

Computer Science Department
(412) 578-2565

Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213
22 November 1982

Greetings!

Enclosed is a copy of the latest draft ("Laser Edition") of the Common LISP Reference Manual. Please remember that this is only a draft and not a final version.

A serious error occurred in the formatting of the document. Through a comedy of errors, the chapter on Arrays was omitted during the final pass through SCRIBE. It was supposed to have been Chapter 17. It is enclosed separately as Chapter 16a, with page numbers 192a through 192j. You are asked to insert it between Chapters 16 and 17 (pages 192 and 193). Cross-references occurring within the manuscript may be in error because of the inadvertent omission.

It is my belief that the Common LISP Group is in essential agreement on the contents of this draft, and that only minor corrections will be needed. A few major language components that have been previously discussed are not included here, such as the proposed "instances" feature, the complex loop macro, and the `LetS` construct, and of course the dispositions of these are not necessarily yet agreed upon.

I would like very much to freeze a version of this document by January 1983 and publish it as a fixed reference to which we can point. This is necessary so that implementations may proceed with confidence. This is not to say that Common LISP will be complete or fixed. It may be expected to continue to evolve and have various aspects extended, included, or modified. However, consideration of looping constructs or package systems should not be allowed to impede the progress of implementations unnecessarily. Presumably a revised Common LISP document will appear later, perhaps in 1984 or 1985. If there is doubt about the design of some feature of the language, it probably should simply be omitted from this frozen version, in the expectation that it can be added the second time around.

It would be helpful if all interested parties would send corrections or comments to me as quickly as possible. The most helpful comments would be corrections of presentation, including indications

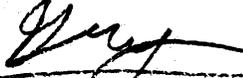
22 November 1982

Page 2

of passages that are inconsistent or unclear, and *minor* improvements to the language. Suggestions for major changes to the language are also welcomed, but I ask that they be included in separate messages or on separate pieces of paper. Of course, if I have failed to reflect in this edition the votes of the committee on any previously balloted issues, this should be drawn to my attention as quickly as possible.

Thank you for your help with and interest in Common LISP.

Sincerely,



Guy L. Steele Jr.

Assistant Professor of Computer Science .

ARPANET: Guy.Steele @ CMU-CS-A

CARNEGIE-MELLON UNIVERSITY

DEPARTMENT OF COMPUTER SCIENCE

SPICE PROJECT

Common Lisp Reference Manual

Guy L. Steele Jr.

16 November 1982

Laser Edition
Supposed to be Completely Coherent

Copyright © 1982 Guy L. Steele Jr.

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.

Would it be wonderful if, under the pressure of all these difficulties, the Convention should have been forced into some deviations from that artificial structure and regular symmetry which an abstract view of the subject might lead an ingenious theorist to bestow on a constitution planned in his closet or in his imagination?

—*James Madison, The Federalist No. 37, January 11, 1788*

Table of Contents

1. Introduction	1
1.1. Purpose	1
1.2. Notational Conventions	3
2. Data Types	9
2.1. Numbers	11
2.1.1. Integers	11
2.1.2. Ratios	12
2.1.3. Floating-point Numbers	13
2.1.4. Complex Numbers	15
2.2. Characters	16
2.3. Symbols	16
2.4. Lists and Conses	18
2.5. Arrays	19
2.5.1. Vectors	20
2.5.2. Strings	21
2.5.3. Bit-vectors	21
2.6. Hash tables	22
2.7. Readtables	22
2.8. Packages	22
2.9. Pathnames	22
2.10. Streams	22
2.11. Random-states	23
2.12. Structures	23
2.13. Functions	23
2.14. Unreadable Data Objects	23
2.15. Overlap, Inclusion, and Disjointness of Types	24
3. Scope and Extent	27
4. Type Specifiers	31
4.1. Type Specifier Symbols	31
4.2. Predicating Type Specifier	31
4.3. Type Specifiers That Combine	32
4.4. Type Specifiers That Specialize	32
4.5. Type Specifiers That Abbreviate	35

4.6. Defining New Type Specifiers	36
4.7. Type Conversion Function	37
4.8. Determining the Type of an Object	38
5. Program Structure	39
5.1. Forms	39
5.1.1. Self-Evaluating Forms	39
5.1.2. Variables	39
5.1.3. Special Forms	40
5.1.4. Macros	41
5.1.5. Function Calls	42
5.2. Functions	42
5.2.1. Named Functions	42
5.2.2. Lambda-Expressions	43
5.3. Top-Level Forms	47
5.3.1. Defining Named Functions	47
5.3.2. Declaring Global Variables and Named Constants	48
5.3.3. Control of Time of Evaluation	49
6. Predicates	51
6.1. Logical Values	51
6.2. Data Type Predicates	52
6.2.1. General Type Predicate	52
6.2.2. Specific Data Type Predicates	52
6.3. Equality Predicates	55
6.4. Logical Operators	58
7. Control Structure	61
7.1. Constants and Variables	62
7.1.1. Reference	62
7.1.2. Assignment	64
7.2. Generalized Variables	65
7.3. Function Invocation	71
7.4. Simple Sequencing	72
7.5. Environment Manipulation	73
7.6. Conditionals	76
7.7. Blocks and Exits	79
7.8. Iteration	80

7.8.1. Simple Iteration	80
7.8.2. General iteration	80
7.8.3. Simple Iteration Constructs	84
7.8.4. Mapping	85
7.8.5. The "Program Feature"	86
7.9. Multiple Values	89
7.9.1. Constructs for Handling Multiple Values	89
7.9.2. Rules for Tail-Recursive Situations	91
7.10. Dynamic Non-local Exits	93
7.10.1. Catch Forms	93
7.10.2. Throw Forms	95
8. Macros	97
8.1. Defining Macros	97
8.2. Expanding Macro Calls	100
9. Declarations	101
9.1. Declaration Syntax	101
9.2. Declaration Forms	103
9.3. Type Declaration for Forms	106
10. Symbols	107
10.1. The Property List	107
10.2. The Print Name	110
10.3. Creating Symbols	111
11. Packages	115
11.1. Built-in Packages	116
11.2. Package System Functions and Variables	116
12. Numbers	121
12.1. Predicates on Numbers	122
12.2. Comparisons on Numbers	122
12.3. Arithmetic Operations	124
12.4. Irrational and Transcendental Functions	127
12.4.1. Exponential and Logarithmic Functions	127
12.4.2. Trigonometric and Related Functions	128

12.4.3. Branch Cuts, Principal Values, and Boundary Conditions in the Complex Plane	130
12.5. Type Conversions and Component Extractions on Numbers	133
12.6. Logical Operations on Numbers	138
12.7. Byte Manipulation Functions	142
12.8. Random Numbers	144
12.9. Implementation Parameters	146
13. Characters	149
13.1. Predicates on Characters	150
13.2. Character Construction and Selection	153
13.3. Character Conversions	155
13.4. Character Control-Bit Functions	156
14. Sequences	159
14.1. Simple Sequence Functions	161
14.2. Catenating, Mapping, and Reducing Sequences	162
14.3. Modifying Sequences	164
14.4. Searching Sequences for Items	167
14.5. Sorting and Merging	169
15. Manipulating List Structure	173
15.1. Conses	173
15.2. Lists	174
15.3. Alteration of List Structure	181
15.4. Substitution of Expressions	182
15.5. Using Lists as Sets	183
15.6. Association Lists	185
16. Hash Tables	189
16.1. Hash Table Functions	190
16.2. Primitive Hash Function	191
17. Strings	193
17.1. String Access	193
17.2. String Comparison	194
17.3. String Construction and Manipulation	195
17.4. Type Conversions on Strings	197

18. Structures	199
18.1. Introduction to Structures	199
18.2. How to Use Defstruct	201
18.3. Using the Automatically Defined Constructor Function	202
18.4. <code>defstruct</code> Slot-Options	203
18.5. Options to <code>defstruct</code>	204
18.6. By-position Constructor Functions	207
19. The Evaluator	209
19.1. Run-Time Evaluation of Forms	209
19.2. The Top-Level Loop	211
20. Streams	213
20.1. Standard Streams	213
20.2. Creating New Streams	214
20.3. Operations on Streams	216
21. Input/Output	219
21.1. Printed Representation of LISP Objects	219
21.1.1. What the <code>read</code> Function Accepts	220
21.1.2. Parsing of Numbers and Symbols	221
21.1.3. Macro Characters	223
21.1.4. Sharp-Sign Abbreviations	229
21.1.5. The Readtable	234
21.1.6. What the <code>print</code> Function Produces	238
21.2. Input Functions	243
21.2.1. Input from ASCII Streams	243
21.2.2. Input from Binary Streams	247
21.3. Output Functions	248
21.3.1. Output to ASCII Streams	248
21.3.2. Output to Binary Streams	250
21.4. Formatted Output	251
21.5. Querying the User	261
22. File System Interface	263
22.1. File Names	263
22.1.1. Pathnames	264
22.1.2. Pathname Functions	266

22.1.3. Defaults and Merging	270
22.1.4. Logical Pathnames	271
22.2. Opening and Closing Files	272
22.3. Renaming, Deleting, and Other Operations	276
22.4. Loading Files	277
22.5. Accessing Directories	279
23. Errors	281
23.1. Signalling Conditions	281
23.2. Establishing Handlers	282
23.3. Error Handlers	283
23.4. Signalling Errors	284
23.5. Standard Condition Names	286
24. Miscellaneous Features	289
24.1. The Compiler	289
24.2. Documentation	290
24.3. Modules	291
24.4. Debugging Tools	292
24.5. Environment Inquiries	295
24.5.1. Time Functions	295
24.5.2. Other Environment Inquiries	297
24.6. Identity Function	299
COMMON LISP Summary	301
Index	317
Index of Concepts	319
Index of Variables and Constants	323
Index of Keywords	327
Index of Functions, Macros, and Special Forms	333

List of Tables

Table 1-1: Sample Function Description	6
Table 1-2: Sample Variable Description	6
Table 1-3: Sample Constant Description	6
Table 1-4: Sample Special Form Description	7
Table 1-5: Sample Macro Description	7
Table 2-1: Minimum Floating-Point Precision and Exponent Size Requirements	14
Table 4-1: Standard Type Specifier Symbols	32
Table 5-1: Names of All COMMON LISP Special Forms	41
Table 21-1: Standard Character Syntax Attributes	222
Table 21-2: Syntax of Numbers	223
Table 21-3: Standard Constituent Character Attributes	224
Table 21-4: Standard Sharp-Sign Macro Character Syntax	230

Acknowledgements

COMMON LISP was designed by a diverse group of people representing many institutions. The many people who have contributed to the design of COMMON LISP are hereby gratefully acknowledged:

Alan Bawden ¹	Richard P. Gabriel ^{5,6}	Jonathan Rees ¹¹
Rodney A. Brooks ¹	Joseph Ginder ³	William L. Scherlis ³
Richard L. Bryan ²	Richard Greenblatt ⁷	Richard M. Stallman ¹
Glenn S. Burke ¹	Martin L. Griss ⁸	Barbara K. Steele ³
Howard I. Cannon ²	Charles L. Hedrick ⁹	William vanMelle ¹⁰
George J. Carrette ¹	Earl A. Killian ⁶	Walter van Roggen ³
David Dill ³	John L. Kulp ²	Allan C. Wechsler ²
Scott E. Fahlman ³	Larry M. Masinter ¹⁰	Daniel L. Weinreb ²
Richard J. Fateman ⁴	John McCarthy ⁵	Jon L White ¹⁰
Neal Feinberg ³	Don Morrison ⁸	Richard Zippel ¹
John Foderaro ⁴	David A. Moon ²	Leonard Zubkoff ³
	Kent M. Pitman ¹	

1. Massachusetts Institute of Technology, 545 Technology Square, Cambridge, Massachusetts 02139
2. Symbolics, Inc., Cambridge, Massachusetts 02139
3. Computer Science Department, Carnegie-Mellon University, Schenley Park, Pittsburgh, Pennsylvania 15213
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
11. Department of Computer Science, Yale University, New Haven, Connecticut 06520

Special thanks go to Jon Bentley, Scott Fahlman {and others???) for extraordinarily careful proofreading of the final drafts of the manuscript.

The organization, typography, and content of this document were inspired in large part by the *MacLISP Reference Manual* by David A. Moon and others [9], and by the *LISP Machine Manual* by Daniel Weinreb and David Moon [15], which in turn acknowledges the efforts of Richard Stallman, Mike McMahon, Alan Bawden, Glenn Burke, and "many people too numerous to list".

