitted, `random-state` returns a *copy* of the current random-number state object (the value of the variable `random-state`). If *state* is a state object, a copy of that state object is returned. If *state* is `t`, then a new state object is returned which has been "randomly" initialized by some means (such as by a time-of-day clock).

12.9. Implementation Parameters

The values of the named constants defined in this section are implementation-dependent. They may be useful for parameterizing code in some situations.

`most-positive-fixnum` [*Constant*]
`most-negative-fixnum` [*Constant*]

The value of `most-positive-fixnum` is that fixnum closest in value to positive infinity provided by the implementation.

The value of `most-negative-fixnum` is that fixnum closest in value to negative infinity provided by the implementation.

`most-positive-short-float` [*Constant*]
`least-positive-short-float` [*Constant*]
`least-negative-short-float` [*Constant*]
`most-negative-short-float` [*Constant*]

The value of `most-positive-short-float` is that short-format floating-point number closest in value to positive infinity provided by the implementation.

The value of `least-positive-short-float` is that positive short-format floating-point number closest in value to zero provided by the implementation.

The value of `least-negative-short-float` is that negative short-format floating-point number closest in value to zero provided by the implementation.

The value of `most-negative-short-float` is that short-format floating-point number closest in value to negative infinity provided by the implementation.

`most-positive-single-float` [*Constant*]
`least-positive-single-float` [*Constant*]
`least-negative-single-float` [*Constant*]
`most-negative-single-float` [*Constant*]
`most-positive-double-float` [*Constant*]
`least-positive-double-float` [*Constant*]
`least-negative-double-float` [*Constant*]
`most-negative-double-float` [*Constant*]
`most-positive-long-float` [*Constant*]
`least-positive-long-float` [*Constant*]
`least-negative-long-float` [*Constant*]
`most-negative-long-float` [*Constant*]

These are analogous to the constants defined above for short-format floating-point numbers.

`short-float-radix` [*Constant*]
`single-float-radix` [*Constant*]
`double-float-radix` [*Constant*]
`long-float-radix` [*Constant*]

These constants indicate, for each floating-point format, the radix used in the floating-point representation. (For most contemporary computers this is 2, but for the IBM 370 it is 16, for example.) See `float-fraction` (page 134).

`short-float-epsilon` [*Constant*]
`single-float-epsilon` [*Constant*]
`double-float-epsilon` [*Constant*]
`long-float-epsilon` [*Constant*]

These constants indicate, for each floating-point format, the smallest positive number *e* of that format such that

```
(not (= (float 1 e) (+ e (float 1 e))))
```

`short-float-negative-epsilon` [*Constant*]
`single-float-negative-epsilon` [*Constant*]
`double-float-negative-epsilon` [*Constant*]
`long-float-negative-epsilon` [*Constant*]

These constants indicate, for each floating-point format, the smallest positive number *e* of that format such that

```
(not (= (float 1 e) (- e (float 1 e))))
```

Chapter 13

Characters

COMMON LISP provides a character data type; objects of this type represent printed symbols such as letters.

Every character has three attributes: code, bits, and font. The code attribute is intended to distinguish among the printed glyphs and formatting functions for characters. The bits attribute allows extra flags to be associated with a character. The font attribute permits a specification of the style of the glyphs (such as italics).

`char-code-limit` [*Constant*]

The value of `char-code-limit` is a non-negative integer which is the upper exclusive bound on values produced by the function `char-code` (page 149), which returns the *code* component of a given character; that is, the values returned by `char-code` are non-negative and strictly less than the value of `char-code-limit`.

Implementation note: For the PERQ, the value will be 256; for the S-1, 512.

`char-font-limit` [*Constant*]

The value of `char-font-limit` is a non-negative integer which is the upper exclusive bound on values produced by the function `char-font` (page 150), which returns the *font* component of a given character; that is, the values returned by `char-font` are non-negative and strictly less than the value of `char-font-limit`.

Implementation note: No COMMON LISP implementation is required to support non-zero font attributes; if it does not, then `char-font-limit` should be `1`. For the PERQ, the value will be `256`; for the S-1, `512`.

`char-bits-limit` [*Constant*]

The value of `char-bits-limit` is a non-negative integer which is the upper exclusive bound on values produced by the function `char-bits` (page 149), which returns the *bits* component of a given character; that is, the values returned by `char-bits` are non-negative and strictly less than the value of `char-bits-limit`. Note that the value of `char-bits-limit` will be a power of two.

Implementation note: No COMMON LISP implementation is required to support non-zero bits attributes; if it does not, then `char-bits-limit` should be `1`. For the PERQ, the value will be `256`; for the S-1, `512`.

13.1. Predicates on Characters

The predicate `characterp` (page 48) may be used to determine whether any LISP object is a character object.

`standard-charp` *char* [*Function*]

The argument *char* must be a character object. `standard-charp` is true if the argument is a "standard character", that is, one of the ninety-five ASCII printing characters or <return>. If the argument is a non-standard character, then `standard-charp` is false.

Note in particular that any character with a non-zero *bits* or *font* attribute is non-standard.

`graphicp` *char* [*Function*]

The argument *char* must be a character object. `graphicp` is true if the argument is a "graphic" (printing) character, and false if it is a "non-graphic" (formatting or control) character. Graphic characters have a standard textual representation as a single glyph, such as "`A`" or "`*`" or "`=`". By convention, the space character is considered to be graphic. Of the standard characters (as defined by `standard-charp`), all but <return> are graphic. If an implementation provides any of the semi-standard characters <backspace>, <tab>, <rubout>, <linefeed>, and <form>, they are not graphic.

Graphic characters of font 0 may be assumed all to be of the same width when printed; programs may depend on this for purposes of columnar formatting. Non-graphic characters and characters of other fonts may be of varying widths.

Any character with a non-zero bits attribute is non-graphic.

`string-charp` *char* [*Function*]

The argument *char* must be a character object. `string-charp` is true if *char* can be stored into a string (see the functions `char` (page 191) and `rplachar` (page 192)), and otherwise is false. Any character which satisfies `standard-charp` and `graphicp` also satisfies `string-charp`; others may also.

`alphap` *char* [*Function*]

The argument *char* must be a character object. `alphap` is true if the argument is an alphabetic character, and otherwise is false.

Of the standard characters (as defined by `standard-charp`), the letters "A" through "Z" and "a" through "z" are alphabetic.

`uppercasep` *char* [*Function*]
`lowercasep` *char* [*Function*]
`bothcasep` *char* [*Function*]

The argument *char* must be a character object. `uppercasep` is true if the argument is an upper-case (majuscule) character, and otherwise is false. `lowercasep` is true if the argument is an lower-case (minuscule) character, and otherwise is false.

`bothcasep` is true if the argument is upper-case *and* there is a corresponding lower-case character (which can be obtained using `char-downcase` (page 150)), or if the argument is lower-case and there is a corresponding upper-case character (which can be obtained using `char-upcase` (page 150)).

If a character is either upper-case or lower-case, it is necessarily alphabetic. However, it is permissible in theory for an alphabetic character to be neither uppercase nor lowercase.

Of the standard characters (as defined by `standard-charp`), the letters "A" through "Z" are upper-case and "a" through "z" are lower-case.

`digitp` *char* `&optional` (*radix* `10.`) [*Function*]

The argument *char* must be a character object, and *radix* must be a non-negative integer. `digitp` is a pseudo-predicate: if *char* is not a digit of the radix specified by *radix*, then it is false; otherwise it returns a non-negative integer which is the "weight" of *char* in that radix.

Digits are necessarily graphic characters.

Of the standard characters (as defined by `standard-charp`), the characters "0" through "9", "A" through "Z", and "a" through "z" are digits. The weights of "0" through "9" are the integers 0 through 9, and of "A" through "Z" (and also "a" through "z") are 10 through 35. `digitp` returns the weight for one of these digits if and only if its weight is strictly less than *radix*. Thus, for example, the digits for radix 16 are "0123456789ABCDEF".

```
(defun convert-string-to-integer (str &optional (radix 10))
  "Given a digit string and optional radix, return an integer."
  (do ((j 0 (+ j 1))
       (n 0 (+ (* n radix)
               (or (digitp (char str j) radix)
                   (ferror "Bad radix-~D digit: ~C"
                           radix
                           (char str i))))))
      ((= j (string-length str)) n)))
```

`alphanumericp` *char* [*Function*]

The argument *char* must be a character object. `alphanumericp` is true if *char* is either alphabetic or numeric. By definition,

```
(alphanumericp x) <=> (or (alphap x) (digitp x))
```

Alphanumeric characters are therefore are necessarily graphic (as defined by `graphicp` (page 146)).

Of the standard characters (as defined by `standard-charp`), the characters "0" through "9", "A" through "Z", and "a" through "z" are alphanumeric.

`char=` *char1 char2* [*Function*]

The arguments *char1* and *char2* must be character objects. `char=` is true if *char1* and *char2* are equivalent character objects, having equivalent attributes, and otherwise is false.

The function `CHAR=` is the finest discriminator of characters available to the programmer. If `(char= c1 c2)` is true, then any function professing to operate on a character must behave the same whether given `c1` or `c2`.

For non-"funny" characters (those not satisfying `funny-charp` (page FUNNY-CHARP-FUN)),

```
(CHAR= C1 C2) <=>
      (AND (= (CHAR-CODE C1) (CHAR-CODE C2))
           (= (CHAR-BITS C1) (CHAR-BITS C2))
           (= (CHAR-FONT C1) (CHAR-FONT C2)))
```

There is no requirement that `(eq c1 c2)` be true merely because `(char= c1 c2)` is true. While `eq` may distinguish two character objects that `char=` does not, it is distinguishing them not as *characters*, but in some sense on the basis of a lower-level implementation characteristic. (Of course, if `(eq c1 c2)` is true then one may expect `(char= c1 c2)` to be true.) However, `eql` (page 49) and `equal` (page 50) compare character objects in the same way that `char=` does.

`char-equal` *char1 char2* [*Function*]

The arguments *char1* and *char2* must be character objects.

The predicate `char-equal` is like `char=`, except that it ignores differences of font and bits attributes and case. By definition,

```
(char-equal c1 c2) <=>
      (char= (char-upcase (character c1))
             (char-upcase (character c2)))
```

For example:

`(char-equal #\A #\a)` is true
`(char= #\A #\a)` is false
`(char-equal #\A (control #\A))` is true

`char<` *char1 char2* [*Function*]
`char>` *char1 char2* [*Function*]

The arguments *char1* abd *char2* must be character objects. The predicate `char<` is true if *char1* precedes *char2* in the (implementation-dependent) total ordering on characters. The predicate `char>` is true if *char1* follows *char2* in the (implementation-dependent) total ordering on characters. Neither is true if the arguments satisfy `char=` (page 148).

The total ordering on characters is guaranteed to have the following properties:

- The alphanumeric characters obey the following partial ordering:

```
A<B<C<D<E<F<G<H<I<J<K<L<M<N<O<P<Q<R<S<T<U<V<W<X<Y<Z
a<b<c<d<e<f<g<h<i<j<k<l<m<n<o<p<q<r<s<t<u<v<w<x<y<z
0<1<2<3<4<5<6<7<8<9
either 9<A or Z<0
either 9<a or z<0
```

  This implies that alphabetic ordering holds, and that the digits as a group are not interleaved with letters, but that the possible interleaving of upper-case letters and lower-case letters is unspecified.

- If two characters have the same bits and font attributes, then their ordering by `char<` is consistent with the numerical ordering by the predicate `<` (page 118) on their code attributes.

`char-lessp` *char1 char2* [*Function*]

`char-greaterp` *char1 char2* [*Function*]

The arguments *char1* and *char2* must be character objects. The predicate `char-lessp` is like `char<`, except that it ignores differences of font and bits attributes and case; similarly `char-greaterp` is like `char>`. By definition,

```
(char-lessp c1 c2) <=>
     (char< (char-upcase (character c1))
            (char-upcase (character c2)))
```

13.2. Character Construction and Selection

`character` *object* [*Function*]

The function `character` coerces its argument to be a character if possible. If the argument is a character, the argument is simply returned. If the argument is a string of length 1, then the sole element of the string is returned. If the argument is a symbol whose print name is of length 1, then the sole element of the print name is returned. If the argument is an integer *n*, then (`int-char` *n*) is returned.

??? **Query:** This definition is more restrictive than the Lisp Machine Lisp version. Should it be loosened?

`char-code` *char* [*Function*]

The argument *char* must be a character object. `char-code` returns the *code* attribute of the character object; this will be a non-negative integer less than the (normal) value of the variable `char-code-limit` (page 145).

`char-bits` *char* [*Function*]

The argument *char* must be a character object. `char-bits` returns the *bits* attribute of the character object; this will be a non-negative integer less than the (normal) value of the variable `char-bits-limit` (page 145).

`char-font` *char* [*Function*]

The argument *char* must be a character object. `char-font` returns the *font* attribute of the character object; this will be a non-negative integer less than the (normal) value of the variable `char-font-limit` (page 145).

`code-char` *code* `&optional` (*bits* `0`) (*font* `0`) [*Function*]

All three arguments must be non-negative integers. If it is possible in the implementation to construct a character object whose code attribute is *code*, whose bits attribute is *bits*, and whose font attribute is *font*, then such an object is returned; otherwise `nil` is returned.

For any integers *c*, *b*, and *f*, if (`code-char` *c b f*) is not `nil` then

```
(char-code (code-char c b f)) => c
(char-bits (code-char c b f)) => b
(char-font (code-char c b f)) => f
```

If the font and bits attributes of a character object *x* are zero, then it is the case that

```
(char= (code-char (char-code c)) c) is true
```

`make-char` *char* `&optional` (*bits* `0`) (*font* `0`) [*Function*]

The argument *char* must be a character, and *bits* and *font* must be non-negative integers. If it is possible in the implementation to construct a character object whose code attribute is that of *char*, whose bits attribute is *bits*, and whose font attribute is *font*, then such an object is returned; otherwise `nil` is returned.

If *bits* and *font* are zero, then `make-char` cannot fail. This implies that for every character object one can "turn off" its bits and font attributes.

13.3. Character Conversions

`char-upcase` *char* [*Function*]

`char-downcase` *char* [*Function*]

The argument *char* must be a character object. `char-upcase` attempts to convert its argument to an upper-case equivalent; `char-downcase` attempts to convert to lower case.

`char-upcase` returns a character object with the same font and bits attributes as *char*, but with possibly a different code attribute. If the code is different from *char*'s, then the predicate `lowercasep` (page 147) is true of *char*, and `uppercasep` (page 147) is true of the result character. Moreover, if (`char=` (`char-upcase` `x`) `x`) is *not* true, then it is true that

```
(char= (char-downcase (char-upcase x)) x)
```

Similarly, `char-downcase` returns a character object with the same font and bits attributes as *char*, but with possibly a different code attribute. If the code is different from *char*'s, then the predicate `uppercasep` (page 147) is true of *char*, and `lowercasep` (page 147) is true of the result character. Moreover, if (`char=` (`char-downcase` `x`) `x`) is *not* true, then it is true that

```
(char= (char-upcase (char-downcase x)) x)
```

`digit-charp` *weight* `&optional` (*radix* `10.`) (*bits* `0`) (*font* `0`) [*Function*]
`digit-weight` *weight* `&optional` (*radix* `10.`) (*bits* `0`) (*font* `0`) [*Function*]

All arguments must be integers. `digit-charp` determines whether or not it is possible to construct a character object whose bits attribute is *bits*, whose font attribute is *font*, and whose *code* is such that the result character has the weight *weight* when considered as a digit of the radix *radix* (see the predicate `digitp` (page 147)). It returns `t` if that is possible, and otherwise returns `nil`.

`digit-charp` cannot return `nil` if *bits* and *font* are zero, *radix* is between 2 and 36 inclusive, and *weight* is non-negative and less than *radix*.

`digit-weight` assumes that its arguments satisfy `digit-charp`, and constructs such a character. If more than one character object can encode such a weight in the given radix, one shall be chosen consistently by any given implementation; moreover, among the standard characters upper-case letters are preferred to lower-case letters).

For example:

```
(digit-char 7) => #\7
(digit-char 12) => nil
(digit-char 12 16) => #\C          ;not #\c
(digit-char 6 2) => nil
(digit-char 1 2) => #\1
```

`char-int` *char* [*Function*]

The argument *char* must be a character object. `char-int` returns a non-negative integer encoding the character object.

If the font and bits attributes of *char* are zero, then `char-int` returns the same integer `char-code` would. Also,

```
(char= c1 c2) <=> (= (char-int c1) (char-int c2))
```

for characters `c1` and *c2*.

This function is provided primarily for the purpose of hashing characters. Also, the function `tyi` (page 239) is defined in terms of `char-int`.

`int-char` *integer* [*Function*]

The argument must be a non-negative integer. `int-char` returns a character object `c` such that `(char-int c)` is equal to *integer*, if possible; otherwise `int-char` is false.

`char-name` *char* [*Function*]

The argument *char* must be a character object. If the character has a name, then that name (a symbol) is returned; otherwise `nil` is returned. All characters which have zero font and bits attributes and which are non-graphic (do not satisfy the predicate `graphicp` (page 146)) have names. Graphic characters may or may not have names.

The standard characters <return> and <space> have the respective names `return` and `space`. The optional characters <tab>, <form>, <rubout>, <linefeed>, and <backspace> have the respective names `tab`, `form`, `rubout`, `linefeed`, and `backspace`.

Characters which have names can be notated as "`#\`" followed by the name: `#\Space`.

`name-char` *sym* [*Function*]

The argument `sym` must be a symbol. If the symbol is the name of a character object, that object is returned; otherwise `nil` is returned.

13.4. Character Control-Bit Functions

COMMON LISP provides explicit names for four bits of the bits attribute: *Control*, *Meta*, *Hyper*, and *Super*. The following definitions are provided for manipulating these. Each COMMON LISP implementation provides these functions for compatibility, even if it does not support any or all of the bits named below.

`char-control-bit` [*Constant*]
`char-meta-bit` [*Constant*]
`char-super-bit` [*Constant*]
`char-hyper-bit` [*Constant*]

The values of these named constants are the "weights" (as integers) for the four named control bits. The weight of the control bit is `1`; of the meta bit, `2`; of the super bit, `4`; and of the hyper bit, `8`.

If a given implementation of COMMON LISP does not support a particular bit, then the corresponding variable is zero instead.

`char-bit` *char name* [*Function*]

`char-bit` takes a character object *char* and the name of a bit, and returns non-`nil` if the bit of that name is set in *char*, or `nil` if the bit is not set in *char*. Valid values for *name* are implementation-dependent, but typically are `:control`, `:meta`, `:hyper`, and `:super`.

For example:

```
(char-bit #\Control-X :control) => true
```

`set-char-bit` *char name newvalue* [*Function*]

`char-bit` takes a character object *char*, the name of a bit, and a flag. A character is returned which is just like *char* except that the named bit is set or reset according to whether *newvalue* is non-`nil` or `nil`. Valid values for *name* are implementation-dependent, but typically are `:control`, `:meta`, `:hyper`, and `:super`.

For example:

```
(set-char-bit #\X :control t) => #\Control-X
(set-char-bit #\Control-X :control t) => #\Control-X
(set-char-bit #\Control-X :control nil) => #\X
```

Chapter 14

Sequences

The type `sequence` encompasses both lists and one-dimensional arrays, including vectors, strings, and bit-vectors. While these are different data structures with different structural properties leading to different algorithmic uses, they do have a common property: each contains contain an ordered set of elements.

There are some operations which are useful on both lists and arrays because they deal with ordered sets of elements. One may ask the number of elements, reverse the ordering, extract a subsequence, and so on. For such purposes COMMON LISP provides a set of generic functions on sequences:

`elt`	`reverse`	`map`	`remove`	`remove-duplicates`
`setelt`	`nreverse`	`some`	`delete`	`delete-duplicates`
`subseq`	`concat`	`every`	`position`	`find`
`copyseq`	`length`	`notany`	`mismatch`	`substitute`
`fill`	`sort`	`notevery`	`maxprefix`	`search`
`replace`	`merge`		`maxsuffix`	`count`

Some of these operations come in more than one version. Such versions are indicated by adding a suffix to the basic name of the operation. In addition, many operations accept one or more optional keyword arguments that can modify the operation in various ways.

If the operation requires testing sequence elements according to some criterion, then the criterion may be specified in one of two ways. The basic operation accepts an item, and elements are tested for being `eql` to that item. (A test other than `eql` can be specified by the `:test` or `:test-not` keyword.)

??? **Query:** Should the default test be `equal` or `eql`? If `eql`, what about `member`, `delet`, and `assoc`?

The variants formed by adding "`-if`" and "`-if-not`" to the basic operation name do not take an item, but instead a one-argument predicate, and elements are tested for satisfying or not satisfying the predicate. As an example,

`(remove` *item sequence*`)`

returns a copy of *sequence* from which all elements `eql` to *item* have been removed;

`(remove` *item sequence* `:test #'equal)`

returns a copy of *sequence* from which all elements `equal` to *item* have been removed;

`(remove-if #'numberp` *sequence*`)`

returns a copy of *sequence* from which all numbers have been removed; and

`(remove-if #'(lambda (x) (fuzzy= x` *number tolerance*`))` *sequence*`)`

returns a copy of *sequence* from which all elements fuzzily equal to *number* to with *tolerance* have been removed.

If an operation tests elements of a sequence in any manner, the keyword argument `:key`, if not `nil`, should be a function of one argument that will extract from an element the part to be tested in place of the whole element. For example, the effect of the MACLISP expression `(assq item seq)` could be obtained by

```
(find item sequence :test #'eq :key #'car)
```

This searches for the first element of *sequence* whose *car* is `eq` to *item*.

For some operations it can be useful to specify the direction in which the sequence is processed. In this case the basic operation normally processes the sequence in the forward direction, and processing in the reverse direction is indicated by a non-`nil` value for the keyword argument `:from-end`.

Many operations allow the specification of a subsequence to be operated upon. Such operations have keyword arguments called `:start` and `:end`. These arguments should be integer indices into the sequence, with *start*≤*end*; they indicate the subsequence starting with and including element *start* and up to but excluding element *end*. The length of the subsequence is therefore *end*−*start*. If *start* is omitted it defaults to zero, and if *end* is omitted or `nil` it defaults to the length of the sequence; therefore if both are omitted the entire sequence is processed by default. For the most part this is permitted purely for the sake of efficiency; one can simply call `subseq` instead to extract the subsequence before operating on it. However, operations which produce indices return indices into the original sequence, not into the subsequence.

```
(position #/b "foobar" :start 2 :end 5) => 3
(position #/b (subseq "foobar" 2 5)) => 1
```

If two sequences are involved, then the `:start` and `:end` values affect *both* sequences. Alternatively, the keyword arguments `:start1`, `:end1`, `:start2`, and `:end2` may be used to specify separate subsequences for each sequence.

For some functions, notably `remove` and `delete`, the keyword argument `:count` is used to specify how many occurrences of the item should be affected. If this is `nil` or is not supplied, all matching items are affected.

In the following function descriptions, an element x of a sequence "satisfies the test" if either of the following holds:

- A basic function was called, *testfn* was specified by the keyword `:test`, and `(funcall` *testfn item* `(`*keyfn* `x))` is true.
- A basic function was called, *testfn* was specified by the keyword `:test-not`, and `(funcall` *testfn item* `(`*keyfn* `x))` is false.
- An "`-if`" function was called, and `(funcall` *predicate* `(`*keyfn* *x*`))` is true.
- An "`-if-not`" function was called, and `(funcall` *predicate* `(`*keyfn* *x*`))` is false.

In each case *keyfn* is the value of the `:key` keyword argument (the default being the identity function). See, for example, `remove` (page 160).

??? **Query:** Again, should the default *testfn* be `eql` or `equal`?

In the following function descriptions, two elements *x* and *y* taken from sequences "match" if either of the following holds:

- *testfn* was specified by the keyword `:test`, and (`funcall` *testfn* (*keyfn x*) (*keyfn y*)) is true.
- *testfn* was specified by the keyword `:test-not`, and (`funcall` *testfn* (*keyfn x*) (*keyfn y*)) is false.

See, for example, `search` (page 164).

14.1. Simple Sequence Functions

`elt` *sequence index* [*Function*]

This returns the element of *sequence* specified by *index*, which must be a non-negative integer less than the length of the *sequence*. The first element of a sequence has index `0`.

`setelt` *sequence index newvalue* [*Function*]

The object *newvalue* is stored into the component of the *sequence* specified by *index*, which must be a non-negative integer less than the length of the *sequence*. The first element of any sequence has index `0`. If *sequence* is a specialized array, then the *newvalue* must be an object which that array can contain. `setelt` returns *newvalue*.

`subseq` *sequence start* `&optional` *end* [*Function*]

This returns the subsequence of *sequence* specified by *start* and *end*. `subseq` *always* allocates a new sequence for a result; it never shares storage with an old sequence. The result subsequence is always of the same type as the argument *sequence*.

`copyseq` *sequence* [*Function*]

A copy is made of the argument *sequence*; the result is `equal` to the argument but not `eq` to it.

```
(copyseq x) <=> (subseq x 0)
```

but the name `copyseq` is more perspicuous when applicable.

`length` *sequence* [*Function*]

The number of elements in *sequence* is returned as a non-negative integer. If the sequence has a fill pointer, the "active length" is returned; that is, `array-active-length` (page 186) is used rather than `array-length` (page ARRAY-LENGTH-FUN).

`reverse` *sequence* [*Function*]

The result is a new sequence of the same kind as *sequence*, containing the same elements but in reverse order. The argument is not modified.

`nreverse` *sequence* [*Function*]

The result is a sequence containing the same elements as *sequence* but in reverse order. The argument may be destroyed and re-used to produce the result. The result may or may not be `eq` to the argument, so it is usually wise to say something like `(setq x (nreverse x))`, because simply `(nreverse x)` is not guaranteed to leave a reversed value in `x`.

14.2. Converting, Catenating, and Mapping Sequences

`to` *result-type sequence* [*Function*]

The *sequence* is converted to be a sequence of type *result-type* and returned. The *result-type* must be a subtype of type `sequence`. If it is specified as simply `array`, for example, then `(array t)` is assumed. If one specifies `sequence`, then `list` is assumed. A specialized type such as `string` or `(vector (complex short-float))` may be specified; of course, the result may be of either that type or some more general type, as determined by the implementation (see Chapter 4).

It is an error if the elements of the *sequence* cannot be put into a sequence of type *result-type*. If the *sequence* is already of the specified type, it may be returned without copying it; in this `(to` *type sequence*`)` differs from `(catenate` *type sequence*`)`, for the latter is required to copy the argument *sequence*.

`catenate` *result-type* `&rest` *sequences* [*Function*]

The result is a new sequence which contains all the elements of all the sequences in order. All of the sequences are copied from; the result does not share any structure with any of the argument sequences (in this `catenate` differs from `append`). The type of the result is specified by *result-type*, which must be a subtype of `sequence`, as for the function `to` (page 158). It must be possible for every element of the argument sequences to be an element of a sequence of type *result-type*.

The implementation must be such that `catenate` is associative, in the sense that the elements of the result sequence are not affected by reassociation (but the type of the result sequence may be affected). If no arguments are provided, `catenate` returns a new empty sequence of type *result-type*.

`map` *result-type function sequence* `&rest` *more-sequences* [*Function*]

The *function* must take as many arguments as there are sequences provided; at least one sequence must be provided. The result of `map` is a sequence such that element *j* is the result of applying *function* to element *j* of each of the argument sequences. The result sequence is as long as the shortest of the input sequences.

If the *function* has side-effects, it can count on being called first on all the elements numbered 0, then on all those numbered 1, and so on.

The type of the result sequence is specified by the argument *result-type*, as for the function `to` (page 158).

Compatibility note: In MACLISP, Lisp Machine LISP, INTERLISP, and indeed even LISP 1.5, the function `map` has always meant a non-value-returning version. In my opinion they blew it. I suggest that for COMMON LISP this should be corrected, as the names `map` and `reduce` have become quite common in the literature, `map` always meaning what in the past LISP people have called `mapcar`. It would simplify things in the future to make the standard (according to the rest of the world) name `map` do the standard thing. Therefore the old `map` function is here renamed `map1` (page 77).

For example:

```
(map 'list #'- '(1 2 3 4)) => (-1 -2 -3 -4)
(map 'bit-vector #'(lambda (x) (if (oddp x) 1 0)) '(1 2 3 4))
   => #"1010"
```

`some` *predicate sequence* `&rest` *more-sequences* [*Function*]
`every` *predicate sequence* `&rest` *more-sequences* [*Function*]
`notany` *predicate sequence* `&rest` *more-sequences* [*Function*]
`notevery` *predicate sequence* `&rest` *more-sequences* [*Function*]

These are all predicates. The *predicate* must take as many arguments as there are sequences provided. The *predicate* is first applied to the elements with index 0 in each of the sequences, and possibly then to the elements with index 1, and so on, until a termination criterion is met or the end of the shortest of the *sequences* is reached.

`some` returns as soon as any invocation of *predicate* returns a non-`nil` value; `some` returns that value. If the end of a sequence is reached, `some` returns `nil`. Thus as a predicate it is true if *some* invocation of *predicate* is true.

`every` returns `nil` as soon as any invocation of *predicate* returns `nil`. If the end of a sequence is reached, `every` returns a non-`nil` value. Thus as a predicate it is true if *every* invocation of *predicate* is true.

`notany` returns `nil` as soon as any invocation of *predicate* returns a non-`nil` value. If the end of a sequence is reached, `notany` returns a non-`nil` value. Thus as a predicate it is true if *no* invocation of *predicate* is true.

`notevery` returns a non-`nil` value as soon as any invocation of *predicate* returns `nil`. If the end of a sequence is reached, `notevery` returns `nil`. Thus as a predicate it is true if *not every* invocation of *predicate* is true.

Compatibility note: The order of the arguments here is not compatible with INTERLISP and Lisp Machine LISP.

This is to stress the similarity of these functions to `map`. The functions are therefore extended here to functions of more than one argument, and multiple sequences.

14.3. Modifying Sequences

`fill` *sequence item* `&key :start :end` [*Function*]

The *sequence* is destructively modified by replacing the elements of the subsequence specified by the `:start` and `:end` parameters with the *item*. The *item* may be any LISP object, but must be a suitable element for the *sequence*. The *item* is stored into all specified components of the *sequence*, beginning at the one specified by the `:start` index (which defaults to zero), and up to but not including the one specified by the `:end` index (which defaults to the length of the sequence). `fill` returns the modified *sequence*.

For example:

```
(setq x (vector 'a 'b 'c 'd 'e)) => #(a b c d e)
(fill x 'z :start 1 :end 3) => #(a z z d e)
  and now x => #(a z z d e)
(fill x 'p) => #(p p p p p)
  and now x => #(p p p p p)
```

`replace` *sequence1 sequence2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]

The sequence *sequence1* is destructively modified by copying successive elements into it from *sequence2*. The elements of *sequence2* must be of a type that may be stored into *sequence1*. The subsequence of *sequence2* specified by `:start2` and `:end2` is copied into the subsequence of *sequence1* specified by `:start1` and `:end1`. (The arguments `:start1` and `:start2` default to `:start`, which defaults to zero. The arguments `:end1` and `:end2` default to `:end`, which defaults to `nil`, meaning the end of the appropriate sequence.) If these subsequences are not of the same length, then the shorter length determines how many elements are copied; the extra elements near the end of the longer subsequence are not involved in the operation. The number of elements copied may be expressed as:

`(min (-` *end1 start1*`) (-` *end2 start2*`))`

The value returned by `replace` is the modified *sequence1*.

If *sequence1* and *sequence2* are the same object and the region being modified overlaps with the region being copied from, then it is as if the entire source region were copied to another place and only then copied back into the target region.

`remove` *item sequence* `&key :from-end :test :test-not :start :end :count :key` [*Function*]

`remove-if` *test sequence* `&key :from-end :start :end :count :key` [*Function*]

`remove-if-not` *test sequence* `&key :from-end :start :end :count :key` [*Function*]

The result is a sequence of the same kind as the argument *sequence*, which has the same elements except that those in the subsequence delimited by `:start` and `:end` and satisfying the test (see

above) have been removed. This is a nondestructive operation; the result is a copy of the input *sequence*, save that some elements are not copied.

The `:count` argument, if supplied, limits the number of elements removed; if more than `:count` elements satisfy the test, only the leftmost `:count` such are removed.

A non-`nil` `:from-end` specification matters only when the `:count` argument is provided; in that case only the rightmost `:count` elements satisfying the test are removed.

For example:

```
(remove 4 '(1 2 4 1 3 4 5)) => (1 2 1 3 5)
(remove 4 '(1 2 4 1 3 4 5) :count 1) => (1 2 1 3 4 5)
(remove 4 '(1 2 4 1 3 4 5) :count 1 :from-end t)
   => (1 2 4 1 3 5)
(remove 3 '(1 2 4 1 3 4 5) :test #'>) => (4 3 4 5)
(remove-if #'oddp '(1 2 4 1 3 4 5)) => (2 4 4)
(remove-if #'evenp '(1 2 4 1 3 4 5) :count 1 :from-end t)
   => (1 2 4 1 3 5)
```

The result of `remove` and related functions may share with the argument *sequence*; a list result may share a tail with an input list, and the result may be `eq` to the input *sequence* if no elements need to be removed.

`delete` *item sequence* `&key :from-end :test :test-not :start :end :count :key` [*Function*]

`delete-if` *test sequence* `&key :from-end :start :end :count :key` [*Function*]

`delete-if-not` *test sequence* `&key :from-end :start :end :count :key` [*Function*]

This is the destructive counterpart to `remove`. The result is a sequence of the same kind as the argument *sequence*, which has the same elements except that those in the subsequence delimited by `:start` and `:end` and satisfying the test (see above) have been deleted. This is a destructive operation. The argument *sequence* may be destroyed and used to construct the result; however, the result may or may not be `eq` to *sequence*.

The `:count` argument, if supplied, limits the number of elements deleted; if more than `:count` elements satisfy the test, only the leftmost `:count` such are deleted.

A non-`nil` `:from-end` specification matters only when the `:count` argument is provided; in that case only the rightmost `:count` elements satisfying the test are deleted.

For example:

```
(delete 4 '(1 2 4 1 3 4 5)) => (1 2 1 3 5)
(delete 4 '(1 2 4 1 3 4 5) :count 1) => (1 2 1 3 4 5)
(delete 4 '(1 2 4 1 3 4 5) :count 1 :from-end t)
   => (1 2 4 1 3 5)
(delete 3 '(1 2 4 1 3 4 5) :test #'>) => (4 3 4 5)
(delete-if #'oddp '(1 2 4 1 3 4 5)) => (2 4 4)
(delete-if #'evenp '(1 2 4 1 3 4 5) :count 1 :from-end t)
   => (1 2 4 1 3 5)
```

`substitute` *newitem olditem sequence* `&key :from-end :test :test-not :start :end :count :key` [*Function*]

`substitute-if` *newitem test sequence* `&key :from-end :start :end :count :key` [*Function*]

`substitute-if-not` *newitem test sequence* `&key :from-end :start :end :count :key` [*Function*]

The result is a sequence of the same kind as the argument *sequence*, which has the same elements except that those in the subsequence delimited by `:start` and `:end` and satisfying the test (see above) have been replaced by *newitem*. This is a nondestructive operation; the result is a copy of the input *sequence*, save that some elements are changed.

The `:count` argument, if supplied, limits the number of elements altered; if more than `:count` elements satisfy the test, only the leftmost `:count` such are replaced.

A non-`nil` `:from-end` specification matters only when the `:count` argument is provided; in that case only the rightmost `:count` elements satisfying the test are removed.

For example:

```
(substitute 9 4 '(1 2 4 1 3 4 5)) => (1 2 9 1 3 9 5)
(substitute 9 4 '(1 2 4 1 3 4 5) :count 1) => (1 2 9 1 3 4 5)
(substitute 9 4 '(1 2 4 1 3 4 5) :count 1 :from-end t)
   => (1 2 4 1 3 9 5)
(substitute 9 3 '(1 2 4 1 3 4 5) :test #'>) => (9 9 4 9 3 4 5)
(substitute-if 9 #'oddp '(1 2 4 1 3 4 5)) => (9 2 4 9 9 4 9)
(substitute-if 9 #'evenp '(1 2 4 1 3 4 5) :count 1 :from-end t)
   => (1 2 4 1 3 9 5)
```

The result of `substitute` and related functions may share with the argument *sequence*; a list result may share a tail with an input list, and the result may be `eq` to the input *sequence* if no elements need to be changed.

`nsubstitute` *newitem olditem sequence* `&key :from-end :test :test-not :start :end :count :key` [*Function*]

`nsubstitute-if` *newitem test sequence* `&key :from-end :start :end :count :key` [*Function*]

`nsubstitute-if-not` *newitem test sequence* `&key :from-end :start :end :count :key` [*Function*]

This is the destructive counterpart to `substitute`. The result is a sequence of the same kind as the argument *sequence*, which has the same elements except that those in the subsequence delimited by `:start` and `:end` and satisfying the test (see above) have been replaced by *newitem*. This is a destructive operation. The argument *sequence* may be destroyed and used to construct the result; however, the result may or may not be `eq` to *sequence*.

14.4. Searching Sequences for Items

`find` *item sequence* `&key :from-end :test :test-not :start :end :key` [*Function*]
`find-if` *test sequence* `&key :from-end :start :end :key` [*Function*]
`find-if-not` *test sequence* `&key :from-end :start :end :key` [*Function*]

If the *sequence* contains an element satisfying the test, then the leftmost such element is returned; otherwise `nil` is returned.

If `:start` and `:end` keyword arguments are given, only the specified subsequence of *sequence* is searched.

If a non-`nil` `:from-end` keyword argument is specified, then the result is the *rightmost* element satisfying the test.

`position` *item sequence* `&key :from-end :test :test-not :start :end :key` [*Function*]
`position-if` *test sequence* `&key :from-end :start :end :key` [*Function*]
`position-if-not` *test sequence* `&key :from-end :start :end :key` [*Function*]

If the *sequence* contains an element satisfying the test, then the index within the sequence of the leftmost such element is returned as a non-negative integer; otherwise `nil` is returned.

If `:start` and `:end` keyword arguments are given, only the specified subsequence of *sequence* is searched. However, the index returned is relative to the entire sequence, not to the subsequence.

If a non-`nil` `:from-end` keyword argument is specified, then the result is the index of the *rightmost* element satisfying the test. (The index returned, however, is an index from the left-hand end, as usual.)

`count` *item sequence* `&key :from-end :test :test-not :start :end :key` [*Function*]
`count-if` *test sequence* `&key :from-end :start :end :key` [*Function*]
`count-if-not` *test sequence* `&key :from-end :start :end :key` [*Function*]

The result is always a non-negative integer, the number of elements in the specified subsequence of *sequence* satisfying the test (see above).

`mismatch` *sequence1 sequence2* `&key :from-end :test :test-not` [*Function*]
`:start :end :start1 :start2 :end1 :end2`

The specified subsequences of *sequence1* and *sequence2* are compared element-wise. If they are of equal length and match in every element, the result is `nil`. Otherwise, the result is a non-negative integer, the index within *sequence1* of the leftmost position at which they fail to match; or, if one is shorter than and a matching prefix of the other, the index within *sequence1* beyond the last position tested is returned.

If a non-`nil` `:from-end` keyword argument is given, then the index of the *rightmost* position in which the sequences differ is returned. The (sub)sequences are aligned at their right-hand ends; the last elements are compared, the penultimate elements, and so on. The index returned is again

an index into *sequence1*.

`maxprefix` *sequence1 sequence2* `&key :from-end :test :test-not` [*Function*]
`:start :end :start1 :start2 :end1 :end2`
`maxsuffix` *sequence1 sequence2* `&key :from-end :test :test-not` [*Function*]
`:start :end :start1 :start2 :end1 :end2`

The arguments *sequence1* and *sequence2* are compared element-wise. The result is a non-negative integer, which for `maxprefix` is the index of the leftmost position at which they fail to match; or, if one is shorter than and a matching prefix of the other, the length of the shorter sequence is returned. If they are of equal length and match in every element, the result is the length of each.

The keyword arguments `:start1` and `:end1` delimit a subsequence of *sequence1* to be matched, and `:start2` and `:end2` delimit a subsequence of *sequence2*. The comparison proceeds by first aligning the left-hand ends of the two subsequences; the index returned is an index into *sequence1*. `maxprefix` is therefore not commutative if `:start1` and `:start2` are not equal.

The `suffix` versions differ in that `1` plus the index of the *rightmost* position in which the sequences differ is returned. The (sub)sequences are aligned at their right-hand ends; the last elements are compared, the penultimate elements, and so on. The index returned is again an index into *sequence1*.

The implementation may choose to match the sequences in any order; there is no guarantee on the number of times the test is made. For example, `maxsuffix` might match lists from left-to-right instead of from right-to-left. Therefore it is a good idea for a user-supplied *predicate* to be free of side-effects.

`search` *sequence1 sequence2* `&key :from-end :test :test-not` [*Function*]
`:start :end :start1 :start2 :end1 :end2`

A search is conducted for a subsequence of *sequence2* that element-wise matches *sequence1*. If there is no such subsequence, the result is `nil`; if there is, the result is the index into *sequence2* of the leftmost element of the leftmost such matching subsequence.

If a non-`nil` `:from-end` keyword argument is given, the index of the leftmost element of the *rightmost* matching subsequence is returned.

The implementation may choose to search the sequence in any order; there is no guarantee on the number of times the test is made. For example, `search-from-end` might search a list from left-to-right instead of from right-to-left. Therefore it is a good idea for a user-supplied predicate be free of side-effects.

`sort` *sequence predicate* `&key :key` [*Function*]
`stable-sort` *sequence predicate* `&key :key` [*Function*]

The *sequence* is destructively sorted according to an ordering determined by the *predicate*. The *predicate* should take two arguments, and return non-`nil` if and only if the first argument is strictly

less than the second (in some appropriate sense). If the first argument is greater than or equal to the second (in the appropriate sense), then the *predicate* should return `nil`.

The `sort` function determines the relationship between two elements by giving keys extracted from the elements to the *predicate*. The function *k*, when applied to an element, should return the key for that element; *k* defaults to the identity function, thereby making the element itself be the key.

The *selector* function should not have any side effects. A useful example of a *selector* function would be a component selector function for a `defstruct` (page 199) structure, for sorting a sequence of structures.

```
(sort a p :key s)
   <=> (sort a #'(lambda (x y) (p (s x) (s y))))
```

While the above two expression are equivalent, the first may be more efficient in some implementations for certain types of arguments. For example, an implementation may choose to apply *k* to each item just once, putting the resulting keys into a separate table, and then sort the parallel tables, as opposed to applying *k* to an item every time just before applying the *predicate*.

If the *k* and *predicate* functions always return, then the sorting operation will always terminate, producing a sequence containing the same elements as the original sequence (that is, the result is a permutation of *sequence*). This is guaranteed even if the *predicate* does not really consistently represent a total order. If the *k* consistently returns meaningful keys, and the *predicate* does reflect some total ordering criterion on those keys, then the elements of the result sequence will conform to that ordering.

The sorting operation performed by `sort` is not guaranteed *stable*, however; elements considered equal by the *predicate* may or may not stay in their original order. The function `stable-sort` guarantees stability, but may be somewhat slower.

The sorting operation may be destructive in all cases. In the case of an array or vector argument, this is accomplished by permuting the elements. In the case of a list, the list is destructively reordered in the same manner as for `nreverse` (page 158). Thus if the argument should not be destroyed, the user must sort a copy of the argument.

Should execution of *k* or *predicate* cause an error, the state of the list or array being sorted is undefined. However, if the error is corrected the sort will, of course, proceed correctly.

Note that since sorting requires many comparisons, and thus many calls to the *predicate*, sorting will be much faster if the *predicate* is a compiled function rather than interpreted.

For example:

```
(defun mostcar (x)
    (if (symbolp x) x (mostcar (car x))))

(sort fooarray #'string-lessp :key #'mostcar)
```

If `fooarray` contained these items before the sort:

```
(Tokens (The lion sleeps tonight))
(Carpenters (Close to you))
((Rolling Stones) (Brown sugar))
((Beach Boys) (I get around))
(Beatles (I want to hold your hand))
```

then after the sort `fooarray` would contain:

```
((Beach Boys) (I get around))
(Beatles (I want to hold your hand))
(Carpenters (Close to you))
((Rolling Stones) (Brown sugar))
(Tokens (The lion sleeps tonight))
```

`merge` *sequence1 sequence2 predicate* `&key :key` [*Function*]

The sequences *sequence1* and *sequence2* are destructively merged according to an ordering determined by the *predicate*. The *predicate* should take two arguments, and return non-`nil` if and only if the first argument is strictly less than the second (in some appropriate sense). If the first argument is greater than or equal to the second (in the appropriate sense), then the *predicate* should return `nil`.

The `merge` function determines the relationship between two elements by giving keys extracted from the elements to the *predicate*. The function *k*, when applied to an element, should return the key for that element; the *k* function defaults to the identity function, thereby making the element itself be the key.

The `:key` function should not have any side effects. A useful example of a `:key` function would be a component selector function for a `defstruct` (page 199) structure, for merging a sequence of structures.

If the *k* and *predicate* functions always return, then the merging operation will always terminate. The result of merging two sequences *x* and *y* is a new sequence *z* such that the length of *z* is the sum of the lengths of *x* and *y*, and *z* contains the all the elements of *x* and *y*. If *x1* and *x2* are two elements of *x*, and *x1* precedes *x2* in *x*, then *x1* precedes *x2* in *z*; similarly for elements of *y*. In other words, *z* is an *interleaving* of *x* and *y*.

Moreover, if *x* and *y* were correctly sorted according to the *predicate*, then *z* will also be correctly sorted. If *x* or *y* is not so sorted, then *z* will not be sorted, but will nevertheless be an interleaving of *x* and *y*.

The merging operation is guaranteed *stable*; if two or more elements are considered equal by the *predicate*, then the elements from *sequence1* will precede those from *sequence2* in the result.

For example:

```
(merge '(1 3 4 6 7) '(2 5 8) #'<) => (1 2 3 4 5 6 7 8)
```

Chapter 15

Manipulating List Structure

A *cons*, or dotted pair, is a compound data object having two components, called the *car* and *cdr*. Each component may be any LISP object. A *list* is a chain of conses linked by *cdr* fields; the chain is terminated by some atom (a non-cons object). An ordinary list is terminated by `nil`, the empty list (also written "`()`"). A list whose *cdr*-chain is terminated by some non-`nil` atom is called a *dotted list*.

The recommended predicate for testing for the end of a list is `endp` (page 168).

15.1. Conses

`car` *x* [*Function*]

Returns the *car* of *x*, which must be a cons or `()`; that is, *x* must satisfy the predicate `listp` (page 47). By definition, the *car* of `()` is `()`. If the cons is regarded as the first cons of a list, then `car` returns the first element of the list.

For example:

```
(car '(a b c)) => a
```

`cdr` *x* [*Function*]

Returns the *cdr* of *x*, which must be a cons or `()`; that is, *x* must satisfy the predicate `listp` (page 47). By definition, the *cdr* of `()` is `()`. If the cons is regarded as the first cons of a list, then `cdr` returns the rest of the list, which is a list with all elements but the first of the original list.

For example:

```
(cdr '(a b c)) => (b c)
```

`c...r` *x* [*Function*]

All of the compositions of up to four *car*'s and *cdr*'s are defined as functions in their own right. The names of these functions begin with "`c`" and end with "`r`", and in between is a sequence of "`a`" and "`d`" letters corresponding to the composition performed by the function.

For example:

`(cddadr x)` is the same as `(cdr (cdr (car (cdr x))))`

If the argument is regarded as a list, then `cadr` returns the second element of the list, `caddr` the third, and `cadddr` the fourth. If the first element of a list is a list, then `caar` is the first element of the sublist, `cdar` is the rest of that sublist, and `cadar` is the second element of the sublist; and so on.

As a matter of style, it is often preferable to define a function or macro to access part of a complicated data structure, rather than to use a long `car/cdr` string:

```
(defmacro lambda-vars (lambda-exp) '(cadr ,lambda-exp))
   ; then use lambda-vars everywhere instead of cadr
```

See also `defstruct` (page 199), which will automatically declare new record data types and access functions for instances of them.

`cons` *x y* [*Function*]

`cons` is the primitive function to create a new *cons*, whose *car* is *x* and whose *cdr* is *y*.

For example:

```
(cons 'a 'b) => (a . b)
(cons 'a (cons 'b (cons 'c '()))) => (a b c)
(cons 'a '(b c d)) => (a b c d)
```

`cons` may be thought of as creating a *cons*, or as adding a new element to the front of a list.

`tree-equal` *x y* [*Function*]

This is a predicate which is true if *x* and *y* are isomorphic trees with identical leaves; that is, if *x* and *y* are `eql`, or if they are both conses and their *car*s are `tree-equal` and their *cdr*s are `tree-equal`. Thus `tree-equal` recursively compares conses (but not any other objects which have components). See `equal` (page 50), which does recursively compare other structured objects.

15.2. Lists

`endp` *object* [*Function*]

The predicate `endp` is the recommended way to test for the end of a list. It is true of conses, false of `nil`, and an error for all other arguments.

Implementation note: Implementations are encouraged to signal an error, especially in the interpreter, for a non-list argument. The `endp` function is defined so as to allow compiled code to perform simply an atom check or a null check if speed is more important than safety.

`list-length` *list* `&optional` *limit* [*Function*]

`list-length` returns, as an integer, the length of *list*. The length of a list is the number of top-level conses in it. If the argument *limit* is supplied, it should be an integer; if the length of the *list* is greater than *limit* (possibly because the *list* is circular!), then *limit* is returned.

For example:

```
(list-length '()) => 0
(list-length '(a b c d)) => 4
(list-length '(a (b c) d)) => 3
(list-length '(a b c d e f g) 4) => 4
```

`list-length` could be implemented by:

```
(defun list-length (x &optional (limit nil limitp))
  (declare (integer limit))
  (do ((n 0 (+ n 1))
       (y x (cdr y)))
      ((endp y) n)
    (when (and limitp (>= n limit))
      (return limit))))
```

See `length` (page 157), which will return the length of any sequence.

`nth` *n list* [*Function*]

(`nth` *n list*) returns the *n*'th element of *list*, where the zeroth element is the *car* of the list. *n* must be a non-negative integer. If the length of the list is not greater than *n*, then the result is ( ), that is, `nil`. (This is consistent with the idea that the *car* and *cdr* of ( ) are each ( ).)

For example:

```
(nth 0 '(foo bar gack)) => foo
(nth 1 '(foo bar gack)) => bar
(nth 3 '(foo bar gack)) => ()
```

Compatibility note: This is not the same as the INTERLISP function called `nth`, which is similar to but not exactly the same as the COMMON LISP function `nthcdr`. This definition of `nth` is compatible with Lisp Machine LISP and NIL. Also, some people have used macros and functions called `nth` of their own in their old MACLISP programs, which may not work the same way; be careful.

`nthcdr` *n list* [*Function*]

(`nthcdr` *n list*) performs the `cdr` operation *n* times on *list*, and returns the result.

For example:

```
(nthcdr 0 '(a b c)) => (a b c)
(nthcdr 2 '(a b c)) => (c)
(nthcdr 4 '(a b c)) => ()
```

In other words, it returns the *n*'th *cdr* of the list.

Compatibility note: This is similar to the INTERLISP function `nth`, except that the INTERLISP function is one-based instead of zero-based.

```
(car (nthcdr n x)) <=> (nth n x)
```

`last` *list* [*Function*]

`last` returns the last cons (*not* the last element!) of *list*. If *list* is ( ), it returns ( ).

For example:

```
(setq x '(a b c d))
(last x) => (d)
(rplacd (last x) '(e f))
x => '(a b c d e f)
(last '(a b c . d)) => (c . d)
```

`list &rest` *args* [*Function*]

`list` constructs and returns a list of its arguments.

For example:

```
(list 3 4 'a (car '(b . c)) (+ 6 -2)) => (3 4 a b 4)
```

`list* arg &rest` *others* [*Function*]

`list*` is like `list` except that the last *cons* of the constructed list is "dotted". The last argument to `list*` is used as the *cdr* of the last cons constructed; this need not be an atom. If it is not an atom, then the effect is to add several new elements to the front of a list.

For example:

```
(list* 'a 'b 'c 'd) => (a b c . d)
```
This is like
```
(cons 'a (cons 'b (cons 'c 'd)))
```
Also:
```
(list* 'a 'b 'c '(d e f)) => (a b c d e f)
(list* x) <=> x
```

`make-list` *size* `&optional` *value* [*Function*]

This creates and returns a list containing *size* elements, each of which is *value* (which defaults to `nil`). *size* should be a non-negative integer.

For example:

```
(make-list 5) => (nil nil nil nil nil)
(make-list 3 'rah) => (rah rah rah)
```

Compatibility note: The Lisp Machine LISP function `make-list` takes arguments *area* and *size*. Areas are not relevant to COMMON LISP. The argument order used here is compatible with NIL.

`append &rest` *lists* [*Function*]

The arguments to `append` are lists. The result is a list which is the concatenation of the arguments. The arguments are not destroyed.

For example:

```
(append '(a b c) '(d e f) '() '(g)) => (a b c d e f g)
```

Note that `append` copies the top-level list structure of each of its arguments *except* the last. The function `catenate` (page 158) can perform a similar operation, but always copies all its arguments. See also `nconc` (page 171), which is like `append` but destroys all arguments but the last.

`(append` *x* `'())` is an idiom once frequently used to copy the list *x*, but the `copylist`

function is more appropriate to this task.

`copylist` *list* [*Function*]

Returns a list which is `equal` to *list*, but not `eq`. Only the top level of list-structure is copied; that is, `copylist` copies in the *cdr* direction but not in the *car* direction. If the list is "dotted", that is, `(cdr (last` *list*`))` is a non-`nil` atom, this will be true of the returned list also. See also `copyseq` (page 157).

`copyalist` *list* [*Function*]

`copyalist` is for copying association lists. The top level of list structure of *list* is copied, just as `copylist` does. In addition, each element of *list* which is a cons is replaced in the copy by a new cons with the same car and cdr.

`copytree` *object* [*Function*]

`copytree` is for copying trees of conses. The argument *object* may be any LISP object. If it is not a cons, it is returned; otherwise the result is a new cons of the results of calling `copytree` on the car and cdr of the argument. In other words, all conses in the tree are copied recursively, stopping only when non-conses are encountered. Circularities and the sharing of substructure are *not* preserved.

`revappend` *x y* [*Function*]

`(revappend` *x y*`)` is exactly the same as `(append (reverse` *x*`)` *y*`)` except that it is more efficient. Both *x* and *y* should be lists. The argument *x* is copied, not destroyed. Compare this with `nreconc` (page 172), which destroys its first argument.

`nconc &rest` *lists* [*Function*]

`nconc` takes lists as arguments. It returns a list which is the arguments concatenated together. The arguments are changed, rather than copied. (Compare this with `append` (page 170), which copies arguments rather than destroying them.)

For example:

```
(setq x '(a b c))
(setq y '(d e f))
(nconc x y) => (a b c d e f)
x => (a b c d e f)
```

Note, in the example, that the value of `x` is now different, since its last cons has been `rplacd`'d to the value of `y`. If one were then to evaluate `(nconc x y)` again, it would yield a piece of "circular" list structure, whose printed representation would be `(a b c d e f d e f d e f . . .)`, repeating forever.

`nreconc` *x y* [*Function*]

(`nreconc` *x y*) is exactly the same as (`nconc` (`nreverse` *x*) *y*) except that it is more efficient. Both *x* and *y* should be lists. The argument *x* is destroyed. Compare this with `revappend` (page 171).

`push` *item place* [*Macro*]

The form *place* should be the name of a generalized variable containing a list; *item* may refer to any LISP object. The *item* is consed onto the front of the list, and the augmented list is stored back into *place* and returned. The form *place* may be any form acceptable as a generalized variable to `setf` (page 60). If the list held in *place* is viewed as a push-down stack, then `push` pushes an element onto the top of the stack.

For example:

```
(setq x '(a (b c) d))
(push 5 (cadr x)) => (5 b c)  and now x => (a (5 b c) d)
```

The effect of (`push` *item place*) is roughly equivalent to

```
(setf place (cons item place))
```

except that the latter would evaluate any subforms of *place* twice, while `push` takes care to evaluate them only once. Moreover, for certain *place* forms `push` may be significantly more efficient than the `setf` version.

`pushnew` *item place* [*Macro*]

The form *place* should be the name of a generalized variable containing a list; *item* may refer to any LISP object. If the *item* is already a member of the list (as determined by `eql` comparisons), then the *item* is consed onto the front of the list, and the augmented list is stored back into *place* and returned; otherwise `nil` is returned. (Thus a `pushnew` form returns a truth value saying whether *item* was new to the list or not.) The form *place* may be any form acceptable as a generalized variable to `setf` (page 60). If the list held in *place* is viewed as a set, then `pushnew` adjoins an element to the set; see `adjoin` (page 177).

For example:

```
(setq x '(a (b c) d))
(pushnew 5 (cadr x)) => (5 b c)  and now x => (a (5 b c) d)
(pushnew 'b (cadr x)) => nil  and x is unchanged
```

The effect of (`pushnew` *item place*) is roughly equivalent to

```
(and (not (member item place))
     (setf place (cons item place)))
```

except that the latter would evaluate *item* twice and any subforms of *place* thrice, while `pushnew` takes care to evaluate them only once each. Moreover, for certain *place* forms `pushnew` may be significantly more efficient than the `setf` version.

??? **Query:** The other way to define `pushnew` is as

```
(setf place (adjoin item place))
```

but that doesn't act as a useful pseudo-predicate. However, it may compile into shorter code. What do people

think?

`pop` *place* [*Macro*]

The form *place* should be the name of a generalized variable containing a list. The result of `pop` is the `car` of the contents of *place*, and as a side-effect the `cdr` of the contents is stored back into *place*. The form *place* may be any form acceptable as a generalized variable to `setf` (page 60). If the list held in *place* is viewed as a push-down stack, then `pop` pops an element from the top of the stack and returns it.

For example:

```
(setq stack '(a b c))
(pop stack) => a  and now stack => (b c).
```

The effect of `(pop` *place*`)` is roughly equivalent to

```
(prog1 (car place) (setf place (cdr place)))
```

except that the latter would evaluate any subforms of *place* thrice, while `pop` takes care to evaluate them only once. Moreover, for certain *place* forms `pop` may be significantly more efficient than the `setf` version.

`butlast` *list* `&optional` *n* [*Function*]

This creates and returns a list with the same elements as *list*, excepting the last *n* elements. *n* defaults to 1. The argument is not destroyed. If the *list* has fewer than *n* elements, then `()` is returned.

For example:

```
(butlast '(a b c d)) => (a b c)
(butlast '((a b) (c d)) => ((a b))
(butlast '(a)) => ()
(butlast nil) => ()
```

The name is from the phrase "all elements but the last".

`nbutlast` *list* `&optional` *n* [*Function*]

This is the destructive version of `butlast`; it changes the *cdr* of the cons $n+1$ from the end of the *list* to `nil`. *n* defaults to 1. If the *list* has fewer than *n* elements, then `nbutlast` returns `()`, and the argument is not modified. (Therefore one normally writes `(setq a (nbutlast a))` rather than simply `(nbutlast a)`.)

For example:

```
(setq foo '(a b c d))
(nbutlast foo) => (a b c)
foo => (a b c)
(nbutlast '(a)) => ()
(nbutlast 'nil) => ()
```

`buttail` *list sublist* [*Function*]

list should be a list, and *sublist* should be a sublist of *list*, i.e., one of the conses that make up *list*. `buttail` (meaning "all but the tail") will return a new list, whose elements are those elements of *list* that appear before *sublist*. If *sublist* is not a tail of *list*, then a copy of *list* is returned. The argument *list* is not destroyed.

For example:

```
(setq x '(a b c d e))
(setq y (cdddr x)) => (d e)
(buttail x y) => (a b c)
```
but
```
(buttail '(a b c d) '(c d)) => (a b c d)
```
since the sublist was not `eq` to any part of the list.

??? **Query:** I realize we voted to change the name from `ldiff` to `buttail`, but it seems senseless to be different from existing INTERLISP and Lisp Machine LISP usage. Can we reconsider?

15.3. Alteration of List Structure

The functions `rplaca` and `rplacd` are used to make alterations in already-existing list structure; that is, to change the cars and cdrs of existing conses.

The structure is not copied but is physically altered; hence caution should be exercised when using these functions, as strange side-effects can occur if portions of list structure become shared unbeknownst to the programmer. The `nconc` (page 171), `nreverse` (page 158), `nreconc` (page 172), and `nbutlast` (page 173) functions already described, and the `delete` (page 161) family described later, have the same property. However, they are normally not used for this side-effect; rather, the list-structure modification is purely for efficiency and compatible non-modifying functions are provided.

`rplaca` *x y* [*Function*]

`(rplaca` *x y*`)` changes the car of *x* to *y* and returns (the modified) *x*. *x* should be a cons, but *y* may be any Lisp object.

For example:

```
(setq g '(a b c))
(rplaca (cdr g) 'd) => (d c)
Now g => (a d c)
```

`rplacd` *x y* [*Function*]

`(rplacd` *x y*`)` changes the cdr of *x* to *y* and returns (the modified) *x*. *x* should be a cons, but *y* may be any Lisp object.

For example:

```
(setq x '(a b c))
(rplacd x 'd) => (a . d)
Now x => (a . d)
```

Compatibility note: In COMMON LISP, as in MACLISP and Lisp Machine LISP, `rplacd` can not be used to set

the property list of a symbol. The `setplist` (page SETPLIST-FUN) function is provided for this purpose.

`setnth` *n list newvalue* [*Function*]

Alters the *n*'th element of *list* to be *newvalue*, where the zeroth element is the *car* of the *list*. *n* must be a non-negative number less than the length of the list. *setnth* returns *newvalue*. See `nth` (page 169).

15.4. Substitution of Expressions

A number of functions are provided for performing substitutions within a tree. All take a tree and a description of old sub-expressions to be replaced by new ones. The functions form a semi-regular collection, according to these properties:

- Whether comparison of items is by `eq` or `equal`.
- Whether substitution is specified by two arguments or by an association list.
- Whether the tree is copied or modified.

These properties may be summarized as follows:

	Accepts two arguments, old and new		Accepts an association list
	Uses `equal`	Uses `eq`	Uses `eq`
Copies	`subst`	`substq`	`sublis`
Modifies	`nsubst`	`nsubstq`	`nsublis`

`subst` *new old tree* [*Function*]

(`subst` *new old tree*) substitutes *new* for all occurrences of *old* in *tree*, and returns the modified copy of *tree*. The original *tree* is unchanged, as `subst` recursively copies all of *tree* replacing elements `equal` to *old* as it goes.

For example:

```
(subst 'Tempest 'Hurricane
       '(Shakespeare wrote (The Hurricane)))
   => (Shakespeare wrote (The Tempest))
```

This function is not "destructive"; that is, it does not change the *car* or *cdr* of any already-existing list structure.

(`subst nil nil x`) is an idiom once frequently used to copy all the conses in a tree, but the `copytree` (page 171) function is more appropriate to the task.

`nsubst` *new old tree* [*Function*]

`nsubst` is a destructive version of `subst`. The list structure of *tree* is altered by replacing each occurrence of *old* with *new*. `equal` is used to decide whether a part of *tree* is the same as *old*.

`substq` *new old tree* [*Function*]

`substq` is just like `subst`, except that `eq`, rather than `equal`, is used to decide whether a part of *tree* is the same as *old*.

`nsubstq` *new old tree* [*Function*]

`nsubstq` is a destructive version of `substq`. `nsubstq` is just like `nsubst`, except that `eq`, rather than `equal`, is used to decide whether a part of *tree* is the same as *old*.

`sublis` *alist tree* [*Function*]

`sublis` makes substitutions for symbols in a tree (a structure of conses). The first argument to `sublis` is an association list. The *car* of each a-list entry should be a symbol. The second argument is the tree in which substitutions are to be made. `sublis` looks at all symbols in the tree; if a symbol appears as a key in the association list occurrences of it are replaced by the object it is associated with. The argument is not modified; new conses are created where necessary and only where necessary, so the newly created structure shares as much of its substructure as possible with the old. For example, if no substitutions are made, the result is `eq` to the old tree.

For example:

```
(sublis '((x . 100) (z . zprime))
        '(plus x (minus g z x p) 4))
  => (plus 100 (minus g zprime 100 p) 4)
```

`nsublis` *alist tree* [*Function*]

`nsublis` is like `sublis` but changes the original list structure instead of copying.

15.5. Using Lists as Sets

COMMON LISP includes functions which allow a list of items to be treated as a *set*. Some of the functions usefully allow the set to be ordered; others specifically support unordered sets. There are functions to add, remove, and search for items in a list, based on various criteria. There are also set union, intersection, and difference functions.

The naming conventions for these functions and for their keyword arguments generally follow the conventions for the generic sequence functions. See Chapter 14.

`member` *item list* `&key :test :test-not :key` [*Function*]
`member-if` *predicate list* `&key :key` [*Function*]
`member-if-not` *predicate list*, `Keys = {[key]` [*Function*]

(`member` *item list*) returns `nil` if *item* is not `eql` to any element in the *list*. Otherwise, it returns the tail of *list* beginning with the first occurrence of *item*. *list* is searched on the top level only. Because `member` returns `nil` if it doesn't find anything, and something non-`nil` if it finds something, it is often used as a predicate.

For example:

```
(member 'snerd '(a b c d)) => nil
(member 'a '(g (a y) c a d e a f)) => (a d e a f)
```

Note that the value returned by `member` is `eq` to the portion of the list beginning with *a*. Thus `rplaca` on the result of `member` may be used, if you first check to make sure `member` did not return `nil`, to alter the found list element.

`mem-if` is like `member`, except that *predicate*, a function of one argument, is used to test elements of *list*.

`mem-if-not` is like `mem-if`, except that the sense of *predicate* is inverted; that is, a test succeeds if `predicate` returns `nil`.

See also `find` (page 163) and `position` (page 163).

`tailp` *sublist list* [*Function*]

This predicate is true if *sublist* is a sublist of *list* (i.e. one of the conses that makes up *list*). Otherwise it is false. Another way to look at this is that `tailp` is true if `(nthcdr` *n list*`)` is *sublist*, for some value of *n*. See `buttail` (page 174).

`adjoin` *item list* `&key :test :test-not` [*Function*]

`adjoin` is used to add an element to a set, provided that it is not already a member. The equality test defaults to `eql`.

`(adjoin` *item list*`)` <=> `(if (member` *item list*`)` *list* `(cons` *item list*`))`

See `pushnew` (page 172).

??? **Query:** To make the tests consistent with the keyword proposal, I had to make `union` and `intersection` take only two list, not *n*. Is this acceptable?

`union` *list1 list2* `&key :test :test-not` [*Function*]

`nunion` *list1 list2* `&key :test :test-not` [*Function*]

`union` takes two lists and returns a new list containing everything that is an element of either of the *lists*. If there is a duplication between two lists, only one of the duplicate instances will be in the result. If either of the arguments has duplicate entries within it, the redundant entries may or may not appear in the result.

For example:

```
(union '(a b c) '(f a d)) => (a b c f d)
```

There is no guarantee that the order of elements in the result will reflect the ordering of the arguments in any particular way. The implementation is therefore free to use any of a variety of strategies.

`nunion` is the destructive version of `union`. It performs the same operation, but may destroy the argument lists, using their cells to construct the result.

`intersection` *list1 list2* `&key :test :test-not` [*Function*]
`nintersection` *list1 list2* `&key :test :test-not` [*Function*]

`intersection` takes two lists and returns a new list containing everything that is an element of both argument lists. If either list has duplicate entries, the redundant entries may or may not appear in the result.

For example:

```
(intersection '(a b c) '(f a d)) => (a)
```

There is no guarantee that the order of elements in the result will reflect the ordering of the arguments in any particular way. The implementation is therefore free to use any of a variety of strategies.

`nintersection` is the destructive version of `intersection`. It performs the same operation, but may destroy *list1* using its cells to construct the result. (The argument *list2* is *not* destroyed.)

`setdifference` *list1 list2* `&key :test :test-not` [*Function*]
`nsetdifference` *list1 list2* `&key :test :test-not` [*Function*]

`setdifference` returns a list of elements of *list1* which do not appear in *list2*. This operation is not destructive.

`nsetdifference` is the destructive version of `setdifference`. This operation may destroy *list1*.

`set-exclusive-or` *list1 list2* `&key :test :test-not` [*Function*]
`nset-exclusive-or` *list1 list2* `&key :test :test-not` [*Function*]

`set-exclusive-or` returns a list of elements which appear in exactly one of *list1* and *list2*. This operation is not destructive.

`nset-exclusive-or` is the destructive version of `set-exclusive-or`. Both lists may be destroyed in producing the result.

`subsetp` *list1 list2* `&key :test :test-not` [*Function*]

`subsetp` is a predicate that is true iff every element of *list1* appears in *list2*.

15.6. Association Lists

An *association list*, or *a-list*, is a data structure used very frequently in LISP. An a-list is a list of pairs (conses); each pair is an association. The *car* of a pair is called the *key*, and the *cdr* is called the *datum*.

An advantage of the a-list representation is that an a-list can be incrementally augmented simply by adding new entries to the front. Moreover, because the searching function `assoc` (page 179) searches the a-list in order, new entries can "shadow" old entries. If an a-list is viewed as a mapping from keys to data, then the mapping can be not only augmented but also altered in a non-destructive manner by adding new entries to

the front of the a-list.

Sometimes an a-list represents a bijective mapping, and it is desirable to retrieve a key given a datum. For this purpose the "reverse" searching function `rassoc` (page 180) is provided. Other variants of a-list searches can be constructed using the function `find` (page 163) or `member` (page 176).

It is permissible to let `nil` be an element of an a-list in place of a pair.

`acons` *key datum a-list* [*Function*]

`acons` constructs a new association list by adding the pair (*key* . *datum*) to the old *a-list*.

```
(acons x y a) <=> (cons (cons x y) a)
```

`pairlis` *keys data* `&optional` *a-list* [*Function*]

`pairlis` takes two lists and makes an association list which associates elements of the first list to corresponding elements of the second list. It is an error if the two lists *keys* and *data* are not of the same length. If the optional argument *a-list* is provided, then the new pairs are added to the front of it.

For example:

```
(pairlis '(beef clams kitty) '(roast fried yu-shiang))
   => ((beef . roast) (clams . fried) (kitty . yu-shiang))
(pairlis '(one two) '(1 2) '((three . 3) (four . 19)))
   => ((one . 1) (two . 2) (three . 3) (four . 19))
```

`assoc` *item a-list* `&key :test :test-not` [*Function*]

(`assoc` *item alist*) looks up *item* in the association list *a-list*. The value is the first pair in the a-list such that *item* and the *car* of the pair satisfy the test, or `nil` if there is none such. (The test defaults to `eql`.)

For example:

```
(assoc 'r '((a . b) (c . d) (r . x) (s . y) (r . z)))
        => (r . x)
(assoc 'goo '((foo . bar) (zoo . goo))) => nil
(assoc '2 '((1 a b c) (2 b c d) (-7 x y z))) => (2 b c d)
```

It is possible to `rplacd` the result of `assoc` *provided* that it is not `nil`, if your intention is to "update" the "table" that was `assoc`'s second argument. (However, it is often better to update an a-list by adding new pairs to the front, rather than altering old pairs.)

For example:

```
(setq values '((x . 100) (y . 200) (z . 50)))
(assoc 'y values) => (y . 200)
(rplacd (assoc 'y values) 201)
(assoc 'y values) => (y . 201) now
```

A typical trick is to say (`cdr (assoc x y)`). Because the *cdr* of `nil` is guaranteed to be `nil`, this yields `nil` if no pair is found *or* if a pair is found whose *cdr* is `nil`. This is useful if `nil` serves

its usual role as a "default value".

Compatibility note: This is of course not compatible with MACLISP, which uses `equal`, not `eql`, as the default comparison test.

```
(assoc item list :test fn)
   <=> (find item list :test fn :key #'car)
```

`rassoc` *item a-list* `&key :test :test-not` [*Function*]

`rassoc` is the reverse form of `assoc`; it compares *item* to the *cdr* of each successive pair in *a-list*, rather than to the *car*.

For example:

```
(rassoc 'a '((a . b) (b . c) (c . a) (z . a))) => (c . a)
```

```
(rassoc item list :test fn)
   <=> (find item list :test fn :key #'cdr)
```

15.7. Hash Tables

A hash table is a LISP object that works something like a property list and something like an association list. Each hash table has a set of *entries*, each of which associates a particular *key* with a particular *value*. The basic functions that deal with hash tables can create entries, delete entries, and find the value that is associated with a given key. Finding the value is very fast even if there are many entries, because hashing is used; this is an important advantage of hash tables over property lists.

A given hash table can only associate one *value* with a given *key*; if you try to add a second *value* it will replace the first. Also, adding a value to a hash table is a destructive operation; the hash table is modified. By contrast, association lists can be augmented non-destructively.

Hash tables come in three kinds, the difference being whether the keys are compared with `eq`, `eql`, or `equal`. In other words, there are hash tables that hash on Lisp *objects* (using `eq` or `eql`) and there are hash tables which hash on abstract *S-expressions* (using `equal`).

Hash tables of the first kind are created with the function `make-hash-table`, which takes various options. New entries are added to hash tables with the `puthash` function. To look up a key and find the associated value, use `gethash`; to remove an entry, use `remhash`. Here is a simple example.

```
(setq a (make-hash-table))
(puthash 'color 'brown a)
(puthash 'name 'fred a)
(gethash 'color a) => brown
(gethash 'name a) => fred
(gethash 'pointy a) => nil
```

In this example, the symbols `color` and `name` are being used as keys, and the symbols `brown` and `fred` are being used as the associated values. The hash table has two items in it, one of which associates from `color` to `brown`, and the other of which associates from `name` to `fred`.

Keys do not have to be symbols; they can be any LISP object. Likewise values can be any LISP object. Hash tables are properly interfaced to the relocating garbage collector so that garbage collection will have no perceptible effect on the functionality of hash tables.

When a hash table is first created, it has a *size*, which is the maximum number of entries it can hold. Usually the actual capacity of the table is somewhat less, since the hashing is not perfectly collision-free. With the maximum possible bad luck, the capacity could be very much less, but this rarely happens. If so many entries are added that the capacity is exceeded, the hash table will automatically grow, and the entries will be *rehashed* (new hash values will be recomputed, and everything will be rearranged so that the fast hash lookup still works). This is transparent to the caller; it all happens automatically.

Compatibility note: This hash table facility is compatible with Lisp Machine LISP. It is similar to the hasharray facility of INTERLISP, and some of the function names are the same. However, it is *not* compatible with INTERLISP. The exact details and the order of arguments are designed to be consistent with the rest of MACLISP rather than with INTERLISP. For instance, the order of arguments to `maphash` is different, there is no "system hash table", and there is not the INTERLISP restriction that keys and values may not be `nil`. Note, however, that the order of arguments to `gethash`, `puthash`, and `remhash` is not consistent with `get`, `putprop`, and `remprop`, either. This is an unfortunate result of the haphazard historical development of Lisp.

15.7.1. Hash Table Functions

This section documents the functions for hash tables, which use *objects* as keys and associate other objects with them.

`make-eq-hash-table &key :size :rehash-size :rehash-threshold` [*Function*]
`make-eql-hash-table &key :size :rehash-size :rehash-threshold` [*Function*]
`make-equal-hash-table &key :size :rehash-size :rehash-threshold` [*Function*]

Calling any of these creates a new hash table; depending on which one is used, the resulting table treats keys as equal if they are `eq`, `eql`, or `equal`, respectively.