Notes on This Edition

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

The chapter on the evaluator does not contain the proposed evaluator code, which is still under review.

The package system is not yet fully designed. The specification given here is likely to be changed.

I deleted the functions `array-length` and `array-active-length`; for one-dimensional arrays, the functions `array-dimension` and `length` may be used instead.

I couldn't think of a good name for the `defstruct array-leader` type option, so I have temporarily flushed it, along with the `:make-array` option.

The function `stream-element-type` (page 227) was added so that you can tell what you got if you specified `:type :default` to `open` (page 283).

The case for a floating-point specifier, apparently mandated to be lower-case by the October 1982 ballot (issue 1), is not specified in this edition. While an upper-case "S" can be confused with the digit "5", so may a lower-case "l" be confused with the digit "1".

The quote data type specifier mandated by the October 1982 ballot (issue 3) has been purposely omitted from this edition. There is some question as to the inconsistency of requiring objects to be quoted in type specifiers that need not be quoted in ordinary executable forms, such as numbers, strings, and keywords. This inconsistency is not easily resolved, because `t` and `nil` mean something as type specifiers other than their quoted selves.

Except for the above two issues, all issues from the October 1982 ballot have more or less been accounted for in this edition. The results from the August 1982 meeting have been incorporated except as noted below.

For August 1982 issue 56, the functions `read-binary-object` (page 257) and `write-binary-object` (page 260) are proposed.

For August 1982 issue 62, I could not find an appropriate definition for the requested new version of `catch-all`. In this edition I have instead fixed up the old definitions of `catch-all` and `unwind-all` (page 93).

For August 1982 issue 78, I am worried that a lexical declaration to shadow a global special declaration will unduly slow down the interpreter. Do we really need this in practice? (You can't detect whether to put one in until you have discovered that you are losing, in which case renaming does the job.)

For August 1982 issue 88, function specs are still under debate and are not included here.

For August 1982 issue 118, a new primitive for getting file information is not yet proposed.

For August 1982 issues 126, 127, and 128, the proposed format floating-point, picture, and metric-prefix directives are omitted here.

For August 1982 issue 153, I am unconvinced that making trace and untrace (page 302) be functions will be more convenient than their present definition as macros. No proposal has yet been made regarding keyword arguments, ways to find out what is currently being traced, and a way to untrace all traced functions. Therefore no change is reflected here.

There are problems with the formatting of the index. These include doubled titles, extraneous blank pages, and failure to indent continuation lines when there are many page references. These should not be too difficult to fix, but I did not want to delay this edition purely to fix formatting bugs.

Chapter 1

Introduction

This manual documents a dialect of LISP called "COMMON LISP", which is a successor to MACLISP [9], influenced strongly by Lisp Machine LISP [15] and also to some extent by SCHEME [12] and INTERLISP [14].

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 goals of COMMON LISP are thus very close to those of STANDARD LISP [8]. COMMON LISP differs from STANDARD LISP primarily in incorporating more features, including a richer and more complicated set of data types and more complex control structures.

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, and sizes of floating-point exponents and fractions, as well as implementation-dependent information such as the nature of the file system, the method of invoking the implementation, and so on.
- The *blue pages* constitutes an implementation guide in the spirit of the INTERLISP virtual machine specification [10]. 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.

1.2. Notational Conventions

In COMMON LISP, as in most LISP dialects, the symbol `nil` (page 51) 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 will be signalled; see `error` (page 294) and `cerror` (page 295).
- 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 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. Purely as a matter of convention, all global variables used by COMMON LISP have names beginning and ending with an asterisk.

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.)

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. Brackets, braces, stars, plus signs, and vertical bars are metasyntactic marks. 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

`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

("+"); 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 summary, 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". These notations are also used for syntactic descriptions expressed as BNF-like productions, for example in Table

`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 79) 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

22-2.

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 253)) 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 22, 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 105₈ (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.

The square brackets, braces, question mark, and exclamation point (that is, "[", "]", "{", "}", "?", and "!") are not used for any purpose in standard COMMON LISP syntax. These characters are explicitly reserved to the user, primarily for use as *macro characters* for user-defined syntax extensions. See section 22.1.3 (page 233).

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 pretend that the colons are somewhat like quote marks, causing such symbols to evaluate to themselves.

Chapter 2

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 constitutes advice from the user that may be useful in gaining efficiency. See `declare` (page 101).)

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 52) may be used to ask the latter question, and the function `type-of` (page 38) to ask the former.

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`. A type called `common` encompasses all the data objects required by the COMMON LISP language. A COMMON LISP implementation is free to provide other data types that are not subtypes of `common`.

COMMON LISP objects may be roughly divided into the following categories: numbers, characters, symbols, lists, arrays, structures, and functions. 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), by causing one to be *displaced* to the other. One-dimensional arrays of any kind are called *vectors*. One-dimensional arrays of characters are called *strings*. One dimensional arrays of bits (that is, of integers whose values are 0 or 1) are called *bit-vectors*.
- *Hash tables* provide an efficient way of mapping any LISP object (a *key*) to an associated object.
- *Readtables* are used to control the built-in expression parser `read` (page 253).
- *Packages* are collections of symbols that serve as name spaces. The parser recognizes symbols by looking up character sequences in the "current package".
- *Pathnames* represent names of files in a fairly implementation-independent manner. They are used to interface to the external file system.
- *Streams* represent sources or sinks of data (typically characters or bytes). They are used to

perform I/O, as well as for internal purposes such as parsing strings.

- *Random-states* are data structures used to encapsulate the state of the built-in random-number generator.
- *Structures* are user-defined record structures, objects that have named components. The `defstruct` (page 211) facility is used to define new structure types. Some COMMON LISP implementations may choose to implement certain system-supplied data types as structures such as *bignums*, *readtables*, *streams*, *hash tables*, 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 *compiled-functions* (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.

These categories are not always mutually exclusive. The required relationships among the various data types are explained in more detail in section 2.15 (page 24).

2.1. Numbers

There are several kinds of numbers defined in COMMON LISP. They are divided into *rational numbers*, consisting of integers and ratios; *floating-point numbers*, with names provided for up to four different precisions; and *complex numbers*.

2.1.1. 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 146) and `most-negative-fixnum` (page 146).

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 ; $e^{\pi i}$ 
15511210043330985984000000. ;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 dddd or #nnR dddd

means the integer in radix-*nn* notation denoted by the digits *dddd*. 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.

```

2.1.2. Ratios

A *ratio* is a number representing the mathematical ratio of two integers. Integers and ratios are collectively called *rational*s. 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 258) 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 “/”. This syntax may be described as follows:

$$\text{ratio} ::= [\text{sign}] \{\text{digit}\}^+ / \{\text{digit}\}^+$$

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 #mR, #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.

2.1.3. Floating-point Numbers

A floating-point number is a (mathematical) rational number of the form $s*f*b^{e-p}$, where s is $+1$ 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 *significand*; and e is an integer, the *exponent*. The value of p and the range of e depends on the implementation and on the type of floating-point number within that implementation. In addition, there is a floating-point zero; depending on the implementation, there may also be a “minus zero”. If there is no minus zero, then “0.0” and “-0.0” are both interpreted as simply a floating-point zero.

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 `float-radix` (page 137).

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 $p \cdot \log_2 b$) 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 2-1.

<u>Format</u>	<u>Minimum Precision</u>	<u>Minimum Exponent Size</u>
Short	13 bits	5 bits
Single	24 bits	8 bits
Double	50 bits	8 bits
Long	50 bits	8 bits

Table 2-1: 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 [6, 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 description of floating-point notation.

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 | 1 | 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 253) 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.

Examples of floating-point numbers:

<code>0.0</code>	; Floating-point zero in default format.
<code>0E0</code>	; Also floating-point zero in default format.
<code>-.0</code>	; This may be a zero or a minus zero, ; depending on the implementation.
<code>0.</code>	; The <i>integer</i> zero, not a floating-point number!
<code>0.0s0</code>	; A floating-point zero in <i>short</i> format.
<code>0s0</code>	; Also a floating-point zero in <i>short</i> format.
<code>3.1415926535897932384d0</code>	; A <i>double</i> -format approximation to π .
<code>6.02E+23</code>	; Avogadro's number, in default format.
<code>602E+21</code>	; Also Avogadro's number, in default format.
<code>3.1010299957f-1</code>	; $\log_{10} 2$, in <i>single</i> format.
<code>-0.000000001s9</code>	; $e^{\pi i}$ in <i>short</i> format, the hard way.

2.1.4. 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 137).)

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.
```

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).

2.2. 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 149).

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-\beta
#\C-M-Return
```

Note that not all COMMON LISP implementations provide for non-zero bits attributes; see `char-bits-limit` (page 149).

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`.

2.3. Symbols

Symbols are LISP data objects that serve several purposes and have several interesting characteristics. Every symbol has a name, called its *print name*. 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:

```

\(  

\+1  

+\1  

\frobboz  

3.14159265\s0  

3.14159265\S0  

3.14159265s0  

APL\360  

ap1\360  

\ (b^2)\ -\ 4*a*c

```

; The symbol whose name is “(”.
; The symbol whose name is “+1”.
; Also the symbol whose name is “+1”.
; The symbol whose name is “fROBBOZ”.
; The symbol whose name is “3.14159265s0”.
; The symbol whose name is “3.14159265S0”.
; A short-format floating-point approximation to π .
; The symbol whose name is “APL\360”.
; Also the symbol whose name is “ap1\360”.
; 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:

```

|"  

|(b^2) - 4*a*c|  

|frobboz|  

|APL\360|  

|APL\360|  

|ap1\360|  

|\|\|

```

; The same as writing \".
; The name is “(b^2) - 4*a*c”.
; The name is “frobboz”, not “FROBBOZ”.
; The name is “APL360”, because
; the “\” quotes the “3”.
; The name is “APL\360”.
; The name is “ap1\360”.
; Same as \|\|: the name is “|”.

2.4. Lists and Conses

A *cons* is a 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)  

(2.0s0 (a 1) #\*)

```

; A list of three symbols.
; 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 need not be surrounded by white space or other delimiters. The dot is required to be delimited in COMMON LISP, as 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.

2.5. 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 is itself a non-negative *fixnum*. The total number of elements in the array is the product of all the dimensions.

An implementation of COMMON LISP may impose a limit on the rank of an array, but this limit may not be

smaller than 63. Therefore, any COMMON LISP program may assume the use of arrays of rank 63 or less.

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 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 196):

```
(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.

In general, arrays can be multi-dimensional, can share their contents with other array objects, and can have their size altered dynamically (either enlarging or shrinking) after creation. A one-dimensional array may also have a *fill pointer*.

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 a matter of semantics; the second, a matter of efficiency.)

2.5.1. Vectors

One-dimensional arrays are called *vectors* in COMMON LISP. Vectors 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.

A vector that is not displaced to another array, has no fill pointer, and is not to have its size adjusted dynamically after creation, is called a *simple* vector. Some implementations can handle simple vectors in an especially efficient manner. The user may provide declarations that certain arrays will be simple vectors. Simple vectors may have a more compact representation than non-simple vectors.

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.
```

When the function `read` parses this syntax, it always constructs a *simple* general vector.

Rationale: Many 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 253), 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 square brackets and braces for any purpose.

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 rather, a special subset of the characters); the one-dimensional instances of this specialization are called *strings*. All implementations are also required to provide specialized arrays of bits, that is, arrays of type `(array bit)`; the one-dimensional instances of this specialization are called *bit-vectors*.

2.5.2. Strings

A string 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. The function `prin1` will print any character vector using this syntax, but the function `read` will always construct a *simple* string from this syntax.

2.5.3. Bit-vectors

A bit-vector can be written as the sequence of bits contained in the string, preceded by “#*”; any delimiter character (such as whitespace) will terminate the bit-vector syntax.

For example:

```
##*10110 ; A five-bit bit-vector; bit 0 is a 1.  
##* ; An empty bit-vector.
```

The function `prin1` will print any bit-vector using this syntax, but the function `read` will always construct a *simple* bit-vector from this syntax.

2.6. Hash tables

Hash tables provide an efficient way of mapping any LISP object (a *key*) to an associated object. They are provided as primitives of COMMON LISP because some implementations may need to use internal storage management strategies that would make it very difficult for the user to implement hash tables himself in a portable fashion. Hash tables are described in chapter 16 (page 189).