The `:size` argument sets the initial size of the hash table, in entries, as a fixnum. The default is 64. (The actual size may be rounded up from the size you specify to the next "good" size, for example to make it a prime number.) You won't necessarily be able to store this many entries into the table before it overflows and becomes bigger; but except in the case of extreme bad luck you will be able to store almost this many.

The `:rehash-size` argument specifies how much to increase the size of the hash table when it becomes full. This can be an integer greater than zero, which is the number of entries to add, or it can be a floating-point number greater than one, which is the ratio of the new size to the old size. The default is `1.3`, which causes the table to be made 30% bigger each time it has to grow.

The `:rehash-threshold` argument specifies how full the hash table can get before it must grow. This can be an integer greater than zero and less than the rehash-size (in which case it will be scaled whenever the table is grown), or it can be a floating-point number between zero and one. The default is `0.8`, which means the table is enlarged when it becomes over 80% full.

For example:

```
(make-hash-table :rehash-size 1.5
                 :size (* number-of-widgets 43))
```

`gethash` *key hash-table* `&optional` *default* [*Function*]

Find the entry in *hash-table* whose key is *key*, and return the associated value. If there is no such entry, return *default*, which is `nil` if not specified.

`gethash` actually returns two values, the second being a predicate value that is true if an entry was found, and false if no entry was found.

`puthash` *key value hash-table* [*Function*]

Create an entry in *hash-table* associating *key* to *value*; if there is already an entry for *key*, then replace the value of that entry with *value*. Returns *value*.

??? **Query:** Should *value* be the last argument? Wouldn't be compatible with Lisp Machine LISP.

`remhash` *key hash-table* [*Function*]

Remove any entry for *key* in *hash-table*. This is a predicate that is true if there was an entry or false if there was not.

`maphash` *function hash-table* [*Function*]

For each entry in *hash-table*, call *function* on two arguments: the key of the entry and the value of the entry. If entries are added to or deleted from the hash table while a `maphash` is in progress, the results are unpredictable. `maphash` returns `nil`.

`clrhash` *hash-table* [*Function*]

Remove all the entries from *hash-table*. Returns the hash table itself.

15.7.2. Primitive Hash Function

`sxhash` *S-expression* [*Function*]

`sxhash` computes a hash code of an S-expression, and returns it as a non-negative fixnum. A property of `sxhash` is that (`equal` *x y*) implies (`=` (`sxhash` *x*) (`sxhash` *y*)).

The manner in which the hash code is computed is implementation-dependent, but is independent of the particular "incarnation" or "core image". Hash values may be written out to files, for example, and read in again into an instance of the same implementation.

Chapter 16

Arrays

16.1. Array Creation

`make-array` *dimensions* `&key :type :initial-value :initial-contents` [*Function*]
`:fill-pointer :displaced-to :displaced-index-offset`

This is the primitive function for making arrays. *dimensions* should be a list of non-negative integers (in fact, fixnums) that are to be the dimensions of the array; the length of the list will be the dimensionality of the array. For convenience when making a one-dimensional array, the single dimension may be provided as an integer rather than a list of one integer.

The `:type` argument should be the name of the type of the elements of the array; an array is constructed of the most specialized type which can nevertheless accommodate elemments of the given type. The type `t` specifies a general array, one whose elements may be any LISP object; this is the default type.

The `:initial-value` argument may be used to initialize each element of the array. The value must be of the type specified by the `:type` option. If the `:initial-value` option is omitted, the initial values of the array elements are undefined (unless the `:initial-contents` or `:displaced-to` option is used). The `:initial-value` option may not be used with the `:initial-contents` or `:displaced-to` option.

The `:initial-contents` argument may be used to initialize the contents of the array. The value is a nested structure of sequences. If the array is zero-dimensional, then the value specifies the single element. Otherwise, the value must be a sequence whose length is equal to the first dimension; each element must be a nested structure for an array whose dimensions are the remaining dimensions, and so on.

For example:

```
(make-array '(4 2 3) :initial-contents
            '(((a b c) (* * *))
              ((d e f) (* * *))
              ((g h i) (* * *))
              ((j k l) (* * *))))
```

The numbers of levels in the structure must equal the rank of the array. Each leaf of the nested structure must be of the type specified by the `:type` option. If the `:initial-contents` option

is omitted, the initial values of the array elements are undefined (unless the `:initial-value` or `:displaced-to` option is used). The `:initial-contents` option may not be used with the `:initial-value` or `:displaced-to` option.

The `:fill-pointer` argument specifies that the array should have a fill pointer. If this option is specified, the array must be one-dimensional. The value is used to initialize the fill pointer for the array. if the value `nil` is specified, the length of the array is used; otherwise the value must be an integer between 0 (inclusive) and the length of the array (inclusive).

The `:displaced-to` argument, if not `nil`, specifies that the array will be a *displaced* array. The argument must then be an array or vector; `make-array` will create an *indirect* or *shared* array which shares its contents with the specified array. In this case the `:displaced-index-offset` option may be useful. The `:displaced-to` option may not be used with the `:initial-value` or `:initial-contents` option.

??? **Query:** A long, extended discussion of displaced arrays is clearly needed here.

The `:displaced-index-offset` argument may be used only in conjunction with the `displaced-to` option. This argument should be a non-negative fixnum (it defaults to zero); it is made to be the index-offset of the created shared array.

For example:

```
;; Create a one-dimensional array of five elements.
(make-array 5)

;; Create a two-dimensional array, 3 by 4, with four-bit elements.
(make-array '(3 4) ':type '(mod 16))

;; Create an array of single-floats.
(make-array 5 ':type ':single-float))

;; Making a shared array.
(setq a (make-array '(4 3)))
(setq b (make-array 8 ':displaced-to a
                      ':displaced-index-offset 2))
;;Now it is the case that:
        (aref b 0) <=> (aref a 0 2)
        (aref b 1) <=> (aref a 1 0)
        (aref b 2) <=> (aref a 1 1)
        (aref b 3) <=> (aref a 1 2)
        (aref b 4) <=> (aref a 2 0)
        (aref b 5) <=> (aref a 2 1)
        (aref b 6) <=> (aref a 2 2)
        (aref b 7) <=> (aref a 3 0)
```

The last example depends on the fact that arrays are, in effect, stored in row-major order for purposes of sharing. Put another way, the sequences of indices for the elements of an array are ordered lexicographically.

Compatibility note: Both Lisp Machine LISP and FORTRAN store arrays in column-major order.

`make-vector` *length* `&key :type :initial-value :initial-contents :fill-pointer` [*Function*]

`make-vector` is like `make-array` (page 183), but guarantees to return a vector. Depending on the implementation, use of a vector (and declaration of such use to the compiler) may result in significantly more efficient code. One may not specify a list of dimensions, but only a single integer, the length. The `:type`, `:initial-value`, `:initial-contents`, and `:fill-pointer` keyword arguments are as for `make-array`.

16.2. Array Access

`aref` *array* `&rest` *subscripts* [*Function*]

This accesses and returns the element of *array* specified by the *subscripts*. The number of subscripts must equal the rank of the array, and each subscript must be a non-negative integer less than the corresponding array dimension.

`aset` *new-value array* `&rest` *subscripts* [*Function*]

This stores *new-value* into the element of *array* specified by the *subscripts*. The number of subscripts must equal the rank of the array, and each subscript must be a non-negative integer less than the corresponding array dimension. The result of `aset` is the value *new-value*.

The argument *new-value* must be of a type suitable for storing into *array* if the *array* is of a specialized type.

??? **Query:** The more I think about it, the more atractive seems the suggestion from RMS simply to flush all these updator functions and use `setf`.

16.3. Array Information

`array-type` *array* [*Function*]

This returns the type of elements of the array. For a general array, this is `t`; for an array of eight-bit integers, `(mod 256)` might be returned. What is returned is the actual type of the array elements, which may be the same as that specified to `make-array`, or may be more general if the implementatation doesn't support arrays of that specific type.

`array-allocated-length` *array* [*Function*]

array may be any array. This returns the total number of elements allocated in *array*. For a one-dimensional array, this is equal to the length of the single axis. (If a fill pointer is in use for the array, however, the function `array-active-length` (page 186) may be more useful.)

`array-active-length` *array* [*Function*]

`array-active-length` returns the fill pointer for the array. This is normally the same as the length of the array unless `array-reset-fill-pointer` (page 189) has been used.

`array-rank` *array* [*Function*]

Returns the number of dimensions (axes) of *array*. This will be a non-negative integer.

Compatibility note: In Lisp Machine LISP this is called `array-#-dims`. This name causes problems in MACLISP because of the # character. The problem is better avoided.

`array-dimension` *axis-number array* [*Function*]

The length of dimension number *axis-number* of the *array* is returned. *array* may be any kind of array, and *axis-number* should be a non-negative integer less than the rank of *array*.

Compatibility note: This is similar to the Lisp Machine LISP function `array-dimension-n`, but is zero-origin for consistency instead of one-origin. Also, in Lisp Machine LISP `(array-dimension-n 0)` returns the length of the array leader.

`array-dimensions` *array* [*Function*]

`array-dimensions` returns a list whose elements are the dimensions of *array*.

`array-in-bounds-p` *array* `&rest` *subscripts* [*Function*]

This predicate checks whether the *subscripts* are all legal subscripts for *array*, and is true if they are; otherwise it is false. The *subscripts* must be integers.

16.4. Functions on Vectors

The functions in this section are equivalent in operation to the corresponding more general functions, but require arguments to be vectors (of general or specialized type). These functions are provided primarily for reasons of efficiency and convenience.

`velt` *vector index* [*Function*]

The element of the *vector* specified by the integer *index* is returned. The *index* must be non-negative and less than the length of the vector. See `elt` (page 157), `aref` (page 185), and `vref` (page 187).

`vsetelt` *vector index newvalue* [*Function*]

The LISP object *newvalue* is stored into the component of the *vector* specified by the integer *index*. The *index* must be non-negative and less than the length of the vector. See `setelt` (page 157), `aset` (page 185), and `vset` (page 187).

16.5. Functions on General Vectors (Vectors of LISP Objects)

The functions in this section are equivalent in operation to the corresponding more general functions, but require arguments to be vectors of type `(vector t)`. These functions are provided primarily for reasons of efficiency and convenience.

`vref` *vector index* [*Function*]

The element of the *vector* specified by the integer *index* is returned. The *index* must be non-negative and less than the length of the vector. See `elt` (page 157), `aref` (page 185), and `velt` (page 186).

`vset` *vector index newvalue* [*Function*]

The LISP object *newvalue* is stored into the component of the *vector* specified by the integer *index*. The *index* must be non-negative and less than the length of the vector. See `setelt` (page 157), `aset` (page 185), and `vsetelt` (page 186).

16.6. Functions on Bit-vectors

`bit` *bit-vector index* [*Function*]

The element of the *bit-vector* specified by the integer *index* is returned. The *index* must be non-negative and less than the length of the vector. The result will always be `0` or `1`. See `elt` (page 157).

`rplacbit` *bit-vector index newbit* [*Function*]

The *newbit* is stored into the component of the *bit-vector* specified by the integer *index*. The *index* must be non-negative and less than the length of the vector. The *newvalue* must be `0` or `1`. See `setelt` (page 157).

`bit-and &rest` *bit-vectors* [*Function*]
`bit-ior &rest` *bit-vectors* [*Function*]
`bit-xor &rest` *bit-vectors* [*Function*]
`bit-eqv &rest` *bit-vectors* [*Function*]
`bit-nand` *bit-vector1 bit-vector2* [*Function*]
`bit-nor` *bit-vector1 bit-vector2* [*Function*]
`bit-andc1` *bit-vector1 bit-vector2* [*Function*]
`bit-andc2` *bit-vector1 bit-vector2* [*Function*]
`bit-orc1` *bit-vector1 bit-vector2* [*Function*]
`bit-orc2` *bit-vector1 bit-vector2* [*Function*]

These functions perform bit-wise logical operations on bit-vectors. All of the arguments to any of these functions must be bit-vectors or one-dimensional arrays of bits, all of the same length. The

result is a bit-vector matching the argument(s) in length, such that bit *j* of the result is produced by operating on bit *j* of each of the arguments. Indeed, if the arguments are in fact bit-vectors of the same length, then

`(bit-`*xxx* . *arguments*`) <=> (map 'bit-vector #'log`*xxx* . *arguments*`)`

That is, each `bit-` function described here is simply a mapping over bit-vectors of a `log` function which applies to integers (and therefore to the bit values 0 and 1). See `logand` (page 135) and friends.

The following table indicates what the result bit is for each operation when two arguments are given. (Those operations which accept an indefinite number of arguments are commutative and associative, and require at least one argument.)

argument1	0	0	1	1	
argument2	0	1	0	1	*Operation name*
`bit-and`	0	0	0	1	and
`bit-ior`	0	1	1	1	inclusive or
`bit-xor`	0	1	1	0	exclusive or
`bit-eqv`	1	0	0	1	equivalence (exclusive nor)
`bit-nand`	1	1	1	0	not-and
`bit-nor`	1	0	0	0	not-or
`bit-andc1`	0	1	0	0	and complement of *argument1* with *argument2*
`bit-andc2`	0	0	1	0	and *argument1* with complement of *argument2*
`bit-orc1`	1	1	0	1	or complement of *argument1* with *argument2*
`bit-orc2`	1	0	1	1	or *argument1* with complement of *argument2*

`bit-not` *bit-vector* [*Function*]

The argument must be a one-dimensional array of bits. A bit-vector containing a copy of the argument with all the bits inverted is returned. That is, bit *j* of the result is 1 iff bit *j* of the argument is zero.

`(bit-not` *bitvec*`) <=> (map 'bit-vector #'lognot` *bitvec*`)`

See `lognot` (page 137).

16.7. Fill Pointers

To make it easy to incrementally fill in the contents of an array, a set of functions for manipulating a *fill pointer* are defined. The fill pointer is a non-negative integer no larger than the total number of elements in the array (as returned by `array-length` (page ARRAY-LENGTH-FUN)); it is the number of "active" or "filled-in" elements in the array. When an array is created, its fill pointer is initialized to the number of elements in the array; the fill pointer should be *reset* before use. The fill pointer constitutes the "active length" of the array. Some functions will ignore elements beyond the fill-pointer index; those that do are so documented.

Multidimensional arrays may have fill pointers; elements are filled in row-major order (last index varies

fastest).

`array-reset-fill-pointer` *array* `&optional` *index* [*Function*]

The fill pointer of *array* is reset to *index*, which defaults to zero. The *index* must be a non-negative integer not greater than the old value of the fill pointer.

`array-push` *array new-element* [*Function*]

array must be a one-dimensional array that has a fill pointer, and *new-element* may be any object. `array-push` attempts to store *new-element* in the element of the array designated by the fill pointer, and increase the fill pointer by one. If the fill pointer does not designate an element of the array (specifically, when it gets too big), it is unaffected and `array-push` returns `nil`. Otherwise, the store and increment take place and `array-push` returns the *former* value of the fill pointer (one less than the one it leaves in the array); thus the value of `array-push` is the index of the new element pushed.

`array-push-extend` *array x* `&optional` *extension* [*Function*]

`array-push-extend` is just like `array-push` except that if the fill pointer gets too large, the array is extended (using `adjust-array-size` (page 189)) so that it can contain more elements; it never "fails" the way `array-push` does, and so never returns `nil`. The optional argument *extension*, which must be a positive integer, is the minimum number of elements to be added to the array if it must be extended.

`array-pop` *array* [*Function*]

array must be a one-dimensional array that has a fill pointer. The fill pointer is decreased by one, and the array element designated by the new value of the fill pointer is returned. If the new value does not designate any element of the array (specifically, if it has reached zero), an error occurs.

16.8. Changing the Size of an Array

`adjust-array-size` *array new-size* `&optional` *new-element* [*Function*]

The *array* is adjusted so that it contains (at least) *new-size* elements. The argument *new-size* must be a non-negative integer.

If *array* is a one-dimensional array, its size is simply changed to be *new-size*, by altering its single dimension. If *array* has more than one dimension, then its *first* dimension is adjusted to the smallest possible value which allows the array to have no fewer than *new-size* elements. There are two degenerate cases, however:

1. If any dimension other than the first is zero, then the array is not changed, and an error occurs if *new-size* is not `0`.

2. If the array has zero dimensions, then the array is not changed, and an error occurs if

new-size is not 0 or 1.

If *array* is made smaller, the extra elements are lost. If *array* is made bigger, the new elements are initialized to *new-element*; if this argument is not provided, then the values of the new elements are undefined.

`adjust-array-size` may, depending on the implementation and the arguments, simply alter the given array or create and return a new one. In the latter case the given array will be altered so as to be displaced to the new array and have the given new dimensions.

If `adjust-array-size` is applied to an array created with the `:displaced-to` (page MAKE-ARRAY-DISPLACED-TO-KWD) option, or to an array used as the argument for the `:displaced-to` option in the creation of another array, then the operation will be performed correctly with respect to the given array, but the effects on the other array will be unpredictable.

Compatibility note: In Lisp Machine LISP, the argument *new-element* is not provided; it would seem useful, however.

Also the Lisp Machine LISP manual is unclear on the precise method of extension for multidimensional arrays. The above definition ties this down.

`array-grow` *array new-element* `&rest` *dimensions* [*Function*]

`array-grow` returns an array of the same type as *array*, with the specified *dimensions*. The number of *dimensions* given must equal the rank of *array*.

Those elements of *array* that are still in bounds appear in the new array. The elements of the new array that are not in the bounds of *array* are initialized to *new-element*; if this argument is not provided, then the initial contents of any new elements are undefined.

`array-grow` may, depending on the implementation and the arguments, simply alter the given array or create and return a new one. In the latter case the given array will be altered so as to be displaced to the new array and have the given new dimensions.

If `array-grow` is applied to an array created with the `:displaced-to` (page MAKE-ARRAY-DISPLACED-TO-KWD) option, or to an array used as the argument for the `:displaced-to` option in the creation of another array, then the operation will be performed correctly with respect to the given array, but the effects on the other array will be unpredictable.

`array-grow` differs from `adjust-array-size` in that it keeps the elements of a multidimensional array in the same logical positions while allowing extension of any or all dimensions, not just the first.

Chapter 17

Strings

A string is a specialized kind of vector whose elements are characters. While, strictly speaking, only vectors of characters are called strings (as opposed to all arrays of characters), the string operations described here will operate properly on any one-dimensional array of characters.

> **Compatibility note:** Lisp Machine LISP allows a fixnum to be coerced into a one-character string whose element is a character whose ASCII value is the fixnum. The net effect is that a single character can be automatically coerced to be a one-character string. It would be inconsistent with adherence to the character standard, and possibly also affect efficiency adversely in some implementations, to remain compatible with this.

As a rule, any string operation will accept a symbol instead of a string as an argument if the operation never modifies that argument; the print-name of the symbol is used. In this respect the string-specific sequence operations are not simply specializations of the generic versions; the generic sequence operations never accept symbols as sequences. This slight inelegance is permitted in COMMON LISP in the name of pragmatic utility. Also, there is a slight non-parallelism in the names of string functions. Where the suffixes `equalp` and `eql` would be more appropriate, for historical compatibility the suffixes `equal` and `=` are used instead to indicate case-insensitive and case-sensitive character comparison, respectively.

Any LISP object may be tested for being a string by the predicate `stringp` (page 48).

Note that strings, like all vectors, may have fill pointers. String operations generally operate only on the active portion of the string (below the fill pointer). See `array-reset-fill-pointer` (page 189) and related functions.

17.1. String Access and Modification

`char` *string index* [*Function*]

The given *index* must be a non-negative integer less than the length of *string*. The character at position *index* of the string is returned as a character object. (This character will necessarily satisfy the predicate `string-charp` (page 146).) As with all sequences in COMMON LISP, indexing is zero-origin.

For example:

```
(char "Floob-Boober-Bab-Boober-Bubs" 0) => #\F
(char "Floob-Boober-Bab-Boober-Bubs" 1) => #\l
```

See `elt` (page 157).

`rplachar` *string index newchar* [*Function*]

The argument *string* must be a string. The given *index* must be a non-negative integer less than the length of the string. The character at position *index* is altered to be *newchar*, which must be a character object which satisfies the predicate `string-charp` (page 146). `rplachar` returns *newchar* as its value. See `setelt` (page 157).

17.2. String Comparison

`string=` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]

`string=` compares two strings, and is true if they are the same (corresponding characters are identical) but is false if they are not. The function `equal` (page 50) calls `string=` if applied to two strings.

The keyword arguments `:start1` and `:start2` are the places in the strings to start the comparison. The arguments `:end1` and `:end2` are the places in the strings to stop comparing; comparison stops just *before* the position specified by a limit. The start arguments default to zero (beginning of string), and the end arguments (if either omitted or `nil`) default to the lengths of the strings (end of string), so that by default the entirety of each string is examined. These arguments are provided so that substrings can be compared efficiently.

`string=` is necessarily false if the (sub)strings being compared are of unequal length; that is, if

```
(not (= (- end1 start1) (- end2 start2)))
```

is true then `string=` is false.

For example:

```
(string= "foo" "foo") is true
(string= "foo" "Foo") is false
(string= "foo" "bar") is false
(string= "together" "frogs" :start1 1 :end1 2 :start2 3 :end2 4
   is true
```

`string-equal` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]

`string-equal` is just like `string=` except that differences in case are ignored; two characters are considered to be the same if `char-equal` (page 148) is true of them.

For example:

```
(string-equal "foo" "Foo") is true
```

`string<` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]
`string>` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]
`string<=` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]
`string>=` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]
`string/=` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]

The two string arguments are compared lexicographically, and the result is `nil` unless *string1* is (less than, greater than, less than or equal to, greater than or equal to, not equal to) *string2*, respectively. If the condition is satisfied, however, then the result is the index within the strings of the first character position at which the strings fail to match; put another way, the result is the length of the longest common prefix of the strings.

A string *a* is less than a string *b* if in the first position in which they differ the character of *a* is less than the corresponding character of *b* according to the function `char<` (page 148), or if string *a* is a proper prefix of string *b* (of shorter length and matching in all the characters of *a*).

The optional arguments *start1* and *start2* are the places in the strings to start the comparison. The optional arguments *end1* and *end2* places in the strings to stop comparing; comparison stops just *before* the position specified by a limit. The *start* arguments default to zero (beginning of string), and the *end* arguments (if either omitted or `nil`) default to the lengths of the strings (end of string), so that by default the entirety of each string is examined. These arguments are provided so that substrings can be compared efficiently. The index returned in case of a mismatch is an index into *string1*.

`string-lessp` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]
`string-greaterp` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]
`string-not-lessp` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]
`string-not-greaterp` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]
`string-not-equal` *string1 string2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]

These are exactly like `string<`, `string>`, `string<=`, `string>=`, and `string<>`, respectively, except that distinctions between upper-case and lower-case letters are ignored. It is if `char-lessp` (page 149) were used instead of `char<` (page 148) for comparing characters.

17.3. String Construction and Manipulation

`make-string` *count* `&optional` *fill-character* [*Function*]

This returns a string of length *count*, each of whose characters has been initialized to the *fill-character*. If *fill-character* is not specified, then the string will be initialized in an implementation-dependent way.

Implementation note: It may be convenient to initialize the string to null characters, or to spaces, or to garbage ("whatever was there").

`string-trim` *character-bag string* [*Function*]
`string-left-trim` *character-bag string* [*Function*]
`string-right-trim` *character-bag string* [*Function*]

`string-trim` returns a substring (in the sense of the function `substring` (page SUBSTRING-FUN)) of *string*, with all characters in *character-bag* stripped off of the beginning and end. The function `string-left-trim` is similar, but strips characters off only the beginning; `string-right-trim` strips off only the end. The argument *character-bag* may be a list of characters or a string.

For example:

```
(string-trim '(#\Space #\Tab #\Return) " garbanzo beans
        ") => "garbanzo beans"
(string-trim " (*)" " ( *three (silly) words* ) ")
   => "three (silly) words"
(string-left-trim " (*)" " ( *three (silly) words* ) ")
   => "three (silly) words* ) "
(string-right-trim " (*)" " ( *three (silly) words* ) ")
   => " ( *three (silly) words"
```

`string-upcase` *string* `&key :start :end` [*Function*]
`string-downcase` *string* `&key :start :end` [*Function*]
`string-capitalize` *string* `&key :start :end` [*Function*]

`string-upcase` returns a string just like *string* with all lower-case alphabetic characters replaced by the corresponding upper-case characters. More precisely, each character of the result string is produced by applying the function `char-upcase` (page 150) to the corresponding character of *string*.

`string-downcase` is similar, except that upper-case characters are converted to lower-case characters (using `char-downcase` (page 150)).

The keyword arguments `:start` and `:end` delimit the portion of the string to be affected.

The argument is not destroyed. However, if no characters in the argument require conversion, the result may be either the argument or a copy of it, at the implementation's discretion.

For example:

```
(string-upcase "Dr. Livingston, I presume?")
   => "DR. LIVINGSTON, I PRESUME?"
(string-downcase "Dr. Livingston, I presume?")
   => "dr. livingston, i presume?"
```

`string-capitalize` produces a copy of *string* such that every word (subsequence of case-modifiable characters delimited by non-case-modifiable characters) has its first character in upper-case and any other letters in lower-case.

For example:

```
(string-capitalize " hello ") => " Hello "
(string-capitalize
    "occlUDeD cASEmenTs FOreSTAll iNADVertent DEFenestraTION")
 => "Occluded Casements Forestall Inadvertent Defenestration"
(string-capitalize 'kludgy-hash-search) => "Kludgy-Hash-Search"
(string-capitalize "DON'T!") => "Don'T!"    ;not "Don't!"
```

17.4. Type Conversions on Strings

`string` *x* [*Function*]

`string` coerces *x* into a string. Most of the string functions apply this to such of their arguments as are supposed to be strings. If *x* is a string, it is returned. If *x* is a symbol, its print-name is returned. If *x* cannot be coerced to be a string, an error occurs.

To get the string representation of a number or any other LISP object, use `prin1string` (page 242), `princstring` (page 242), or `format` (page 244).

Chapter 18

Structures

COMMON LISP provides a facility for creating named record structures with named components. In effect, the user can declare a new data type; every data structure of that type has components with specified names. Constructor, access, and assignment constructs are automatically defined when the data type is declared.

This chapter is divided into two parts. The first part discusses the basics of the structure facility, which is very simple and allows the user to take advantage of the type-checking, modularity, and convenience of user-defined record data types. The second part discusses a number of specialized features of the facility which have advanced applications. These features are completely optional, and you needn't even know they exist in order to take advantage of the basics.

> **Rationale:** It is important not to scare the novice away from `defstruct` with a multiplicity of features. The basic idea is very simple, and we should encourage its use by providing a very simple description. The hairy stuff, including all options, is shoved to the end of the chapter.

18.1. Introduction to Structures

The structure facility is embodied in the `defstruct` macro, which allows the user to create and use aggregate datatypes with named elements. These are like "structures" in PL/I, or "records" in PASCAL.

As an example, assume you are writing a LISP program that deals with space ships in a two-dimensional plane. In your program, you need to represent a space ship by a LISP object of some kind. The interesting things about a space ship, as far as your program is concerned, are its position (represented as x and y coordinates), velocity (represented as components along the x and y axes), and mass.

A ship might therefore be represented as a record structure with five components: x-position, y-position, x-velocity, y-velocity, and mass. This structure could in turn be implemented as a LISP object in a number of ways. It could be a list of five elements; the x-position could be the *car*, the y-position the *cadr*, and so on. Equally well it could be a vector of five elements: the x-position could be element 0, the y-position element 1, and so on. The problem with either of these representations is that the components occupy places in the object which are quite arbitrary and hard to remember. Someone looking at `(cadddr ship1)` or `(vref ship1 3)` in a piece of code might find it difficult to determine that this is accessing the y-velocity component of `ship1`. Moreover, if the representation of a ship should have to be changed, it would be very

difficult top find all the places in the code to be changed to match (not all occurrences of `cadddr` are intended to extract the *y*-velocity from a ship).

Ideally components of record structures should have names. One would like to write something like `(ship-y-velocity ship1)` instead of `(cadddr ship1)`. One would also like a more mnemonic way to create a ship than this:

```
(list 0 0 0 0 0)
```

Indeed, one would like `ship` to be a new data type, just like other LISP data types, that one could test with `typep` (page 46), for example. The `defstruct` facility provides all of this.

`defstruct` itself is a macro which defines a structure. For the space ship example one we might define the structure by saying:

```
(defstruct ship
  x-position
  y-position
  x-velocity
  y-velocity
  mass)
```

This declares that every `ship` is an object with five named components. The evaluation of this form does several things:

- It defines `ship-x-position` to be a function of one argument, a ship, which returns its *x*-position; `ship-y-position` and the other components are given similar function definitions. These functions are called the *access functions*, as they are used to access elements of the structure.

- The symbol `ship` becomes the name of a data type, of which instances of ships are elements. This name becomes acceptable to `typep` (page 46), for example; `(typep x 'ship)` is true iff `x` is a ship. Moreover, all ships are instances of the type `structure`, because `ship` is a subtype of `structure`.

- A function named `ship-p` of one argument is defined; it is a predicate which is true if its argument is a ship, and is false otherwise.

- A function called `make-ship` is defined which, when invoked, will create a data structure with five components, suitable for use with the access functions. Thus executing

  ```
  (setq ship2 (make-ship))
  ```

  sets `ship2` to a newly-created `ship` object. One can specify the initial values of any desired component in the call to `make-ship` in this way:

  ```
  (setq ship2 (make-ship :mass *default-ship-mass*
                         :x-position 0
                         :y-position 0))
  ```

  This constructs a new ship and initializes three of its components. This function is called the *constructor function*, because it constructs a new structure.

  ??? **Query:** It seems desirable to make the `defstruct`-produced "`make-`" constructs as similar as possible to `make-array`, `make-symbol`, etc. Toward this end I here suggest that keywords be used to specify components, the keywords being formed simply by interning the slot names in the keyword package. Once

this is done, all arguments are evaluable (keywords being self-evaluating), and so the constructor macro might as well be a keyword-accepting function, like `make-array`.

- Two ways are provided to alter components of a `ship`. One way is to use the macro `setf` (page 60) in conjunction with an access function (because `defstruct` effectively performs an appropriate `defsetf` (page DEFSETF-FUN)):

```
(setf (ship-x-position ship2) 100)
```

This alters the *x*-position of *ship2* to be `100`. This works because `defstruct` generates an appropriate `defsetf` (page DEFSETF-FUN) form for each access function.

The other way is to use the special *alterant macro*, which allows alteration of several components at once in parallel:

```
(alter-ship enterprise                ;Counter-clockwise inter-quadrant warp!
            :x-position (- (ship-y-position enterprise))
            :y-position (ship-x-position enterprise))
```

Besides allowing parallel updating of several components, use of the alterant macro may be more efficient in certain cases.

??? **Query:** I'd really like to get rid of alterant macros. They're harder to understand than `setf`, and clutter up the language. I suspect that the gained efficiency is minuscule except in certain odd cases involving backed bytes. Is it worth it?

This simple example illustrates the power of `defstruct` to provide abstract record structures in a convenient manner. `defstruct` has many other features as well for specialized purposes.

18.2. How to Use Defstruct

`defstruct` *name-and-options* {*slot-description*}+ [*Macro*]

Defines a record-structure data type. A general call to `defstruct` looks like this:

```
(defstruct (name option-1 option-2 ...)
           slot-description-1
           slot-description-2
           ...)
```

name must be a symbol; it becomes the name of a new data type consisting of all instances of the structure. The function `typep` (page 46) will accept and use this name as appropriate.

Usually no options are needed at all. If no options are specified, then one may write simply *name* instead of (*name*) after the word `defstruct`. The syntax of options and the options provided are discussed in section ???.

Each *slot-description-j* is of the form

```
(slot-name default-init
     slot-option-name-1 slot-option-value-1
     slot-option-name-2 slot-option-value-2
     ...)
```

Each *slot-name* must be a symbol; an access function is defined for each slot. If no options and no

default-init are specified, then one may write simply *slot-name* instead of (*slot-name*) as the slot description. The *default-init* is a form which is evaluated *each time* a structure is to be constructed; the value is used as the initial value of the slot. If no *default-init* is specified, then the initial contents of the slot are undefined and implementation-dependent. The available slot-options are described in Section 18.4.

Compatibility note: Slot-options are not currently provided in Lisp Machine LISP, but this is an upward-compatible extension.

Besides defining an access function for each slot, `defstruct` arranges for `setf` to work properly on such access functions, defines a predicate named *name*`-p`, and defines constructor and alterant macros named `make-`*name* and `alter-`*name*, respectively. All names of automatically created functions and macros are symbols of the same package (if any) to which the structure *name* itself belongs.

Because evaluation of a `defstruct` form causes many functions and macros to be defined, one must take care that two `defstruct` forms do not define the same name (just as one must take care not to use `defun` to define two distinct functions of the same name). For this reason, as well as for clarity in the code, it is conventional to prefix the names of all of the slots with some text which identifies the structure. In the example above, all the slot names start with "`ship-`". The `:conc-name` (page 202)option can be used to provide such prefixes automatically.

18.3. Using the Automatically Defined Macros

After you have defined a new structure with `defstruct`, you can create instances of this structure by using the constructor macro, and alter the values of its slots by using the alterant macro. By default, `defstruct` defines these macros automatically, forming their names by adding prefixes to the name of the structure; for a structure named `foo`, the respective macro names would be `make-foo` and `alter-foo`. You can specify the names yourself by giving the name you want to use as the argument to the `:constructor` (page 204) and `:alterant` (page 204) options, or specify that you don't want a macro created at all by using `nil` as the argument.

18.3.1. Constructor Functions

A call to a constructor function, in general, has the form

(*name-of-constructor-macro*
slot-keyword-1 form-1
slot-keyword-2 form-2
...)

All arguments are keyword arguments. Each *slot-keyword* should be a keyword whose name matches the name of a slot of the structure (`defstruct` determines the possible keywords simply by interning each slot-name in the keyword package). All the *keywords* and *forms* are evaluated.

If *slot-keyword-j* names a slot, then that element of the created structure will be initialized to the value of

form-j. If no *slot-keyword-j*/*form-j* pair is present for a given slot, then the slot will be initialized by evaluating the *default-init* form specified for that slot in the call to `defstruct`. (In other words, the initialization specified in the `defstruct` defers to any specified in a call to the constructor macro.) If the default initialization form is used, it is evaluated at construction time, but in the lexical environment of the `defstruct` form in which it appeared. If the `defstruct` itself also did not specify any initialization, the element's initial value is undefined. You should always specify the initialization, either in the `defstruct` or in the call to the constructor function, if you care about the initial value of the slot.

Compatibility note: The Lisp Machine LISP documentation is slightly unclear about when the initialization specified in the `defstruct` form gets evaluated: at `defstruct` evaluation time, or at constructor time? The code reveals that it is at constructor time, which causes problems with referential transparency with respect to lexical variables (which currently don't exist officially in Lisp Machine LISP anyway). The above remark concerning the lexical environment in effect requires that the initialization form is treated as a thunk; it is evaluated at constructor time, but in the environment where it was written (the `defstruct` environment). Most of the time this makes no difference anyway, as the initialization form is typically a quoted constant or refers only to special variables. The requirement is imposed here for uniformity, and to ensure that what look like special variable references in the initialization form are in fact always treated as such.

The order of evaluation of the initialization forms is *not* necessarily the same as the order in which they appear in the constructor call or in the `defstruct` form; code should not depend on the order of evaluation. The initialization forms *are* re-evaluated on every constructor-macro call, so that if, for example, the form `(gensym)` were used as an initialization form, either in the constructor-macro call or as the default form in the `defstruct` declaration, then every call to the constructor macro would call `gensym` once to generate a new symbol.

18.3.2. Alterant Macros

A call to the alterant macro, in general, has the form

```
(name-of-alterant-macro  instance-form
      slot-keyword-1  form-1
      slot-keyword-2  form-2
      ...)
```

instance-form is evaluated, and should return an instance of the structure. Each *form-j* is evaluated, and the corresponding slot named by *slot-keyword-j* is changed to have the result as its new value. The assignments are parallel; that is, the slots are altered *after* all the *forms* have been evaluated, so you can exchange the values of two slots, as follows:

```
(alter-ship enterprise
    ship-x-position (ship-y-position enterprise)
    ship-y-position (ship-x-position enterprise))
```

As with the constructor macro, the order of evaluation of the *forms* is undefined.

Single slots can also be altered by using `setf` (page 60). Using the alterant macro may produce more efficient code than using consecutive `setf` forms.

18.4. `defstruct` Slot-Options

Each *slot-description* in a `defstruct` form may specify one or more slot-options. A slot-option consists of a pair of a keyword and a value (which is not a form to be evaluated, but the value itself).

For example:

```
(defstruct ship
  (ship-x-position 0.0 :type short-float)
  (ship-y-position 0.0 :type short-float)
  (ship-x-velocity 0.0 :type short-float)
  (ship-y-velocity 0.0 :type short-float)
  (ship-mass *default-ship-mass* :type short-float :read-only t))
```

This specifies that the first four slots will always contain short-format floating-point numbers, that the last three slots are "invisible" (will not ordinarily be shown when a ship is printed), and that the last slot may not be altered once a ship is constructed.

The available slot-options are:

`:type` The option (`:type` *type*) specifies that the contents of the slot will always be of the specified data type. This is entirely analogous to the declaration of a variable or function; indeed, it effectively declares the result type of the access function. An implementation may or may not choose to check the type of the new object when initializing or assigning to a slot.

`:invisible` The option `:invisible` specifies that the contents of this slot should not be printed when an instance of the structure is printed.

`:read-only` The option `:read-only` specifies that this slot may not be altered; it will always contain the value specified at construction time. The alterant macro will not accept the name of this slot, and `setf` (page 60) will not accept the access function for this slot.

18.5. Options to `defstruct`

The preceding description of `defstruct` is all that the average user will need (or want) to know in order to use structures. The remainder of this chapter discusses more complex features of the `defstruct` facility.