2.7. Readtables

A readtable is a data structure that maps characters into syntax types for the LISP expression parser. In particular, a readtable indicates for each character with syntax *macro character* what its macro definition is. This is a mechanism by which the user may reprogram the parser to a limited but useful extent. See section 22.1.5 (page 244).

2.8. Packages

Packages are collections of symbols that serve as name spaces. The parser recognizes symbols by looking up character sequences in the "current package". Packages can be used to hide names internal to a module from other code. Mechanisms are provided for exporting symbols from a given package to the primary "user" package. See chapter 11 (page 115).

2.9. Pathnames

Pathnames are the means by which a COMMON LISP program can interface to an external file system in a reasonably implementation-independent manner. See section 23.1.1 (page 274).

2.10. Streams

A stream is a source or sink of data, typically characters or bytes. Nearly all function that perform I/O do so with respect to a specified stream. The function `open` (page 283) takes a pathname and returns a stream connected to the file specified by the pathname. There are a number of standard streams that are used by default for various purposes. See chapter 21 (223).

2.11. Random-states

For information about random-state objects and the random-number generator, see section 12.8 (page 144).

2.12. Structures

Structures are instances of user-defined data types that have a fixed number of named components. They are analogous to records in PASCAL. Structures are declared using the `defstruct` (page 211) construct; `defstruct` automatically defines access and constructor functions for the new data type.

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 `:print-function` (page 217) option to `defstruct`). 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.

2.13. Functions

A *function* is anything that may be correctly given to the `funcall` (page 71) or `apply` (page 71) function, to be executed as code when arguments are supplied.

A *compiled-function* is a compiled code object.

A list whose *car* is `lambda` 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 `symbol-function` (page 63).

2.14. Unreadable Data Objects

Some objects 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 “#<”.

2.15. Overlap, Inclusion, and Disjointness of Types

The COMMON LISP data type hierarchy is tangled, and purposely left somewhat open-ended so that implementors may experiment with new data types as extensions to the language. This section states explicitly all the defined relationships between types, including subtype/supertype relationships, disjointness, and exhaustive partitioning. The user of COMMON LISP should not depend on any relationships not explicitly stated here. For example, it is not valid to assume that because a number is not complex and not rational that it must be a `float`, because implementations are permitted to provide yet other kinds of numbers.

First we need some terminology. If x is a supertype of y , then any object of type y is also of type x , and y is said to be a subtype of x . If types x and y are disjoint, then no object (in any implementation) may be both of type x and of type y . Types a_1 through a_n are an *exhaustive union* of type x if each a_j is a subtype of x , and any object of type x is necessarily of at least one of the types a_j ; a_1 through a_n are furthermore an *exhaustive partition* if they are also pairwise disjoint.

- The type `t` is a supertype of every type whatsoever. Every object belongs to type `t`.
- The type `nil` is a subtype of every type whatsoever. No object belongs to type `nil`.
- The types `cons`, `symbol`, `array`, `number`, and `character` are pairwise disjoint.
- The types `rational`, `float`, and `complex` are pairwise disjoint subtypes of `number`.
- The types `integer` and `ratio` are disjoint subtypes of `rational`.
- The types `fixnum` and `bignum` are disjoint subtypes of `integer`.
- The types `short-float`, `single-float`, `double-float`, and `long-float` are subtypes of `float`. Any two of them must be either disjoint or identical; if identical, then any other types between them in the above ordering must also be identical to them (for example, if `single-float` and `long-float` are identical types, then `double-float` must be identical to them also).
- The type `null` is a subtype of `symbol`; the only object of type `null` is `nil`.
- The types `cons` and `null` form an exhaustive partition of the type `list`.
- The type `standard-char` is a subtype of `string-char`; `string-char` is a subtype of `character`.

- The type `string` is a subtype of `vector`, for `string` means `(vector string-char)`.
- The type `bit-vector` is a subtype of `vector`, for `bit-vector` means `(vector bit)`.
- The type `vector` is a subtype of `array`; for all types `x`, the type `(vector x)` is a subtype of the type `(array x (*))`, the set of all one-dimensional arrays.
- The types `hash-table`, `readtable`, `package`, `pathname`, `stream`, and `random-state` are pairwise disjoint.
- Any two types created by `defstruct` (page 211) are disjoint unless one is a supertype of the other by virtue of the `:include` (page 215) option.
- An exhaustive union for the type `common` is formed by the types `cons`, `symbol`, `(array x)` where `x` is a subtype of `common`, `fixnum`, `bignum`, `ratio`, `short-float`, `single-float`, `double-float`, `long-float`, `(complex x y)` where `x` and `y` are subtypes of `common`, `standard-char`, `hash-table`, `readtable`, `package`, `pathname`, `stream`, `random-state`, and all types created by `defstruct`. No data type not mentioned in this document may be a subtype of `common`. Note that a type such as `number` or `array` may or may not be a subtype of `common`, depending on whether or not the given implementation has extended the set of objects of that type.

Chapter 3

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 copy-cell (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 `copy-cell`) 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.

- *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 286) 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 destroy the entity if it can prove that reference to it is no longer possible.)

Example: most COMMON LISP data objects have indefinite extent.

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 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 73) 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 95) 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 93), `catch-all` (page 93), `unwind-all` (page 93), or `unwind-protect` (page 94) special form has dynamic scope.
- An exit point established by a `block` (page 79) construct has lexical scope and dynamic extent. (Such exit points are also established by `do` (page 80), `prog` (page 87), and other iteration constructs.)
- The tags established by a `prog` (page 87) and referenced by `go` (page 89) have lexical scope and dynamic extent.

- Named constants such as `nil` (page 51) and `pi` (page 130) have indefinite scope and indefinite 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 79) exit points outstanding, each apparently named `here`. However, the `return-from` (page 79) 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 62) construct (here abbreviated by the `#'` syntax) was executed.

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 211), 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.

4.2. Predicating Type Specifier

A type specifier list (`satisfies predicate-name`) denotes the set of all objects that satisfy the predicate named by *predicate-name*, which must be a symbol whose global function definition is a one-argument predicate. (A name is required; lambda-expressions are not allowed in order to avoid scoping problems.) For example, the type `(satisfies numberp)` is the same as the type `number`. The call `(typep x '(satisfies p))` results in applying `p` to `x` and returning `t` if the result is true and `nil` if the result is false.

As an example, the type `string-char` could be defined as

```
(deftype string-char () (and character (satisfies string-charp)))
```

See `deftype` (page 36).

As a rule, a predicate appearing in a `satisfies` type specifier should not cause any side effects when invoked.

4.3. Type Specifiers That Combine

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

t	nil	common	
null	cons	list	symbol
array	vector	bit-vector	string
sequence	simple-vector	simple-bit-vector	simple-string
function	compiled-function	pathname	character
number	rational	float	string-char
integer	ratio	short-float	standard-char
fixnum	complex	single-float	package
bignum	random-state	double-float	stream
bit	readtable	long-float	hash-table

Table 4-1: Standard Type Specifier Symbols

(member *object1 object2 ...*)

This denotes the set containing precisely those objects named. An object is of this type if and only if it is `eq1` (page 56) to one of the specified objects.

Compatibility note: This is approximately equivalent to what the INTERLISP DECL package calls `memq`.

(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 168) 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`.

??? Query: Should `or` and `and` type specifiers guarantee the order in which the types are examined? This matters if a `satisfies` type specifier has side effects or if it relies on previous type restrictions, as in writing `(typep x '(and number (satisfies primep)))`.

4.4. 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 an array for the express purpose of holding only short-float objects, you may optionally specify to `make-array` (page 193) 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 *element-type dimensions*)

This denotes the set of specialized arrays whose elements are all members of the type *element-type* and whose dimensions match *dimensions*. For declaration purposes, this type encompasses those arrays that can result by specifying *element-type* as the element type to the function `make-array` (page 193); this may be different from what the type means for discrimination purposes. *element-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-format
                           ;   floating-point numbers.
```

Note that (array t) is a proper subset of (array *).

(vector *element-type size*)

This denotes the set of specialized one-dimensional arrays whose elements are all of type *element-type* and whose lengths match *size*. This is entirely equivalent to `(array element-type (size))`.

For example:

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

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

`(simple-vector element-type size)`

This is the same as `(vector element-type size)` except that it includes only simple vectors.

`(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 137); 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 ...) value-type)`

This type may be used only for declaration and not for discrimination; `typep` (page 52) 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 a value that is a member of the types specified by the *value-type* form. The `&optional`, `&rest`, and `&key` keywords may appear in the list of argument types. The *value-type* may be a `values` type specifier, to indicate the types of multiple values.

As an example, the function `cons` (page 174) 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). The function `truncate` (page 135) is of type `(function (number number) (values number number))`, as well as of type `(function (integer (mod 8)) integer)`.

`(values value1-type value2-type ...)`

This type specifier is extremely restricted: it may be used *only* as the *value-type* in a function type specifier or in a `the` (page 106) declaration. It is used to specify individual types when multiple values are involved. The `&optional`, `&rest`, and `&key` keywords may appear in the *value-type* list; they thereby indicate the parameter list of a function that, when given to `multiple-value-call` (page 90) along with the values, would be suitable for receiving those values.

4.5. 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, `member`).

`(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 -2s-1 2s-1-1)`. Simply `signed-byte` or `(signed-byte *)` is the same as `integer`.

`(unsigned-byte s)`

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

`(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 (array string-char (*size*)): the set of strings of the indicated size. One may also use the name simple-string to include only simple strings.

(bit-vector *size*)

This means the same as (array bit (*size*)): the set of bit-vectors of the indicated size. One may also use the name simple-bit-vector to include only simple bit-vectors.

4.6. Defining New Type Specifiers

New type specifiers can come into existence in two ways. First, defining a new structure type with `defstruct` (page 211) automatically causes the name of the structure to be a new type specifier symbol. Second, the `deftype` special form can be used to define new type-specifier abbreviations.

```
deftype name lambda-list {declaration | doc-string}* {form}* [Macro]
```

This is very similar to a `defmacro` (page 99) 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.

If the optional documentation string *doc-string* is present, then it is attached to the *name* as a documentation string of type *type*; see *documentation* (page 301).

For example:

```
(deftype mod (n) '(integer 0 (,n)))
(deftype list () '(or null cons))
(deftype square-matrix (&optional type size)
  "SQUARE-MATRIX includes all square two-dimensional arrays."
  '(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. A better definition is:

```
(defun equidimensional (a)
  (or (< (array-rank a) 2)
      (apply #'= (array-dimensions a))))

(deftype square-matrix (&optional type size)
  '(and (array ,type (,size ,size))
        (satisfies equidimensional)))
```

4.7. Type Conversion Function

coerce *object* *result-type*

[Function]

The *result-type* must be a type specifier; the *object* is converted to an “equivalent” object of the specified type. As a rule, if *object* is already of the specified type, as determined by *typep* (page 52), then it is simply returned. It is not generally possible to convert any object to be of any type whatsoever; only certain conversions are permitted:

- Any sequence type may be converted to any other sequence type, provided that the new sequence can contain all actual elements of the old sequence (it is an error if it cannot). If the *result-type* is specified as simply *array*, for example, then `(array t)` 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. If the *sequence* is already of the specified type, it may be returned without copying it; in this `(coerce type sequence)` differs from `(concatenate type sequence)`, for the latter is required to copy the argument *sequence*. In particular, if one specifies *sequence*, then the argument may simply be returned, if it already is a sequence.

```
(coerce '(a b c) 'vector) => #(a b c)
```

- Some strings, symbols, and integers may be converted to characters. If *object* is a string of length 1, then the sole element of the string is returned. If *object* is a symbol whose print name is of length 1, then the sole element of the print name is returned. If *object* is an integer *n*, then (int-char *n*) is returned. See character (page 154).

```
(coerce "a" 'character) => #\a
```

- Any non-complex number can be converted to be a short-float, single-float, double-float, or long-float. If simply float is specified, and *object* is not already a float of some kind, then the object is converted to be a single-float.

```
(coerce 0 'short-float) => 0.0S0
(coerce 3.5L0 'float) => 3.5L0
(coerce 7/2 'float) => 3.5
```

- Any number can be converted to be a complex number. If the number is not already complex, then a zero imaginary part is provided by coercing the integer zero to the type of the given real part.

```
(coerce 4.5s0 'complex) => #C(4.5S0 0.0S0)
(coerce 7/2 'complex) => #C(7/2 0)
(coerce #C(7/2 0) '(complex double-float))
=> #C(3.5D0 0.0D0)
```

Coercions from floating-point numbers to rationals and from ratios to integers are purposely *not* provided, because of rounding problems. The functions rational (page 134), rationalize, floor (page 135), ceiling, truncate, and round may be used for such purposes.