This section explains each of the options that can be given to `defstruct`. As with slot-options, a `defstruct` option may be either a keyword or a list of a keyword and arguments for that keyword.

??? **Query:** Suppose we could standardize on keyword-value pairs, as everywhere else?

`:conc-name` This provides for automatic prefixing of names of access functions. It is conventional to begin the names of all the access functions of a structure with a specific prefix, the name of the structure followed by a hyphen. This is the default behavior.

The argument to the `:conc-name` option specifies an alternate prefix to be used. (If a hyphen is to be used as a separator, it must be specified as part of the prefix.) If `nil` is

specified as an argument, then *no* prefix is used; then the names of the access functions are the same as the slot names, and it is up to the user to name the slots reasonably.

Note that no matter what is specified for `:conc-name`, with constructor functions and alterant macros one uses slot keywords that match the slot names, with no prefix attached. On the other hand, one uses the access-function name when using `setf`. Here is an example:

```
(defstruct (door (:conc-name nil))
  knob-color width material)
(setq my-door (make-door :knob-color 'red :width 5.0))
(door-knob-color my-door) ==> red
(alter-door my-door :knob-color 'green :material 'wood)
(door-material my-door) => wood
(setf (door-width my-door) 43.7)
(door-width my-door) => 43.7
```

`:type` The `:type` option specifies what kind of LISP object will be used to implement the structure. It takes one argument, which must be one of the types enumerated below. If the `:type` option is not provided, the type defaults to `:vector`, and the `:named` option is assumed unless `:unnamed` is explicitly specified.

Rationale: Making a structure be `:unnamed` mostly just saves space. It is probably better to protect the novice by providing by default a named vector, since that provides maximal features, nice printing, reasonable use of space (better than lists or arrays in most implementations), etc.

`vector` Use a general vector, storing components as vector elements. This is normally `:named`.

`array` Use a one-dimensional array, storing components in the body of the array. By default this is `:named`.

`(array` *type*`)` A specialized array may be used, in which case every component must be of a type which can be stored in such an array. The array must be one-dimensional.

Compatibility note: This is a suggested feature not yet in Lisp Machine LISP.

`array-leader` Make an object which is an array, and can be indexed as one, but which additionally has hidden `defstruct` components. By default this is `:named`. (See the option `:make-array` (page 206), described below.)

??? **Query:** This has to be renamed. But to what?

`list` Use a list. A structure of this type cannot be distinguished by `typep`, even if the `:named` option is used. By default this is `:unnamed`.

`integer` This unusual type implements the structure as a single integer. The structure may only have one slot. This is only useful with the byte field feature (see page DEFSTRUCT-BYTE-FIELD); it lets you store several small numbers within fields of an integer, giving the fields names. This cannot be `:named`.

Compatibility note: The `:integer` option is a suggested feature not yet in Lisp Machine Lisp. It is similar to the `fixnum` option.

Compatibility note: All the "`named-`" types such as `:named-array` from Lisp Machine Lisp have been omitted here, as they tend to multiply. An implementation may provide them, but they are not required here. The `:named` and `:unnamed` options may be used separately to get the same effect.

`:named` The `:named` option specifies that the structure is "named"; this option takes no argument. A named structure has an associated predicate for determining whether a given Lisp object is a structure of that name. Some named structures in addition can be distinguished by the predicate `typep` (page 46). If neither `:named` nor `:unnamed` is specified, then the default depends on the `:type` option.

`:unnamed` The `:unnamed` option specifies that the structure is not named; this option takes no argument.

`:constructor` This option takes one argument, a symbol, which specifies the name of the constructor function. If the argument is not provided or if the option itself is not provided, the name of the constructor is produced by concatenating the string "`make-`" and the name of the structure, putting the name in the same package as the structure name. If the argument is provided and is `nil`, no constructor function is defined.

This option actually has a more general syntax which is explained in ???.

`:alterant` This option takes one argument, which specifies the name of the alterant macro. If the argument is not provided or if the option itself is not provided, the name of the alterant macro is made by concatenating the string "`alter-`" to the name of the structure, putting the name in the same package as the structure name. If the argument is provided and is `nil`, no alterant macro is defined. Use of the alterant macro is explained on ???.

`:predicate` This option takes one argument, which specifies the name of the type predicate. If the argument is not provided or if the option itself is not provided, the name of the predicate is made by concatenating the name of the structure to the string "`-p`", putting the name in the same package as the structure name. If the argument is provided and is `nil`, no predicate is defined. A predicate can be defined only if the structure is `:named` (page 204).

`:include` This option is used for building a new structure definition as an extension of an old structure definition. As an example, suppose you have a structure called `person` that looks like this:

```
(defstruct person name age sex)
```

Now suppose you want to make a new structure to represent an astronaut. Since astronauts are people too, you would like them to also have the attributes of name, age, and sex, and you would like Lisp functions that operate on `person` structures to operate just as well on `astronaut` structures. You can do this by defining `astronaut` with the `:include` option, as follows:

```
(defstruct (astronaut (:include person))
   helmet-size
   (favorite-beverage 'tang))
```

The :include option causes the structure being defined to have the same slots as the included structure, in such a way that the access functions and alterant macro for the included structure will also work on the structure being defined. In this example, an astronaut will therefore have five slots: the three defined in person, and the two defined in astronaut itself. The access functions defined by the person structure can be applied to instances of the astronaut structure, and they will work correctly.

??? Query: RPG asks: should astronaut-name be defined, as well as letting person-name operate on astronauts?

The following examples illustrate how you can use astronaut structures:

```
(setq x (make-astronaut name 'buzz
                        age 45.
                        sex t
                        helmet-size 17.5))

(person-name x) => buzz
(favorite-beverage x) => tang
```

Note that the :conc-name (page 202) option was *not* inherited from the included structure; it only applies to the names of the access functions of person and not to those of astronaut.

The argument to the :include option is required, and must be the name of some previously defined structure. The included structure must be of the same :type as this structure. The structure name of the including structure definition becomes the name of a data type, of course; moreover, it becomes a subtype of the included structure. In the above example, astronaut is a subtype of person; hence

```
(typep (make-astronaut) 'person)
```

is true, indicating that all operations on persons will work on astronauts.

The following is an advanced feature of the :include option. Sometimes, when one structure includes another, the default values or slot-options for the slots that came from the included structure are not what you want. The new structure can specify default values or slot-options for the included slots different from those the included structure specifies, by giving the :include option as:

(:include *name slot-description-1 slot-description-2 ...*)

Each *slot-description-j* must have a *slot-name* or *slot-keyword* which is the same as that of some slot in the included structure. If *slot-description-j* has no *default-init*, then in the new structure the slot will have no initial value. Otherwise its initial value form will be replaced by the *default-init* in *slot-description-j*. A normally writable slot may be made read-only, and a normally visible slot may be made invisible in the defined structure. If a slot is invisible or read-only in the included structure, then it must also be so in the including structure. If a type is specified for a slot, it must be a the same as or a subtype of the type specified in the included structure. If it is a strict subtype, the implementation may or may not choose to error-check assignments.

For example, if we had wanted to define astronaut so that the default age for an astronaut is 45, then we could have said:

```
(defstruct (astronaut (:include person (age 45)))
  helmet-size
  (favorite-beverage 'tang))
```

`:make-array` If an array is used to represent the structure being defined (the `:type` (page 203) option is `:array` or `:array-leader`), this option allows you to control those aspects of the array used to implement the structure that are not otherwise constrained by `defstruct`. For example, if you are creating a structure of type `:array-leader`, you almost certainly want to specify the dimensions of the array to be created, and you may want to specify the type of the array.

The argument to the `:make-array` option should be a list of alternating keyword symbols to the function `make-array` (page 183) and forms whose values are the arguments to those keywords.

`defstruct` may need to specify some arguments to `make-array` for its own purposes. If these conflict with the specifications given to the `:make-array` keyword, an error is signalled.

> **Compatibility note:** This is more robust than the current Lisp Machine LISP specification that `defstruct` quietly overrides what you specify.

Constructor macros for structures implemented as arrays all allow the keyword `:make-array` to be supplied. Attributes supplied therein override any `:make-array` option attributes supplied in the original `defstruct` form. If some attribute appears in neither the invocation of the constructor nor in the `:make-array` option to `defstruct`, then the constructor will chose appropriate defaults.

If a structure is of type `:array-leader`, you probably want to specify the dimensions of the array. The dimensions of an array are given to `:make-array` as a position argument rather than a keyword argument, so there is no way to specify them in the above syntax. To solve this problem, you may use the special keyword `:dimensions` or `:length` (they mean the same thing), with a value that is anything acceptable as `make-array`'s first argument.

`:print-function`

The argument to this option should be a function of four arguments which is to be used to print structures of this type. When a structure of this type is to be printed, the function is called on the structure to be printed, a stream to print to, an integer indicating the current depth (to be compared against `prinlevel` (page 236)), and a flag which is true for `prin1`-style printout and false for `princ`-style printout. This option can be used only with `:named` structures.

> **Compatibility note:** This is suggested merely to provide a simple way to set up the printing function in a central place and in an implementation-independent manner. In Lisp Machine LISP this would presumably set up an invoke handler for the type. There needs to be a good way to interface to the grinder, too.

`:initial-offset`

This allows you to tell `defstruct` to skip over a certain number of slots before it starts allocating the slots described in the body. This option requires an argument (which must

be a non-negative integer) which is the number of slots you want `defstruct` to skip. To make use of this option requires that you have some familiarity with how `defstruct` is implementing your structure; otherwise, you will be unable to make use of the slots that `defstruct` has left unused.

`:eval-when` Normally the macros defined by `defstruct` are defined at eval time, compile time, and load time. This option allows the user to control this behavior. The argument to the `:eval-when` option is just like the list that is the first subform of an `eval-when` (page EVAL-WHEN-FUN) special form. For example,

```
(:eval-when (:eval :compile))
```

will cause the macros to be defined only when the code is running interpreted or inside the compiler.

18.6. By-position Constructor Functions

If the `:constructor` (page 204) option is given as (`:constructor` *name arglist*), then instead of making a keyword driven constructor function, `defstruct` defines a "positional" constructor function, taking arguments whose meaning is determined by the argument's position rather than by a keyword. The *arglist* is used to describe what the arguments to the constructor will be. In the simplest case something like `(:constructor make-foo (a b c))` defines `make-foo` to be a three-argument constructor function whose arguments are used to initialize the slots named `a`, `b`, and `c`.

In addition, the keywords `&optional`, `&rest`, and `&aux` are recognized in the argument list. They work in the way you might expect, but there are a few fine points worthy of explanation.

For example:

```
(:constructor create-foo
        (a &optional b (c 'sea) &rest d &aux e (f 'eff)))
```

This defines `create-foo` to be a constructor of one or more arguments. The first argument is used to initialize the `a` slot. The second argument is used to initialize the `b` slot. If there isn't any second argument, then the default value given in the body of the `defstruct` (if given) is used instead. The third argument is used to initialize the `c` slot. If there isn't any third argument, then the symbol `sea` is used instead. Any arguments following the third argument are collected into a list and used to initialize the `d` slot. If there are three or fewer arguments, then `nil` is placed in the `d` slot. The `e` slot *is not initialized*; its initial value is undefined. Finally, the `f` slot is initialized to contain the symbol `eff`.

The actions taken in the `b` and `e` cases were carefully chosen to allow the user to specify all possible behaviors. Note that the `&aux` "variables" can be used to completely override the default initializations given in the body.

With this definition, one can write

```
(create-foo 1 2)
```

instead of

```
(make-foo a 1 b 2)
```

and of course `create-foo` provides defaulting different from that of `make-foo`.

It is permissible to use the `:constructor` option more than once, so that you can define several different constructor functions, each taking different parameters.

If you write the keyword `:make-array` in place of a variable name, then the corresponding argument will specify the `:make-array` option at construction time, just as for an ordinary constructor function.

Because a constructor of this type operates By Order of Arguments, it is sometimes known as a BOA constructor.

Chapter 19

The Evaluator

19.1. Run-Time Evaluation of Forms

`eval` [*Function*]

???

`evalhook` [*Variable*]

???

`evalhook` [*Function*]

???

19.2. The Top-Level Loop

Where do `describe` and `inspect` go???

Chapter 20

Streams

Streams are objects that serve as sources or sinks of data. Character streams produce or absorb characters; binary streams produce or absorb integers. The normal action of a COMMON LISP system is to read characters from a character input stream, parse the characters into successive S-expressions, evaluate each S-expression in turn, and print the results to an output character stream.

Typically streams are connected to files or to an interactive terminal. Streams, being LISP objects, serve as the ambassadors of external devices by which input/output is accomplished.

A stream may be input-only, output-only, or bidirectional. What operations may be performed on a stream depends on which of the three types of stream it is.

20.1. Standard Streams

There are several variables whose values are streams used by many functions in the LISP system. These variables and their uses are listed here. By convention, variables which are expected to hold a stream capable of input have names ending with "`-input`", and similarly "`-output`" for output streams. Those expected to hold a bidirectional stream have names ending with "`-io`".

`standard-input` [*Variable*]

In the normal LISP top-level loop, input is read from `standard-input` (that is, whatever stream is the value of the global variable `standard-input`). Many input functions, including `read` (page 237) and `inch` (page 239), take a stream argument which defaults to `standard-input`.

`standard-output` [*Variable*]

In the normal LISP top-level loop, output is sent to `standard-output` (that is, whatever stream is the value of the global variable `standard-output`). Many output functions, including `print` (page 242) and `ouch` (page 243), take a stream argument which defaults to `standard-output`.

`error-output` [*Variable*]

The value of `error-output` is a stream to which error messages should be sent. Normally this is the same as `standard-output`, but `standard-output` might be bound to a file and `error-output` left going to the terminal or a separate file of error messages.

`query-io` [*Variable*]

The value of `query-io` is a stream to be used when asking questions of the user. The question should be output to this stream, and the answer read from it. When the normal input to a program may be coming from a file, questions such as "Do you really want to delete all of the files in your directory??" should be sent directly to the user, and the answer should come from the user, not from the data file. `query-io` is used by such functions as `yes-or-no-p` (page 254).

`terminal-io` [*Variable*]

The value of `terminal-io` is ordinarily the stream which connects to the user's console.

`trace-output` [*Variable*]

The value of `trace-output` is the stream on which the `trace` (page TRACE-FUN) function prints its output.

`standard-input`, `standard-output`, `error-output`, `trace-output`, and `query-io` are initially bound to synonym streams which pass all operations on to the stream which is the value of `terminal-io`. (See `make-synonym-stream` (page 212).) Thus any operations performed on those streams will go to the terminal.

No user program should ever change the value of `terminal-io`. A program which wants (for example) to divert output to a file should do so by binding the value of `standard-output`; that way error messages sent to `error-output` can still get to the user by going through `terminal-io`, which is usually what is desired.

20.2. Creating New Streams

Perhaps the most important constructs for creating new streams are those that open files; see `with-open-file` (page 267) and `open` (page 268). The following functions construct streams without reference to a file system.

`make-synonym-stream` *symbol* [*Function*]

`make-synonym-stream` creates and returns a "synonym stream". Any operations on the new stream will be performed on the stream which is then the value of the dynamic variable named by the *symbol*. If the value of the variable should change or be bound, then the synonym stream will operate on the new stream.

??? Query: In Lisp Machine Lisp this is called `make-syn-stream`. The documentor found it necessary to explain that "syn" meant "synonym"; it certainly isn't obvious. The abbreviation "syn" could be mistaken for any number of other things, such as "synchronous" or "syntactic" or "synthetic" ... Here this confusion is eliminated.

`make-broadcast-stream &rest` *streams* [*Function*]

Returns a stream which only works in the output direction. Any output sent to this stream will be sent to all of the streams given. The set of operations which may be performed on the new stream is the intersection of those for the given streams. The results returned by a stream operation are the values returned by the last stream in *streams*; the results of performing the operation on all preceding streams are discarded.

`make-concatenated-stream &rest` *streams* [*Function*]

Returns a stream which only works in the input direction. Input is taken from the first of the *streams* until it reaches end-of-file; then that stream is discarded, and input is taken from the next of the *streams*, and so on. If no arguments are given, the result is a stream with no content; any input attempt will result in end-of-file.

`make-io-stream` *input-stream output-stream* [*Function*]

Returns a bidirectional stream which gets its input from *input-stream* and sends its output to *output-stream*.

`make-echo-stream` *input-stream output-stream* [*Function*]

Returns a bidirectional stream which gets its input from *input-stream* and sends its output to *output-stream*. In addition, all input taken from *input-stream* is echoed to *output-stream*.

`make-string-input-stream` *string* `&optional` *start end* [*Function*]

Returns an input stream which will supply the characters the substring of *string* delimited by *start* and *end* in order and then signal end-of-file.

`make-string-output-stream &optional` *line-length* [*Function*]

Returns an output stream which will accumulate all output given it for the benefit of the function `get-output-stream-string`.

`get-output-stream-string` *string-output-stream* [*Function*]

Given a stream produced by `make-string-output-stream`, this returns a string containing all the characters output to the stream so far. The stream is then reset; thus each call to `get-output-stream-string` gets only the characters since the last such call (or the creation of the stream, if no such previous call has been made).

20.3. Operations on Streams

`streamp` *object* [*Function*]

`streamp` is true if its argument is a stream, and otherwise is false.

```
(streamp x) <=> (typep x 'stream)
```

`input-stream-p` *stream* [*Function*]

This predicate is true if its argument (a stream) can handle input operations, and otherwise is false.

`output-stream-p` *stream* [*Function*]

This predicate is true if its argument (a stream) can handle output operations, and otherwise is false.

`close` *stream* `&optional` *abort-flag* [*Function*]

The stream is closed. No further input/output operations may be performed on it. However, certain inquiry operations may still be performed, and it is permissible to close an already-closed stream.

If *abort-flag* is not `nil` (it defaults to `nil`), it indicates an abnormal termination of the use of the stream. An attempt is made to clean up any side effects of having created the stream in the first place. For example, if the stream performs output to a file, the file is deleted and any previously existing file is not superseded.

`charpos`, `linenum`, and so on?

Chapter 21

Input/Output

21.1. Printed Representation of LISP Objects

LISP objects are not normally thought of as being text strings; they have very different properties from text strings as a consequence of their internal representation. However, to make it possible to get at and talk about LISP objects, LISP provides a representation of objects in the form of printed text; this is called the *printed representation*, which is used for input/output purposes and in the examples throughout this manual. Functions such as `print` (page 242) take a LISP object and send the characters of its printed representation to a stream. The collection of routines that does this is known as the (LISP) *printer*. The `read` function takes characters from a stream, interprets them as a printed representation of a LISP object, builds a corresponding object, and returns it; the collection of routines that does this is called the (LISP) *reader*.

Ideally, one could print a LISP object and then read the printed representation back in, and so obtain the same identical object. In practice this is difficult, and for some purposes not even desirable. Instead, reading a printed representation produces an object that is (with obscure technical exceptions) `equal` (page 50) to the originally printed object.

Most LISP objects have more than one possible printed representation. For example, the integer twenty-seven can be written in any of these ways:

```
27    27.    #o33    #x1B    #b11011    #.(* 3 3 3)
```

A list of two symbols `A` and `B` can be printed in many, many ways:

```
(A B)    (a b)    (  a  b )    (\A |B|)
    (|\A|
  B
)
```

The last example, which is spread over three lines, may be ugly, but it is legitimate. In general, wherever whitespace is permissible in a printed representation, any number of spaces, tab characters, and newlines may appear.

When `print` produces a printed representation, it must choose arbitrarily from among many possible printed representations. It attempts to choose one that is readable. There are a number of global variables that can be used to control the actions of `print`, and a number of different printing functions.

This section describes in detail what is the standard printed representation for any Lisp object, and also describes how `read` operates.

21.1.1. What the `read` Function Accepts

The purpose of the reader LISP is to accept characters, interpret them as the printed representation of a LISP object, and construct and return such an object. The reader cannot accept everything that the printer produces; for example, the printed representations of compiled code objects and closures cannot be read in. However, the reader has many features that are not used by the output of the printer at all, such as comments, alternative representations, and convenient abbreviations for frequently-used unwieldy constructs. The reader is also parameterized in such a way that it can be used as a lexical analyzer for a more general user-written parser.

When the reader is invoked, it reads a character from the input stream and dispatches according to the attributes of that character. Every character that can appear in the input stream can have one of the following attributes: *whitespace*, *constituent*, *escape character*, or *macro character*. In addition, a macro character may be *terminating* or *non-terminating* (of tokens).

Supposing that the first character has been read; call it x. The reader then performs the following actions:

- If x is a *whitespace* character, then discard it and start over, reading another character.
- If x is a *macro character*, then execute the function associated with that character. The function may return zero values or one value (see `values` (page 82)). If one value is returned, that object is returned by the reader. If zero values are returned, the reader starts anew, reading a character from the input stream and dispatching. The function may of course read characters from the input stream; if it does, it will see those characters following the macro character.
- If x is an *escape character*, then read the next character and pretend it is a *constituent*, ignoring its usual syntax. Drop into the following case.
- If x is a *constituent*, then it begins an extended token, representing a symbol or a number. The reader reads more characters, accumulating them until a *whitespace* character or a *macro character* that is *terminating* is found, or until end-of-file is reached. However, whenever an *escape character* is found during the accumulation, the character after that is treated as a pure *constituent* and also accumulated, no matter what its usual syntax is. Similarly, any *non-terminating macro character* is simply accumulated as if it were a constituent. Call the eventually found *whitespace* character or *macro character* y. All characters beginning with x up to but not including y form a single extended token. (If end-of-file was encountered, the characters beginning with x up to the end of the file form the extended token.) This token is then interpreted as a number if possible, and otherwise as a symbol. The number or symbol is then returned by the reader.

Compatibility note: Note that characters of type *single* are not provided for (what MACLISP calls a "single character object"). They can be viewed as simply a kind of macro character. That is,

```
(setsyntax '$ 'single nil)
  <=> (setsyntax '$ 'macro #'(lambda (ignore ignore) '$))
```

which is easy enough to do oneself. After all, one might prefer to see a character rather than a symbol.

<tab> *whitespace*	<form> *whitespace*	<return> *whitespace*
<space> *whitespace*	@ *constituent*	*terminating macro character*
! *constituent*	A *constituent*	a *constituent*
" *terminating macro character*	B *constituent*	b *constituent*
# *terminating macro character*	C *constituent*	c *constituent*
$ *constituent*	D *constituent*	d *constituent*
% *constituent*	E *constituent*	e *constituent*
& *constituent*	F *constituent*	f *constituent*
' *terminating macro character*	G *constituent*	g *constituent*
(*terminating macro character*	H *constituent*	h *constituent*
) *terminating macro character*	I *constituent*	i *constituent*
* *constituent*	J *constituent*	j *constituent*
+ *constituent*	K *constituent*	k *constituent*
, *terminating macro character*	L *constituent*	l *constituent*
- *constituent*	M *constituent*	m *constituent*
. *constituent*	N *constituent*	n *constituent*
/ *constituent*	O *constituent*	o *constituent*
0 *constituent*	P *constituent*	p *constituent*
1 *constituent*	Q *constituent*	q *constituent*
2 *constituent*	R *constituent*	r *constituent*
3 *constituent*	S *constituent*	s *constituent*
4 *constituent*	T *constituent*	t *constituent*
5 *constituent*	U *constituent*	u *constituent*
6 *constituent*	V *constituent*	v *constituent*
7 *constituent*	W *constituent*	w *constituent*
8 *constituent*	X *constituent*	x *constituent*
9 *constituent*	Y *constituent*	y *constituent*
: *constituent*	Z *constituent*	z *constituent*
; *terminating macro character*	[*constituent*	{ *constituent*
< *constituent*	\ *escape character*	| *terminating macro character*
= *constituent*	] *constituent*	} *constituent*
> *constituent*	^ *constituent*	~ *constituent*
? *constituent*	_ *constituent*	<rubout> *constituent*
<backspace> *constituent*	<linefeed> *whitespace*	

Table 21-1: Standard Character Syntax Attributes

The characters of the standard character set initially have the attributes shown in Table 21-1.

21.1.2. Parsing of Numbers and Symbols

When an extended token is read, it is interpreted as a number or symbol. As a rule, letters not preceded by escape characters are converted to upper case. If the token can be interpreted as a number according to the BNF syntax in Table 21-2, then a number object of the appropriate type is constructed and returned. It should be noted that in a given implementation it may be that not all tokens conforming to the syntax for numbers

number ::= *integer* | *ratio* | *floating-point-number*
integer ::= [sign] {digit}+ [.]
ratio ::= [sign] {digit}+ / {digit}+
floating-point-number ::= [*sign*] {*digit*}* . {*digit*}+ [*exponent*] | [*sign*] {*digit*}+ . {*digit*}* *exponent*
sign ::= + | -
digit ::= 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
exponent ::= *exponent-marker* [sign] {digit}+
exponent-marker ::= e | s | f | d | l | b | E | F | D | S | L | B

The notation "{*x*}*" means zero or more occurrences of "*x*", the notation "{*x*}+" means one or more occurrences of "*x*", and the notation "[*x*]" means zero or one occurrences of "*x*".

Table 21-2: Syntax of Numbers

can actually be converted into number objects. For example, specifying too large or too small an exponent for a floating-point number may make the number impossible to represent in the implementation. Similarly, a ratio with denominator zero (such as "-35/000") cannot be represented in *any* implementation. The exponent markers "b" and "B" are undefined, but are reserved for future extension of the floating-point type. In any such circumstance where a token with the syntax of a number cannot be converted to an internal number object, an error is signalled. (On the other hand, an error cannot be signalled for specifying too many significant digits for a floating-point number.)

Note that a token representing a number may not contain any escape characters. An escape character robs the following character of all syntactic qualities, forcing it to be strictly alphabetic.

If the token consists solely of dots (with no escape characters), then an error is signalled, except in one circumstance: if the token is a single dot, and occurs in a situation appropriate to "dotted list" syntax, then it is accepted as a part of such syntax. (Signalling an error catches not only misplaced dots in dotted list syntax, but also lists that were truncated by `prinlength` (page 236) cutoff.)

In all other cases the token is construed to be the name of a symbol. If there are any package markers (colons) in the token, they divide the token into pieces used to control creation of the symbol. The *last* colon is used to divide the token into two parts. The first part specifies a package. A null first part indicates the `keyword` package; otherwise it is recursively interpreted as the name of a symbol, and that symbol must name a package. The second part is the name of the symbol, except that if the second part is null (implying that the last colon was the last character of the token) then `read` is used to read the object after the token, and it is read in such a way that the package specified by the first part will be the default package for reading symbols. This allows one to combine the colon and vertical-bar syntaxes, for example: "`editor:|odd(function)name|`".

??? **Query:** This opens up an entire pervasive-package-syntax can of worms. OK?

If a symbol token contains no package markers, then the entire token is the name of the symbol. The

symbol is looked up relative to the default package (see `package` (page 112)).

<tab>	alphabetic*	{	alphabetic
<linefeed>	alphabetic*	|	alphabetic*
<form>	alphabetic*	}	alphabetic
<return>	alphabetic*	~	alphabetic*
<space>	alphabetic*	@	alphabetic
!	alphabetic	A, a	alphabetic, superdigit
"	alphabetic*	B, b	alphabetic, superdigit, reserved exponent
#	alphabetic*	C, c	alphabetic, superdigit
$	alphabetic	D, d	alphabetic, superdigit, double-float exponent
%	alphabetic	E, e	alphabetic, superdigit, float exponent
&	alphabetic	F, f	alphabetic, superdigit, single-float exponent
'	alphabetic*	G, g	alphabetic, superdigit
(	alphabetic*	H, h	alphabetic, superdigit
)	alphabetic*	I, i	alphabetic, superdigit
*	alphabetic	J, j	alphabetic, superdigit
+	alphabetic, plus sign	K, k	alphabetic, superdigit
,	alphabetic*	L, l	alphabetic, superdigit, long-float exponent
-	alphabetic, minus sign	M, m	alphabetic, superdigit
.	alphabetic, dot, decimal point	N, n	alphabetic, superdigit
/	alphabetic, ratio marker	O, o	alphabetic, superdigit
0	digit	P, p	alphabetic, superdigit
1	digit	Q, q	alphabetic, superdigit
2	digit	R, r	alphabetic, superdigit
3	digit	S, s	alphabetic, superdigit, short-float exponent
4	digit	T, t	alphabetic, superdigit
5	digit	U, u	alphabetic, superdigit
6	digit	V, v	alphabetic, superdigit
7	digit	W, w	alphabetic, superdigit
8	digit	X, x	alphabetic, superdigit
9	digit	Y, y	alphabetic, superdigit
:	package marker	Z, z	alphabetic, superdigit
;	alphabetic*	[	alphabetic
<	alphabetic	\	alphabetic*
=	alphabetic	]	alphabetic
>	alphabetic	^	alphabetic
?	alphabetic	_	alphabetic
<rubout>	alphabetic	~	alphabetic
<backspace>	alphabetic		

* The interpretations in this table apply only to characters determined to have the *constituent* attribute. Entries marked with an asterisk are normally shadowed because the indicated characters have *whitespace*, *macro character*, or *escape character* syntax.

Table 21-3: Standard Constituent Character Attributes

The interpretation of standard characters within extended tokens is shown in Table 21-3. These interpretations can be used, of course, only for characters defined to be *constituent* characters. For characters of type *whitespace*, *macro character*, or *escape character*, the interpretations in Table 21-3 are effectively shadowed. (The interpretation of "superdigits" is relevant to the reading of rational numbers in a radix greater than ten.)

21.1.3. Macro Characters

If the reader encounters a macro character, then the function associated with that macro character is called, and may produce an object to be returned. This function may read following characters in the stream in whatever syntax it likes (it may even call `read` recursively) and returns the object represented by that syntax. Macro characters may not be recognized, of course, when read as part of other special syntaxes (such as for strings).

The reader is therefore organized into two parts: the basic dispatch loop, which also distinguishes symbols and numbers, and the collection of macro characters. Any character can be reprogrammed as a macro character; this is a means by which the reader can be extended. The macro characters normally defined are:

(The left parenthesis character initiates reading of a pair or list. The function `read` (page 237) is called recursively to read successive objects, until a right parenthesis is found to be next in the input stream. A list of the objects read is returned. Thus

```
(a b c)
```

is read as a list of three objects (the symbols `a`, `b`, and `c`). The right parenthesis need not follow the printed representation of the last object immediately; whitespace characters may precede it. This can be useful for putting one object on each line and making it easy to add new objects:

```
(defun traffic-light (color)
  (caseq color
    (green)
    (red (stop))
    (amber (accelerate))          ; Insert more colors after this line.
    ))
```

It may be that *no* objects precede the right parenthesis, as in "`()`" or "`( )`"; this reads as a list of zero objects (the empty list).

If a token is read between objects that is just a dot "`.`", not preceded by an escape character, then exactly one more object must follow (possibly followed by whitespace), and then the right parenthesis:

```
(a b c . d)
```

This means that the *cdr* of the last pair in the list is not `nil`, but rather the object whose representation followed the dot. The above example might have been the result of evaluating

```
(cons 'a (cons 'b (cons 'c 'd))) => (a b c . d)
```

Similarly, we have

```
(cons 'znets 'wolq-zorbitan) => (znets . wolq-zorbitan)
```

It is permissible for the object following the dot to be a list:

```
(a b c d . (e f . (g)))   is the same as   (a b c d e f g)
```

but this is a non-standard form that `print` will never produce.

) The right-parenthesis character is part of various constructs (such as the syntax for lists) using the left-parenthesis character, and is invalid except when used in such a construct.

' The single-quote (accent acute) character provides an abbreviation to make it easier to put constants in programs. '*foo* reads the same as (`quote` *foo*): a list of the symbol `quote` and *foo*.

; Semicolon is used to write comments. The semicolon and everything up through the next newline are ignored. Thus a comment can be put at the end of any line without affecting the reader (except that semicolon, being a macro character and therefore a delimiter, will terminate a token, and so cannot be put in the middle of a number or symbol).

For example:

```
;;;; COMMENT-EXAMPLE and related nonsense.
;;; This function is useless except to demonstrate comments.
;;; Notice that there are several kinds of comments.

(defun comment-example (x y)    ;X is anything; Y is an a-list.
  (cond ((listp x) x)           ;If X is a list, use that.
        ;; X is now not a list.  There are two other cases.
        ((symbolp x)
         ;; Look up a symbol in the a-list.
         (cdr (assq x y)))      ;Remember, (cdr nil) is nil.
        ;; Do this when all else fails:
        (t (cons x              ;Add x to a default list.
                 '((lisp t)     ;LISP is okay.
                   (fortran nil)     ;FORTRAN is not.
                   (pl/i -500)  ;Note that you can put comments in
                   (ada .001)   ; "data" as well as in "programs".
                   ;; COBOL??
                   (teco -1.0e9))))))
```

This example illustrates a few conventions for comments in common use. Comments may begin with one to four semicolons.

- Single-semicolon comments are all aligned to the same column at the right; usually each comments about only the line it is on. Occasionally two or three contain a single sentence together; this is indicated by indenting all but the first by a space.
- Double-semicolon comments are aligned to the level of indentation of the code. A space follows the two semicolons. Usually each describes the state of the program at that point, or describes the section that follows.
- Triple-semicolon comments are aligned to the left margin. Usually they are not used within S-expressions, but precede them in large blocks.
- Quadruple-semicolon comments are interpreted as subheadings by some software such as the ATSIGN listing program.

" The double-quote character begins the printed representation of a string. Characters are read from the input stream and accumulated until another double-quote is encountered, except that if an *escape character* is seen, it is discarded, the next character is accumulated, and accumulation continues. When a matching double-quote is seen, all the accumulated characters up to but not including the matching

double-quote are made into a string and returned.

The vertical-bar character begins one printed representation of a symbol. Characters are read from the input stream and accumulated until another vertical-bar is encountered, except that if an *escape character* is seen, it is discarded, the next character is accumulated, and accumulation continues. When a matching vertical-bar is seen, all the accumulated characters up to but not including the matching vertical-bar are made into a symbol and returned. In this syntax, no characters are ever converted to upper case; the name of the symbol is precisely those characters between the vertical bars (allowing for any escape characters).

The backquote (accent grave) character makes it easier to write programs to construct complex data structures by using a template. As an example, writing

```
(cond ((numberp ,x) ,@y) (t (print ,x) ,@y))
```

is roughly equivalent to writing

```
(list 'cond
      (cons (list 'numberp x) y)
      (list* 't (list 'print x) y))
```

The general idea is that the backquote is followed by a template, a picture of a data structure to be built. This template is copied, except that within the template commas can appear. Where a comma occurs, the form following the comma is to be evaluated to produce an object to be inserted at that point. Assume `b` has the value 3, for example, then evaluating the form denoted by "`(a b ,b ,(+ b 1) b)`" produces the result `(a b 3 4 b)`.

If a comma is immediately followed by an at-sign ("`@`"), then the form following the at-sign is evaluated to produce a *list* of objects. These objects are then "spliced" into place in the template. For example, if `x` has the value `(a b c)`, then

```
(x ,x ,@x foo ,(cadr x) bar ,(cdr x) baz ,@(cdr x))
   => (x (a b c) a b c foo b bar (b c) baz b c)
```

The backquote syntax can be summarized formally as follows. For each of several situations in which backquote can be used, a possible interpretation of that situation as an equivalent form is given. Note that the form is equivalent only in the sense that when it is evaluated it will calculate the correct result. An implementation is quite free to interpret backquote in any way such that a backquoted form, when evaluated, will produce a result `equal` to that produced by the interpretation shown here.

- *simple* is the same as '*simple*, that is, (`quote` *simple*), for any form *simple* that is not a list or a general vector.
- ,`@i[form]` is the same as *form*, for any *form*, provided that the representation of *form* does not begin with "`@`" or "`.`". (A similar caveat holds for all occurrences of a form after a comma.)
- ,`@`*form* is an error.
- (*x1 x2 x3 ... xn . atom*) may be interpreted to mean (`append` *x1 x2 x3 ... xn* (`quote` *atom*)), where the underscore indicates a transformation of an *xj* as follows:
  - *form* is interpreted as (`list` *form*), which contains a backquoted form that must then be further interpreted.

- ,*form* is interpreted as (`list` *form*).
- ,@*form* is interpreted simply as *form*.

- (*x1 x2 x3 ... xn*) may be interpreted to mean the same as (*x1 x2 x3 ... xn* . `nil`).
- (*x1 x2 x3 ... xn* . ,*form*) may be interpreted to mean (`append` *x1 x2 x3 ... xn form*), where the underscore indicates a transformation of an *xj* as above.
- (*x1 x2 x3 ... xn* . ,@*form*) is an error.
- `#`(*x1 x2 x3 ... xn*) may be interpreted to mean (`make-vector nil :initial-contents` (*x1 x2 x3 ... xn*)).

No other uses of comma are permitted; in particular, it may not appear within the `#A` or `#S` syntax.

Anywhere ",@" may be used, the syntax ",." may be used instead to indicate that it is permissible to destroy the list produced by the form following the ",."; this may permit more efficient code, using `nconc` (page 171) instead of `append` (page 170), for example.

If the backquote syntax is nested, the innermost backquoted form should be expanded first. This means that if several commas occur in a row, the leftmost one belongs to the innermost backquote.

Once again, it is emphasized that an implementation is free to interpret a backquoted form as any form that, when evaluated, will produce a result that is `equal` to the result implied by the above definition. In particular, no guarantees are made as to whether the constructed copy of the template will or will not share list structure with the template itself. As an example, the above definition implies that `((,a b) ,c ,@d)` will be interpreted as if it were

```
(append (list (append (list a) (list 'b) 'nil)) (list c) d 'nil)
```

but it could also be legitimately interpreted to mean any of the following:

```
(append (list (append (list a) (list 'b))) (list c) d)
(append (list (append (list a) '(b))) (list c) d)
(append (list (cons a '(b))) (list c) d)
(list* (cons a '(b)) c d)
(list* (cons a (list 'b)) c d)
(list* (cons a '(b)) c (copylist d))
```

(There is no good reason why `copylist` should be performed, but it is not prohibited.)

, The comma character is part of the backquote syntax and is invalid if used other than inside the body of a backquote construction as described above.

The sharp-sign character is a *dispatching* macro character. It reads an optional digit string and then one more character, and uses that character to select a function to run as a macro-character function. See the next section for predefined sharp-sign macro characters.

21.1.4. Sharp-Sign Abbreviations

The standard syntax includes forms introduced by a sharp sign ("#"). These take the general form of a sharp sign, a second character that identifies the syntax, and following arguments in some form. If the second character is a letter, then case is not important; #O and #o are considered to be equivalent, for example.

Certain sharp-sign forms allow an unsigned decimal number to appear between the sharp sign and the second character; some other forms even require it.

#<tab> signals error	#<form> signals error	#<return> signals error
#<space> signals error	#@ undefined	# undefined
#! undefined	#A array	#a array
#" bit-vector	#B binary rational	#b binary rational
## reference to label	#C complex number	#c complex number
#$ undefined	#D undefined	#d undefined
#% undefined	#E undefined	#e undefined
#& undefined	#F undefined	#f undefined
#' `function` abbreviation	#G undefined	#g undefined
#(general vector	#H undefined	#h undefined
#) signals error	#I undefined	#i undefined
#* undefined	#J undefined	#j undefined
#+ read-time conditional	#K undefined	#k undefined
#, load-time evaluation	#L undefined	#l undefined
#- read-time conditional	#M undefined	#m undefined
#. read-time evaluation	#N undefined	#n undefined
#/ undefined	#O octal rational	#o octal rational
#0 (infix argument)	#P undefined	#p undefined
#1 (infix argument)	#Q undefined	#q undefined
#2 (infix argument)	#R radix-*n* rational	#r radix-*n* rational
#3 (infix argument)	#S structure	#s structure
#4 (infix argument)	#T undefined	#t undefined
#5 (infix argument)	#U undefined	#u undefined
#6 (infix argument)	#V undefined	#v undefined
#7 (infix argument)	#W undefined	#w undefined
#8 (infix argument)	#X hexadecimal rational	#x hexadecimal rational
#9 (infix argument)	#Y undefined	#y undefined
#: undefined	#Z undefined	#z undefined
#; undefined	#[undefined	#{ undefined
#< signals error	#\ named character	#| undefined
#= labels LISP object	#] undefined	#} undefined
#> undefined	#^ undefined	#~ undefined
#? undefined	#_ undefined	#<rubout> undefined
#<backspace> undefined	#<backspace> signals error	

Table 21-4: Standard Sharp-Sign Macro Character Syntax

The currently-defined sharp-sign constructs are described below and summarized in Table 21-4; more are likely to be added in the future. However, the constructs "#!", "#?", "#[", "#]", "#{", and "#}" are explicitly reserved for the user and will never be defined by the COMMON LISP standard.

#\ #*x* reads in as a character object that represents the character *x*. Also, #*name* reads in as the character object whose name is *name*. Note that the backslash "\" allows this construct to be parsed easily by EMACS-like editors.

In the single-character case, the character *x* must be followed by a non-constituent character, lest a *name* appear to follow the "#\". A good model of what happens is that after "#\" is read, the reader backs up over the "\" and then reads an extended token, treating the initial "\" as an escape character (whether it really is or not in the current readtable).

Upper-case and lower-case letters are distinguished after "#\"; "#\A" and "#\a" denote different character objects. Any character works after #\, even those that are normally special to `read`, such as parentheses. Non-printing characters may be used after #\, although for them names are generally preferred.

#*name* reads in as a character object whose name is *name* (actually, whose name is (`string-upcase` *name*); therefore the syntax is case-insensitive). The following names are standard across all implementations:

`return`	The carriage return or newline character.
`space`	The space or blank character.

The following names are semi-standard; if an implementation supports them, they should be used for the described characters and no others.

`rubout`	The rubout or delete character.
`form`	The formfeed or page-separator character.
`tab`	The tabulate character.
`backspace`	The backspace character.
`linefeed`	The line feed character.

The *name* should have the syntax of a symbol.

When the LISP printer types out the name of a special character, it uses the same table as the #\ reader; therefore any character name you see typed out is acceptable as input (in that implementation). Standard names are always preferred over non-standard names for printing.

The following convention is used in implementations that support non-zero bits attributes for character objects. If a name after #\ is longer than one character and has a hyphen in it, then it may be split into the two parts preceding and following the first hyphen; the first part (actually, `string-upcase` of the first part) may then be interpreted as the name or initial of a bit, and the second part as the name of the character (which may in turn contain a hyphen and be subject to further splitting).

For example:

```
#\Control-Space                 #\Control-Meta-Tab
#\C-M-Return                    #\H-S-M-C-Rubout
```

If the character name consists of a single character, then that character is used. Another "\" may be necessary to quote the character.

```
#\Control-@                    #\Control-Meta-\"
#\Control-\a                    #\Meta->
```

If an unsigned decimal integer appears between the "#" and "\", it is interpreted as a font number, to become the `char-font` (page 150) of the character object.

> **Compatibility note:** Formerly, Lisp Machine LISP and MACLISP used #\ to mean only the #*name* version of this syntax, using #/ for the #*x* version. Lisp Machine LISP has recently changed to allow #/ to handle both syntaxes. The incompatibility is a result of the general exchange of the / and \ characters.
>
> Also, MACLISP and Lisp Machine LISP define #\ and #/ to be a syntax for *numbers*, integers that represent characters. Here they are a syntax for character objects. Code conforming to the "Character Standard for LISP" will not depend on this distinction; but non-conforming code (such as that which does arithmetic on bare character values) may not be compatible.

#' #'*foo* is an abbreviation for (`function` *foo*). *foo* may be the printed representation of any LISP object. This abbreviation may be remembered by analogy with the ' macro-character, since the `function` and `quote` special forms are similar in form.

#(A series of representations of objects enclosed by "#(" and ")" is read as a general vector of those objects. This is analogous to the notation for lists.

If an unsigned decimal integer appears between the "#" and "(", it specifies explicitly the length of the vector. In that case, it is an error if too many objects are specified before the closing ")", and if too few are specified the last one is used to fill all remaining elements of the vector.

For example:

```
#(a b c c c c)
#6(a b c c c c)
#6(a b c)
#6(a b c c)
```

all mean the same thing: a vector of length 6 with elements `a`, `b`, and four instances of `c`.

#" A series of binary digits (`0` and `1`) enclosed by "#"" and "" is read as a bit-vector of those objects. This is analogous to the notation for strings.

If an unsigned decimal integer appears between the "#" and "", it specifies explicitly the length of the bit-vector. In that case, it is an error if too many bits are specified before the closing "", and if too few are specified the last one is used to fill all remaining elements of the bit-vector.

For example:

```
#"101000"
#6"101000"
#6"1010"
#6"10100"
```

all mean the same thing.

#. #.*foo* is read as the object resulting from the evaluation of the LISP object represented by *foo*, which may be the printed representation of any LISP object. The evaluation is done during the `read` process, when the `#.` construct is encountered. This, therefore, performs a "read-time" evaluation of *foo*. By contrast, #, (see below) performs a "load-time" evaluation.

This allows you, for example, to include in your code complex list-structure constants that cannot be written with `quote`. Note that the reader does not put `quote` around the result of the evaluation. You must do this yourself if you want it, typically by using the ' macro-character. An example of a

case where you do not want `quote` around it is when this object is an element of a constant list.

`#,` `#,`*foo* is read as the object resulting from the evaluation of the LISP object represented by *foo*, which may be the printed representation of any LISP object. The evaluation is done during the `read` process, *unless* unless the compiler is doing the reading, in which case it is arranged that *foo* will be evaluated when the file of compiled code is loaded. This, therefore, performs a "load-time" evaluation of *foo*. By contrast, `#.` (see above) performs a "read-time" evaluation. In a sense, `#,` is like specifying `(eval load)` to `eval-when` (page EVAL-WHEN-FUN), while `#.` is more like specifying `(eval compile)`. It makes no difference when loading interpreted code, but when code is to be compiled, `#.` specifies compile-time evaluation and `#,` specifies load-time evaluation.

`#,` allows you, for example, to include in your code complex list-structure constants that cannot be written with `quote`. Note that the reader does not put `quote` around the result of the evaluation. You must do this yourself if you want it, typically by using the ' macro-character. An example of a case where you do not want `quote` around it is when this object is an element of a constant list.

`#B` `#b`*rational* reads *rational* in binary (radix 2).

`#O` `#o`*rational* reads *rational* in octal (radix 8).

`#X` `#x`*rational* reads *rational* in hexadecimal (radix 16). The digits above `9` are the letters `A` through `F` (the lower-case letters `a` through `f` are also acceptable).

`#`*n*`R` `#`*radix*`r`*rational* reads *rational* in radix *radix*. *radix* must consist of only digits, and it is read in decimal; its value must be between 2 and 36 (inclusive).

For example, `#3r102` is another way of writing `11`, and `#11R32` is another way of writing `35`. For radices larger than 10, letters of the alphabet are used in order for the digits after `9`.

`#S` The syntax `#s(`*name slot1 value1 slot2 value2 ...*`)` denotes a structure. This is legal only if *name* is the name of a structure already defined by `defstruct` (page 199), and if the structure has a standard constructor macro, which it normally will. Let *cm* stand for the name of this constructor macro; then this syntax is equivalent to

#.(cm slot1 'value1 slot2 'value2 ...)

That is, the constructor macro is called, with the specified slots having the specified values (note that one does not write quote-marks in the `#s` syntax). Whatever object the constructor macro returns is returned by the `#s` syntax.

If *name* is `vector` or `array`, however, the syntax is instead `#s(`*name dimension-info key1 value1 key2 value2 ...*`)`, and is treated as equivalent to

#.(cf key1 'value1 key2 'value2 ...)

where *cf* is `make-array` or `make-vector`, as appropriate.

`#`*n*`=` The syntax `#`*n*`=`*object* reads as whatever LISP object has *object* as its printed representation. However, that object is labelled by *n*, a required unsigned decimal integer, for possible reference by the syntax `#`*n*`#` (below). The scope of the label is the S-expression being read by the outermost call to `read`. Within this S-expression the same label may not appear twice.

??? **Query:** Should we require that a label occur textually before any references?

##

The syntax *#n#*, where *n* is a required unsigned decimal integer, serves as a reference to some object labelled by *#n=*; that is, *#n#* represents a pointer to the same identical (`eq`) object labelled by *#n=*. This permits notation of structures with shared or circular substructure. For example, a structure created in the variable y by this code:

```
(setq x (list 'p 'q))
(setq y (list (list 'a 'b) x 'foo x))
(rplacd (last x) (cdr x))
```

could be represented in this way:

```
((a b) . #1=(#2=(p q) foo #2# . #1#))
```

Without this notation, but with `prinlength` (page 236) set to `10`, the structure would print in this way:

```
((a b) (p q) foo (p q) (p q) foo (p q) (p q) foo (p q) ...)
```

#+

The *#+* syntax provides a read-time conditionalization facility. The general syntax is "*#+feature form*". If *feature* is "true", then this syntax represents a LISP object whose printed representation is *form*. If *feature* is "false", then this syntax is effectively whitespace; it is as if it did not appear.

The *feature* should be the printed representation of a symbol or list. If *feature* is a symbol, then it is true iff it is a member of the list that is the value of the global variable `features` (page FEATURES-VAR).

> **Compatibility note:** MACLISP uses the `status` special form for this purpose, and Lisp Machine LISP duplicates `status` essentially only for the sake of `(status features)`. The use of a variable allows one to bind the features list, for example when compiling.

Otherwise, *feature* should be a boolean expression composed of `and`, `or`, and `not` operators on (recursive) *feature* expressions.

For example, suppose that in implementation A the features `spice` and `perq` are true, and in implementation B the feature `lispm` is true. Then the expressions on the left below are read the same as those on the right in implementation A:

```
(cons #+spice "Spice" #+lispm "Lispm" x)   (cons "Spice" x)
(setq a '(1 2 #+perq 43 #+(not perq) 27))  (setq a '(1 2 43))
(let ((a 3) #+(or spice lispm) (b 3))      (let ((a 3) (b 3))
  (foo a))                                   (foo a))
```

In implementation B, however, they are read in this way:

```
(cons #+spice "Spice" #+lispm "Lispm" x)   (cons "Lispm" x)
(setq a '(1 2 #+perq 43 #+(not perq) 27))  (setq a '(1 2 27))
(let ((a 3) #+(or spice lispm) (b 3))      (let ((a 3) (b 3))
  (foo a))                                   (foo a))
```

The *#+* construction must be used judiciously if unreadable code is not to result. The user should make a careful choice between read-time conditionalization and run-time conditionalization. See the macros named `if-for` (page IF-FOR-FUN) and `if-in` (page IF-IN-FUN).

#-

#-feature form is equivalent to *#+*`(not` *feature*`)` *form*.

#<

This is not legal reader syntax. It is used in the printed representation of objects that cannot be read back in. Attempting to read a *#<* will cause an error. (More precisely, it *is* legal syntax, but the macro-character function for it signals an error.)

#<space>, #<tab>, #<return>, #<form>

A # followed by a standard whitespace character is not legal reader syntax. This is so that abbreviated forms produced via `prinlevel` (page 236) cutoff will not read in again; this serves as a safeguard against losing information. (More precisely, it *is* legal syntax, but the macro-character function for it signals an error.)

#)

This is not legal reader syntax. This is so that abbreviated forms produced via `prinlevel` (page 236) cutoff will not read in again; this serves as a safeguard against losing information. (More precisely, it *is* legal syntax, but the macro-character function for it signals an error.)

21.1.5. The Readtable

Previous sections have described the standard syntax accepted by the `read` function. This section discusses the advanced topic of altering the standard syntax, either to provide extended syntax for LISP objects or to aid the writing of other parsers.

There is a data structure called the *readtable* that is used to control the reader. It contains information about the syntax of each character equivalent to that in Table 21-1. Initially it is set up exactly as in Table 21-1 to give the standard COMMON LISP meanings to all the characters, but the user can change the meanings of characters to alter and customize the syntax of characters. It is also possible to have several readtables describing different syntaxes and to switch from one to another by binding the symbol `readtable`.

Even if an implementation supports characters with non-zero *bits* and *font* attributes, it need not (but may) allow for such characters to have syntax descriptions in the readtable. However, every character of type `string-char` must be represented in the readtable.

`readtable` [*Variable*]

The value of `readtable` is the current readtable. The initial value of this is a readtable set up for standard COMMON LISP syntax. You can bind this variable to temporarily change the readtable being used.

To program the reader for a different syntax, a set of functions are provided for manipulating readtables. Normally, you should begin with a copy of the standard COMMON LISP readtable and then customize the individual characters within that copy.

`copy-readtable &optional` *from-readtable to-readtable* [*Function*]

A copy is made of *from-readtable*, which defaults to the current readtable (the value of the global variable `readtable`). If *from-readtable* is `nil`, then a copy of a standard COMMON LISP readtable is made; for example,

```
(setq readtable (copy-readtable nil))
```

will restore the input syntax to standard COMMON LISP syntax, even if the original readtable has been clobbered (assuming it is not so badly clobbered that you cannot type in the above expression!).

If *to-readtable* is unsupplied or `nil`, a fresh copy is made. Otherwise *to-readtable* must be a readtable, which is clobbered with the copy.

`set-syntax-from-char` *to-char from-char* `&optional` *to-readtable from-readtable* [*Function*]

Makes the syntax of *to-char* in *to-readtable* be the same as the syntax of *from-char* in *from-readtable*. The *to-readtable* defaults to the current readtable (the value of the global variable `readtable` (page 229)), and *from-readtable* defaults to `nil`, meaning to use the syntaxes from the standard LISP readtable.

Only attributes as shown in Table 21-1 are copied; moreover, if a *macro character* is copied, the macro definition function is copied also. However, attributes as shown in Table 21-3 are not copied; they are "hard-wired" into the extended-token parser. For example, if the definition of "S" is copied to "`*`", then "`*`" will become a *constituent*, but will be simply *alphabetic* and cannot be used as an exponent indicator for short-format floating-point number syntax.

It "works" to copy a macro definition from a character such as "`|`" to another character; the standard definition for "`|`" looks for another character that is the same as the character that invoked it. It doesn't "work" to copy the definition of "`(`" to "`{`", for example; it can be done, but it lets one write lists in the form "`{a b c)`", not "`{a b c}`", because the definition always looks for a closing "`)`". See the function `read-delimited-list` (page 238), which is useful in this connection.

`set-macro-character` *char function* `&optional` *non-terminating-p readtable* [*Function*]
`get-macro-character` *char* `&optional` *readtable* [*Function*]

`set-macro-character` causes *char* to be a macro character that when seen by `read` causes *function* to be called. If *non-terminating-p* is not `nil` (it defaults to `nil`), then it will be a non-terminating macro character: it may be embedded within extended tokens. `get-macro-character` returns the function associated with *char*, and as a second value returns the *non-terminating-p* flag; it returns `nil` if *char* does not have macro-character syntax. In each case, *readtable* defaults to the current readtable.

function is called with two arguments, *stream* and *char*. The *stream* is the input stream, and *char* is the macro-character itself. In the simplest case, *function* may return a LISP object. This object is taken to be that whose printed representation was the macro character and any following characters read by the *function*. As an example, a plausible definition of the standard single-quote character is:

```
(defun single-quote-reader (stream ignore)
  (list 'quote (read stream)))
(set-macro-character #\' #'single-quote-reader)
```

The function reads an object following the single-quote and returns a list of the symbol `quote` and that object. The *char* argument is ignored.

The function may choose instead to return *zero* values (for example, by using `(values)` as the return expression). In this case the macro character and whatever it may have read contribute

nothing to the object being read. As an example, here is a plausible definition for the standard semicolon (comment) character:

```
(defun semicolon-reader (stream ignore)
  (do () ((char= (inch stream) #\Return)))  ;Eat rest of line.
  (values))                                 ;Return no values.
(set-macro-character #\; #'semicolon-reader)
```

The *function* should not have any side-effects other than on the *stream*. Front ends (such as editors and rubout handlers) to the reader may cause *function* to be called repeatedly during the reading of a single expression in which the macro character only appears once, because of backtracking and restarting of the `read` operation.

`make-dispatch-macro-character` *char* `&optional` *non-terminating-p readtable* [*Function*]

This causes the character *char* to be a dispatching macro character in *readtable* (which defaults to the current readtable). If *non-terminating-p* is not `nil` (it defaults to `nil`), then it will be a non-terminating macro character: it may be embedded within extended tokens.

Initially every character in the dispatch table has a character-macro function that signals an error. Use `set-dispatch-macro-character` to define entries in the dispatch table.

`set-dispatch-macro-character` *disp-char sub-char function* `&optional` *readtable* [*Function*]
`get-dispatch-macro-character` *disp-char sub-char* `&optional` *readtable* [*Function*]

`set-dispatch-macro-character` causes *function* to be called when the *disp-char* followed by *sub-char* is read. The *readtable* defaults to the current readtable. The arguments and return values for *function* are the same as for normal macro characters, documented above under `set-macro-character` (page 230), except that *function* gets *sub-char* as its second argument, and also receives a third argument that is the non-negative integer whose decimal representation appeared between *disp-char* and *sub-char*, or `nil` if there was none. The *sub-char* may not be one of the ten decimal digits; they are always reserved for specifying an infix integer argument.

`get-dispatch-macro-character` returns the macro-character function for *sub-char* under *disp-char*.

As an example, suppose one would like #$*foo* to be read as if it were `(dollars` *foo*`)`. One might say:

```
(defun sharp-dollar-reader (stream ignore ignore)
  (list 'dollars (read stream)))
(set-dispatch-macro-character #\# #\$ #'sharp-dollar-reader)
```

Compatibility note: This macro-character mechanism is different from those in MACLISP, INTERLISP, and Lisp Machine LISP. Recently LISP systems have implemented very general readers, even readers so programmable that they can parse arbitrary compiled BNF grammars. Unfortunately, these readers can be complicated to use. This design is an attempt to make the reader as *simple* as possible to understand, use, and implement. Splicing macros have been eliminated; a recent informal poll indicates that no one uses them to produce other than zero or one value. The ability to access parts of the object preceding the macro character have been eliminated. The single-character-object feature has been eliminated, because it is seldom used and trivially obtainable by defining a macro.

The user is encouraged to turn off most macro characters, turn others into single-character-object macros, and then use

`read` purely as a lexical analyzer on top of which to build a parser. It is unnecessary, however, to cater to more complex lexical analysis or parsing than that needed for COMMON LISP.

21.1.6. What the `print` Function Produces

The COMMON LISP printer is controlled by a number of special variables. Foremost among these is `prinescape`.

`prinescape` [*Variable*]

When this flag is `nil`, then escape characters are not output when an S-expression is printed. In particular, a symbol is printed by simply printing the characters of its print name. The function `princ` (page 242) effectively binds `prinescape` to `nil`.

When this flag is not `nil`, then an attempt is made to print an S-expression in such a way that it can be read again to produce an `equal` structure. The function `prin1` (page 242) effectively binds `prinescape` to `t`.

Compatibility note: This flag controlkes what was called *slashification* in MACLISP.

The initial value of this variable is `t`.

`prinpretty` [*Variable*]

When this flag is `nil`, then only a small amount of whitespace is output when printing an expression, as described below.

When this flag is not `nil`, then the printer will endeavor to insert extra whitespace where appropriate to make the expression more readable.

`princircle` [*Variable*]

When this flag is `nil` (the default), then the printing process proceeds by recursive descent; an attempt to print a circular structure may lead to looping behavior and failure to terminate.

When this flag is not `nil`, then the printer will endeavor to detect cycles in the structure to be printed, and to use *#n=* and *#n#* syntax to indicate the circularities.

How an expression is printed depends on its data type.

Integers. If appropriate, a radix specifier may be printed; see `prinradix` below. If an integer is negative, a minus sign is printed and then the absolute value of the integer is printed. Non-negative integers are printed in the radix specified by `base` in the usual positional notation, most significant digit first. The number zero is represented by the single digit `0`, and never has a sign. A decimal point may then be printed.

`base` [*Variable*]

The value of `base` determines in what radix the printer will print rationals. This may be any integer from 2 to 36, inclusive; the default value is 10 (decimal radix). For radices above 10, letters of the alphabet are used to represent digits above "9".

Compatibility note: MACLISP and Lisp Machine LISP have a default `base` of 8.

Floating-point numbers are always printed in decimal, no matter what the value of `base`.

`prinradix` [*Variable*]

If the variable `prinradix` is non-`nil`, the printer will print a radix specifier to indicate the radix in which it is printing a rational number. For example, if the current base is twenty-four (decimal), the decimal integer twenty-three would print as "`#24RN`". If `base` is 2, 8, or 16, then the radix specifier used is `#B`, `#O`, or `#X`. For integers, base ten is indicated by a trailing decimal point, instead of using a leading radix specifier; for ratios, "`#10R`" is used. The default value of `prinradix` is `nil`.

Ratios. If appropriate, a radix specifier may be printed; see `prinradix`. If the ratio is negative, a minus sign is printed. Then the absolute value of the numerator is printed, as for an integer; then a "`/`"; then the denominator. The numerator and denominator are both printed in the radix specified by `base`.

Floating-point numbers. Floating point numbers are printed in one of two ways. If the floating point number is between 10^{-3} (inclusive) and 10^{7} (exclusive), it may be printed as the integer part of the number, then a decimal point, followed by the fractional part of the number; there is always at least one digit on each side of the decimal point. Outside of that range, it will be printed in "computerized scientific notation", with the exponent character indicating the precision of the number. For example, Avogadro's number as a short-format floating-point number would be printed as "`6.02S23`". If the format of the number matches that specified by `read-default-float-format` (page 237), however, then the exponent marker "`E`" is used.

Characters. When `prinescape` (page PRINESCAPE-FUN) is `nil`, a character prints as itself; it is sent directly to the output stream. When `prinescape` is not `nil`, then `#\` syntax is used. For example, the printed representation of the character `#\a` with control and meta bits on would be `#\CONTROL-META-\a`.

Symbols. When `prinescape` (page PRINESCAPE-FUN) is `nil`, the only characters of the print name of the symbol are output. When `prinescape` is not `nil`, backslashes "`\`" and vertical bars "`|`" are included as required, and package prefixes may be printed (using colon "`:`" syntax) if necessary. As a special case, `nil` may sometimes be printed as "`()`" instead, when `prinescape` and `prinpretty` are each not `nil`.

Strings. The characters of the string are output in order. If `prinescape` (page 232) is not `nil`, a double quote "`"`" is output beforehand and afterward, and all and double quotes and escape characters are preceded by "`\`".

Conses. Wherever possible, list notation is preferred over dot notation. Therefore the following algorithm

is used:

1. Print an open parenthesis "(".
2. Print the *car* of the cons.
3. If the *cdr* is a cons, make is the current cons, print a space, and go to step 2.
4. If the *cdr* is not null, print a space, a dot ".", a space, and the *cdr*.
5. Print a close parenthesis "(".

This form of printing is clearer than showing each individual cons cell. Although the two S-expressions below are equivalent, and the reader will accept either one and produce the same data structure, the printer will always print such a data structure in the second form.

```
(a . (b . ((c . (d . nil)) . (e . nil))))
(a b (c d) e)
```

The printing of conses is affected by the variables `prinlevel` (page 236) and `prinlength` (page 236).

General Vectors. The printed representation of a zero length vector is "`#()`". The printed representation of a non-zero length vector begins with a `#(`. Following the `#(` is printed the first element of the vector. If there are any other elements, they are printed in turn, with a space printed before each additional element. A close parenthesis after the last element terminates the printed representation of the vector. The printing of vectors is affected by the variables `prinlevel` (page 236) and `prinlength` (page 236).

Bit-vectors. A bit vector is printed as "`#"`", then all the bits in the vector as "0" and "1" characters, then a closing double-quote "`"`". The empty bit vector is therefore printed as "`#""`".

Arrays. ???

Structures defined by `defstruct` (page 199) are printed under the control of the `:printer` option to `defstruct`.

Any other types are printed in an implementation-dependent manner. It is recommended that printed representations of all such objects begin with the characters "#<" and end with ">" so that the reader will catch such objects and not permit them to be read under normal circumstances.

When debugging or when frequently dealing with large or deep objects at toplevel, the user may wish to restrict the printer from printing large amounts of information. The variables `prinlevel` and `prinlength` allow the user to control how deep the printer will print, and how many elements at a given level the printer will print. Thus the user can see enough of the object to identify it without having to wade through the entire expression.

`prinlevel` [*Variable*]
`prinlength` [*Variable*]

The `prinlevel` variable controls how many levels deep a nested data object will print. If `prinlevel` is `nil` (the initial value), then no control is exercised. Otherwise the value should be

an integer, indicating the maximum level to be printed. An object to be printed is at level `0`; its components (as of a list or vector) are at level `1`; and so on. If an object to be recursively printed has components and is at a level equal or greater to the value of `prinlevel`, then the object is printed as simply "#".

The `prinlength` variable controls how many elements at a given level are printed. A value of `nil` (the initial value) indicates that there be no limit to the number of components printed. Otherwise the value of `prinlength` should be an integer. Should the number of elements of a data object exceed the value `prinlength`, the printer will print three dots "..." in place of those elements beyond the number specified by `prinlength`. (In the case of a dotted list, if the list contains exactly as many elements as the value of `prinlength`, and in addition has the non-null atom terminating it, that terminating atom is printed, rather than printing "...".)

As an example, here are the ways the object

```
(if (member x items) (+ (car x) 3) '(foo . #(a b c d "Baz")))
```

would be printed for various values of `prinlevel` = *v* and `prinlength` = *n*.

```
v  n
0  1  #
1  1  (if ...)
1  2  (if # ...)
1  3  (if # # ...)
1  4  (if # # #)
2  1  (if ...)
2  2  (if (member x ...) ...)
2  3  (if (member x items) (+ # 3) ...)
3  2  (if (member x items) ...)
3  3  (if (member x items) (+ (car x) 3) ...)
3  4  (if (member x items) (+ (car x) 3) '(foo . #(a b c d ...
```

21.2. Input Functions

21.2.1. Input from ASCII Streams

Many input functions take optional arguments called *input-stream* and *eof-value*. The *input-stream* argument is the stream from which to obtain input; if unsupplied or `nil` it defaults to the value of the special variable `standard-input` (page 211). One may also specify `t` as a stream, meaning the value of the special variable `terminal-io` (page 212).

Rationale: Allowing the use of `t` provides some semblance of MACLISP compatibility.

The *eof-value* argument controls what happens if input is from a file (or any other input source that has a definite end) and the end of the file is reached. If no *eof-value* argument is supplied, an error will be signalled at end of file. If there is an *eof-value*, it is the value to be returned. Note that an *eof-value* of `nil` means to return `nil` if the end of the file is reached; it is *not* equivalent to supplying no *eof-value*. The *eof-value* argument is always evaluated; the resulting value is used, however, only when end of file is encountered.

Functions such as `read` (page 237) that read an "object" rather than a single character will always signal an error, regardless of *eof-value*, if the file ends in the middle of an object. For example, if a file does not contain enough right parentheses to balance the left parentheses in it, `read` will complain. If a file ends in a symbol or a number immediately followed by end-of-file, `read` will read the symbol or number successfully and when called again will see the end-of-file and return *eof-value*. Similarly, the function `readline` (page 239) will successfully read the last line of a file even if that line is terminated by end-of-file rather than the newline character. If a file contains ignorable text at the end, such as blank lines and comments, `read` will not consider it to end in the middle of an object and will return *eof-value*.

??? **Query:** Should `nil` as an *eof-value* be reserved to mean the same thing as omitting the *eof-value*?

Compatibility note: These end-of-file conventions are compatible with Lisp Machine LISP, but not completely compatible with Maclisp. Maclisp's deviations from this are generally considered to be bugs rather than features.

The MACLISP "feature" of letting *input-stream* and *eof-value* appear in either order is not supported.

Note that all of these functions will echo their input if used on an interactive stream. The functions that input more than one character at a time allow the input to be edited. The function `inchpeek` (page 240) echoes all of the characters that are skipped over (if any) if `inch` would have echoed them; the character not removed from the stream is not echoed either.

`read &optional` *input-stream eof-value* [*Function*]

`read` reads in the printed representation of a LISP object from *input-stream*, builds a corresponding LISP object, and returns the object. The details are explained above.

`read-default-float-format` [*Variable*]

The value of this variable must be one of `short`, `single` (the initial value), `double`, or `long`. It indicates the floating-point format to be used for reading floating-point numbers that have no exponent marker or have "`e`" or "`E`" for an exponent marker. (Other exponent markers explicitly prescribe the floating-point format to be used.) The printer also uses this variable to guide the choice of exponent markers when printing floating-point numbers.

`read-preserving-whitespace &optional` *input-stream eof-value* [*Function*]

Certain printed representations given to `read`, notably those of symbols and numbers, require a delimiting character after them. (Lists do not, because the close parenthesis marks the end of the list.) Normally `read` will throw away the delimiting character if it is a white-space character, but will preserve it (using `untyi` (page 239)) if the character is syntactically meaningful, since it may be the start of the next expression.

The function `read-preserving-whitespace` is provided for some specialized situations where it is desirable to determine precisely what character terminated the extended token.

As an example, consider this macro-character definition:

```
(defun slash-reader (stream ignore)
  (do ((path (list (read-preserving-whitespace stream))
             (cons (progn (inch stream)
                          (read-preserving-whitespace stream))
                   path)))
      ((not (char= (inchpeek stream) #\/))
       (cons 'pathname (nreverse path)))))
(set-macro-character #\/ #'slash-reader)
```

Consider now calling `read` on this expression:

```
(zyedh /usr/games/zork /usr/games/boggle)
```

The "/" macro reads objects separated by more "/" characters; thus `/usr/games/zork` is intended to read as `(pathname usr games zork)`. The entire example expression should therefore be read as

```
(zyedh (pathname usr games zork) (pathname usr games boggle))
```

However, if `read` had been used instead of `read-preserving-whitespace`, then after the reading of the symbol `zork`, the following space would be discarded, and then the next call to `inchpeek` would see the following "/", and the loop would continue, producing this interpretation:

```
(zyedh (pathname usr games zork usr games boggle))
```

On the other hand, there are times when whitespace *should* be discarded. If one has a command interpreter that takes single-character commands, but occasionally reads a LISP object, then if the whitespace after a symbol were not discarded it might be interpreted as a command some time later after the symbol had been read.

`read-delimited-list` *char* `&optional` *input-stream* [*Function*]

This reads objects from *stream* until the next character after an object's representation (ignoring whitespace characters) is *char*. (The *char* should not have whitespace syntax in the current readtable.) A list of the objects read is returned.

This function is particularly useful for defining new macro-characters. Suppose one were to want "`#{`*a b c ... z*`}`" to read as a list of all pairs of the elements *a*, *b*, *c*, . . . , *z*; for example:

`#{p q z a}` reads as `((p q) (p z) (p a) (q z) (q a) (z a))`

This can be done by specifying a macro-character definition for "#{" that does two things: read in all the items up to the "}", and construct the pairs. `read-delimited-list` performs the first task.

```
(defun sharp-leftbrace-reader (stream ignore ignore)
  (mapcon #'(lambda (x)
              (mapcar #'(lambda (y) (list x y)) (cdr x)))
          (read-delimited-list #\} stream)))
(set-dispatch-macro-character #\# #\{ #'sharp-leftbrace-reader)
```

Note that `read-delimited-list` does not take an *eof-value* argument. The reason for this is that it is always an error to hit end-of-file during the operation of `read-delimited-list`.

`readline` `&optional` *input-stream eof-value* [*Function*]

`readline` reads in a line of text, terminated by the implementation's usual way for indicating end-of-line (typically a <return> character). It returns the line as a character string (*without* the <return> character). This function is usually used to get a line of input from the user. A second returned value is a flag that is false if the line was terminated normally, or true if end-of-file terminated the (non-empty) line.

`inch` `&optional` *input-stream eof-value* [*Function*]

`tyi` `&optional` *input-stream eof-value* [*Function*]

`inch` inputs one character from *input-stream* and returns it as a character object. The character is echoed if *input-stream* is interactive.

`tyi` is similar to `inch`, but returns the character as an integer; it is as if `inch` were used, and `char-int` (page 151) applied to the result.

It is almost always preferable to use `inch` rather than `tyi`, if only for reasons of portability.

`uninch` *character* `&optional` *input-stream* [*Function*]

`untyi` *integer* `&optional` *input-stream* [*Function*]

`uninch` puts the *character* onto the front of *input-stream*. The *character* must be the same character that was most recently read from the *input-stream*. The *input-stream* "backs up" over this character; when a character is next read from *input-stream*, it will be the specified character, followed by the previous contents of *input-stream*. `uninch` returns `nil`.

`untyi` is similar to `uninch`, but takes an integer rather than a character object. It is as if `uninch` were used after applying `int-char` (page 151) to the first argument. It is almost always preferable to use `uninch` rather than `untyi`, if only for reasons of portability.

One may only apply `uninch` or `untyi` to the character most recently read from input-*stream*; moreover, one may not invoke `uninch` or `untyi` twice consecutively without an intervening `inch` or `tyi` operation. The result is that one may back up only by one character, and one may not insert any characters into the input stream that were not already there.

> **Rationale:** This is not intended to be a general mechanism, but rather an efficient mechanism for allowing the LISP reader and other parsers to perform one-character lookahead in the input stream. This protocol admits a wide variety of efficient implementations, such as simply decrementing a buffer pointer. To have to specify the character in the call to `uninch` is admittedly redundant, since there at any given time is only one character that may be legally specified. The redundancy is intentional, again to give the implementation latitude.

`inchpeek` `&optional` *peek-type input-stream eof-value* [*Function*]

`tyipeek` `&optional` *peek-type input-stream eof-value* [*Function*]

What `inchpeek` does depends on the *peek-type*, which defaults to `nil`. With a *peek-type* of `nil`, `inchpeek` returns the next character to be read from *input-stream*, without actually removing it from the input stream. The next time input is done from *input-stream* the character will still be there. It is as if one had called `inch` and then `uninch` in succession.

If *peek-type* is `t`, then `inchpeek` skips over whitespace characters, and then performs the peeking operation on the next character. This is useful for finding the (possible) beginning of the next printed representation of a Lisp object. As above, the last character (the one that starts an object) is not removed from the input stream.

If *peek-type* is a character object, then `inchpeek` skips over input characters until a character that is `char=` (page 148) to that object is found; that character is left in the input stream.

Characters passed over by `inchpeek` are echoed if *input-stream* is interactive.

`tyipeek` is similar to `inchpeek`, but returns an integer rather than a character object; it is as if `inchpeek` were used, and `char-int` (page 151) applied to the result. (If, however, an *eof-value* is provided and returned, `char-int` is not applied!) `tyipeek` also requires an integer instead of a character as the *peek-type*.

It is almost always preferable to use `inchpeek` rather than `tyipeek`, if only for reasons of portability.

`listen &optional` *input-stream* [*Function*]

The predicate `listen` is true if there is a character immediately available from *input-stream*, and is false if not. This is particularly useful when the stream obtains characters from an interactive device such as a keyboard; a call to `inch` (page 238) would simply wait until a character was available, but `listen` can sense whether or not input is available and allow the program to decide whether or not to attempt input. On a non-interactive stream, the general rule is that `listen` is true except when at end-of-file.

`inch-no-hang &optional` *input-stream eof-value* [*Function*]

`tyi-no-hang &optional` *input-stream eof-value* [*Function*]

These functions are exactly like `inch` (page 238) and `tyi` (page 238), except that if it would be necessary to wait in order to get a character (as from a keyboard), `nil` is immediately returned without waiting. This allows one efficiently to check for input being available and get the input if it is. This is different from the `listen` (page 240) operation in two ways. First, these functions potentially actually read a character, while `listen` never inputs a character. Second, `listen` does not distinguish between end-of-file and no input being available, while these functions do make that distinction, returning *eof-value* at end-of-file (or signalling an error if no *eof-value* was given), but always returning `nil` if no input is available.