4.8. Determining the Type of an Object

type-of *object*

[Function]

(type-of *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. If the argument is a user-defined named structure created by defstruct then type-of will return the type name of that structure. Because the result is implementation-dependent, it is usually better to use type-of of one argument primarily for debugging purposes; however, there are a few situations where portable code requires the use of type-of, such as when the result is to be given to the coerce (page 37) or map (page 163) function. On the other hand, often the typep (page 52) function or the typecase construct is more appropriate for some purpose than type-of.

Compatibility note: In MACLISP this function is called typep, and anomalously so, for it is not a predicate.

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. (Any COMMON LISP data object not explicitly defined to be a valid form is not a valid form, and attempting to evaluate such an object will cause an error to be signalled.)

5.1.1. Self-Evaluating Forms

All numbers, characters, strings, and bit-vectors are *self-evaluating* forms. When such an object is evaluated, 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

64), or *bound*, as by `let` (page 73). 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 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 `defconstant` (page 48) 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 `defconstant`. 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 or values, but it may instead call for a non-local exit; see `return-from` (page 79), `go` (page 89), and `throw` (page 95).

The set of special forms is fixed in COMMON LISP; no way is provided for the user to define more. The

and	(page 58)	return-from	(page 79)
or	(page 59)	tagbody	(page 87)
quote	(page 62)	go	(page 89)
function	(page 62)	multiple-value-call	(page 90)
setq	(page 64)	multiple-value-prog1	(page 90)
progn	(page 72)	catch	(page 93)
let*	(page 74)	catch-all	(page 93)
progv	(page 75)	unwind-all	(page 93)
flet	(page 75)	unwind-protect	(page 94)
labels	(page 75)	throw	(page 95)
macrolet	(page 75)	declare	(page 101)
if	(page 77)	the	(page 106)
block	(page 79)		

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

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

user can create new syntactic constructs, however, by defining macros.

The set of special forms in COMMON LISP is purposely kept very small, because any program-analyzing program must have special knowledge about every type of special form. Such a program needs no special knowledge about macros, because it is simple to expand the macro and operate on the resulting expansion. (This is not to say that many such programs, particularly compilers, will not have such special knowledge. A compiler may be able to produce much better code if it recognizes such constructs as `typecase` and `multiple-value-bind` and gives them customized treatment.)

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 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 `return` (page 79) will take a form such as `(return x)` and from that form compute a new form `(return-from nil x)`. 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 99).

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 95). A function that does return may produce no value or several values; see `values` (page 89). If and when the function returns, whatever values it returns become the values 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`. Therefore 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 62) 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 47) special form. A *local name* can be given to a function by using the `labels` (page 75) 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 47) and `labels` (page 75) 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]])}]*
```

```
          [&allow-other-keys]
```

```
        [&aux {var | (var [initform])}]*)
```

```
        {declaration | documentation-string}*  
        {form}*)
```

Each element of a lambda-list is either a *parameter specifier* or a *lambda-list keyword*; lambda-list keywords begin with “&”. (Note that lambda-list keywords are not keywords in the usual sense; they do not belong to the keyword package. They are ordinary symbols whose name begins with an ampersand.)

In all cases a *var* must be a symbol, the name of a variable, and similarly for *svar* also; each *keyword* must be a keyword symbol, such as “:start”. An *initform* may be any form.

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

- Specifiers for the *required* parameters. These are all the parameter specifiers up to the first lambda-list keyword; if there is no such lambda-list keyword, then all the specifiers are for required parameters.
- Specifiers for *optional* parameters. If the lambda-list keyword `&optional` is present, the *optional* parameter specifiers are those following the lambda-list keyword `&optional` up to the next lambda-list keyword or the end of the list.
- A specifier for a *rest* parameter. The lambda-list keyword `&rest`, if present, must be followed by a single *rest* parameter specifier, which in turn must be followed by another lambda-list keyword or the end of the lambda-list.
- Specifiers for *keyword* parameters. If the lambda-list keyword `&key` is present, all specifiers up to the next lambda-list keyword or the end of the list are *keyword* parameter specifiers. The keyword parameter specifiers may optionally be followed by the lambda-list keyword

`&allow-other-keys`.

- Specifiers for *aux* variables. These are not really parameters. If the lambda-list keyword `&aux` is present, all specifiers after it are *auxiliary variable* specifiers.

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 101)).

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, then 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 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.) If there is no *rest* parameter and there are no *keyword* parameters, then there should be no unprocessed arguments (it is an error if there are).

Next any *keyword* parameters are processed. For this purpose the same arguments are processed that would be made into a list for a *rest* parameter. (Indeed, it is permitted to specify both `&rest` and `&key`; in this case the arguments are used for both purposes. This is the only situation in which an argument is used in the processing of more than one parameter specifier.) If `&key` is specified, there must remain an even number of arguments; these are considered as pairs, the first argument in each pair being interpreted as a keyword name and the second as the corresponding value. It is an error for the first object of each pair to be anything but a keyword.

Rationale: This last restriction is imposed so that a compiler may issue warnings about malformed calls to functions that take keyword arguments.

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 to indicate the keyword name, 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)) ...)
```

The keyword parameter specifiers are, like all parameter specifiers, effectively processed from left to right. For each keyword parameter specifier, if there is an argument pair whose keyword name matches that specifier's keyword name (that is, the names are eq), then the parameter variable for that specifier is bound to the second item (the value) of that argument pair. If more than one such argument pair matches, it is not an error; the leftmost argument pair is used. 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, unless `&allow-other-keys` was specified, in which case the argument pair is simply ignored (but such an argument pair is accessible through the `&rest` parameter if one was specified).

After all parameter specifiers have been processed, the auxiliary variable specifiers (those following the lambda-list keyword `&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 73):

```
(lambda (x y &aux (a (car x)) (b 2) c) ...)
<=> (lambda (x y) (let ((a (car x)) (b 2) c) ...))
```

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).

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 a lexical binding can later be reinstated 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)

```

Examples of mixtures:

```

((lambda (a &optional (b 3) &rest x &key c (d a))
 (list a b c d x))
 1) => (1 3 () nil 1)

((lambda (a &optional (b 3) &rest x &key c (d a))
 (list a b c d x))
 1 2) => (1 2 () nil 1)

((lambda (a &optional (b 3) &rest x &key c (d a))
 (list a b c d x))
 :c 7) => (:c 7 () nil :c)

((lambda (a &optional (b 3) &rest x &key c (d a))
 (list a b c d x))
 1 6 :c 7) => (1 6 (:c 7) 7 1)

((lambda (a &optional (b 3) &rest x &key c (d a))
 (list a b c d x))
 1 6 :d 8) => (1 6 (:d 8) nil 8)

((lambda (a &optional (b 3) &rest x &key c (d a))
 (list a b c d x))
 1 6 :d 8 :c 9 :d 10) => (1 6 (:d 8 :c 9 :d 10) 9 8)

```

All lambda-list keywords are permitted, but not terribly useful, in lambda-expressions appearing explicitly as the first element of a function-call form, as shown in the examples above. They are extremely useful, however, in functions given global names by `defun` (page 47).

All symbols whose names begin with “&” are conventionally reserved for use as lambda-list keywords and should not be used as variable names. Implementations of COMMON LISP are free to provide additional lambda-list keywords.

`lambda-list-keywords`

[*Constant*]

The value of `lambda-list-keywords` is a list of all the lambda-list keywords used in the implementation, including the additional ones used only by `defmacro` (page 99). It must contain at least the symbols `&optional`, `&rest`, `&key`, `&allow-other-keys`, `&aux`, `&body`, and `&whole`.

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 forms 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.

It is not illegal to use these forms at other than top level, but whether it is meaningful to do so depends on context. Compilers, for example, may not recognize these forms properly in other than top-level contexts. (As a special case, however, if a `progn` (page 72) form appears at top level, then all forms within that `progn` are considered by the compiler to be top-level forms.)

Compatibility note: In MACLISP, a top-level `progn` is considered to contain top-level forms only if the first form is “(quote compile)”. This odd marker is unnecessary in COMMON LISP.

Macros are usually defined by using the special form `defmacro` (page 99). This facility is fairly complicated, and is described in Chapter 8.

5.3.1. Defining Named Functions

`defun name lambda-list {declaration | doc-string}* {form}*`

[*Macro*]

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

`(lambda lambda-list {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 (if not followed by a declaration, 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 attached to the *name* as a documentation string of type `function`; see `documentation` (page 301).

The body of the defined function is implicitly enclosed in a `block` (page 79) construct whose name is the same as the *name* of the function. Therefore `return` (page 79) and `return-from` (page 79) 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!).

5.3.2. Declaring Global Variables and Named Constants

```
defvar name [initial-value [documentation]] [Macro]
defparameter name initial-value [documentation] [Macro]
defconstant name initial-value [documentation] [Macro]
```

`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 101)), 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. The *initial-value* form 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).

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

`defparameter` is similar to `defvar`, but requires an *initial-value* form, and always evaluates it and assigns the result to the variable. The semantic distinction is that `defvar` is intended to declare a variable changed by the program, whereas `defparameter` is intended to declare a variable that is normally constant, but can be changed (possibly at run time), considered as a change *to* the program. `defparameter` therefore does not indicate that the quantity *never* changes; in particular, it does not license the compiler to build assumptions about the value into programs being compiled.

`defconstant` is like `defparameter`, but *does* assert that the value of the variable *name* is fixed, and does license the compiler to build assumptions about the value into programs being compiled. It is an error if there are any special bindings of the variable at the time the `defconstant` form is executed (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 `equalp` (page 57) to the specified *initial-value*.

Once a name has been declared by `defconstant` to be constant, any further assignment to or binding of that special variable is an error. This is the case for such system-supplied constants as `t` (page 51) and `most-positive-fixnum` (page 146). A compiler may also choose to issue warnings about bindings of the lexical variable of the same name.

For any of these constructs, the documentation should be a string. It is attached to the name of the variable, parameter, or constant under the `variable` documentation type; see `documentation` (page 301).

5.3.3. Control of Time of Evaluation.

`eval-when` (*{situation}**) *{form}** [Function]

The body of an `eval-when` form is processed as an implicit `progn`, but only in the situations listed. A *situation* may be `compile`, `load`, or `eval`.

`eval` specifies that the interpreter should process the body. `compile` specifies that the compiler should evaluate the body *at compile time* in the compilation context. `load` specifies that the compiler should arrange to evaluate the forms in the body when the compiled file containing the `eval-when` form is loaded.

The default interpretation is that top-level forms are effectively processed in `eval` and `load` situations. `eval-when` is occasionally useful to get different effects. For example, if the compiler is to be able to read a file properly that uses user-defined reader macro characters, it is necessary to write

```
(eval-when (compile load eval)
  (set-macro-character #\$ #'(lambda (stream char)
    (declare (ignore char))
    (list 'dollar (read stream))))))
```


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 76), `if` (page 77), `when` (page 77), and `unless` (page 77) 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 183).

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 type* [Function]

`typep` 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`. A specifier of the form `(satisfies fn)` is handled simply by applying *fn* to *object*; the *object* is considered to be of the specified type if the result is not `nil`.

`subtypep` *type1 type2* [Function]

The arguments must be type specifiers that are acceptable to `typep` (page 52). The two type specifiers are compared; this predicate is true if *type1* is definitely a (not necessarily proper) subtype of *type2*. If the result is `nil`, however, then *type1* may or may not be a subtype of *type2* (sometimes it is impossible to tell, especially when `satisfies` type specifiers are involved). A second returned value indicates the certainty of the result; if it is true, then the first value is an accurate indication of the subtype relationship. Thus there are three possible result combinations:

<code>t t</code>	<i>type1</i> is definitely a subtype of <i>type2</i>
<code>nil t</code>	<i>type1</i> is definitely not a subtype of <i>type2</i>
<code>[nil nil]</code>	<code>subtypep</code> could not determine the relationship

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 58); 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. 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. Note that the empty list is not a cons, so (*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))
```

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 "integer" or "fixnum", 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)
- bit-vector-p object* [Function]
bit-vector-p is true if its argument is a bit-vector, and otherwise is false.
(bit-vector-p x) <=> (typep x 'bit-vector)
- vectorp object* [Function]
vectorp is true if its argument is a vector, and otherwise is false.
(vectorp x) <=> (typep x 'vector)
- simple-string-p object* [Function]
simple-string-p is true if its argument is a simple string, and otherwise is false.
(simple-string-p x) <=> (typep x 'simple-string)
- simple-bit-vector-p object* [Function]
simple-bit-vector-p is true if its argument is a simple bit-vector, and otherwise is false.
(simple-bit-vector-p x) <=> (typep x 'simple-bit-vector)
- simple-vector-p object* [Function]
simple-vector-p is true if its argument is a simple vector, and otherwise is false.
(simple-vector-p x) <=> (typep x 'simple-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.