`clear-input &optional` *input-stream* [*Function*]

This clears any buffered input associated with *input-stream*. It is primarily useful for clearing type-ahead from keyboards when some kind of asynchronous error has occurred. If this operation doesn't make sense for the stream involved, when `clear-input` does nothing. `clear-input` returns `nil`.

`read-from-string` *string* `&optional` *start end preserve-whitespace-p eof-value* [*Function*]

The characters of *string* are given successively to the LISP reader, and the LISP object built by the reader is returned. Macro characters and so on will all take effect.

The arguments *start* and *end* delimit a substring of *string* beginning at the character indexed by *start* and up to but not including the character indexed by *end*. By default *start* is 0 (the beginning of the string) and *end* is (`length` *string*). This is as for other string functions.

The flag *preserve-delimiters-p*, if provided and not `nil`, indicates that the operation should preserve whitespace as for `read-preserving-whitespace` (page 236).

The *eof-value* is what to return if the end of the (sub)string is reached before the operation is completed, as with other reading functions.

`read-from-string` returns two values; the first is the object read and the second is the index of the first character in the string not read. If the entire string was read, this will be either the length of the string or one greater than the length of the string. The parameter *preserve-whitespace-p* may affect this second value.

For example:

```
(read-from-string "(a b c)") => (a b c) and 7
```

`parse-number` *string* `&optional` *start end radix no-junk-allowed* [*Function*]

This function examines the substring of *string* delimited by *start* and *end* (which default to the beginning and end of the string). It skips over whitespace characters and then attempts to parse a number, in the syntax for <number> given in Table 21-2. The *radix* defaults to 10, and must be an integer between 2 and 36. If the *radix* is not 10, then floating-point numbers will not be permitted by the parse.

If *no-junk-allowed* is `nil` (the default), then the first value returned is the number parsed, or `nil` if no syntactically correct number was seen. The second value is the index into the string of the delimiter that terminated the parse, or the index beyond the substring if the parse terminated at the end of the substring.

If *no-junk-allowed* is not `nil`, then the entire substring is scanned. An error is signalled if the substring does not consist entirely of the representation of a number, possibly surrounded on either side by whitespace characters. The returned value is the number parsed, or 0 if no number was found (the substring was blank).

21.2.2. Input from Binary Streams

`in` *binary-input-stream* `&optional` *eof-value* [*Function*]

`in` reads one byte from the *binary-input-stream* and returns it in the form of a non-negative integer.

21.3. Output Functions

21.3.1. Output to ASCII Streams

These functions all take an optional argument called *output-stream*, which is where to send the output. If unsupplied or `nil`, *output-stream* defaults to the value of the variable `standard-output` (page 211). If it is `t`, the value of the variable `terminal-io` (page 212) is used.

`prin1` *object* `&optional` *output-stream* [*Function*]
`print` *object* `&optional` *output-stream* [*Function*]
`princ` *object* `&optional` *output-stream* [*Function*]

`prin1` outputs the printed representation of *object* to *output-stream*, using escape characters. As a rule, the output from `prin1` is suitable for input to the function `read` (page 236); see ???. `prin1` returns *object*.

`print` is just like `prin1` except that the printed representation of *object* is preceded by a <return> character and followed by a <space>. `print` returns *object*.

`princ` is just like `prin1` except that the output has no escape characters. A symbol is printed as simply the characters of its print-name; a string is printed without surrounding double-quotes; and there may be differences for other data types as well. The general rule is that output from `princ` is intended to look good to people, while output from `prin1` is intended to be acceptable to the function `read` (page 236). `princ` returns *object*.

The output from these functions is affected by the values of the variables `base` (page 233), `prinlevel` (page 234), and `prinlength` (page 234).

Compatibility note: In MACLISP, these three functions return `t`, not the argument *object*. There is some old code that depends on the value being non-`nil`, such as in:

```
(and condition (print x) (print y) (print z))
```

which should have been written as

```
(cond (condition (print x) (print y) (print z)))
```

but someone was too lazy to do it that way (`when` didn't exist in those days). Ugh. COMMON LISP does not support this bad style.

`prin1string` *object* [*Function*]
`princstring` *object* [*Function*]

The object is effectively printed, as by `prin1` or `princ`, and the characters that would be output are made into a string and returned.

`ouch` *character* `&optional` *output-stream* [*Function*]
`tyo` *integer* `&optional` *output-stream* [*Function*]

`ouch` outputs the *character* to *output-stream*.

`tyo` is similar, but takes an integer instead of a character; it is as if `int-char` were applied to the

first argument and then `ouch` were called.

It is almost always preferable to use `ouch` rather than `tyo`, if only for reasons of portability.

Both functions return `t`.

`terpri &optional` *output-stream* [*Function*]
`fresh-line &optional` *output-stream* [*Function*]

`terpri` outputs a newline to *output-stream*; this may be simply a carriage-return character, a return-linefeed sequence, or whatever else is appropriate for the stream. `terpri` returns `nil`.

`fresh-line` is similar to `terpri`, but outputs a newline only if the stream is not already at the start of a line. (If for some reason this cannot be determined, then a newline is output anyway.) This guarantees that the stream will be on a "fresh line" while consuming as little vertical distance as possible. `fresh-line` is a (side-effecting) predicate that is true if it output a newline, and otherwise false.

`force-output &optional` *output-stream* [*Function*]
`clear-output &optional` *output-stream* [*Function*]

Some streams may be implemented in an asynchronous or buffered manner. The function `force-output` attempts to ensure that all output sent to *output-stream* has reached its destination, and only then returns `nil`.

The function `clear-output`, on the other hand, attempts to abort any outstanding output operation in progress, to allow as little output as possible to continue to the destination. This is useful, for example, to abort a lengthy output to the terminal when an asynchronous error occurs. `clear-output` returns `nil`.

The function `format` (page 244) is very useful for producing nicely formatted text. It can do anything any of the above functions can do, and it makes it easy to produce good looking messages and such. `format` can generate a string or output to a stream.

The function `pprint` (page PPRINT-FUN) is useful for printing LISP objects "prettily" in an indented format. Also, `grindef` (page GRINDEF-FUN) is useful for formatting LISP programs.

21.3.2. Output to Binary Streams

`out` *integer binary-output-stream* [*Function*]

`out` writes one byte, the value of *integer* (which must be non-negative and smaller than the largest valid byte value) to the *binary-output-stream*.

??? **Query:** Should this limitation on the argument be enforced? Should it quietly grab the low *n* bits? What about writing signed bytes to a file? What about writing floating-point numbers or characters to a binary file?

21.4. Formatted Output

`format` *destination control-string* `&rest` *arguments* [*Function*]

`format` is used to produce formatted output. `format` outputs the characters of *control-string*, except that a tilde ("~") introduces a directive. The character after the tilde, possibly preceded by prefix parameters and modifiers, specifies what kind of formatting is desired. Most directives use one or more elements of *args* to create their output; the typical directive puts the next element of *args* into the output, formatted in some special way.

The output is sent to *destination*. If *destination* is `nil`, a string is created that contains the output; this string is returned as the value of the call to `format`. In all other cases `format` returns `nil`, performing output to *destination* as a side effect. If *destination* is a stream, the output is sent to it. If *destination* is `t`, the output is sent to the stream that is the value of the variable `standard-output` (page 211).

A `format` directive consists of a tilde ("~"), optional prefix parameters separated by commas, optional colon (":") and atsign ("@") modifiers, and a single character indicating what kind of directive this is. The alphabetic case of the directive character is ignored. The prefix parameters are generally decimal numbers. Examples of control strings:

```
"~S"        ;This is an S directive with no parameters or modifiers.
"~3,4:@s"   ;This is an S directive with two parameters, 3 and 4,
            ;    and both the colon and atsign flags.
"~,4S"      ;Here the first prefix parameter is omitted and takes
            ;    on its default value, while the second parameter is 4.
```

The `format` function includes some extremely complicated and specialized features. It is not necessary to understand all or even most of its features to use `format` effectively. The beginner should skip over anything in the following documentation that is not immediately useful or clear. The more sophisticated features are there for the convenience of programs with complicated formatting requirements.

Sometimes a prefix parameter is used to specify a character, for instance the padding character in a right- or left-justifying operation. In this case a single quote (" ' ") followed by the desired character may be used as a prefix parameter, so that you don't have to know the decimal numeric values of characters in the character set. For example, you can use "~5,'0d" to print a decimal number in five columns with leading zeros, or "~5,'*d" to get leading asterisks.

In place of a prefix parameter to a directive, you can put the letter "V", which takes an argument from *arguments* as a parameter to the directive. Normally this should be an integer (but in general it doesn't really have to be). This feature allows variable column-widths and the like. Also, you can use the character "#" in place of a parameter; it represents the number of arguments remaining to be processed.

Here are some relatively simple examples to give you the general flavor of how `format` is used.

```
(format nil "foo") => "foo"
(setq x 5)
(format nil "The answer is ~D." x) => "The answer is 5."
(format nil "The answer is ~3D." x) => "The answer is   5."
(format nil "The answer is ~3,'0D." x) => "The answer is 005."
(format nil "The answer is ~:D." (expt 47 x))
                                  => "The answer is 229,345,007."

(setq y "elephant")
(format nil "Look at the ~A!" y) => "Look at the elephant!"
(format nil "Type ~:C to ~A." (control #\D) "delete all your files")
     => "Type Control-D to delete all your files."

(setq n 3)
(format nil "~D item~:P found." n) => "3 items found."
(format nil "~R dog~:[s are~; is~] here." n (= n 1))
     => "three dogs are here."
(format nil "~R dog~:*~[~1; is~:;s are~] here." n)
     => "three dogs are here."
(format nil "Here ~[~1;is~:;are~] ~:*~R pupp~:@P." n)
     => "Here are three puppies."
```

The directives will now be described. The term *arg* in general refers to the next item of the set of *arguments* to be processed. The word or phrase at the beginning of each description is a mnemonic word for the directive.

`~A` *Ascii.* An *arg*, any LISP object, is printed without escape characters (as by `princ` (page 241)). In particular, if *arg* is a string, its characters will be output verbatim. Normally all occurrences of `nil` in the printed object will be printed as "`nil`", but the colon modifier (`~:A`) will cause them to be printed as "`()`".

~mincol`A` inserts spaces on the right, if necessary, to make the width at least *mincol* columns. The `@` modifier causes the spaces to be inserted on the left rather than the right.

~mincol, colinc, minpad, padchar`A` is the full form of `~A`, which allows elaborate control of the padding. The string is padded on the right with at least *minpad* copies of *padchar*; padding characters are then inserted *colinc* characters at a time until the total width is at least *mincol*. The defaults are `0` for *mincol* and *minpad*, `1` for *colinc*, and the space character for *padchar*.

`~S` *S-expression.* This is just like `~A`, but *arg* is printed *with* escape characters (as by `prin1` (page 241) rather than `princ`). The output is therefore suitable for input to `read` (page 236). `~S` can accept all the arguments and modifiers that `~A` can.

`~D` *Decimal.* An *arg*, which should be an integer, is printed in decimal radix. `~D` will never put a decimal point after the number.

~mincol`D` uses a column width of *mincol*; spaces are inserted on the left if the number requires fewer than *mincol* columns for its digits and sign. If the number doesn't fit in *mincol* columns, additional columns are used as needed.

~mincol, padchar`D` uses *padchar* as the pad character instead of space.

If *arg* is not an integer, it is printed in `~A` format and decimal base.

The `@` modifier causes the number's sign to be printed always; the default is only to print it

if the number is negative. The : modifier causes commas to be printed between groups of three digits; the third prefix parameter may be used to change the character used as the comma. Thus the most general form of ~D is ~*mincol*,*padchar*,*commachar*D.

~B *Binary.* This is just like ~D but prints in binary radix (radix 2) instead of decimal. The full form is therefore ~*mincol*,*padchar*,*commachar*B.

~O *Octal.* This is just like ~D but prints in octal radix (radix 8) instead of decimal. The full form is therefore ~*mincol*,*padchar*,*commachar*O.

~X *Hexadecimal.* This is just like ~D but prints in hexadecimal radix (radix 16) instead of decimal. The full form is therefore ~*mincol*,*padchar*,*commachar*X.

~R *Radix.* ~*n*R prints *arg* in radix *n*. The modifier flags and any remaining parameters are used as for the ~D directive. Indeed, ~D is the same as ~10R. The full form here is therefore ~*radix*,*mincol*,*padchar*,*commachar*R.

If no arguments are given to ~R, then an entirely different interpretation is given. The argument should be an integer; suppose it is 4.

- ~R prints *arg* as a cardinal English number: "four".
- ~:R prints *arg* as an ordinal English number: "fourth".
- ~@R prints *arg* as a Roman numeral: "IV".
- ~:@R prints *arg* as an old Roman numeral: "IIII".

~P *Plural.* If *arg* is not eql to the integer 1, a lower-case "s" is printed; if *arg* is eql to 1, nothing is printed. (Notice that if *arg* is a floating-point 1.0, the "s" *is* printed.)

~:P does the same thing, after doing a ~:* to back up one argument; that is, it prints a lower-case "s" if the *last* argument was not 1. This is useful after printing a number using ~D.

~@P prints "y" if the argument is 1, or "ies" if it is not. ~:@P does the same thing, but backs up first.

```
(format nil "~D tr~:@P/~D win~:P" 7 1) => "7 tries/1 win"
(format nil "~D tr~:@P/~D win~:P" 1 0) => "1 try/0 wins"
(format nil "~D tr~:@P/~D win~:P" 1 3) => "1 try/3 wins"
```

~F *Floating-point.*

??? **Query:** Is this really what we want?

arg is printed in floating point. ~*n*F rounds *arg* to a precision of *n* digits. The minimum value of *n* is 2, since a decimal point is always printed. If the magnitude of *arg* is too large or too small, it is printed in exponential notation. If *arg* is not a number, it is printed in ~A format. Note that the prefix parameter *n* is not *mincol*; it is the number of digits of precision desired. Examples:

```
(format nil "~2F" 5)  =>  "5.0"
(format nil "~4F" 5)  =>  "5.0"
(format nil "~4F" 1.5) => "1.5"
(format nil "~4F" 3.14159265)  =>  "3.142"
(format nil "~3F" 1e10)  =>  "1.0e10"
```

~E *Exponential.*

??? **Query:** Is this the right thing Study PL/I, FORTRAN.

arg is printed in exponential notation. This is identical to ~F, including the use of a prefix parameter to specify the number of digits, except that the number is always printed with a trailing exponent, even if it is within a reasonable range.

~C *Character.* The next *arg* should be a character; it is printed according to the modifier flags.

~C prints the character in an implementation-dependent abbreviated format. This format should be culturally compatible with the host environment.

Implementation note: In Lisp Machine LISP, the following format is used. If the character has any control bits set, and the output stream can represent the necessary Greek characters, then the control bits are output as alpha (α) for Control, beta (β) for Meta, lambda (λ) for Hyper, and pi (π) for Super. If the character itself is alpha, beta, lambda, pi, or equivalence-sign (≡), then it is preceded by an equivalence-sign to quote it. After all this, the base character itself is output.

Implementations which do not have Greek characters may well choose to represent control characters by initials and hyphens thus:

```
C-A              C-M-$              H-S-C-#
```

This has the advantage of staying within the standard character set.

~:C spells out the names of the control bits, and represents non-printing characters by their names: "`Control-Meta-F`", "`Control-Return`", "`Space`". This is a "pretty" format for printing characters.

~:@C prints what ~:C would, and then if the character requires unusual shift keys on the keyboard to type it, this fact is mentioned: "`Control-∂ (Top-F)`". This is the format used for telling the user about a key he is expected to type, for instance in prompt messages. The precise output may depend not only on the implementation, but on the particular I/O devices in use.

~@C prints the character in a way that the LISP reader can understand, using "#\" syntax.

Rationale: In some implementations the ~S directive would accomplish this also, but the ~C directive is compatible with LISP dialects which do not have a character data type.

~% Outputs a newline (see `terpri` (page 242)). ~*n*% outputs *n* newlines. No *arg* is used. Simply putting a newline in the control string would work, but ~% is often used because it makes the control string look nicer in the middle of a LISP program.

~& Unless the stream knows that it is already at the beginning of a line, this outputs a newline (see `fresh-line` (page 242)). ~*n*& does a `:fresh-line` operation and then outputs *n*−*1* newlines.

~| Outputs a page separator character, if possible. ~*n*| does this *n* times. With a : modifier, if the output stream supports the `clear-screen` (page CLEAR-SCREEN-FUN) operation this directive clears the screen; otherwise it outputs

page separator character(s) as if no : modifier were present. | is vertical bar, not capital I.

~~ *Tilde.* Outputs a tilde. ~*n*~ outputs *n* tildes.

~<return> Tilde immediately followed by a <return> ignores the <return> and any following non-<return> whitespace. With a :, the <return> is ignored but any following whitespace is left in place. With an @, the <return> is left in place but any following whitespace is ignored. This directive is typically used when a format control string is too long to fit nicely into one line of the program:

```
(defun pet-rock-warning (rock friend amount)
  (unless (equalp rock friend)
          (format t "~&Warning!  Your pet rock ~A just ~
                     bit your friend ~A,~%  and ~
                     ~:[he~;she~] is suing you for $~D!"
                  rock friend (femalep friend) amount)))
(pet-rock-warning "Fred" "Susan" 500)  prints:
Warning:  Your pet rock Fred just bit your friend Susan,
  and she is suing you for $500!
```

~T *Tabulate.* Spaces over to a given column. ~*colnum*,*colinc*T will output sufficient spaces to move the cursor to column *colnum*. If the cursor is already past column *colnum*, it will output spaces to move it to column *colnum*+*k***colinc*, for the smallest non-negative integer *k* possible. *colnum* and *colinc* default to 1.

~:T is like ~T, but *colnum* and *colinc* are in units of pixels, not characters; this makes sense only for streams which can set the cursor position in pixel units.

If for some reason the current column position cannot be determined or set, any ~T operation will simply output two spaces. When `format` is creating a string, ~T will work, assuming that the first character in the string is at the left margin (column 0).

~@T performs *relative* tabulation. ~*colrel*,*colinc*@T is equivalent to ~*curcol*+*colrel*,*colinc*@T where *curcol* is the current output column. If the current output column cannot be determined, however this outputs *colrel* spaces, not two spaces.

~:@T performs relative tabulation in units of pixels instead of columns.

~* The next *arg* is ignored. ~*n** ignores the next *n* arguments.

~:* "ignores backwards"; that is, it backs up in the list of arguments so that the argument last processed will be processed again. ~*n*:* backs up *n* arguments.

When within a ~{ construct (see below), the ignoring (in either direction) is relative to the list of arguments being processed by the iteration.

This is a "relative goto"; for an "absolute goto", see ~G.

~*n*G *Goto.* Goes to the *n*th *arg*, where 0 means the first one. Directives after a ~*n*G will take arguments in sequence beginning with the one gone to.

When within a ~{ construct, the "goto" is relative to the list of arguments being processed by the iteration.

This is an "absolute goto"; for a "relative goto", see ~*.

The format directives after this point are much more complicated than the foregoing; they constitute "control structures" which can perform conditional selection, iteration, justification, and non-local exits. Used with restraint, they can perform powerful tasks. Used with wild abandon, they can produce unreadable and unmaintainable spaghetti with goulash on top.

~[*str0*~;*str1*~;...~;*strn*~]

Conditional expression. This is a set of control strings, called *clauses*, one of which is chosen and used. The clauses are separated by ~; and the construct is terminated by ~]. For example,

```
"~[Siamese~;Manx~;Persian~;Tortoise-Shell~] Cat"
```

The *arg*th clause is selected, where the first clause is number 0. If a prefix parameter is given (as ~*n*[), then the parameter is used instead of an argument (this is useful only if the parameter is specified by "#"). If *arg* is out of range then no clause is selected. After the selected alternative has been processed, the control string continues after the ~].

~[*str0*~;*str1*~;...~;*strn*~:;*default*~] has a default case. If the *last* "~;" used to separate clauses is instead "~:;", then the last clause is an "else" clause, which is performed if no other clause is selected. For example:

```
"~[Siamese~;Manx~;Persian~;Tortoise-Shell~:;Alley~] Cat"
```

~[~*tag00*,*tag01*,...;*str0*~*tag10*,*tag11*,...;*str1*...~] allows the clauses to have explicit tags. The parameters to each ~; are numeric tags for the clause which follows it. That clause is processed which has a tag matching the argument. If ~*a1*,*a2*,*b1*,*b2*,...:; (note the colon) is used, then the following clause is tagged not by single values but by ranges of values *a1* through *a2* (inclusive), *b1* through *b2*, etc. ~:; with no parameters may be used at the end to denote a default clause. For example:

```
"~[~'+,'-,'*,'/;operator ~'A,'Z,'a,'z:;letter ~
  ~'0,'9:;digit ~:;other ~]"
```

~:[*false*~;*true*~] selects the *false* control string if *arg* is `nil`, and selects the *true* control string otherwise.

~@[*true*~] tests the argument. If it is not `nil`, then the argument is not used up by the ~@[command, but remains as the next one to be processed, and the one clause *true* is processed. If the *arg* is `nil`, then the argument is used up, and the clause is not processed. The clause therefore should normally use exactly one argument, and may expect it to be non-`nil`. For example:

```
(setq prinlevel nil prinlength 5)
(format nil "~@[ PRINLEVEL=~D~]~@[ PRINLENGTH=~D~]"
        prinlevel prinlength)
  =>  " PRINLENGTH=5"
```

The combination of ~[and # is useful, for example, for dealing with English conventions for printing lists:

```
(setq foo "Items:~#[ none~; ~S~; ~S and ~
           ~S~:;~@{~#[~1; and~] ~S~^,~}~].")
(format nil foo)
        =>  "Items: none."
(format nil foo 'foo)
        =>  "Items: FOO."
(format nil foo 'foo 'bar)
        =>  "Items: FOO and BAR."
(format nil foo 'foo 'bar 'baz)
        =>  "Items: FOO, BAR, and BAZ."
(format nil foo 'foo 'bar 'baz 'quux)
        =>  "Items: FOO, BAR, BAZ, and QUUX."
```

~; Separates clauses in ~[and ~< constructions. It is undefined elsewhere.

~] Terminates a ~[. It is undefined elsewhere.

~{*str*~} *Iteration.* This is an iteration construct. The argument should be a list, which is used as a set of arguments as if for a recursive call to `format`. The string *str* is used repeatedly as the control string. Each iteration can absorb as many elements of the list as it likes as arguments; if *str* uses up two arguments by itself, then two elements of the list will get used up each time around the loop. If before any iteration step the list is empty, then the iteration is terminated. Also, if a prefix parameter *n* is given, then there will be at most *n* repetitions of processing of *str*. Finally, the ~^ directive can be used to terminate the iteration prematurely.

Here are some simple examples:

```
(format nil "The winners are:~{ ~S~}." '(fred harry jill)
    => "The winners are: FRED HARRY JILL."
(format nil "Pairs:~{ <~S,~S>~}." '(a 1 b 2 c 3))
     => "Pairs: <A,1> <B,2> <C,3>."
```

~:{*str*~} is similar, but the argument should be a list of sublists. At each repetition step one sublist is used as the set of arguments for processing *str*; on the next repetition a new sublist is used, whether or not all of the last sublist had been processed. Example:

```
(format nil "Pairs:~:{ <~S,~S>~}."
            '((a 1) (b 2) (c 3)))
     => "Pairs: <A,1> <B,2> <C,3>."
```

~@{*str*~} is similar to ~{*str*~}, but instead of using one argument which is a list, all the remaining arguments are used as the list of arguments for the iteration. Example:

```
(format nil "Pairs:~@{ <~S,~S>~}."
            'a 1 'b 2 'c 3)
     => "Pairs: <A,1> <B,2> <C,3>."
```

~:@{*str*~} combines the features of ~:{*str*~} and ~@{*str*~}. All the remaining arguments are used, and each one must be a list. On each iteration the next argument is used as a list of arguments to *str*. Example:

```
(format nil "Pairs:~:@{ <~S,~S>~}."
            '(a 1) '(b 2) '(c 3))
     => "Pairs: <A,1> <B,2> <C,3>."
```

Terminating the repetition construct with ~:} instead of ~} forces *str* to be processed at least once even if the initial list of arguments is null (however, it will not override an explicit prefix parameter of zero).

If *str* is empty, then an argument is used as *str*. It must be a string, and precedes any arguments processed by the iteration. As an example, the following are equivalent:

```
(funcall* #'format stream string args)
(format stream "~1{~:}" string args)
```

This will use `string` as a formatting string. The `~1{` says it will be processed at most once, and the `~:}` says it will be processed at least once. Therefore it is processed exactly once, using `args` as the arguments.

As another (rather sophisticated) example, the `format` function itself uses `format-error` (a routine internal to the `format` package) to signal error messages, which in turn uses `ferror`, which uses `format` recursively. Now `format-error` takes a string and arguments, just like `format`, but also prints the control string to `format` (which at this point is available in the variable `ctl-string`) and a little arrow showing where in the processing of the control string the error occurred. The variable `ctl-index` points one character after the place of the error.

```
(defun format-error (string &rest args)
  (ferror nil "~1{~:}~%~VT↓~%~3X\"~A\"~%"
          string args (+ ctl-index 3) ctl-string))
```

This first processes the given string and arguments using `~1{~:}`, then goes to a new line, tabs a variable amount for printing the down-arrow, and prints the control string between double-quotes. The effect is something like this:

```
(format t "The item is a ~[Foo~;Bar~;Loser~]." 'quux)
>>ERROR: The argument to the FORMAT "~[" command
         must be a number.
                    ↓
   "The item is a ~[Foo~;Bar~;Loser~]."
...
```

~} Terminates a ~{. It is undefined elsewhere.

~*mincol*, *colinc*, *minpad*, *padchar*<*str*~>

Justification. This justifies the text produced by processing *str* within a field at least *mincol* columns wide. *str* may be divided up into segments with ~;, in which case the spacing is evenly divided between the text segments.

With no modifiers, the leftmost text segment is left justified in the field, and the rightmost text segment right justified; if there is only one, as a special case, it is right justified. The : modifier causes spacing to be introduced before the first text segment; the @ modifier causes spacing to be added after the last. The *minpad* parameter (default 0) is the minimum number of padding characters to be output between each segment. The padding character is specified by *padchar*, which defaults to the space character. If the total width needed to satisfy these constraints is greater than *mincol*, then the width used is *mincol*+*k***colinc* for the smallest possible non-negative integer value *k*; *colinc* defaults to 1, and *mincol* defaults to 0.

Examples:

```
(format nil "~10<foo~;bar~>")          => "foo    bar"
(format nil "~10:<foo~;bar~>")         => "  foo  bar"
(format nil "~10:@<foo~;bar~>")        => "  foo bar "
(format nil "~10<foobar~>")            => "    foobar"
(format nil "~10:<foobar~>")           => "    foobar"
(format nil "~10@<foobar~>")           => "foobar    "
(format nil "~10:@<foobar~>")          => "  foobar  "
(format nil "$~10,,,'*<~3F~>" 2.59023) => "$******2.59"
```

Note that *str* may include `format` directives. All the clauses in *str* are processed in order; it is the resulting pieces of text that are justified. The last example illustrates how the ~< directive can be combined with the ~F directive to provide more advanced control over the formatting of numbers.

??? **Query:** Unfortunately, the ~F command as defined above isn't really flexible enough?

The ~^ directive may be used to terminate processing of the clauses prematurely, in which case only the completely processed clauses are justified.

If the first clause of a ~< is terminated with ~:; instead of ~;, then it is used in a special way. All of the clauses are processed (subject to ~^, of course), but the first one is not used in performing the spacing and padding. When the padded result has been determined, then if it will fit on the current line of output, it is output, and the text for the first clause is discarded. If, however, the padded text will not fit on the current line, then the text segment for the first clause is output before the padded text. The first clause ought to contain a newline (such as a ~% directive). The first clause is always processed, and so any arguments it refers to will be used; the decision is whether to use the resulting segment of text, not whether to process the first clause. If the ~:; has a prefix parameter *n*, then the padded text must fit on the current line with *n* character positions to spare to avoid outputting the first clause's text. For example, the control string

```
"~%;; ~{~<~%;; ~1:; ~S~>~^,~}.~%"
```

can be used to print a list of items separated by commas, without breaking items over line boundaries, and beginning each line with ";; ". The prefix parameter 1 in ~1:; accounts for the width of the comma which will follow the justified item if it is not the last element in the list, or the period if it is. If ~:; has a second prefix parameter, then it is used as the width of the line, thus overriding the natural line width of the output stream. To make the preceding example use a line width of 50, one would write

```
"~%;; ~{~<~%;; ~1,50:; ~S~>~^,~}.~%"
```

If the second argument is not specified, then `format` uses the line width of the output stream. If this cannot be determined (for example, when producing a string result), then `format` uses 72 as the line length.

~> Terminates a ~<. It is undefined elsewhere.

~^ *Up and out.* This is an escape construct. If there are no more arguments remaining to be processed, then the immediately enclosing ~{ or ~< construct is terminated. If there is no such enclosing construct, then the entire formatting operation is terminated. In the ~< case, the formatting *is* performed, but no more segments are processed before doing the justification. The ~^ should appear only at the *beginning* of a ~< clause, because it aborts the entire clause it appears in (as well as all following clauses). ~^ may appear anywhere in a ~{ construct.

```
(setq donestr "Done.~^  ~D warning~:P.~^  ~D error~:P.")
(format nil donestr) => "Done."
(format nil donestr 3) => "Done.  3 warnings."
(format nil donestr 1 5) => "Done.  1 warning.  5 errors.
```

If a prefix parameter is given, then termination occurs if the parameter is zero. (Hence ~^ is equivalent to ~#^.) If two parameters are given, termination occurs if they are equal. If three are given, termination occurs if the second is between the other two in ascending order. Of course, this is useless if all the prefix parameters are constants; at least one of them should be a # or a V parameter.

If ~^ is used within a ~:{ construct, then it merely terminates the current iteration step (because in the standard case it tests for remaining arguments of the current step only); the next iteration step commences immediately. To terminate the entire iteration process, use ~:^.

Here are some examples of the use of ~^ within a ~< construct.

```
(format nil "~15<~S~;~^~S~;~^~S~>" 'foo)
        =>  "             FOO"
(format nil "~15<~S~;~^~S~;~^~S~>" 'foo 'bar)
        =>  "FOO          BAR"
(format nil "~15<~S~;~^~S~;~^~S~>" 'foo 'bar 'baz)
        =>  "FOO   BAR   BAZ"
```

Compatibility note: The ~Q directive and user-defined directives have been omitted here, as well as control lists (as opposed to strings), which are rumored to be changing in meaning.

21.5. Querying the User

The following functions provide a convenient and consistent interface for asking questions of the user. Questions are printed and the answers are read using the stream `query-io`, which normally is synonymous with `terminal-io` but can be rebound to another stream for special applications.

We describe first two simple functions for asking yes-or-no questions, and then the general function `fquery` on which all querying is built.

`y-or-n-p &optional` *message stream* [*Function*]

This predicate is for asking the user a question whose answer is either "yes" or "no". It types out *message* (if supplied and not `nil`), reads an answer in some implementation-dependent manner (intended to be short and simple, like reading a single character such as "Y"" or "N"), and is true if the answer was "yes" or false if the answer was "no".

If the *message* argument is supplied and not `nil`, it will be printed on a fresh line (see `fresh-line` (page 242)). Otherwise it is assumed that a message has already been printed. If you want a question mark and/or a space at the end of the message, you must put it there yourself; `y-or-n-p` will not add it. *stream* defaults to the value of the global variable `query-io` (page 212).

For example:

```
(y-or-n-p "Cannot establish connection.  Retry?  ")
```

`y-or-n-p` should only be used for questions which the user knows are coming. If the user is unlikely to anticipate the question, or if the consequences of the answer might be grave and irreparable, then `y-or-n-p` should not be used, because the user might type ahead and thereby accidentally answer the question. For such questions as "Shall I delete all of your files?", it is better to use `yes-or-no-p`.

`yes-or-no-p &optional` *message stream* [*Function*]

This predicate, like `y-or-n-p`, is for asking the user a question whose answer is either "Yes" or "No". It types out *message* (if supplied and not `nil`), attracts the user's attention, and reads a reply in some implementation-dependent manner. It is intended that some thought have to go into the reply, such as typing the full word "`yes`" or "`no`" followed by a <return>.

If the *message* argument is supplied, it will be printed on a fresh line (see `fresh-line` (page 242)). Otherwise the caller is assumed to have printed the message already. If you want a question mark and/or a space at the end of the message, you must put it there yourself; `yes-or-no-p` will not add it. *stream* defaults to the value of the global variable `query-io` (page 212).

To allow the user to answer a yes-or-no question with a single character, use `y-or-n-p`. `yes-or-no-p` should be used for unanticipated or momentous questions; this is why it attracts attention and why it requires thought to answer it.

??? **Query:** Maybe `fquery` should be developed inot a more abstract menu sort of interface function? Maybe it belongs in yellow pages?

The preceding two functions allow the asking of simple yes-or-no questions. More complicated questions can be asked using `fquery`, described below. `fquery` is quite general and complicated. It is best to write some interface function for each particular kind of question, using `fquery` in the definition. In this way the complicated arguments to `fquery` need be written in only a few places.

`fquery` *options format-string* `&rest` *format-args* [*Function*]

This asks a question, printed by executing

```
(format query-io format-string format-args...)
```

and returns the answer. `fquery` takes care of checking for valid answers, reprinting the question when the user clears the screen, giving help, and so forth.

options is a list of alternating keywords and values, used to select among a variety of features. Most callers will have a constant list to pass as *options* (rather than consing up a different list each time).

`:type` The expected form of the answer. The types currently defined are:

- `:inch` A single character, as read by `inch` (page 238). This the default.
- `:tyi` This is similar to `inch`; the answer is a single character, but the result is an integer, as if read by `tyi` (page 238).

`:readline` A string, typed as a line terminated by a carriage return, as read by `readline` (page 238).

`:choices` Defines the allowed answers. The allowed forms of choices are complicated and explained below. The default is the same set of choices as for `y-or-n-p` (page 252), if `:type` is `:inch` or `:tyi`, or the same as for `yes-or-no-p`, if `:type` is `:readline`. Note that the `:type` and `:choices` options should be consistent with each other.

Compatibility note: In Lisp Machine LISP, `:choices` always defaults to `y-or-n-p` choices, even if `:type` is `:readline`. This is clearly bogus.

`:list-choices`
If true, the allowed choices are listed (in parentheses) after the question. The default is true; supplying false causes the choices not to be listed unless the user tries to give an answer which is not one of the allowed choices.

`:help-function`
Specifies a function to be called if the user types "?". (Note that other implementation-dependent non-standard keyboard characters might trigger the help function as well, or other actions.) The default help-function simply lists the available choices. Specifying `nil` disables the special treatment of "?". Specifying a function of three arguments (the stream, the value of the `:choices` option, and the type-function) allows smarter help processing. The type-function is a function selected by the `:type` option; it does `inch`, `tyi`, or `readline`, but with additional processing. Often it can be ignored by the help-function.

`:fresh-line` If true (the default), `query-io` is advanced to a fresh line before asking the question. If false, the question is printed wherever the cursor was left by previous typeout.

`:beep` If true, `fquery` beeps to attract the user's attention to the question. The default is false, which means not to beep unless the user tries to give an answer which is not one of the allowed choices.

`:clear-input` If true, `fquery` throws away type-ahead before reading the user's response to the question. Use this for unexpected questions. The default is false, which means not to throw away typeahead unless the user tries to give an answer which is not one of the allowed choices. In that case, type-ahead is discarded since the user probably wasn't expecting the question.

The argument to the `:choices` option is a list each of whose elements is a *choice*. The cdr of a choice is a list of the user inputs which correspond to that choice. These should be characters if the `:type` is `:inch`, integers corresponding to characters for `:tyi`, or strings for `:readline`. The car of a choice is either an atom which `fquery` should return if the user answers with that choice (in which case nothing is echoed), or a list whose first element is such an atom and whose second element is the string to be echoed when the user selects the choice.

Compatibility note: In Lisp Machine LISP the choice-value is specified to be a symbol. To allow `nil` to be returned, or even integers, atoms (non-lists) are specified here.

In most cases a `:type` of `:readline` would use the first format, since the user's input has already been echoed, and `:inch` or `:tyi` would use the second format, since the input has not been

echoed and furthermore is a single character, which would not be mnemonic to see on the display.

As an example, here is a definition of the function `y-or-n-p` in terms of `fquery`:

```
(defun y-or-n-p (&optional message (stream query-io))
  (let ((query-io stream))
    (fquery '(:fresh-line nil
              :list-choices nil
              :choices
                  (((t "Yes.") #\y #\t #\Space)
                   ((nil "No.") #\n #\Rubout)))
            (if message "~&~a" "~*")
            message)))
```

As another example, here is a definition of `yes-or-no-p`:

```
(defun y-or-n-p (&optional message (stream query-io))
  (let ((query-io stream))
    (fquery '(:fresh-line nil
              :list-choices nil
              :beep t
              :type :readline
              :choices ((t "Yes") (nil "No")))
            (if message "~&~a" "~*")
            message)))
```

As a third example, this function allows more complex choices. One may type `P`, `Q`, `R`, or `D`, in which respective cases the symbol `proceed`, `quit`, `retry`, or `debug` is returned. Space or rubout may be typed instead of `P` or `Q`, respectively.

```
(defun i/o-error-query &optional message (stream query-io))
  (let ((query-io stream))
    (fquery '(:fresh-line nil
              :list-choices t
              :choices
                  (((proceed "Proceed") #\p #\Space)
                   ((quit "Quit") #\q #\Rubout)
                   ((retry "Retry") #\r)
                   ((debug "Debug") #\d)))
            (if message "~&~a" "~*")
            message)))
```

Chapter 22

File System Interface

A frequent use of streams is to communicate with a *file system* to which groups of data (files) can be written and from which files can be retrieved.

COMMON LISP defines a standard interface for dealing with such a file system. This interface is designed to be simple and general enough to accommodate the facilities provided by "typical" operating system environments within which COMMON LISP is likely to be implemented. The goal is to make COMMON LISP programs that perform only simple operations on files reasonably portable.

To this end COMMON LISP assumes that files are named, that given a name one can construct a stream connected to a file of that name, and that the names can be fit into a certain canonical, implementation-independent form called a *pathname.*

Facilities are provided for manipulating pathnames, for creating streams connected to files, and for manipulating the file system through pathnames and streams.

22.1. File Names

COMMON LISP programs need to use names to designate files. The main difficulty in dealing with names of files is that different file systems have different naming formats for files. For example, here is a table of several file systems (actually, operating systems that provide file systems) and what the "same" file name might look like for each one:

System	File name
TOPS-20	`<LISPIO>FORMAT.FASL.13`
TOPS-10	`FORMAT.FAS[1,4]`
ITS	`LISPIO;FORMAT FASL`
MULTICS	`>udd>LispIO>format.fasl`
TENEX	`<LISPIO>FORMAT.FASL;13`
UNIX	`/usr/lispio/format.fasl`

It would be impossible for each program that deals with file names to know about each different file name format that exists; a new COMMON LISP implementation might use a format different from any of its predecessors. Therefore COMMON LISP provides *two* ways to represent file names: *namestrings*, which are

strings in the implementation-dependent form customary for the file system, and *pathnames*, which are special data objects that represent file names in an implementation-independent way. Functions are provided to convert between these two representations, and all manipulations of files can be expressed in machine-independent terms by using pathnames.

In order to allow COMMON LISP programs to operate in a network environment that may have more than one kind of file system, the pathname facility allows a file name to specify which file system is to be used. In this context, each file system is called a *host*, in keeping with the usual networking terminology.

22.1.1. Pathnames

All file systems dealt with by COMMON LISP are forced into a common framework, in which files are named by a LISP data object of type `pathname`.

A pathname always has six components, described below. These components are the common interface that allows programs to work the same way with different file systems; the mapping of the pathname components into the concepts peculiar to each file system is taken care of by the COMMON LISP implementation.

host	The name of the file system on which the file resides.
device	Corresponds to the "device" or "file structure" concept in many host file systems: the name of a (logical or physical) device containing files.
directory	Corresponds to the "directory" concept in many host file systems: the name of a group of related files (typically those belonging to a single user or project).
name	The name of a group of files which can be thought of as conceptually the "same" file.
type	Corresponds to the "filetype" or "extension" concept in many host file systems. This says what kind of file this is. Files with the same name but different type are usually related in some specific way, such as one being a source file, another the compiled form of that source, and a third the listing of errors messages from the compiler.
version	Corresponds to the "version number" concept in many host file systems. typically a number that is incremented every time the file is modified.

Note that a pathname is not necessarily the name of a specific file. Rather, it is a specification (possibly only a partial specification) of how to access a file. A pathname need not correspond to any file that actually exists, and more than one pathname can refer to the same file. For example, the pathname with a version of "newest" may refer to the same file as a pathname with the same components except a certain number as the version. Indeed, a pathname with version "newest" may refer to different files as time passes, because the meaning of such a pathname depends on the state of the file system. In file systems with such facilities as "links", multiple file names, logical devices, and so on, two pathnames that look quite different may turn out to address the same file. To access a file given a pathname one must do a file system operation such as `open`

(page 268).

Two important operations involving pathnames are *parsing* and *merging*. Parsing is the conversion of a namestring (which might be something supplied interactively by the user when asked to supply the name of a file) into a pathname object. This operation is implementation-dependent, because the format of namestrings is implementation-dependent. Merging takes a pathname with missing components and supplies values for those components from a source of defaults.

Not all of the components of a pathname need to be specified. If a component of a pathname is missing, its value is `nil`. Before the file system interface can do anything interesting with a file, such as opening the file, all the missing components of a pathname must be filled in (typically from a set of defaults). Pathnames with missing components may used internally for various purposes; in particular, parsing a namestring that does not specify certain components will result in a pathname with missing components. However, the host component is not allowed to be missing from any pathname; since the behavior of a pathname is host-dependent to some extent.

??? **Query:** Is `:unspecific` really needed over and above `nil`?

A component of a pathname can also be the keyword `:unspecific`. This means that the component has been explicitly determined not to be there, as opposed to being missing. One way this can occur is with *generic* pathnames, which refer not to a file but to a whole family of files. The version, and usually the type, of a generic pathname are `:unspecific`. Another way `:unspecific` is used is to represent components that are simply not supported by a file system. When a pathname is converted to a namestring, `nil` and `:unspecific` both cause the component not to appear in the string. When merging, however, a `nil` value for a component will be replaced with the default for that component, while `:unspecific` will be left alone.

A component of a pathname can also be the special symbol `:wild`. This is only useful when the pathname is being used with a directory-manipulating operation, where it means that the pathname component matches anything. The printed representation of a pathname typically designates `:wild` by an asterisk; however, this is host-dependent.

What values are allowed for components of a pathname depends, in general, on the pathname's host. However, in order for pathnames to be usable in a system-independent way certain global conventions are adhered to. These conventions are stronger for the type and version than for the other components, since the type and version are explicitly manipulated by many programs, while the other components are usually treated as something supplied by the user which just needs to be remembered and copied from place to place.

The type is always a string or `nil`, `:unspecific`, or `:wild`. Many programs that deal with files have an idea of what type they want to use.

The version is either a positive integer or a special symbol. The meanings of `nil`, `:unspecific`, and `:wild` have been explained above. The keyword `:newest` refers to the largest version number that already exists in the file system when reading a file, or that number plus one when writing a new file. The keyword

`:oldest` refers to the smallest version number that exists. Some COMMON LISP implementations may choose to define other special version symbols, such as `:installed`, for example, if the file system for that implementation will support them.

The host may be a string, indicating a file system, or a list of strings, of which the first names the file system and the rest may be used for such a purpose as inter-network routing.

The device, directory, and name also can each be a simple string (with host-dependent rules on allowed characters and length) or a list of strings (in which case such a component is said to be *structured*). Structured components are used to handle such file system features as hierarchical directories. COMMON LISP programs do not need to know about structured components unless they do host-dependent operations. Specifying a string as a pathname component for a host that requires a structured value will cause conversion of the string to the appropriate form. Specifying a structured component for a host that does not provide for that component to be structured causes conversion to a string by the simple expedient of taking the first element of the list and ignoring the rest.

Some host file systems have features that do not fit into this pathname model. For instance, directories might be accessible as files, there might be complicated structure in the directories or names, or there might be relative directories, such as the "<" syntax in MULTICS or the special ".." file name of UNIX. Such features are not allowed for by the standard COMMON LISP file system interface. An implementation is free to accommodate such features in its pathname representation and provide a parser that can process such specifications in namestrings; such features are then likely to work within that single implementation. However, note that once your program depends explicitly on any such features, it will not be portable.

22.1.2. Pathname Functions

These functions are what programs use to parse and default file names that have been typed in or otherwise supplied by the user.

As a rule, any argument called *pathname* may actually be a pathname, a string or symbol, or a stream, and any argument called *defaults* may be a pathname, a string or symbol, a stream, or a *pathname defaults a-list.*

In the examples, it is assumed that the host named CMUC runs the TOPS-20 operating system, and therefore uses TOPS-20 file system syntax; furthermore, an explicit host name is indicated by following it with a double colon. Remember, however, that namestring syntax is implementation-dependent, and this syntax is used purely for the sake of examples.

`pathname` *thing* [*Function*]
`truename` *thing* [*Function*]

The `pathname` function converts its argument to be a pathname. The argument may be a pathname, a string or symbol, or a stream.

The `truename` function behaves identically to `pathname`, with one exception. If the argument is

a stream connected to a file in a file system, then the pathname returned by `truename` reflects the "true" name of the file according to the file system, as opposed to the name originally given to the file system to specify the file (which is what `pathname` will return). Thus the `truename` function may be used to account for any file-name translations performed by the file system, as opposed to logical-pathname translations performed by COMMON LISP (see `translated-pathname` (page 266)).

For example, suppose that "`DOC:`" is a TOPS-20 logical device name that is translated by the TOPS-20 file system to be "`PS:<DOCUMENTATION>`".

```
(setq file (open "CMUC::DOC:DUMPER.HLP"))
(namestring (pathname file)) => "CMUC::DOC:DUMPER.HLP"
(namestring (truename file))
   => "CMUC::PS:<DOCUMENTATION>DUMPER.HLP.13"
```

`parse-namestring` *thing* `&optional` *convention defaults break-characters start end* [*Function*]

This turns *thing* into a pathname. The *thing* is usually a string (that is, a namestring), but it may be a symbol (in which case the print name is used) or a pathname or stream (in which case no parsing is needed, but an error check may be made for matching hosts).

This function does *not* do defaulting; it only does parsing. The *convention* and *defaults* arguments are present because in some implementations it may be that a namestring can only be parsed with reference to a particular file name syntax of several available in the implementation. If *thing* does not contain a manifest host name, then if *convention* is non-`nil`, it must be a string naming the file name syntax (using a host name will indicate that the conventions peculiar to that host should be used if that is meaningful), or a list of strings, of which the first is used. If *host* is `nil` then the host name is obtained from the default pathname in *defaults* and used to determine the syntax convention.

For a string (or symbol) argument, `parse-namestring` parses a file name within it in the range delimited by *start* and *end* (which are integer indices into *string*, defaulting to the beginning and end of the string). Parsing is terminated upon reaching the end of the specified substring or upon reaching a character in *break-characters*, which may be a string or a list of characters; this defaults to an empty set of characters.

Two values are returned by `parse-namestring`. If the parsing is successful, then the first value is a pathname object for the parsed file name, and otherwise the first value is `nil`. The second value is an integer, the index into *string* one beyond the last character processed. This will be equal to *end* if processing was terminated by hitting the end of the substring; it will be the index of a break character if such was the reason for termination; it will be the index of an illegal character if that was what caused processing to (unsuccessfully) terminate. If *thing* is not a string or symbol, then *start* is always returned as the second value.

Parsing an empty string always succeeds, producing a pathname with all components (except the host) `:unspecific`.

Note that if *host* is specified and not `nil`, and *thing* contains a manifest host name, an error is

signalled if they are not the same host.

`merge-pathname-defaults` *pathname* `&optional` *defaults default-type default-version* [*Function*]

This is the function that most programs should call to process a file name supplied by the user. It fills in unspecified components of *pathname* from the defaults, and returns a new pathname. *pathname* can be a pathname, string, or symbol. The returned value will always be a pathname.

defaults defaults to the value of `default-pathname-defaults` (page 265). *default-type* defaults to `:unspecific`. *default-version* defaults to `:newest`.

The rules for merging can be rather complicated in some situations; they are described in detail in section ???. An approximate rule of thumb is simply that any components missing in the pathname are filled in from the defaults.

For example:

```
(merge-pathname-defaults "CMUC::FORMAT"
                         "CMUC::PS:<LISPIO>"
                         "FASL")
  => a pathname object which re-expressed as a namestring would be
     "CMUC::PS:<LISPIO>FORMAT.FASL.0"
```

`make-pathname &key :host :device :directory :name` [*Function*]
`:type :version :defaults`

Given some components, `make-pathname` constructs and returns a pathname. Missing components default to `nil`, except the host (all pathnames must have a host). The `:defaults` option specifies what defaults to get the *host* from if the `:host` option is `nil` or not specified; however, no other components are supplied from the `:defaults`. The default value of the `:defaults` option is the value of `default-pathname-defaults` (page 265). All other keywords specify components for the pathname.

Whenever a pathname is constructed, whether by `make-pathname` or some other function, the components may be canonicalized if appropriate. For example, if a file system is insensitive to case, then alphabetic characters may be forced to upper case or lower case by the implementation.

`pathnamep` *object* [*Function*]

This predicate is true if *object* is a pathname, and otherwise is false.

```
(pathnamep x) <=> (typep x 'pathname)
```

`pathname-host` *pathname* [*Function*]
`pathname-device` *pathname* [*Function*]
`pathname-directory` *pathname* [*Function*]
`pathname-name` *pathname* [*Function*]
`pathname-type` *pathname* [*Function*]
`pathname-version` *pathname* [*Function*]

These return the components of the argument *pathname*, which may be a pathname, string, or symbol. The returned values can be strings, special symbols, or lists of strings in the case of structured components. The type will always be a string or a symbol. The version will always be a number or a symbol.

`pathname-plist` *pathname* [*Function*]

These return the property list of the argument *pathname*, which may be a pathname, string, or symbol (see `plist` (page 103)).

`namestring` *pathname* [*Function*]
`file-namestring` *pathname* [*Function*]
`directory-namestring` *pathname* [*Function*]
`host-namestring` *pathname* [*Function*]
`enough-namestring` *pathname* `&optional` *defaults* [*Function*]

The *pathname* argument may be a namelist, a namestring, or a stream which is or was open to a file. The name represented by *pathname* is returned as a namelist in canonical form.

If *pathname* is a stream, the name returned represents the name used to *open* the file, which may not be the *actual* name of the file (see `truename` (page 260)).

`namestring` returns the full form of the *pathname* as a string. `file-namestring` returns a string representing just the *name*, *type*, and *version* components of the *pathname*; the result of `directory-namestring` represents just the *directory-name* portion; and `host-namestring` returns a string for just the *host-name* portion. Note that a valid namestring cannot necessarily be constructed simply by concatenating some of the three shorter strings in some order.

`enough-namestring` takes another argument, *defaults*. It returns an abbreviated namestring which is just sufficient to identify the file named by *pathname* when considered relative to the *defaults* (which defaults to the value of `default-pathname-defaults` (page 265)). That is,

```
(merge-pathname-defaults (enough-namestring pathname defaults)
                         defaults)
  <=> (parse-pathname pathname)
```

These functions return useful information.

`user-homedir-pathname` `&optional` *host* [*Function*]

Returns a pathname for the user's "home directory" on *host*, which defaults in some appropriate implementation-dependent manner. . The concept of "home directory" is itself somewhat implementation-dependent, but from the point of view os COMMON LISP it is the directory where the user keeps personal files such as initialization files and mail. This function returns a pathname without any name, type, or version component (those components are all `nil`).

`init-file-pathname` *program-name* `&optional` *host* [*Function*]

Returns the pathname of the user's init file for the program *program-name* (a string), on the *host*, which defaults in some appropriate implementation-dependent manner. Programs that load init files containing user customizations call this function to determine where to look for the file, so that they need not know the separate init file name conventions of each host operating system.

22.1.3. Defaults and Merging

Defaulting of pathname components is done by filling in components taken from another pathname; this filling-in is called *merging*. This is especially useful for cases such as a program that has an input file and an output file, and asks the user for the name of both, letting the unsupplied components of one name default from the other. Unspecified components of the output pathname will come from the input pathname, except that the type should default not to the type of the input but to the appropriate default type for output from this program.

The pathname merging operation takes as input a given pathname, a defaults pathname a default type, and a default version, and returns a new pathname. Basically, the missing components in the given pathname are filled in from the defaults pathname, except that if no type is specified the default type is used, and if no version is specified the default version is used. By default, the default type is `:unspecific`, meaning that if the input pathname has no type, the user really wants a file with no type. Programs that have a default type for the files they manipulate will supply it to the merging operation. The default version is usually `:newest`; if no version is specified the newest version in existence should be used. The default type and version can be `nil`, to preserve the information that they were missing in the input pathname.

The full details of the merging rules are as follows. First, if the given pathname explicitly specifies a host and does not supply a device, then the device will be the default file device for that host. Next, if the given pathname does not specify a host, device, directory, or name, each such component is copied from the defaults.

The merging rules for the type and version are more complicated, and depend on whether the pathname specifies a name. If the pathname doesn't specify a name, then the type and version, if not provided, will come from the defaults, just like the other components. However, if the pathname does specify a name, then the type and version are not affected by the defaults. The reason for this is that the type and version "belong to" some other filename, and are unlikely to have anything to do with the new one. Finally, if this process leaves the type or version missing, the default type or default version is used (these were inputs to the merging operation).

The effect of all this is that if the user supplies just a name, the host, device, and directory will come from the defaults, but the type and version will come from the default type and default version arguments to the merging operation. If the user supplies nothing, or just a directory, the name, type, and version will come over from the defaults together. If the host's file name syntax provides a way to input a type or version without a name, the user can let the name default but supply a different type or version than the one in the defaults.

The following special variables are parts of the pathname interface that are relevant to defaults.

`default-pathname-defaults` [*Variable*]

This is the default pathname defaults pathname; if any pathname primitive that needs a set of defaults is not given one, it uses this one. As a general rule, however, each program should have its own pathname defaults a-list rather than using this one.

`load-pathname-defaults` [*Variable*]

This is the pathname defaults pathname for the `load` (page 270) and `comfile` (page 280) functions. Other functions may share these defaults if they deem that to be an appropriate user interface.

22.1.4. Logical Pathnames

Logical pathnames, unlike ordinary pathnames, do not correspond to any particular file server. Like every pathname, however, a logical pathname must have a host, in this case called a "logical" host. Every logical pathname can be translated into a corresponding "actual" pathname; there is a mapping from logical hosts into actual hosts used to effect this translation.

The reason for having logical pathnames is to make it easy to keep bodies of software on more than one file system. A program may need to have a suite of files at its disposal, but different file systems may have different conventions about what directories may be used to store such files. Ideally, it should be easy to write a program in such a way that it will work correctly no matter which site it is run at. This is easily done by writing the program to use a logical name; this logical name can then be provided with a customized translation for each implementation, thereby centralizing the implementation dependency.

Here is how translation is done. For each logical host, there is a mapping that takes a directory name and produces a corresponding actual host name, device name, and directory name. To translate a logical pathname, the system finds the mapping for that pathname's host and looks up that pathname's directory in the mapping. If the directory is found, a new pathname is created whose host is the actual host, and whose device and directory names come from the mapping. The other components of the new pathname taken from the old pathname. There is also, for each logical host, a "default device". If the directory is not found in the mapping, then the new pathname will have the same directory name as the old one, and its device will be the default device for the logical host.

This means that when you invent a new logical device for a certain set of files, you also make up a set of logical directory names, one for each of the directories that the set of files is stored in. Now when you create the mappings at particular sites, you can choose any actual host for the files to reside on, and for each of your logical directory names, you can specify the actual directory name to use on the actual host. This gives you flexibility in setting up your directory names; if you used a logical directory name called `fred` and you want to move your set of files to a new file server that already has a directory called `fred`, being used by someone else, you can translate `fred` to some other name and so avoid getting in the way of the existing directory.

Furthermore, you can set up your directories on each host to conform to the local naming conventions of that host.

`add-logical-pathname-host` *logical-host actual-host default-device translations* [*Function*]

This creates a new logical host named *logical-host*. Its corresponding actual host (that is, the host to which it will forward most operations) is named by *actual-host*. *logical-host* and *actual-host* should both be strings. The *default-device* should be a string naming the default device for the logical host. The *translations* should be a list of translation specifications. Each translation specification should be a list of two items. The first should be a string naming a directory for the logical host. The second is a pathname (or string, symbol, or stream) whose device component and directory component provide the translation for the logical directory.

Compatibility note: Lisp Machine LISP does this: "The default device for the logical host will be the device of the first translation specification." This seems a bit of a crock.

`translated-pathname` *pathname* [*Function*]

This converts a logical pathname to an actual pathname. If the *pathname* already refers to an actual host rather than to a logical host, the argument is simply returned.

`back-translated-pathname` *logical-pathname actual-pathname* [*Function*]

This converts an actual pathname to a logical pathname. *actual-pathname* should be a pathname whose host is the actual host corresponding to the logical host of *logical-pathname*. This returns a pathname whose host is the logical host and whose translation (as by `translated-pathname` (page 266)) is *pathname*.

An example of how this would be used is in connection with truenames. Given a stream `s` that was obtained by opening a logical pathname,

```
(pathname s)
```

returns the logical pathname that was opened;

```
(truename s)
```

returns the true name of the file that is open, which of course is a pathname on the actual host. To get this in the form of a logical pathname, one would do

```
(back-translated-pathname (pathname s) (truename s))
```

If the argument *logical-pathname* is actually an actual pathname, then the argument *actual-pathname* is simply returned. Thus the above example will work no matter what kind of pathname was opened to create the stream.

The namestring corresponding to a logical pathname is, like all namestrings, of implementation-dependent format. As a rule, however, there is no way to specify a device; parsing a logical pathname always returns a pathname whose device component is `:unspecific`.

22.2. Opening and Closing Files

When a file is *opened*, a stream object is constructed to serve as the file system's ambassador to the LISP environment; operations on the stream are reflected by operations on the file in the file system. The act of *closing* the file (actually, the stream) ends the association; the transaction with the file system is terminated, and input/output may no longer be performed on the stream. The stream function `close` (page 214) may be used to close a file; the functions described below may be used to open them.

`with-open-file` *bindspec* {*form*}* [*Special form*]

`(with-open-file (`*stream filename . options*`) . `*body*`)` evaluates the forms of *body* (an implicit `progn`) with the variable *stream* bound to a stream which reads or writes the file named by the value of *filename*. The *options* should evaluate to a keyword or list of keywords.

When control leaves the body, either normally or abnormally (such as by use of `throw` (page 87)), the file is closed. If a new output file is being written, and control leaves abnormally, the file is aborted and it is, so far as possible, as if it had never been opened. Because `with-open-file` always closes the file, even when an error exit is taken, it is preferred over `open` for most applications.

filename is the name of the file to be opened; it can be a namelist or namestring. If an error occurs (such as "File Not Found"), the user is asked to supply an alternate pathname, unless this is overridden by *options*. At that point the user can quit or enter the error handler if the error was not due to a misspelled pathname after all.

options is either a single keyword or a (possibly empty) list of options, where an option is either a keyword or a list of a keyword and arguments to that keyword. (If a keyword with an argument is to be used, then *options* must be a list of options and not a single option.)

Compatibility note: Lisp Machine LISP uses a format where the argument to a keyword simply follows the keyword in the list. This is not compatible with other keyword formats, for example that of `defstruct`. It only makes a difference here in the case of `:byte-size`. It seems worthwhile to minimize the number of keyword formats in COMMON LISP.

??? **Query:** Could we get this to conform to the standard keyword-pairs format?

Valid keywords are:

`:in` or `:input` or `:read`
: Open file for input. This is the default.

`:out` or `:output` or `:write` or `:print`
: Open file for output; a new file is to be created.

`:append`
: Open an existing file for output, arranging that output to the resulting stream should be appended to the previous contents of the file.

Compatibility note: Not all file systems can support this operation. An implementation may choose to simulate it by copying the old file into a new one and then continuing to write the new one.

Compatibility note: The Lisp Machine LISP implementation appears not to support this, but MACLISP does in the `open` function.

`:read-alter` Open a file in read-alter mode; the result is a stream which can perform both input and output on a random-access file.

Compatibility note: Not all file systems can support this operation.

`:character` or `:ascii`
The unit of transaction is a character; the file is a text file. This is the default.

`:binary` or `:fixnum`
The unit of transaction is a small unsigned integer. The `:byte-size` option may be used to specify the number of binary bits in the transaction unit. This precise way in which this interacts with the file system is implementation-dependent.

`:byte-size` This keyword takes an argument, an integer specifying the number of bits per transaction unit; this is used in conjunction with the `:binary` option. If the `:binary` keyword is specified but the `:byte-size` keyword is not, then an implementation-dependent "natural" byte size is used.

`:echo` This keyword requires an argument, an output stream, and is valid only when opening a stream for input. The result stream will echo everything read from it onto the output stream.

`:probe` This keyword specifies that the file is not being opened to do I/O, but only to find out information about it. A stream is returned, but it cannot handle I/O transactions; it is as if the stream were immediately closed after opening it. `:probe` implies `:noerror` (see below).

??? Query: In Lisp Machine Lisp, `:probe` also implies `:fixnum`. Why??

`:noerror` If the file cannot be opened, then instead of returning a stream, a string containing the error message is returned. If `:noerror` is not specified, then an error is signalled using the error message, and the user is asked to supply a different filename.

`open` *filename* `&optional` *options* [*Function*]

Returns a stream which is connected to the file specified by *filename*. The *options* argument is as for `with-open-file` (page 267). If an error occurs, such as "File Not Found", the user is asked to supply an alternate pathname, unless the `:noerror` (page 268) option is used, in which case the error message is returned as a string.

When the caller is finished with the stream, it should close the file by using the `close` (page 214) function. The `with-open-file` (page 267) special form does this automatically, and so is preferred for most purposes. `open` should be used only when the control structure of the program necessitates opening and closing of a file in some way more complex than provided by `with-open-file`. It is suggested that any program which uses `open` directly should use the special form `unwind-protect` (page 86) to close the file if an abnormal exit occurs.

Implementation note: While `with-open-file` tries to automatically close the stream on exit from the construct, for robustness it is helpful if the garbage collector can detect discarded streams and automatically close them.

22.3. Renaming, Deleting, and Other Operations

Compatibility note: The MACLISP/Lisp Machine LISP names `renamef`, `deletef`, etc., are explicitly avoided here because they are not sufficiently mnemonic and because the trailing-`f` convention conflicts with a similar convention for forms related to `setf` (page 60).

`rename-file` *file new-name* `&optional` *error-p* [*Function*]

file can be a filename or a stream which is open to a file. The specified file is renamed to *new-name* (a filename). If *error-p* is true (the default), then if a file-system error occurs it will be signalled as a LISP error. If *error-p* is false and an error occurs, the error message will be returned as a string. If no error occurs, `renamef` returns `nil`.

`delete-file` *file* `&optional` *error-p* [*Function*]

file can be a filename or a stream which is open to a file. The specified file is deleted. If *error-p* is true (the default), then if a file-system error occurs it will be signalled as a LISP error. If *error-p* is false and an error occurs, the error message will be returned as a string. If no error occurs, `deletef` returns `nil`.

`probe-file` *filename* [*Function*]

This pseudo-predicate is false if there is no file named *filename*, and otherwise returns a filename which is the true name of the file (which may be different from *filename* because of file links, version numbers, or other artifacts of the file system; see `truename` (page 260)).

`file-creation-date` *file* [*Function*]

file can be a filename or a stream which is open to a file. This returns the creation date of the file as an integer in universal time format, or `nil` if this cannot be determined.

`file-author` *file* [*Function*]

file can be a filename or a stream which is open to a file. This returns the name of the author of the file as a string, or `nil` if this cannot be determined.

`filepos` *file-stream* `&optional` *position* [*Function*]

`filepos` returns or sets the current position within a random-access file.

(`filepos` *file-stream*) returns a non-negative integer indicating the current position within the *file-stream*, or `nil` if this cannot be determined. Normally, the position is zero when the stream is first created. For a `:character` (page 268) stream, the position is in units of characters; for a `:binary` (page 268) file, the position is in bytes.

(`filepos` *file-stream position*) sets the position within *file-stream* to be *position*. The *position* may be an integer, or `nil` for the beginning of the stream, or `t` for the end of the stream. If the integer is too large, an error occurs (the `file-length` (page 270) function returns the length

beyond which `filepos` may not access). With two arguments, `filepos` is a (side-effecting) predicate that is true if it actually performed the operation, or false if it could not.

`file-length` *file-stream* [*Function*]

file-stream must be a stream which is open to a file. The length of the file is returned as a non-negative integer, or `nil` if the length cannot be determined. For a `:character` (page 268) stream, the position is in units of characters; for a `:binary` (page 268) file, the position is in bytes.

22.4. Loading Files

To *load* a file is to read through the file, evaluating each form in it. Programs are typically stored in files; the expressions in the file are mostly special forms such as `defun` (page 42), `defmacro` (page 91), and `defvar` (page 43) which define the functions and variables of the program.

Loading a compiled ("fasload") file is similar, except that the file does not contain text, but rather pre-digested expressions created by the compiler which can be loaded more quickly.

`load` *filename* `&key :package :verbose :noerror :preserve-defaults` [*Function*]

This function loads the file named by *filename* into the Lisp environment. It is assumed that a text (character file) can be automatically distinguished from an object (binary) file by some appropriate implementation-dependent means, possibly by the file type. If the *filename* does not explicitly specify a type, and both text and object types of the file are available in the file system, `load` should try to select the more appropriate file by some implementation-dependent means.

The `:package` keyword argument can be used to specify the package into which to load the file; it can be either a package or the name of a package as a string or a symbol. If `:noerror` is specified and not `nil`, `load` just returns `nil` if the file cannot be opened. If the file is successfully loaded, `load` always returns a non-`nil` value.

`load` maintains a default filename in the variable `load-pathname-defaults` (page 265), used to default missing components of the *filename* argument; thus `(load "")` will load the same file previously loaded. Normally `load` updates the filename defaults from *filename*, but if `:preserve-defaults` is specified and not `nil` this is suppressed.

22.5. Accessing Directories

*** still missing ***

Chapter 23

Errors

*** This chapter still needs a lot of work! ***

*** right now this is completely wrong ***

COMMON LISP handles errors through a system of *conditions*. One may establish *handlers* which gain control which conditions occur, and *signal* a condition when an error actually occurs. When the system or a user function detects an error it signals an appropriately named condition and some handler established for that condition will deal with it.

The condition mechanism is completely general and could be used for purposes other than "error" handling. There are some functions supplied in COMMON LISP which make use of the condition mechanism to handle errors in a convenient way.

23.1. Signalling Conditions

Condition handlers are associated with conditions (see next section). Every condition is essentially a name, which is a symbol (possibly `nil`). When an unusual situation occurs, such as an error, a condition is signalled. The system (essentially) searches up the stack of nested function invocations looking for a handler establised to handle that condition. The handler is a function which gets called to deal with the condition.

`signal` *condition-name* `&rest` *args* [*Function*]

This searches through all currently established condition handlers, perhaps twice, starting with the most recent. If it finds one which was established to handle `nil` or *condition-name*, then it calls that handler with a first argument of *condition-name* and with *args* as the rest of the arguments. If the first value returned by the handler is `nil`, `signal` will continue searching for another handler; otherwise it will return the first two values returned by the handler. If *condition-name* is not `t`, and if no handler was willing to handle the condition, then a second pass of the established condition handlers is made, searching for any handler established to handle `t`. If one is found that is used in the same manner as in the first pass search. If there is still no willing handler found then `signal` returns `nil`.

Thus a handler set up to handle condition `nil` will handle *all* conditions which are not handled by a more recently established handler. This is intended to make it easy to set up a debugger which intercepts all errors and handles them itself. Note that such a handler doesn't have to actually handle all conditions; it will be offered the chance to do so but can return `nil` to refuse any condition which it doesn't wish to handle.

Conditions established to handle condition `t` will handle *any* condition for which a more specific willing handler can't be found. This makes it easy to set up, at any time, a handler which which will be given a chance to handle all conditions that no one else wants.

23.2. Establishing Handlers

Condition handlers are established through the `condition-bind` or `condition-setq` special forms. These have behaviors somewhat analagous to `let` and `setq`. They make use of the ordinary variable binding mechanism, so that if a `condition-bind` is thrown through the handlers get disestablished. It also means that in multiple stack group implementations of COMMON LISP the handlers are established only in the current stack group.

`condition-bind` *bindings* {*form*}* [*Special form*]

This is used to establish handlers for conditions then perform the *body* in that established handler environment.

For example:

For example:

```
(condition-bind ((cond1 hand1)
                 (cond2 hand2)
                 ...
                 (condm handm))
  form1
  form2
  ...
  formn)
```

Each *condj* is either the name of a condition or a list of names of conditions. Each *handj* is a form which is evaluated to produce a handler function. No guarantee is made on the order of evaluation of these forms, but the conditions are established in sequential order, so that *cond1* will be looked at first. The expressions *formj* are then evaluated in order; the values of all but the last are discarded (that is, the body of the `condition-bind` form is an implicit `progn`). The value of the `condition-bind` form is the value of *formn* (if the body is empty, `nil` is the value). The established conditions become disestablished when the `condition-bind` form is exited.

As an example consider:

For example:

```
(condition-bind ((:wrong-type-argument 'some-wta-handler)
                 ((silliness-1 silliness-2) silliness-handler))
  (format msgfiles '"Yodle-eh-eh-ho!")
  (+ 23 nil))
```

This sets up the function `some-wta-handler` to handle the `:wrong-type-argument` condition. The value of the symbol `silliness-handler` is set up to handle both the `silliness-1` and `silliness-2` conditions. With these handlers set up, it outputs a message and then causes a `:worng-type-argument` error by trying to add 23 to `nil`, which is not a number. The condition handler `some-wta-handler` will be given a chance to handle the error.

`condition-setq` {*spec*}* [*Special form*]

The `condition-setq` form is used to establish condition handlers as a side effect of some operation -- for instance loading a file which contains condition handlers and a `condition-setq` form to establish them.

It takes the form:

For example:

```
(condition-setq cond1 hand1
                cond2 hand2
                ...
                condn handn)
```

Each *condj* is either the name of a condition or a list of names of conditions. Each *handj* is a form which is evaluated to produce a handler function. No guarantee is made on the order of evaluation of these forms, but the conditions are established in sequential order, so that *cond1* will be looked at first.

The conditions established by `condition-setq` remain established until execution is unwound (either normally or by being thrown) past the most recent `condition-bind`. Multiple uses of `condition-setq` cause the most recently established handler to be tried first when a condition is signalled. For example, consider:

For example:

```
(condition-setq :wrong-type-argument 'default-wta-handler)
(+ 23 nil)
(condition-setq :worng-type-argument 'hairy-wta-handler)
(+ 105 nil)
```

When the first `:wrong-type-argument` error is signalled (because of the attempt to add 23 to `nil`) the function `default-wta-handler` will be given first chance at handling the error. When the second error is signalled (because of the attempt to add 105 to `nil`) the function `hairy-wta-handler` will be given first chance. If it declines (by returning `nil` as its first result) then `default-wta-handler` will be given a chance.

23.3. Error Handlers

When `signal` (page 271) invokes a condition handler it passes it the *condition-name* along with whatever other arguments were handed to `signal`. Condition handlers set up to handle errors can safely assume certain things about those arguments for all errors signalled by the system or signalled by user code via the functions `ferror` (page 275) and `cerror` (page 275).

An error handler can expect to be invoked as

For example:

```
(funcall* error-handler
          condition-name
          control-string
          proceedable-flag
          restartable-flag
          function
          params)
```

where *params* may vary in length. Handlers for particular condition names may expect certain parameters to always be included in the *params* list. The parameters supplied by the system for certain standard conditions are given in ???section-ref "standard condition names"???. The program siganalling the condition is free to pass any extra parameters. All error handlers should therefore be written with `&rest` arguments.

The *condition-name* is the name of the condition signalled.

control-string should be a string which when given to `format` (page 243) as a control string, along with *params* as additional arguments, produces some reasonable explanation of the error. It is up to the handling function whether it makes use of that control string.

The third and fourth arguments are flags. If the *proceedable-flag* is non-`nil` then the error is said to be *proceedable*. If the *restartable-flag* is non-`nil` then the error is said to be *restartable*. The values of these flags may be used by the signallers and handlers to pass more information than a single bit. It is up to the user how these are used. For instance, a set of signallers and handlers may pass information concerning the values expected from the handler when an error is proceeded.

function is the name of the function which initiated the signalling of the error, or `nil` if the signaller can't determine it.

An error handler can do some processing and then make one of four respones to the error (assuming the error was signalled with `ferror` (page 275) or `cerror` (page 275)). It can return `nil` to decline handling the error, it can *proceed*, it can *restart* or it can *throw*.

Throwing simply consists of using the function `throw` (page 87) to some tag outside the scope of the error.

Proceeding and *Restarting* are achieved by returning from the error handler with multiple values. The first

value should be one of the following:

`:return` This means to *proceed* the error. If the error was signalled by `cerror` and the error was proceedable then the second value returned with `:return` is returned as the value of `cerror`. If the error was not proceedable (always the case for errors signalled by `ferror`, then the system forces a `break` (page 277).

`:error-restart`
This means to *restart* the error. If the error was signalled by `cerror` and the error was restartable then the second value returned with `:error-restart` is thrown to the catch tag `error-restart`. If the error was not restartable (always the case for errors signalled by `ferror`, then the system forces a `break` (page 277). An error may also be simply restarted from the handler by throwing directly from there to a catch tag of `error-restart`, but that is not as bullet proof if the error wasn't in fact restartable.

No other values are legal as the first values returned by error handlers. For errors signalled by `ferror` or `cerror` illegal values will force a `break`.

23.4. Signalling Errors

LISP programs can signal errors by using one of the functions `ferror` (for *f*atal *error*) or `cerror` (for *c*ontinuable *error*).

`ferror` *condition-name control-string* `&rest` *params* [*Function*]
`ferror` signals the error condition *condtition-name*. The associated error message is obtainable by calling `format` (page 243) on *control-string* and *params*. The error is neither proceedable nor restartable. Function `ferror` never returns. It can be thrown through however. A usual COMMON LISP environment will have some sort of error handler established for condition name `t`. Thus the user can get at least minimal error handling with `ferror` using a null *condition-name* knowing that the error will at least be signalled to the user console.

`cerror` *proceedable-flag restartable-flag condition-name control-string* `&rest` *params* [*Function*]
`cerror` is similar to `ferror` (see above) except for *proceedable-flag* and *restartable-flag*. These are passed through to the eventual error handler. If `cerror` is called with a non-`nil` *proceedable-flag* the caller should be prepared to accept the returned value of `cerror` and use it to retry the operation that failed. If `cerror` is passed a non-`nil` *restartable-flag* then there should be a catch for taf `error-restart` somewhere above the caller.