`compiled-function-p` *object* [Function]

`compiled-function-p` is true if its argument is any compiled code object, and otherwise is false.

```
(compiled-function-p x) <=> (typep x 'compiled-function)
```

`commonp` *object* [Function]

`commonp` is true if its argument is any common data type, and otherwise is false.

```
(commonp x) <=> (typep x 'common)
```

See also `standard-charp` (page 150), `string-charp` (page 150), `streamp` (page 227), `packagep` (page 117), `random-state-p` (page 146), `readtablep` (page 245), `hash-table-p` (page 191), and `pathnamep` (page 278).

6.3. Equality Predicates

COMMON LISP provides a spectrum of predicates for testing for equality of two objects: `eq` (the most specific), `eq1`, `equal`, and `equalp` (the most general). `eq` and `equal` have the meanings traditional in LISP. `eq1` 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 117) 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`. In some implementations, 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.

`eq1 x y`

[Function]

The `eq1` 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 122)), or if they are character objects that represent the same character (that is, they are `char=` (page 152)).

For example:

```

(eq1 'a 'b) is false
(eq1 'a 'a) is true
(eq1 3 3) is true
(eq1 3 3.0) is false
(eq1 (cons 'a 'b) (cons 'a 'c)) is false
(eq1 (cons 'a 'b) (cons 'a 'b)) is false
(setq x '(a . b)) (eq1 x x) is true
(eq1 #\A #\A) is true
(eq1 "Foo" "Foo") is false
(eq1 "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 `eq1`. 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 (only if uninterned symbols are involved).

Most 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 may 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` only if they are `eq`, with one exception: strings and bit-vectors are compared element-by-element. Upper-case and lower-case letters in strings are considered to be

distinct by `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 in Lisp Machine LISP are represented as `fixnums`.

Two pathname objects 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 compare a tree of conses, using `eq1` (or any other desired predicate) on the leaves, use `tree-equal` (page 174).

`equalp x y`

[Function]

Two objects are `equalp` if they are `equal`; if they are characters and satisfy `char-equal` (page 153), which ignores alphabetic case and certain other attributes of characters; 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`.

- 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 may 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, and the corresponding components are `equalp`. The specializations need not match; for example, a string and a general array that happens to contain the same characters will be `equalp` (though definitely not `equal`).

Two symbols can be `equalp` only if they are `eq`, that is, the same identical object.

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 52) 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`, the value `nil` is immediately returned 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-simple-vector))
        (eq (vref a-simple-vector n) 'foo))
    (princ "Foo!"))
```

The above expression prints "Foo!" if element `n` of `a-simple-vector` is the symbol `foo`, provided also that `n` is indeed a valid index for `a-simple-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-simple-vector))
      (eq (vref a-simple-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 77) and `when` (page 77), which are sometimes stylistically more appropriate than `and` for conditional purposes.

From the general definition, one can deduce that $(\text{and } x) \Leftrightarrow x$. Also, `(and)` is true, which is an identity for this operation.

`and` can be defined in terms of `cond` (page 76) 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* other than the last evaluates to something other than `nil`, or immediately returns that non-`nil` value 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 77) and `unless` (page 77), which are sometimes stylistically more appropriate than `or` for conditional purposes.

From the general definition, one can deduce that $(\text{or } x) \Leftrightarrow x$. Also, `(or)` is false, which is the identity for this operation.

`or` can be defined in terms of `cond` (page 76) 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 that produce side-effects, and consequently requires constructs for controlling the sequencing of side-effects. The construct `progn` (page 72), 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 72) and `prog2` (page 72).

For looping, COMMON LISP provides the general iteration facility `do` (page 80), 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 `case`. 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 79), `return` (page 79), `catch` (page 93), and `throw` (page 95).

The multiple-value constructs provide an efficient way for a function to return more than one value; see `values` (page 89).

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 by `read` (page 253) 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, 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 by `read` (page 253) to mean

```
(remove-if (function numberp) '(1 a b 3))
```

`symbol-value` *symbol* [*Function*]

`symbol-value` returns the current value of the dynamic (special) variable named by *symbol*. An error occurs if the symbol has no value; see `boundp` (page 63) and `makunbound` (page 65). Note that constant symbols are really variables that cannot be changed, and so `symbol-value` may be used to get the value of a named constant. In particular, `symbol-value` of a keyword will

(normally) return that keyword.

`symbol-value` cannot access the value of a lexical variable.

This function is particularly useful for implementing interpreters for languages embedded in LISP. The corresponding assignment primitive is `set` (page 64).

`symbol-function` *symbol* [Function]

`symbol-function` returns the current global function definition named by *symbol*. An error occurs if the symbol has no function definition; see `fboundp` (page 63). Note that the definition may be a function, or may be an object representing a special form or macro. See `macro-p` (page 63) and `special-form-p` (page 63).

`symbol-function` cannot access the value of a lexical function name produced by `fllet` (page 75) or `labels` (page 75).

This function is particularly useful for implementing interpreters for languages embedded in LISP. The corresponding assignment primitive is `fset` (page 65).

`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 64), `fset` (page 65), `makunbound` (page 65), and `fmakunbound` (page 65).

`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, the macro-call form) is returned; otherwise `nil` is returned. (The function `macroexpand` (page 100) is the best way to invoke the expansion function.)

The function `special-form-p` also takes a symbol. If the symbol globally names a special form (example: `quote` (page 62)), 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 is possible because an implementation is permitted to implement any macro also as a special form for speed. On the other hand, the macro definition must 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 66), which is the "general assignment statement", capable of assigning to variables, array elements, and other locations.

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

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. (If several variables are to be assigned to in parallel in the context of a loop, the `do` (page 80) construct may be appropriate.)

`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 64) 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 *prog* (page 75), a construct that 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 *plet* (page 75) or *labels* (page 75). *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 that implement the conceptual *access* and *update* operations. For a variable, merely mentioning the name of the variable accesses it, while the *setq* (page 64) special form can be used to update it. The function *car* (page 173) accesses the *car* of a cons, and the function *rp1aca* (page 181) updates it. The function *symbol-value* (page 62) accesses the

dynamic value of a variable named by a given symbol, and the function `set` (page 64) 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`). Now, rather than having to remember two functions for each kind of generalized variable (having to remember, for example, that `rplaca` corresponds to `car`), 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 that accesses it) into a form that updates it. The uniformity of this approach may be seen from the following table:

Access function	Update function	Update using <code>setf</code>
<code>x</code>	<code>(setq x newvalue)</code>	<code>(setf x newvalue)</code>
<code>(car x)</code>	<code>(rplaca x newvalue)</code>	<code>(setf (car x) newvalue)</code>
<code>(symbol-value x)</code>	<code>(set x newvalue)</code>	<code>(setf (symbol-value x) newvalue)</code>

`setf` is actually a macro that examines an access form and produces a call to the corresponding update function.

Given the existence of `setf` in COMMON LISP, it is not necessary to have `setq`, `rplaca`, and `set` as well; they are redundant. They are retained because of their historical importance in LISP. However, most other update functions (such as `putprop`, the update function for `get` (page 108)) have been eliminated in the expectation that `setf` be uniformly used in their place.

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

`(setf place newvalue)` 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*.

If more than one *place-newvalue* pair is specified, the pairs are processed sequentially:

```
(setf place1 newvalue1
      place2 newvalue2)
...
      placen newvaluen)
```

is precisely equivalent to

```
(progn (setf place1 newvalue1)
       (setf place2 newvalue2)
       ...
       (setf placen newvaluen))
```

For consistency, it is legal to write `(setf)`, which simply returns `nil`.

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 173)	caaar	(page 174)	caddr	(page 174)
cdr	(page 173)	cdaaar	(page 174)	cdddr	(page 174)
caar	(page 174)	cadaar	(page 174)	elt	(page 161)
cdar	(page 174)	cddaar	(page 174)	aref	(page 196)
cadr	(page 174)	caadar	(page 174)	svref	(page 197)
cddr	(page 174)	cdadar	(page 174)	sgvref	(page 197)
caaar	(page 174)	caddar	(page 174)	symbol-value	(page 62)
cdaar	(page 174)	cdddar	(page 174)	symbol-function	(page 63)
cadar	(page 174)	caaddr	(page 174)	get	(page 108)
cddar	(page 174)	cdaadr	(page 174)	symbol-plist	(page 109)
caadr	(page 174)	cadadr	(page 174)	gethash	(page 191)
cdadr	(page 174)	cddadr	(page 174)	documentation	(page 301)
caddr	(page 174)	caaddr	(page 174)	nth	(page 175)
cdddr	(page 174)	cdaddr	(page 174)		

- A function call form whose first element is the name of a selector function constructed by `defstruct` (page 211).
- 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):

<u>Function name</u>	<u>Required type</u>
char (page 203)	string-char
bit (page 197)	(mod 2)
subseq (page 161)	sequence

In the case of `subseq`, the replacement value must be a sequence whose elements may be contained by the sequence argument to `subseq`.

- 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):

<u>Function name</u>	<u>Argument that is a <i>place</i></u>	<u>Update function used</u>
char-bit	(page 157) First	set-char-bit (page 157)
ldb	(page 143) Second	dpb (page 143)
mask-field	(page 143) Second	deposit-field (page 144)

- A call on `getf` (page 109), in which case `(setf (getf x y) z)` expands into `(putf x y z)`.
- A `the` (page 106) 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 `setf` expands the macro call and then analyzes the resulting form.

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

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 70).

`psetf` *{place newvalue}**

[Macro]

`psetf` is like `setf` except that if more than one *place-newvalue* pair is specified then the assignments of new values to places is done in parallel. More precisely, all subforms that are to be evaluated are evaluated from left to right; after all evaluations have been performed, all of the assignments are performed.

`psetf` always returns `nil`.

`shiftf` *place {place}* newvalue*

[Macro]

Each *place* form may be any form acceptable as a generalized variable to `setf` (page 66). In the form `(shiftf place1 place2 ... placen newvalue)`, the values in *place1* through *placen* are accessed and saved, and *newvalue* is evaluated, for a total of $n+1$ values in all. Values 2 through $n+1$ are then stored into *place1* through *placen*, and value 1 (the original value of *place1*) is returned. It is as if all the places form a shift register; the *newvalue* is shifted in from the right, all values shift over to the left one place, and the value shifted out of *place1* is returned.

For example:

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

The effect of `(shiftf place1 place2 ... placen newvalue)` is roughly equivalent to

```
(prog1 place1
  (setf place1 place2)
  (setf place2 place3)
  ...
  (setf placen newvalue))
```

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

For example:

```
(setq n 0)
(setq x '(a b c d))
(shiftf (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 `shiftf` may be significantly more efficient than the `prog1` version.

Rationale: `shiftf` and `rotatef` (below) have been included in COMMON LISP as generalizations of two-argument versions formerly called `swapf` and `exchf`. The two-argument versions have been found to be very useful, but the names were easily confused. The generalization to many argument forms and the change of names were both inspired by the work of Suzuki [13], which indicates that use of these primitives can make certain complex pointer-manipulation programs clearer and easier to prove correct.

`rotatef {place}*`

[Macro]

Each *place* form may be any form acceptable as a generalized variable to `setf` (page 66). In the form `(rotatef place1 place2 ... placen)`, the values in *place1* through *placen* are accessed and saved. Values 2 through *n* and value 1 are then stored into *place1* through *placen*. It is as if all the places form an end-around shift register that is rotated one place to the left, with the value of *place1* being shifted around the end to *placen*. Note that `(rotatef place1 place2)` exchanges the contents of *place* and *place2*.

The effect of `(rotatef place1 place2 ... placen newvalue)` is roughly equivalent to

```
(psetf place1 place2
  place2 place3
  ...
  placen place1)
```

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

rotatef always returns nil.

Other macros that manipulate generalized variables include *getf* (page 109), *putf* (page 109), *remf* (page 110), *incf* (page 126), *decf* (page 126), *push* (page 179), and *pop* (page 180).

```
defsetf access-fn {update-fn [doc-string] |
                  lambda-list lambda-list {declaration | doc-string}* {form}*} [Macro]
```

A *defsetf* declaration may take one of two forms. In either form, *access-fn* must be a symbol, the name of a function or macro for which a *setf*-inverse is to be defined.

The simple form of *defsetf* is

```
(defsetf access-fn update-fn [doc-string])
```

The *update-fn* must name a function or macro that takes one more argument than *access-fn* does. When *setf* (page 66) is given a *place* that is a call on *access-fn*, it expands into a call on the *update-fn* that is given all the arguments to the *access-fn* and also, as the last argument, the new value. For example, after

```
(defsetf getfrob putfrob)
```

the form `(setf (getfrob 'a 3) foo)` would expand into `(putfrob 'a 3 foo)`.

The complex form of *defsetf* has the same form as *defmacro* (page 99) except that there are *two* lambda-lists, the first representing the argument forms to the *access-fn* and the second representing the value(s) of the *newvalue* form given to *setf*. The body of the *defsetf* definition must then compute a replacement form for the *setf* form, just as for any other macro.

The body is responsible for ensuring that the expansion causes subforms to be evaluated exactly once each and in the correct (left-to-right) order.

If the second lambda-list specifies other than a single required argument, *setf* will effectively arrange to use *multiple-value-call* (page 90) to receive the values from the *newvalue* form. For example, consider this simple function:

```
(defun uncons (cell) (values (car cell) (cdr cell)))
```

An appropriate *defsetf* definition would be:

```
(defsetf uncons (cell) (a &optional (d 'foo))
  '(values (setf (car ,cell) ,a) (setf (cdr ,cell) ,d)))
```

The result of expanding `(setf (uncons (reckon q)) (floor 5 3))` would then be something like:

```
(let ((G0001 (reckon q)))
  (multiple-value-call #'(lambda (G0002 &optional (G0003 'foo))
                          (values (setf (car G0001) G0002)
                                  (setf (cdr G0001) G0003))))
  (floor 5 3)))
```

Note that the values of parameters in the second lambda-list will be names of variables by means of which the expansion may refer to the values returned by *newvalue*.

7.3. Function Invocation

The most primitive form for function invocation in LISP of course has no name; any list that 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 arg &rest more-args* [Function]

This applies *function* to a list of arguments. *function* may be a compiled-code object, or a lambda-expression, or a symbol; in the latter case the global functional value of that symbol is used (but it is illegal for the symbol to be the name of a macro or special form). The arguments for the *function* consists of the last argument to `apply` appended to a list of all the other arguments to `apply` but the *function* itself.

For example:

```
(setq f '+) (apply f '(1 2)) => 3
(setq f '-') (apply f '(1 2)) => -1
(apply #'max 3 5 '(2 7 3)) => 7
(apply 'cons '((+ 2 3) 4)) =>
  ((+ 2 3) . 4)    not (5 . 4)
```

After the *function* argument there may be any number of individual arguments (possibly none) followed by a list of all the rest of the arguments. If no individual arguments are specified and the final list argument is empty, then the function receives no arguments. Note that if the function takes keyword arguments, the keywords as well as the corresponding values must appear in the argument list:

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

This can be very useful in conjunction with the `&allow-other-keys` feature:

```
(defun foo (size &rest keys &key double &allow-other-keys)
  (let ((v (apply #'make-simple-vector size keys)))
    (if double (concatenate v v) v)))

(foo 4 :initial-contents '(a b c d) :double t)
=> #(a b c d a b c d)
```

??? Query: The above example looks like the right thing, but conflicts with the specification that it is an error to pass an incorrect keyword to a function. What shall we do to preserve the utility of `&allow-other-keys`?

`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 (symbol-function '+))
(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*`.

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 that are to be evaluated sequentially, discarding the results of all forms but the last, and returning the results of the last form.

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}**

[Macro]

`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 that 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 `multiple-value-prog1` (page 90).

`prog2` *first second* *{form}**

[Macro]

`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 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 64)) 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 that 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 that produces the corresponding value, thereby improving program readability.

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

`let*` is similar to `let` (page 73), 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 it 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 64)) 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 that produces values for the bindings.

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

When executed by the LISP interpreter, `compiler-let` behaves exactly like `let` (page 73) with all the variable bindings implicitly declared `special`. When the compiler processes this form, however, no code is compiled for the bindings; instead, the processing of the body by the compiler is done with the special variables bound to the indicated values *in the execution context of the compiler*. This is primarily useful for communication among complicated macros.

`progv` *symbols values* {*form*}*

[*Special form*]

`progv` is a special form that 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 and then made to have no value; see `makunbound` (page 65). 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 73), 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* {*declaration* | *doc-string*}* {*form*}*)}*) {*form*}* [Special form]

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

`macrolet` ({(*name varlist* {*declaration* | *doc-string*}* {*form*}*)}*) {*form*}* [Special form]

`flet` may be used to define locally named functions. Within the body of the `flet` form, function names matching those defined 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 defined. Each definition is similar in format to a `defun` (page 47) 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 defined 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 99).

7.6. Conditionals

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

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 ...)		else if q then ...
(r ...)	roughly	else if r then ...
...	corresponds	...
(t ...))	to	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 58) construct may be stylistically preferable, depending on the context. If the value is not important, but only the effect, then the *when* (page 77) construct may be stylistically preferable.

when *pred {form}**

[Macro]

(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 58) or *if* (page 77) may be stylistically more appropriate.

unless *pred {form}**

[Macro]

(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 59) or *if* (page 77) may be stylistically more appropriate.

`case` *keyform* {{{*key*}*} *form**})* [Macro]

`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 76), 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 `eq1` 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 `eq1` 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**})* [Macro]

`typecase` is a conditional that 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 76) or `case` (page 78), 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 `typecase` 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 72); the latter has nothing to do with `return`.

The *name* is not evaluated; it must be a symbol. 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 47) 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-from` to return prematurely from a function defined by `defun`.

`return-from name [result]` [Special form]

`return [result]` [Macro]

`return-from` is used to return from a `block` or from such constructs as `do` and `prog` that implicitly establish a `block`. The *name* is not evaluated, and must be a symbol. A `block` construct with the same name must lexically enclose the occurrence of `return-from`; whatever the evaluation of *result* produces is immediately returned from the block. (If the *result* form is omitted, it defaults to `nil`. As a matter of style, this form ought to be used to indicate that the

particular value returned doesn't matter.)

The `return-from` form itself never returns, and cannot have a value; it causes results to be returned from a block construct. If the evaluation of *result* produces multiple values, those multiple values are returned by the construct exited.

(`return form`) is identical in meaning to (`return-from nil form`); it returns from a block named `nil`. As a rule, blocks established implicitly by iteration constructs such as `do` are named `nil`, so that `return` will exit properly from such a construct.

7.8. Iteration

COMMON LISP provides a number of iteration constructs. The `loop` (page 80) construct provides a trivial iteration facility; it is little more than a `progn` (page 72) with a branch from the bottom back to the top. The `do` (page 80) and `do*` (page 80) constructs provide a general iteration facility for controlling the variation of several variables on each cycle. For specialized iterations over the elements of a list or *n* consecutive integers, `dolist` (page 84) and `dotimes` (page 84) are provided. The `tagbody` (page 87) construct is the most general, permitting arbitrary `go` (page 89) statements within it. (The traditional `prog` (page 87) construct is a synthesis of `tagbody`, `block` (page 79), and `let` (page 73).) All of the iteration constructs permit statically defined non-local exits in the form of the `return-from` (page 79) and `return` statements.

7.8.1. Simple Iteration

`loop {form}*` [Macro]

Each *form* is evaluated in turn, from left to right. When the last *form* has been evaluated, then the first *form* is evaluated again, and so on, in a never-ending cycle. The `loop` construct never returns a value. It must be explicitly terminated, for example by establishing a `block` (page 79) around it and using a `return-from` statement, or by using `throw` (page 95).

`loop` does *not* establish an implicit block named `nil`.

Rationale: This construct is included primarily as a primitive building block for more complicated iteration macros that is perhaps more easily understood by a compiler than a full-blown `tagbody` (page 87).

A `loop` construct has this meaning only if every *form* is non-atomic (a list). The case where one or more than one *form* is a symbol is reserved for future extensions.

7.8.2. General iteration

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

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

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)
    {declaration}*
    . tagbody)
```

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 64)).

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 72), 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 76) 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 *setq* (page 64)) to the respective index variables. Any variable that 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 *loop* (page 80) construct performs a “do forever” that steps no variables.) The infinite loop can be terminated by the use of *return* (page 79), *return-from* (page 79), *go* (page 89) to an outer level, or *throw* (page 95).

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 an implicit `tagbody` (page 87). Tags may appear within the body of a `do` loop for use by `go` (page 89) statements appearing in the body (but such `go` statements may not appear in the variable specifiers, the *end-test*, or the *result* forms). When the end of a `do` body is reached, the next iteration cycle (beginning with the evaluation of *step* forms) occurs.

An implicit `block` (page 79) named `nil` surrounds the entire `do` form. A `return` (page 79) statement may be used at any point to exit the loop immediately.

`declare` (page 101) 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.

Here are some examples of the use of `do`:

```
(do ((i 0 (+ i 1)) ; Sets every null element of a-vector to zero.
    (n (array-dimension a-vector 0)))
    ((= i n))
    (when (null (aref a-vector i))
      (setf (aref a-vector i) 0)))
```

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 `reverse` (page 162) to put an accumulated `do` loop result into the correct order is a standard idiom.

Other examples:

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

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

Note the use of `endp` (page 175) rather than `null` (page 52) to test for the end of a list in the above two examples. This results in more robust code.

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.)

A `do` loop may be explained in terms of the more primitive constructs `block` (page 79), `return` (page 79), `let` (page 73), `loop` (page 80), `tagbody` (page 87), and `psetq` (page 64) as follows:

```
(block nil
  (let ((var1 init1)
      (var2 init2)
      ...
      (varn initn))
    {declaration}*
    (loop (when end-test (return (progn . result)))
          (tagbody . tagbody)
          (psetq var1 step1
                var2 step2
                ...
                varn stepn))))
```

`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. It is as if, in the above explanation, `let` were replaced by `let*` (page 74) and `psetq`

were replaced by `setq` (page 64).

7.8.3. 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 that 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 79) statement may be used to return immediately from a `dolist` or `dotimes` form, discarding any following iterations that might have been performed; in effect, a block named `nil` surrounds the construct. The body of the loop is implicitly a `tagbody` (page 87) construct; it may contain tags to serve as the targets of `go` (page 89) statements. Declarations may appear before the body of the loop.

`dolist` (*var* *listform* [*resultform*]) {*declaration*}* {*tag* | *statement*}* [Macro]

`dolist` provides straightforward iteration over the elements of a list. First `dolist` evaluates the form *listform*, which should produce a list. It then executes the body 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. (When the *resultform* is evaluated, the control variable *var* is still bound, and has the value `nil`.) 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*]) {*declaration*}* {*tag* | *statement*}* [Macro]

`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. (When the *resultform* is evaluated, the control variable *var* is still bound, and has as its value the number of times the body was executed.) If *resultform* is omitted, the result is `nil`.

Altering the value of *var* in the body of the loop (by using `setq` (page 64), 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 119) and related constructs.

7.8.4. 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 163) may be used to map over any kind of sequence. The following functions operate only on lists.

<code>mapcar</code> <i>function list &rest more-lists</i>	[Function]
<code>maplist</code> <i>function list &rest more-lists</i>	[Function]
<code>mapc</code> <i>function list &rest more-lists</i>	[Function]
<code>mapl</code> <i>function list &rest more-lists</i>	[Function]
<code>mapcan</code> <i>function list &rest more-lists</i>	[Function]
<code>mapcon</code> <i>function list &rest more-lists</i>	[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 178) 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 165) 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 that have `&optional` and `&rest` parameters.