`error-restart` {*form*}* [*Macro*]
`error-restart` can be used to denote a section of a program that can be restarted if certain errors occure during its execution. The form of an `error-restart` is:

For example:

```
(error-restart
  form1
  form2
  ...
  formn)
```

The expressions *formj* are evaluated in order; the values of all but the last are discarded (that is, the body of the `error-restart` form is an implicit `progn`). The value of the `error-restart` form is the value of *formn* (if the body is empty, `nil` is the value). If a *restaribale* error occurs during the evaluation of one of the *formj*'s, and the handler responds by forcing a a restart, then the forms, beginning with *form1* will be re-evaluated in order. The only way a restartable error can occur is if `cerror` is called with a *restartable-flag* which is non-`nil`.

`error-restart` is implemented as a macro which expands into:

For example:

```
(prog ()
  loop (*catch 'error-restart
         (return (progn form1
                        form2
                        ...
                        formn)))
       (go loop))
```

`check-arg` *var-name predicate description type-name* [*Macro*]

The `check-arg` macro is used to check arguments to make sure they are valid, signal a `:wrong-type-argument` condition if they are not, and use the value returned by the handler to replace the invalid value.

var-name is the name of the variable whose value is being checked to be of the correct type. If the error is proceeded this variable will be `setq`'ed to a replacement value. *predicate* is a test for whether the variable is of the correct type. It can either be a symbol whose function definition takes one argument and returns non-`nil` if the type is correct, or it can be a non-atomic form which is evaluated to check the type - usually such a form would contain a reference to the variable *var-name*. *description* is a string which expresses *predicate* in English. It is used in error messages. *type-name* gets passed to the `:wrong-type-argument` handler as the first required parameter of that class of error handlers (see section ??? <<<section ref "standard condition names">>>).

Thus `check-arg` has what consistitutes a valid argument specified to it in three ways. *predicate* is executable, *description* is human understandable and *type-name* is program understandable.

`check-arg` uses *predicate* to determine whether the value of the variable is of the correct type. If it is not a `:wrong-type-argument` condition is signalled with four parameters - *type-name*, the bad value, the symbol *var-name* and *description*. If the error is proceeded, the variable is set to the value returned and `check-arg` repeats the process. Only the first of these two parameters are defined for `:wrong-type-argument` handlers, and so theyt should not depend on the meanin of more than these two.

Consider for example:

For example:

```
(check-arg foo
           (and (fixnump foo) (< foo 0.))
           "a negative fixnum"
           negative-fixnum)
```

If `foo` is not of the right type an error will be signalled and a `:wrong-type-argument` which makes use of the *control-string* and *parameters* passed to it will print the message (at least):

```
Argument foo was 33, which is not a negative fixnum.
```

23.5. Break-points

Often error handlers want to pass control to the user's terminal. The user can then examine variable bindings and respond to the error, or perhaps just start some new computation. Control is passed by using the special form `break`.

`break` *tag* [*conditional-form*] [*Special form*]

This enters a *breakpoint* loop, which is similar to a LISP top level loop. (`break` *tag*) will always enter the loop; (`break` *tag conditional-form*) will evaluate *conditional-form* and only enter the break loop if it returns non-`nil`. If the break loop is entered, `break` prints out

For example:

```
;bkpt tag
```

and then enters a loop reading, evaluating, and printing forms. After reading a form `break` checks for the following special cases. If the symbol `<altmode>p` is typed, `break` return `nil`. If the the symbol `<altmode>r` is typed, `break` throws to a catch tag `error-restart`. If the symbol `<altmode>g` is typed, `break` throws back to the LISP top level. If (`return` *form*) is typed, `break` evaluates *form* and returns the result. In other respects a break loop appears very simialr to a top level loop.

23.6. Standard Condition Names

Some condition names are used by the COMMON LISP system itself. They are listed below along with the arguments they expect and the conventions followed in use of these conditions. All error condition handlers expect at least four arugments: *condition-name*, *control-string*, *proceedbale-flag*, and *restartable-flag*. In addition some condition names expect particular values for the fifth and subsequent arguments. These are included in the list below. It is always permissible to supply even more arguments than those required.

*** this list is not yet complete ***

`:wrong-type-argument`

Requires *type-name* and *value*, where the first is a symbol indicating what type of value is

required, and the second is the bad value supplied to the fnction signalling the error. If the error is proceeded, the value returned by the handler (that is, the second value returned; the first would be `:return`) should be a new value for the argument to be used instead of the one which was of the wrong type.

`:inconsistent-arguments`

Requires *list-of-inconsistent-argument-values*. This condition is signalled when the arguments to a function are inconsistent with each other, but the fault does not lie with any particular one of them. If the error is proceeded, the value returned by the handler should be returned by the function whose arguments were inconsistent.

Chapter 24

The Compiler

The compiler is a program which makes code run faster, by translating programs into an implementation-dependent form (subrs) which can be executed more efficiently by the computer. Most of the time you can write programs without worrying about the compiler; compiling a file of code should produce an equivalent but more efficient program. When doing more esoteric things, one may need to think carefully about what happens at "compile time" and what happens at "load time". Then the difference between the syntaxes "`#.`" and "`#,`" becomes important, and the `eval-when` (page EVAL-WHEN-FUN) construct becomes particularly useful.

Most declarations are not used by the COMMON LISP interpreter; they may be used to give advice to the compiler. The compiler may attempt to check your advice and warn you if it is inconsistent.

Unlike most other LISP dialects, COMMON LISP recognizes `special` declarations in interpreted code as well as compiled code. This potential source of incompatibility between interpreted and compiled code is thereby *eliminated* in COMMON LISP.

The internal workings of a compiler will of course be highly implementation-dependent. The following functions provide a standard interface to the compiler, however.

`compile` *name* `&optional` *definition* [*Function*]

If *definition* is supplied, it should be a lambda-expression, the interpreted function to be compiled. If it is not supplied, then *name* should be a symbol with a definition that is a lambda expression or select expression; that definition is compiled and the resulting compiled code is put back into the symbol as its function definition.

The definition is compiled and a subr object produced. If *name* is a symbol, then the subr object is installed as the global function definition of the symbol and the symbol is returned. If *name* is `nil`, then the subr object itself is returned.

For example:

```
(defun foo ...) => foo                          ;A function definition.
(compile 'foo) => foo                           ;Compile it.
        ;Now foo runs faster.
(compile nil '(lambda (a b c) (- (* b b) (* 4 a c))))
  => a compiled function of three arguments which computes b^2-4ac
```

`comfile` *input-filespec* `&optional` *output-filespec* [*Function*]

Each argument should be a valid file name specifier for `with-open-file` (page 267). The file should be a LISP source file; its contents are compiled and written as a binary object ("`FASL`") file to *output-filespec*. The *output-filespec* defaults in a manner appropriate to the implementation's file system conventions.

`disassemble` *name-or-subr* [*Function*]

The argument should be either a function object, a lambda-expression, or a symbol with a function definition. If the relevant function is not a compiled function, it is first compiled. In any case, the compiled code is then "reverse-assembled" and printed out in a symbolic format. This is primarily useful for debugging the compiler, but also often of use to the novice who wishes to understand the workings of compiled code.

Implementation note: Implementors are encouraged to make the output readable, preferably with helpful comments.

References

[1] Brooks, Rodney A.; Gabriel, Richard P.; and Steele, Guy L. Jr.
An Optimizing Compiler for Lexically Scoped LISP.
In *Proceedings of the 1982 Symposium on Compiler Construction*, pages 261-275. ACM SIGPLAN, Boston, June, 1982.
Proceedings published as *ACM SIGPLAN Notices* 17, 6 (June 1982).

[2] Coonen, Jerome T.
An Implementation Guide to a Proposed Standard for Floating-Point Arithmetic.
Computer 13(1):68-79, January, 1980.
Errata for this paper appeared as [3].

[3] Coonen, Jerome T.
Errata for 'An Implementation Guide to a Proposed Standard for Floating-Point Arithmetic'.
Computer 14(3):62, March, 1981.
These are errata for [2].

[4] Fateman, Richard J.
Reply to an Editorial.
ACM SIGSAM Bulletin 25:9-11, March, 1973.

[5] IEEE Computer Society Standard Committee, Microprocessor Standards Subcommittee, Floating-Point Working Group.
A Proposed Standard for Binary Floating-Point Arithmetic.
Computer 14(3):51-62, March, 1981.

[6] Moon, David.
MacLISP Reference Manual, Revision 0.
M.I.T. Project MAC, Cambridge, Massachusetts, April 1974.

[7] Moore, J. Strother II.
The InterLISP Virtual Machine Specification.
Technical Report CSL 76-5, Xerox Palo Alto Research Center, Palo Alto, California, September, 1976.

[8] Penfield, Paul, Jr.
Principal Values and Branch Cuts in Complex APL.
In *APL 81 Conference Proceedings*, pages 248-256. ACM SIGAPL, San Francisco, September, 1981.
Proceedings published as *APL Quote Quad* 12, 1 (September 1981).

[9] Steele, Guy Lewis Jr., and Sussman, Gerald Jay.
The Revised Report on SCHEME: A Dialect of LISP.
AI Memo 452, Massachusetts Institute of Technology Artificial Intelligence Laboratory, Cambridge, Massachusetts, January, 1978.

[10] Teitelman, Warren, et al.
InterLISP Reference Manual
Xerox Palo Alto Research Center, Palo Alto, California, 1978.
Third revision.

[11] Weinreb, Daniel, and Moon, David.
LISP Machine Manual, Fourth Edition.
Massachusetts Institute of Technology Artificial Intelligence Laboratory, Cambridge, Massachusetts, July 1981.

COMMON LISP Summary

`sample-function` *arg1 arg2* `&optional` *arg3 arg4* [*Function*]

`sample-variable` [*Variable*]

`sample-constant` [*Constant*]

`sample-special-form` [*name*] ({*var*}*) {*form*}+ [*Special form*]

`sample-macro` *var* {*tag* | *statement*}* [*Macro*]

`deftype` *name varlist* {*form*}* [*Special form*]

`defun` *name lambda-list* {(`declare` {*declaration*}*)}* [*doc-string*] {*form*}* [*Special form*]

`defselect` *name* [*doc-string*] {(*keys lambda-list* {(`declare` {*declaration*}*)}* {*form*}*)}* [*Special form*]

`defvar` *name* [*initial-value* [*documentation*]] [*Special form*]

`defconst` *name initial-value* [*documentation*] [*Special form*]

`nil` [*Constant*]

`t` [*Constant*]

`typep` *object* `&optional` *type* [*Function*]

`subtypep` *type1 type2* [*Function*]

`null` *object* [*Function*]

`symbolp` *object* [*Function*]

`atom` *object* [*Function*]

`consp` *object* [*Function*]

`listp` *object* [*Function*]

`numberp` *object* [*Function*]

`integerp` *object* [*Function*]

`rationalp` *object* [*Function*]

`floatp` *object* [*Function*]

`complexp` *object* [*Function*]

`characterp` *object* [*Function*]

`stringp` *object* [*Function*]

`vectorp` *object* [*Function*]

`arrayp` *object* [*Function*]

`functionp` *object* [*Function*]

`subrp` *object* [*Function*]

`closurep` *object* [*Function*]

`eq` *x y* [*Function*]

`eql` *x y* [*Function*]

`equal` *x y* [*Function*]

`equalp` *x y* `&optional` *fuzz* [*Function*]

`not` *x* [*Function*]

`and` {*form*}* [*Special form*]

`or` {*form*}* [*Special form*]

`quote` *object* [*Special form*]

`function` *fn* [*Special form*]

`closure` *varlist function* [*Function*]
`symeval` *symbol* [*Function*]
`fsymeval` *symbol* [*Function*]
`boundp` *symbol* [*Function*]
`fboundp` *symbol* [*Function*]
`macro-p` *symbol* [*Function*]
`special-form-p` *symbol* [*Function*]
`setq` {*var form*}* [*Special form*]
`psetq` {*var form*}* [*Special form*]
`set` *symbol value* [*Function*]
`fset` *symbol value* [*Function*]
`makunbound` *symbol* [*Function*]
`fmakunbound` *symbol* [*Function*]
`setf` *place newvalue* [*Macro*]
`swapf` *place newvalue* [*Macro*]
`exchf` *place1 place2* [*Macro*]
`apply` *function arglist* [*Function*]
`funcall` *fn* `&rest` *arguments* [*Function*]
`funcall*` *f* `&rest` *args* [*Function*]
`progn` {*form*}* [*Special form*]
`prog1` *first* {*form*}* [*Special form*]
`prog2` *first second* {*form*}* [*Special form*]
`let` ({*var* | (*var value*)}*) {*form*}* [*Macro*]
`let*` ({*var* | (*var value*)}*) {*form*}* [*Special form*]
`progv` *symbols values* {*form*}* [*Special form*]
`flet` ({(*name lambda-list* {*declare-form*}* [*doc-string*] {*form*}*)}*) {*form*}* [*Special form*]
`labels` ({(*name lambda-list* {*declare-form*}* [*doc-string*] {*form*}*)}*) {*form*}* [*Special form*]
`macrolet` ({(*name varlist* {*form*}*)}*) {*form*}* [*Special form*]
`cond` {(*test* {*form*}*)}* [*Special form*]
`if` *pred then* [*else*] [*Special form*]
`when` *pred* {*form*}* [*Special form*]
`unless` *pred* {*form*}* [*Special form*]
`case` *keyform* {(({*key*}*) {*form*}*)}* [*Special form*]
`typecase` *keyform* {(*type* {*form*}*)}* [*Special form*]
`block` *name* {*form*}* [*Special form*]
`return` *result* [*Special form*]
`return-from` *blockname result* [*Special form*]
`do` ({(*var* [*init* [*step*]])}*) (*end-test* {*form*}*) {*tag* | *statement*}* [*Special form*]
`do*` *bindspecs endtest* {*form*}* [*Special form*]
`dolist` (*var listform* [*resultform*]) {*tag* | *statement*}* [*Special form*]
`dotimes` (*var countform* [*resultform*]) {*tag* | *statement*}* [*Special form*]
`mapcar` *function list* `&rest` *more-lists* [*Function*]
`maplist` *function list* `&rest` *more-lists* [*Function*]

mapc *function list* &rest *more-lists* [*Function*]
mapl *function list* &rest *more-lists* [*Function*]
mapcan *function list* &rest *more-lists* [*Function*]
mapcon *function list* &rest *more-lists* [*Function*]
prog ({*var* | (*var* [*init*])}*) {*tag* | *statement*}* [*Special form*]
prog* [*Special form*]
go *tag* [*Special form*]
values &rest *args* [*Function*]
values-list *list* [*Function*]
multiple-value-list *form* [*Special form*]
mvcall *function* {*form*}* [*Special form*]
mvprog1 *form* {*form*}* [*Special form*]
multiple-value-bind ({*var*}*) *values-form* {*form*}* [*Special form*]
multiple-value *variables form* [*Special form*]
catch *tag* {*form*}* [*Special form*]
catch-all *catch-function* {*form*}* [*Special form*]
unwind-all *catch-function* {*form*}* [*Special form*]
unwind-protect *protected-form* {*cleanup-form*}* [*Special form*]
throw *tag result* [*Special form*]
macro *name* (*var*) {*form*}* [*Macro*]
defmacro *name varlist* {*form*}* [*Macro*]
defmacro-check-args [*Variable*]
defmacro-maybe-displace [*Variable*]
macro-expansion-hook [*Variable*]
displace *macro-call expansion* [*Function*]
macroexpand *form* [*Function*]
macroexpand-1 *form* [*Function*]
declare {*declaration*}* [*Special form*]
locally {*declare-form*}* {*form*}* [*Special form*]
the *type form* [*Special form*]
getpr *symbol indicator* &optional *default* [*Function*]
putpr *symbol indicator newvalue* [*Function*]
rempr *symbol indicator* [*Function*]
plist *symbol* [*Function*]
getf *place indicator* &optional *default* [*Function*]
putf *place indicator newvalue* [*Macro*]
remf *place indicator* [*Macro*]
get-properties *place indicator-list* [*Function*]
map-properties *function place* [*Function*]
get-pname *sym* [*Function*]
samepnamep *sym1 sym2* [*Function*]
make-symbol *pname* [*Function*]
copysymbol *sym* &optional *copy-props* [*Function*]

gensym &optional *x* [*Function*]

gentemp *prefix* &optional *package* [*Function*]

symbol-package *sym* [*Function*]

make-package *package-name* &optional *copy-from* [*Function*]

package [*Variable*]

package *package* [*Function*]

package-name *package* [*Function*]

begin-package *package-name* [*Function*]

end-package *package-name* [*Function*]

intern *string-or-symbol* &optional *package* [*Function*]

remob *string-or-symbol* &optional *package* [*Function*]

internedp *string-or-symbol* &optional *package* [*Function*]

externalp *string-or-symbol* &optional *package* [*Function*]

export *symbols* &optional *package* [*Function*]

unexport *symbols* &optional *package* [*Function*]

import *symbols* &optional *to-package* [*Function*]

shadow *symbols* &optional *to-package* [*Function*]

use *from-package* &optional *to-package ignore-list force-list* [*Function*]

provide *package* [*Function*]

require *package* &optional *pathname* [*Function*]

package-use-conflicts *from-package* &optional *to-package* [*Function*]

do-symbols (*var* [*package*] [*result-form*]) {*tag* | *statement*}* [*Special form*]

do-external-symbols (*var* [*package*] [*result-form*]) {*tag* | *statement*}* [*Special form*]

do-internal-symbols (*var* [*package*] [*result-form*]) {*tag* | *statement*}* [*Special form*]

do-all-symbols (*var* [*result-form*]) {*tag* | *statement*}* [*Special form*]

zerop *number* [*Function*]

plusp *number* [*Function*]

minusp *number* [*Function*]

oddp *integer* [*Function*]

evenp *integer* [*Function*]

= *number* &rest *more-numbers* [*Function*]

/= *number* &rest *more-numbers* [*Function*]

< *number* &rest *more-numbers* [*Function*]

> *number* &rest *more-numbers* [*Function*]

<= *number* &rest *more-numbers* [*Function*]

>= *number* &rest *more-numbers* [*Function*]

max *number* &rest *more-numbers* [*Function*]

min *number* &rest *more-numbers* [*Function*]

fuzzy= *number1 number2* &optional *fuzz* [*Function*]

fuzziness *number* [*Function*]

+ &rest *numbers* [*Function*]

- *number* &rest *more-numbers* [*Function*]

* &rest *numbers* [*Function*]

`/` *number* `&rest` *more-numbers* [*Function*]

`1+` *number* [*Function*]

`1-` *number* [*Function*]

`incf` *place* [*delta*] [*Macro*]

`decf` *place* [*delta*] [*Macro*]

`conjugate` *number* [*Function*]

`gcd` `&rest` *rationals* [*Function*]

`lcm` *rational* `&rest` *more-rationals* [*Function*]

`exp` *number* [*Function*]

`expt` *base-number power-number* [*Function*]

`log` *number* `&optional` *base* [*Function*]

`sqrt` *number* [*Function*]

`isqrt` *integer* [*Function*]

`abs` *number* [*Function*]

`phase` *number* [*Function*]

`signum` *number* [*Function*]

`sin` *radians* [*Function*]

`cos` *radians* [*Function*]

`tan` *radians* [*Function*]

`cis` *radians* [*Function*]

`asin` *number* [*Function*]

`acos` *number* [*Function*]

`atan` *y* `&optional` *x* [*Function*]

`pi` [*Variable*]

`sinh` *number* [*Function*]

`cosh` *number* [*Function*]

`tanh` *number* [*Function*]

`asinh` *number* [*Function*]

`acosh` *number* [*Function*]

`atanh` *number* [*Function*]

`float` *number* `&optional` *other* [*Function*]

`rational` *number* [*Function*]

`rationalize` *number* `&optional` *tolerance* [*Function*]

`numerator` *rational* [*Function*]

`denominator` *rational* [*Function*]

`floor` *number* `&optional` *divisor* [*Function*]

`ceil` *number* `&optional` *divisor* [*Function*]

`trunc` *number* `&optional` *divisor* [*Function*]

`round` *number* `&optional` *divisor* [*Function*]

`mod` *number* `&optional` *divisor tolerance* [*Function*]

`rem` *number* `&optional` *divisor tolerance* [*Function*]

`ffloor` *number* `&optional` *divisor* [*Function*]

`fceil` *number* `&optional` *divisor* [*Function*]

ftrunc *number* &optional *divisor* [*Function*]

fround *number* &optional *divisor* [*Function*]

float-fraction *float* [*Function*]

float-exponent *float* [*Function*]

scale-float *float integer* [*Function*]

complex *realpart* &optional *imagpart* [*Function*]

realpart *number* [*Function*]

imagpart *number* [*Function*]

logior &rest *integers* [*Function*]

logxor &rest *integers* [*Function*]

logand &rest *integers* [*Function*]

logeqv &rest *integers* [*Function*]

lognand *integer1 integer2* [*Function*]

lognor *integer1 integer2* [*Function*]

logandc1 *integer1 integer2* [*Function*]

logandc2 *integer1 integer2* [*Function*]

logorc1 *integer1 integer2* [*Function*]

logorc2 *integer1 integer2* [*Function*]

boole *op integer1 integer2* [*Function*]

boole-clr [*Variable*]

boole-set [*Variable*]

boole-1 [*Variable*]

boole-2 [*Variable*]

boole-c1 [*Variable*]

boole-c2 [*Variable*]

boole-and [*Variable*]

boole-ior [*Variable*]

boole-xor [*Variable*]

boole-eqv [*Variable*]

boole-nand [*Variable*]

boole-nor [*Variable*]

boole-andc1 [*Variable*]

boole-andc2 [*Variable*]

boole-orc1 [*Variable*]

boole-orc2 [*Variable*]

lognot *integer* [*Function*]

logtest *integer1 integer2* [*Function*]

logbitp *index integer* [*Function*]

ash *integer count* [*Function*]

logcount *integer* [*Function*]

haulong *integer* [*Function*]

haipart *integer count* [*Function*]

byte *size position* [*Function*]

standard-charp *char* [*Function*]

graphicp *char* [*Function*]

string-charp *char* [*Function*]

alphap *char* [*Function*]

uppercasep *char* [*Function*]

lowercasep *char* [*Function*]

bothcasep *char* [*Function*]

digitp *char* &optional (*radix* 10.) [*Function*]

alphanumericp *char* [*Function*]

char= *char1 char2* [*Function*]

char-equal *char1 char2* [*Function*]

char< *char1 char2* [*Function*]

char> *char1 char2* [*Function*]

char-lessp *char1 char2* [*Function*]

char-greaterp *char1 char2* [*Function*]

character *object* [*Function*]

char-code *char* [*Function*]

char-bits *char* [*Function*]

char-font *char* [*Function*]

code-char *code* &optional (*bits* 0) (*font* 0) [*Function*]

make-char *char* &optional (*bits* 0) (*font* 0) [*Function*]

char-upcase *char* [*Function*]

char-downcase *char* [*Function*]

digit-charp *weight* &optional (*radix* 10.) (*bits* 0) (*font* 0) [*Function*]

digit-weight *weight* &optional (*radix* 10.) (*bits* 0) (*font* 0) [*Function*]

char-int *char* [*Function*]

int-char *integer* [*Function*]

char-name *char* [*Function*]

name-char *sym* [*Function*]

char-control-bit [*Constant*]

char-meta-bit [*Constant*]

char-super-bit [*Constant*]

char-hyper-bit [*Constant*]

char-bit *char name* [*Function*]

set-char-bit *char name newvalue* [*Function*]

elt *sequence index* [*Function*]

setelt *sequence index newvalue* [*Function*]

subseq *sequence start* &optional *end* [*Function*]

copyseq *sequence* [*Function*]

length *sequence* [*Function*]

reverse *sequence* [*Function*]

nreverse *sequence* [*Function*]

to *result-type sequence* [*Function*]

`catenate` *result-type* `&rest` *sequences* [*Function*]

`map` *result-type function sequence* `&rest` *more-sequences* [*Function*]

`some` *predicate sequence* `&rest` *more-sequences* [*Function*]

`every` *predicate sequence* `&rest` *more-sequences* [*Function*]

`notany` *predicate sequence* `&rest` *more-sequences* [*Function*]

`notevery` *predicate sequence* `&rest` *more-sequences* [*Function*]

`fill` *sequence item* `&key :start :end` [*Function*]

`replace` *sequence1 sequence2* `&key :start :end :start1 :end1 :start2 :end2` [*Function*]

`remove` *item sequence* `&key :from-end :test :test-not :start :end :count :key` [*Function*]

`remove-if` *test sequence* `&key :from-end :start :end :count :key` [*Function*]

`remove-if-not` *test sequence* `&key :from-end :start :end :count :key` [*Function*]

`delete` *item sequence* `&key :from-end :test :test-not :start :end :count :key` [*Function*]

`delete-if` *test sequence* `&key :from-end :start :end :count :key` [*Function*]

`delete-if-not` *test sequence* `&key :from-end :start :end :count :key` [*Function*]

`substitute` *newitem olditem sequence* `&key :from-end :test :test-not :start :end :count :key` [*Function*]

`substitute-if` *newitem test sequence* `&key :from-end :start :end :count :key` [*Function*]

`substitute-if-not` *newitem test sequence* `&key :from-end :start :end :count :key` [*Function*]

`nsubstitute` *newitem olditem sequence* `&key :from-end :test :test-not :start :end :count :key` [*Function*]

`nsubstitute-if` *newitem test sequence* `&key :from-end :start :end :count :key` [*Function*]

`nsubstitute-if-not` *newitem test sequence* `&key :from-end :start :end :count :key` [*Function*]

`find` *item sequence* `&key :from-end :test :test-not :start :end :key` [*Function*]

`find-if` *test sequence* `&key :from-end :start :end :key` [*Function*]

`find-if-not` *test sequence* `&key :from-end :start :end :key` [*Function*]

`position` *item sequence* `&key :from-end :test :test-not :start :end :key` [*Function*]

`position-if` *test sequence* `&key :from-end :start :end :key` [*Function*]

`position-if-not` *test sequence* `&key :from-end :start :end :key` [*Function*]

`count` *item sequence* `&key :from-end :test :test-not :start :end :key` [*Function*]

`count-if` *test sequence* `&key :from-end :start :end :key` [*Function*]

`count-if-not` *test sequence* `&key :from-end :start :end :key` [*Function*]

`mismatch` *sequence1 sequence2* `&key :from-end :test :test-not :start :end :start1 :start2 :end1 :end2` [*Function*]

`maxprefix` *sequence1 sequence2* `&key :from-end :test :test-not :start :end :start1 :start2 :end1 :end2` [*Function*]

`maxsuffix` *sequence1 sequence2* `&key :from-end :test :test-not :start :end :start1 :start2 :end1 :end2` [*Function*]

search *sequence1 sequence2* &key :from-end :test :test-not [*Function*]
:start :end :start1 :start2 :end1 :end2

sort *sequence predicate* &key :key [*Function*]

stable-sort *sequence predicate* &key :key [*Function*]

merge *sequence1 sequence2 predicate* &key :key [*Function*]

car *x* [*Function*]

cdr *x* [*Function*]

c...r *x* [*Function*]

cons *x y* [*Function*]

tree-equal *x y* [*Function*]

endp *object* [*Function*]

list-length *list* &optional *limit* [*Function*]

nth *n list* [*Function*]

nthcdr *n list* [*Function*]

last *list* [*Function*]

list &rest *args* [*Function*]

list* arg &rest *others* [*Function*]

make-list *size* &optional *value* [*Function*]

append &rest *lists* [*Function*]

copylist *list* [*Function*]

copyalist *list* [*Function*]

copytree *object* [*Function*]

revappend *x y* [*Function*]

nconc &rest *lists* [*Function*]

nreconc *x y* [*Function*]

push *item place* [*Macro*]

pushnew *item place* [*Macro*]

pop *place* [*Macro*]

butlast *list* &optional *n* [*Function*]

nbutlast *list* &optional *n* [*Function*]

buttail *list sublist* [*Function*]

rplaca *x y* [*Function*]

rplacd *x y* [*Function*]

setnth *n list newvalue* [*Function*]

subst *new old tree* [*Function*]

nsubst *new old tree* [*Function*]

substq *new old tree* [*Function*]

nsubstq *new old tree* [*Function*]

sublis *alist tree* [*Function*]

nsublis *alist tree* [*Function*]

member *item list* &key :test :test-not :key [*Function*]

member-if *predicate list* &key :key [*Function*]

member-if-not *predicate list*, Keys = {[key] [*Function*]

`tailp` *sublist list* [*Function*]

`adjoin` *item list* `&key :test :test-not` [*Function*]

`union` *list1 list2* `&key :test :test-not` [*Function*]

`nunion` *list1 list2* `&key :test :test-not` [*Function*]

`intersection` *list1 list2* `&key :test :test-not` [*Function*]

`nintersection` *list1 list2* `&key :test :test-not` [*Function*]

`setdifference` *list1 list2* `&key :test :test-not` [*Function*]

`nsetdifference` *list1 list2* `&key :test :test-not` [*Function*]

`set-exclusive-or` *list1 list2* `&key :test :test-not` [*Function*]

`nset-exclusive-or` *list1 list2* `&key :test :test-not` [*Function*]

`subsetp` *list1 list2* `&key :test :test-not` [*Function*]

`acons` *key datum a-list* [*Function*]

`pairlis` *keys data* `&optional` *a-list* [*Function*]

`assoc` *item a-list* `&key :test :test-not` [*Function*]

`rassoc` *item a-list* `&key :test :test-not` [*Function*]

`make-eq-hash-table &key :size :rehash-size :rehash-threshold` [*Function*]

`make-eql-hash-table &key :size :rehash-size :rehash-threshold` [*Function*]

`make-equal-hash-table &key :size :rehash-size :rehash-threshold` [*Function*]

`gethash` *key hash-table* `&optional` *default* [*Function*]

`puthash` *key value hash-table* [*Function*]

`remhash` *key hash-table* [*Function*]

`maphash` *function hash-table* [*Function*]

`clrhash` *hash-table* [*Function*]

`sxhash` *S-expression* [*Function*]

`make-array` *dimensions* `&key :type :initial-value :initial-contents`
`:fill-pointer :displaced-to :displaced-index-offset` [*Function*]

`make-vector` *length* `&key :type :initial-value :initial-contents`
`:fill-pointer` [*Function*]

`aref` *array* `&rest` *subscripts* [*Function*]

`aset` *new-value array* `&rest` *subscripts* [*Function*]

`array-type` *array* [*Function*]

`array-allocated-length` *array* [*Function*]

`array-active-length` *array* [*Function*]

`array-rank` *array* [*Function*]

`array-dimension` *axis-number array* [*Function*]

`array-dimensions` *array* [*Function*]

`array-in-bounds-p` *array* `&rest` *subscripts* [*Function*]

`velt` *vector index* [*Function*]

`vsetelt` *vector index newvalue* [*Function*]

`vref` *vector index* [*Function*]

`vset` *vector index newvalue* [*Function*]

`bit` *bit-vector index* [*Function*]

`rplacbit` *bit-vector index newbit* [*Function*]

bit-and &rest *bit-vectors* [*Function*]

bit-ior &rest *bit-vectors* [*Function*]

bit-xor &rest *bit-vectors* [*Function*]

bit-eqv &rest *bit-vectors* [*Function*]

bit-nand *bit-vector1 bit-vector2* [*Function*]

bit-nor *bit-vector1 bit-vector2* [*Function*]

bit-andc1 *bit-vector1 bit-vector2* [*Function*]

bit-andc2 *bit-vector1 bit-vector2* [*Function*]

bit-orc1 *bit-vector1 bit-vector2* [*Function*]

bit-orc2 *bit-vector1 bit-vector2* [*Function*]

bit-not *bit-vector* [*Function*]

array-reset-fill-pointer *array* &optional *index* [*Function*]

array-push *array new-element* [*Function*]

array-push-extend *array x* &optional *extension* [*Function*]

array-pop *array* [*Function*]

adjust-array-size *array new-size* &optional *new-element* [*Function*]

array-grow *array new-element* &rest *dimensions* [*Function*]

char *string index* [*Function*]

rplachar *string index newchar* [*Function*]

string= *string1 string2* &key :start :end :start1 :end1 :start2 :end2 [*Function*]

string-equal *string1 string2* &key :start :end :start1 :end1 :start2 :end2 [*Function*]

string< *string1 string2* &key :start :end :start1 :end1 :start2 :end2 [*Function*]

string> *string1 string2* &key :start :end :start1 :end1 :start2 :end2 [*Function*]

string<= *string1 string2* &key :start :end :start1 :end1 :start2 :end2 [*Function*]

string>= *string1 string2* &key :start :end :start1 :end1 :start2 :end2 [*Function*]

string/= *string1 string2* &key :start :end :start1 :end1 :start2 :end2 [*Function*]

string-lessp *string1 string2* &key :start :end [*Function*]
:start1 :end1 :start2 :end2

string-greaterp *string1 string2* &key :start :end [*Function*]
:start1 :end1 :start2 :end2

string-not-lessp *string1 string2* &key :start :end [*Function*]
:start1 :end1 :start2 :end2

string-not-greaterp *string1 string2* &key :start :end [*Function*]
:start1 :end1 :start2 :end2

string-not-equal *string1 string2* &key :start :end [*Function*]
:start1 :end1 :start2 :end2

make-string *count* &optional *fill-character* [*Function*]

string-trim *character-bag string* [*Function*]

string-left-trim *character-bag string* [*Function*]

string-right-trim *character-bag string* [*Function*]

string-upcase *string* &key :start :end [*Function*]

string-downcase *string* &key :start :end [*Function*]

string-capitalize *string* &key :start :end [*Function*]

`inch` &optional *input-stream eof-value* [*Function*]

`tyi` &optional *input-stream eof-value* [*Function*]

`uninch` *character* &optional *input-stream* [*Function*]

`untyi` *integer* &optional *input-stream* [*Function*]

`inchpeek` &optional *peek-type input-stream eof-value* [*Function*]

`tyipeek` &optional *peek-type input-stream eof-value* [*Function*]

`listen` &optional *input-stream* [*Function*]

`inch-no-hang` &optional *input-stream eof-value* [*Function*]

`tyi-no-hang` &optional *input-stream eof-value* [*Function*]

`clear-input` &optional *input-stream* [*Function*]

`read-from-string` *string* &optional *start end preserve-whitespace-p eof-value* [*Function*]

`parse-number` *string* &optional *start end radix no-junk-allowed* [*Function*]

`in` *binary-input-stream* &optional *eof-value* [*Function*]

`prin1` *object* &optional *output-stream* [*Function*]

`print` *object* &optional *output-stream* [*Function*]

`princ` *object* &optional *output-stream* [*Function*]

`prin1string` *object* [*Function*]

`princstring` *object* [*Function*]

`ouch` *character* &optional *output-stream* [*Function*]

`tyo` *integer* &optional *output-stream* [*Function*]

`terpri` &optional *output-stream* [*Function*]

`fresh-line` &optional *output-stream* [*Function*]

`force-output` &optional *output-stream* [*Function*]

`clear-output` &optional *output-stream* [*Function*]

`out` *integer binary-output-stream* [*Function*]

`format` *destination control-string* &rest *arguments* [*Function*]

`y-or-n-p` &optional *message stream* [*Function*]

`yes-or-no-p` &optional *message stream* [*Function*]

`fquery` *options format-string* &rest *format-args* [*Function*]

`pathname` *thing* [*Function*]

`truename` *thing* [*Function*]

`parse-namestring` *thing* &optional *convention defaults break-characters start end* [*Function*]

`merge-pathname-defaults` *pathname* &optional *defaults default-type default-version* [*Function*]

`make-pathname` &key `:host :device :directory :name`
`:type :version :defaults` [*Function*]

`pathnamep` *object* [*Function*]

`pathname-host` *pathname* [*Function*]

`pathname-device` *pathname* [*Function*]

`pathname-directory` *pathname* [*Function*]

`pathname-name` *pathname* [*Function*]

`pathname-type` *pathname* [*Function*]

`pathname-version` *pathname* [*Function*]

`pathname-plist` *pathname* [*Function*]

`namestring` *pathname* [*Function*]
`file-namestring` *pathname* [*Function*]
`directory-namestring` *pathname* [*Function*]
`host-namestring` *pathname* [*Function*]
`enough-namestring` *pathname* `&optional` *defaults* [*Function*]
`user-homedir-pathname` `&optional` *host* [*Function*]
`init-file-pathname` *program-name* `&optional` *host* [*Function*]
`default-pathname-defaults` [*Variable*]
`load-pathname-defaults` [*Variable*]
`add-logical-pathname-host` *logical-host actual-host default-device translations* [*Function*]
`translated-pathname` *pathname* [*Function*]
`back-translated-pathname` *logical-pathname actual-pathname* [*Function*]
`with-open-file` *bindspec* {*form*}* [*Special form*]
`open` *filename* `&optional` *options* [*Function*]
`rename-file` *file new-name* `&optional` *error-p* [*Function*]
`delete-file` *file* `&optional` *error-p* [*Function*]
`probe-file` *filename* [*Function*]
`file-creation-date` *file* [*Function*]
`file-author` *file* [*Function*]
`filepos` *file-stream* `&optional` *position* [*Function*]
`file-length` *file-stream* [*Function*]
`load` *filename* `&key :package :verbose :noerror :preserve-defaults` [*Function*]
`signal` *condition-name* `&rest` *args* [*Function*]
`condition-bind` *bindings* {*form*}* [*Special form*]
`condition-setq` {*spec*}* [*Special form*]
`ferror` *condition-name control-string* `&rest` *params* [*Function*]
`cerror` *proceedable-flag restartable-flag condition-name control-string* `&rest` *params* [*Function*]
`error-restart` {*form*}* [*Macro*]
`check-arg` *var-name predicate description*
type-name [*Macro*]
`break` *tag* [*conditional-form*] [*Special form*]
`compile` *name* `&optional` *definition* [*Function*]
`comfile` *input-filespec* `&optional` *output-filespec* [*Function*]
`disassemble` *name-or-subr* [*Function*]
(End of COMMON LISP summary.)

Index

Index of Concepts

Index of Variables and Constants

Index of Keywords

Index of Functions, Macros, and Special Forms