7.8.5. 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 80), have bodies with the characteristics of a `prog`.

`prog` actually performs three distinct operations: it binds local variables, it permits use of the `return` statement, and it permits use of the `go` statement. In COMMON LISP, these three operations have been separated into three distinct constructs: `let` (page 73), `block` (page 79), and `tagbody` (page 87). These three constructs may be used independently as building blocks for other types of constructs.

`tagbody` *{tag | statement}* [Special form]*

The part of a `prog` after the variable list is called the *body*. An item in the body may be a symbol or an integer, in which case it is called a *tag*, or a list, in which case it is called a *statement*.

Each element of the body is processed from left to right. A *tag* are ignored; a *statement* is evaluated, and its results are discarded. If the end of the body is reached, the `tagbody` returns `nil`.

If (`go tag`) is evaluated, control jumps to the part of the body labelled with the *tag*. The `go 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.

The scope of the tags established by a `tagbody` is lexical, and the extent is dynamic. Once a `tagbody` construct has been exited, it is no longer legal to `go` to a *tag* in its body. It is permissible for a `go` to jump to a `tagbody` that is not the innermost `tagbody` construct containing that `go`; the tags established by a `tagbody` will only shadow other tags of like name in an outer `tagbody`.

`prog` *{var | (var [init])}* {declaration}* {tag | statement}* [Macro]*
`prog*` *{var | (var [init])}* {declaration}* {tag | statement}* [Macro]*

A typical `prog` looks like:

```
(prog (var1 var2 (var3 init3) var4 (var5 init5))
      {declaration}*
      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 73) 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. Any *declaration* appearing in the `prog` is used as if appearing at the top of the `let` body.

The body of the `prog` is executed as if it were a `tagbody` (page 87) construct; the `go` (page 89) statement may be used to transfer control to a *tag*.

A `prog` implicitly establishes a `block` (page 79) named `nil` around the entire `prog` construct, so that `return` (page 79) may be used at any time to exit from the `prog` construct.

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))))
```

The `prog` construct may be explained in terms of the simpler constructs `block` (page 79), `let` (page 73), and `tagbody` (page 87) as follows:

```
(prog variable-list {declaration}* . body)
<=> (block nil (let variable-list {declaration}* (tagbody . body)))
```

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 74) is to `let` (page 73).

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 `tagbody` (page 87) construct. The *tag* must be a symbol or an integer; the *tag* is not evaluated. `go` transfers control to the point in the body labelled by a tag `eq1` to the one given. If there is no such tag in the body, the bodies of lexically containing `tagbody` constructs (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.

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 79). If the use of `go` seems to be unavoidable, perhaps the control structure implemented by `go` should be packaged up as a macro definition.

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 objects 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 89), 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 135)) return multiple values; those that do are so documented.

The special forms for receiving multiple values are `multiple-value-bind` (page 90), `multiple-value` (page 91), and `multiple-value-list` (page 90). 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* [Macro]

`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)
```

`multiple-value-call` *function* *{form}** [Special form]

`multiple-value-call` 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 `multiple-value-call` is whatever is returned by the function.

For example:

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

`multiple-value-prog1` *form* *{form}** [Special form]

`multiple-value-prog1` 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 `multiple-value-prog1`. See `prog1` (page 72), which always returns a single value.

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

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)
```

`multiple-value variables form`

[Macro]

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, then 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 219) returns multiple values if the form given it to evaluate produces multiple values.
- `apply` (page 71), `funcall` (page 71), and `multiple-value-call` (page 90), pass back multiple values from the function applied or called.
- When a lambda-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 72) itself passes back multiple values from its last subform, as does any construct some part of which is defined to be an "implicit `progn`"; these include `progv` (page 75), `let` (page 73), `let*` (page 74), `when` (page 77), `unless` (page 77), `case` (page 78), `typecase` (page 78), `multiple-value-bind` (page 90), `multiple-value` (page 91), `catch` (page 93), and `catch-all` (page 93).
- `multiple-value-prog1` (page 90) passes back multiple values from its first subform. However, `prog1` (page 72) always returns a single value.
- `unwind-protect` (page 94) returns multiple values if the form it protects does.

- `catch` (page 93) returns multiple values if the result form in a `throw` (page 95) exiting from such a `catch` produces multiple values.
- `cond` (page 76) passes back multiple values from the last subform 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 77) passes back multiple values from whichever subform is selected (the *then* form or the *else* form).
- `and` (page 58) and `or` (page 59) pass back multiple values from the last subform, but not from subforms other than the last.
- `block` (page 79) passes back multiple values from its last subform when it exits normally. If `return-from` (page 79) is used to terminate the `block` prematurely, then `return-from` passes back multiple values from its subform as the values of the terminated `block`. Other constructs that create implicit blocks, such as `do` (page 80), `dolist` (page 84), `dotimes` (page 84), `prog` (page 87), and `prog*` (page 87), also pass back multiple values specified by `return-from` (or `return` (page 79)). 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. Similarly, `dolist` and `dotimes` pass back multiple values from the *resultform* if that is executed.

Among special forms that *never* pass back multiple values are `setq` (page 64), `multiple-value` (page 91), and `prog1` (page 72). 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 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 `(multiple-value-call #'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 that 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 79) and `return` (page 79), 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 and characters 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.

`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
                (win (values-list results)) ; Check tag.
                (lose (cleanup)           ; If win, return results.
                     (ferror "Lose lose!")) ; If lose, clean up
                (otherwise                 ; and signal an error.
                 (throw tag (values-list results)))) ; Otherwise relay throw.
              (determine-win-or-lose)))
```

Note that an attempt to use `go` (page 89) or `return-from` (page 79) to exit from a `catch-all` will also be trapped by the *catch-function*; the tag given to the catch function will be some internal implementation-dependent object that can nevertheless be given to `throw` to continue the `go` or `return-from` operation.

`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.

Compatibility note: In `MACLISP`, a `go` from within a `MACLISP` `catchall` (note the different spelling) quietly breaks up the `catchall` frame without invoking the `catchall` function, which means that it catches all throws but not all exits! In `COMMON LISP`, `catch-all` traps *all* attempts to exit.

`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 that 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 that 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 should typically make two passes over the control stack. In the first pass it 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` constructs 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 219) is given a list whose *car* is a symbol, it looks for local definitions of that symbol (by `let` (page 75), `labels` (page 75), and `macrolet` (page 75)); if that fails, it looks for a global definition. If the definition is a macro definition, then the original list is said to be a *macro call*. Associated with the definition 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 71), `funcall` (page 71), or `map` (page 163); 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) {declaration}* {form}* [Macro]`

The primitive special form for establishing a global macro definition is `macro`. Note, however, that the use of `macro` is often very awkward, and it is preferable to use `defmacro` (page 99) in almost all circumstances. A call to `macro` has the following form:

`(macro name (var) . body)`

This is very similar 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 77) construct could be defined in terms of `cond` (page 76) as a macro:

```
(macro if (call-form)
  '(cond (,(cadr call-form) ,(caddr call-form))
        (t ,(caddr call-form))))
```

If the above form is executed by the interpreter, it will cause the definition of the symbol `if` to be a macro associated with which is a one-argument expansion function equivalent to:

```
(lambda (calling-form)
  (list 'cond
        (list (cadr calling-form) (caddr calling-form))
        (list 't (caddr 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 99) alleviates this problem. It should also be clear that the backquote facility (page 237) 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 49) construct:

```
(eval-when (compile)
  (macro name (var)
    body))
```

`defmacro` *name* *lambda-list* [*declaration* | *doc-string*]* *form** [Macro]

`defmacro` is a macro-defining macro that, unlike `macro`, decomposes the calling form in a more elegant and useful way. `defmacro` has the same syntax as `defun` (page 47): *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, in effect the expander function and the list of (unevaluated) argument forms is given to `apply` (page 71). The parameter specifiers are processed as for any lambda-expression, using the macro-call argument forms as the arguments. 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.

If the optional documentation string *doc-string* is present (if not followed by a declaration, 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 attached to the *name* as a documentation string of type `function`; see `documentation` (page 301).

Like the lambda-list in a `defun`, a `defmacro` *lambda-list* may contain the lambda-list keywords `&optional`, `&rest`, `&key`, `&allow-other-keys`, and `&aux`. For `&optional` and `&key` 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 treated as a body, 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 MACLISP and 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 `defmacro`, a definition for three-argument `if` in terms of `cond` 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`.

See also `macrolet` (page 75), which establishes macro definitions over a restricted lexical scope.

8.2. Expanding Macro Calls

`macroexpand form &rest environment` [Function]

`macroexpand-1 form &rest environment` [Function]

If *form* is a macro call, then `macroexpand-1` will expand the macro call *once* and return two values: the expansion and `t`. If *form* is not a macro call, then the two values *form* and `nil` are returned.

A *form* is considered to be a macro call only if it is a cons whose *car* is a symbol that names a macro. The *environment* is similar to that used within the evaluator and made visible via `evalhook` (page 220); any local macro definitions established within the *environment* by `macrolet` (page 75) will be considered. If only *form* is given as an argument, the *environment* is null, and only global macro definitions (as established by `defmacro` (page 99)) will be considered.

Macro expansion is carried out as follows. Once `macroexpand-1` has determined that a symbol names a macro, it obtains the expansion function for that macro. The value of the variable `*macroexpand-hook*` (page 100) is then called as a function of two arguments: the expansion function and the *form*. The value returned from this call is taken to be the expansion of the macro call. The initial value of `*macroexpand-hook*` is `funcall` (page 71), and the net effect is to invoke the expansion function, giving it the *form* as its single argument. (The purpose of `*macroexpand-hook*` is to facilitate various techniques for improving interpretation speed by caching macro expansions.)

`macroexpand` is similar to `macroexpand-1`, but repeatedly expands *form* until it is no longer a macro call. (In effect, `macroexpand` simply calls `macroexpand-1` repeatedly until the second value returned is `nil`.) A second value of `t` or `nil` is returned as for `macroexpand-1`, indicating whether the original *form* was a macro call.

`*macroexpand-hook*` [Variable]

The value of `*macroexpand-hook*` is used as the expansion interface hook by `macroexpand-1` (page 100).

Chapter 9

Declarations

Declarations allow you to specify extra information about your program to the LISP system. All declarations are completely optional and correct declarations 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.

Note that it is considered an error for a program to violate a declaration (such as a type declaration), but an implementation is not required to detect such errors (though such detection, where feasible, is to be encouraged).

9.1. Declaration Syntax

`declare` {*declaration-form*}* [*Special form*]

A `declare` form is known as a *declaration*. Declarations may occur only at top level, or at the beginning of the bodies of certain special forms; that is, a declaration not at top level may occur only as a statement of such a special form, and all statements preceding it (if any) must also be `declare` forms (or possibly documentation strings, in some cases). Declarations may occur in lambda-expressions, and in the following forms:

<code>deftype</code>	(page 36)	<code>dolist</code>	(page 84)
<code>defun</code>	(page 47)	<code>dotimes</code>	(page 84)
<code>defsetf</code>	(page 70)	<code>prog</code>	(page 87)
<code>let</code>	(page 73)	<code>prog*</code>	(page 87)
<code>let*</code>	(page 74)	<code>multiple-value-bind</code>	(page 90)
<code>flet</code>	(page 75)	<code>macro</code>	(page 97)
<code>labels</code>	(page 75)	<code>defmacro</code>	(page 99)
<code>macrolet</code>	(page 75)	<code>locally</code>	(page 103)
<code>do</code>	(page 80)	<code>do-symbols</code>	(page 119)
<code>do*</code>	(page 80)	<code>do-all-symbols</code>	(page 119)

If a declaration is found anywhere else an error will be signalled.

It is permissible for a macro call to expand into a declaration and be recognized as such, provided

that the macro call appears where a declaration may legitimately appear.

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 that 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. As an example of a pervasive declaration,

```
(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.

As a rule, code in any initialization forms used to compute initial values for bound variables is *not* affected by pervasive declarations in that special form, with one exception: lambda-list initialization forms, which appear in lambda-expressions as well as *defun* (page 47), *defmacro* (page 99), *macro* (page 97), *deftype* (page 36), and *defsetf* (page 70) forms, are affected by special declarations for variables not bound by that form; so is every initialization form in a *let** *except* for the first initialization form.

For example:

```
(defun foo (x (y *princircle*))      ;This reference to *princircle* is
  (declare (special *princircle*)) ; special because of this declaration.
  ...)
```

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 49).

`locally` {*declaration*}* {*form*}*

[Macro]

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))
         (optimize space)
         (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.

For example:

```
(defun hack (thing *mod*)
  (declare (special *mod*)) ;The binding of *mod* is visible to
  (hack1 (car thing)))      ; hack1, but not that of thing.

(defun hack1 (arg)
  (declare (special *mod*)) ;Declare that references to *mod*
                             ; within hack1 are special.
  (if (atom arg) *mod*
      (cons (hack1 (car arg)) (hack1 (cdr arg)))))
```

Note that it is conventional, though not required, to give special variables names that begin

and end with an asterisk.

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. In particular, values assigned to the variables by *setq* (page 64), as well as the initial values of the variables, must be 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 32).

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 47): 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.
- optimize** (`optimize (quality1 value1) (quality2 value2) ...`) advises the compiler that each *quality* should be given attention according to the specified corresponding *value*. A quality is a symbol; standard qualities include `speed` (of the object code), `space` (both code size and run-time space), `safety` (run-time error checking), and `compilation-speed` (speed of the compilation process). Other qualities may be recognized by particular implementations. A *value* should be a non-negative integer, normally in the range 0 to 3. The value 0 means that the quality is totally unimportant, and 3 that the quality is extremely important; 1 and 2 are intermediate values, with 1 the "normal" or "usual" value. One may abbreviate "`(quality 3)`" to simply "`quality`". This declaration is pervasive.

For example:

```
(defun often-used-subroutine (x y)
  (declare (optimize (safety 2)))
  (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))))))
```

- declaration** (`declaration name1 name2 ...`) advises the compiler that each *namej* is a valid but non-standard declaration name. The purpose of this is to tell one compiler not to issue warnings for declarations meant for another compiler or other program processor.

For example:

```
(declare (declaration author target-language
          target-machine))
(declare (target-language ada) (target-machine IBM-650))
(declare (author "Harry Tweeker"))
```

An implementation is free to support other (implementation-dependent) declaration forms as well. On the other hand, a COMMON LISP compiler is free to ignore entire classes of declaration forms (for example, implementation-dependent declaration forms not supported by that compiler's implementation!), except for the `declaration` declaration form. Compiler implementors are encouraged, however, to program the compiler by default to issue a warning if the compiler finds a declaration form of a kind it never uses. Such a warning is required if a declaration form is not one of those defined above and has not been declared in a `declaration` declaration.

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. For this purpose the special form is defined: (`the type form`) means that the value of `form` is declared to be of type `type`.

the *value-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 *value-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 (concatenate x y)) ; concatenate will produce 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))
```

The `values` type specifier may be used to indicate the types of multiple values:

```
(the (values integer integer) (floor x y))
(the (values string t) (gethash the-key the-string-table))
```

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 that has three user-visible components:

- The *property list* is a list that 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.

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 with a symbol (called the *indicator*) 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 186) and `pairlis` (page 186)).

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 66) 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 using `get` (page 108) within a `setf` (page 66) form.

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 109), `putf` (page 109), and `remf` (page 110).

`get` *symbol indicator* &optional *default* [Function]

`get` 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*.

```
(get x y) <=> (getf (symbol-plist x) y)
```

Suppose that the property list of `foo` is `(bar t baz 3 hunoz "Huh?")`. Then, for example:

```
(get 'foo 'baz) => 3
(get 'foo 'hunoz) => "Huh?"
(get 'foo 'zoo) => nil
```

??? Query: In MACLISP, `get` of a non-symbol quietly returns `nil`. What about COMMON LISP?

`setf` (page 66) may be used with `get` to create a new property-value pair, possibly replacing an old pair with the same property name.

For example:

```
(get 'clyde 'species) => nil
(setf (get 'clyde 'species) 'elephant) => elephant
and now (get 'clyde 'species) => elephant
```

`remprop` *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.

```
(remprop x y) <=> (remf (symbol-plist x) y)
```

For example:

```
If the property list of foo was
  (color blue height 6.3 near-to bar)
then
  (remprop 'foo 'height) => t
and foo's property list would have been altered to be
  (color blue near-to bar)
```

`symbol-plist` *symbol*

[Function]

This returns the list that 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 `symbol-plist` does *not* work. One must give the symbol itself to `get`, or use the function `getf` (page 109).

`setf` (page 66) may be used with `symbol-plist` to destructively replace the entire property list of a symbol. Care must be taken that the new property list is in fact a list of even length.

`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*. Often *place* is computed from a generalized variable acceptable to `setf` (page 66). See `get` (page 108).

`setf` (page 66) may be used with `getf`, in which case the *place* must indeed be acceptable as a *place* to `setf`. The effect is to perform a `putf` operation.

`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 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 66).

`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 66). See `remprop` (page 109).

`get-properties place indicator-list`

[Function]

`get-properties` is like `getf` (page 109), 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 66).

`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.

10.2. The Print Name

Every symbol has an associated string called the *print-name*. 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 section 22.1.1; page 230) and *printer* (see section 22.1.6, page 248).

`symbol-print-name sym`

[Function]

This returns the print-name of the symbol *sym*.

For example:

```
(symbol-print-name '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 253) 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 (make-symbol "XYZ")) is true
(samepnamep 'xyz (make-symbol "WXY")) is false
```

10.3. Creating Symbols

Symbols can be used in two rather different ways. An *interned* symbol is one that 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 253), 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 that are cataloged and the only LISP objects that 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 117), 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, see chapter 11 (page 115).

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` *print-name*

[Function]

`(make-symbol print-name)` creates a new uninterned symbol, whose print-name is the string *print-name*. 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 *print-name* or may be a copy of it, at the implementation's discretion. The user should not assume that `(symbol-print-name (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.

`copy-symbol` *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 that 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 112) may be more appropriate to use.

`gentemp` *prefix* &optional *package* [Function]

`gentemp`, like `gensym` (page 112), creates and returns a new symbol. `gentemp` differs from `gensym` in that it interns the symbol (see `intern` (page 117)) in the *package* (which defaults to the current package; see `*package*` (page 117)). `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`.

`keywordp` *symbol* [Function]

The argument must be a symbol. The predicate `keywordp` is true if the symbol is a keyword (that is, belongs to the keyword package).

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 [15]. The proper design of a package system for LISP is still a subject of research; COMMON LISP therefore defines only a minimal facility purposely designed to accommodate experimentation with extensions by implementors. Certain desirable features and facilities have been omitted. What is defined here is intended to be just enough to allow modules in the yellow pages library to hide most internal symbols and make names of external functions and variables visible to the user, without making any commitment on such issues as nested packages and hierarchical inheritance.

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.

The value of the special variable **package** (page 117) 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*, and is said to *own* the symbol. (Some symbols are not owned by any package; they are said to be *uninterned*.)

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 named `keyword` contains all keyword symbols. Because keyword symbols are used so frequently, COMMON LISP permits "keyword:foo" to be abbreviated to simply ":foo". (The keyword package is also treated specially in that whenever a symbol is added to the keyword package, the symbol is automatically declared to be a constant and is made to have itself as its value. This is why every keyword evaluates to itself.)

All other uses of colons within names of symbols are not defined by COMMON LISP, but are reserved for experimentation by implementors; this includes names that end in a colon, contain two or more colons, or consist of just a colon.

Each symbol contains a package slot that is used to record the home package of the symbol. When the symbol is printed, if it is in the keyword package then it is printed with a preceding colon; otherwise, if it is present in the current package, it is printed without any qualification; otherwise, it is printed with the name of the home package as the qualifier. A symbol that is uninterned (has no home package) is printed preceded by "#:".

11.1. Built-in Packages

The following packages are built into the COMMON LISP system:

<code>user</code>	The user package is, by default, the current package at the time a COMMON LISP system starts up. The standard symbols used by COMMON LISP as function names, variable names, and for other purposes are available in this package.
<code>keyword</code>	This package contains all of the keywords used by built-in or user-defined LISP functions.
<code>si</code>	This package name is reserved to the implementation. (The name is an abbreviation for "system internals".) Normally this is used to contain names of functions and variables that are needed to implement the user-level COMMON LISP facilities.

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. The *package-name* should be a symbol or a string.

If a package of this name already exists, a correctable error is signalled. The *copy-from* argument may specify another package of which the new one will initially be a copy; that is, the new package will logically contain (but not own) all symbols in the *copy-from* 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.

packagep object [Function]

packagep is true if its argument is a package, and otherwise is false.

`(packagep x) <=> (typep x 'package)`

package package [Function]

This converts its argument to be a package object. If the argument is already a package, it is 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.

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.

If *package* is the keyword *package* and a symbol of the specified name is not already in the package, then as the symbol is installed in the keyword package it will be given itself as its value; see *symbol-value* (page 62). If the argument is a symbol rather than a string, then it must not already have a home package other than the keyword package.

`unintern` *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. `unintern` returns `t` if it actually removed a symbol, and `nil` otherwise.

Compatibility note: The equivalent of this in MACLISP is `remob`.

`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.

`export` *symbols* [Function]

The argument should be a list of symbols or strings, or possibly a single symbol or string. It is arranged for symbols of the specified names to be available in both the current package and the user package; for each name, the *same* symbol must be in both packages. If a name is in neither package, a symbol is created and interned in both. If a name is in just one, it is interned in the other. If the two packages have different symbols of the same name, the one in the current package is first removed by using `unintern` (page 118). In any case, the owner of the symbol is changed to be the user package. `export` returns `t`.

By convention, a call to `export` listing all exported symbols is placed near the start of a file to advertise which of the symbols used mentioned the file are intended to be used by other programs.

`shadow` *symbols* [Function]

The argument should be a list of symbols or strings, or possibly a single symbol or string. For each specified name, it is arranged that the current package, if it is not the user package, will contain a symbol of that name that is different from the symbol of that name in the user package. If the current package and the user package share a symbol whose name has been specified, that symbol is first removed from the current package.

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 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 175). One might write:

```

(provide 'interlisp)
(export '(masterscope helpsys dwimify ...))
(shadow '(nth ...))
(require 'odd-utilities)
...
(defun nth (x n)          ;InterLISP NTH function.
  ...)

```

shadow returns t.

`do-symbols` (*var* [*package*] [*result-form*]) {*declaration*}* {*tag* | *statement*}* [Macro]

`do-symbols` 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 `do-symbols` form. (When the *resultform* is evaluated, the control variable *var* is still bound, and has the value `nil`.) If *resultform* is omitted, the result is `nil`. `return` (page 79) may be used to terminate the iteration prematurely. 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* by using `unintern` (page 118), the effects are unpredictable.

`do-all-symbols` (*var* [*result-form*]) {*declaration*}* {*tag* | *statement*}* [Macro]

This is similar to `do-symbols`, but executes the body once for every symbol contained in "every" package. (This may not get all symbols whatsoever, depending on the implementation.) 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 that accept only certain kinds of numbers are so documented below.

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 `eq1` (page 56) is probably appropriate; if both objects are known to be numbers, then `=` (page 122) 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. For instance, 3.142857142857142857 is a more precise approximation to π than 3.14159, but the latter is more accurate. 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 53), `rationalp` (page 53), `floatp` (page 54), `complexp` (page 54), and `numberp` (page 53).

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 2 3) is false	(/= 3 2 3) is false
(< 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 $j \leq k$

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 eq1 (page 56).

max *number &rest more-numbers*

[Function]

The arguments may be any non-complex numbers. max returns the argument that 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 that 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.

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, successively subtracts from the first argument all the others, and returns the result. For example, *(- 3 4 5)* => -6.

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, successively 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
(/ 3 4 5) => 3/20
```

To divide one integer by another producing an integer result, use one of the functions `floor`, `ceiling`, `truncate`, or `round` (page 135).

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 `truncate` (page 135) 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 arguments 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", `truncate` (page 135), `floor` (page 135), `ceiling` (page 135), and `round` (page 135) 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.

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 66). 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 *integers* [Function]

Returns the greatest common divisor of all the arguments, which must be integers. The result is always a non-negative integer. If no arguments are given, `gcd` returns 0, which is an identity for this operation.

For example:

```
(gcd 91 -49) => 7
```

`lcm` *integer* &rest *more-integers* [Function]

This returns the least common multiple of its arguments, which must be integers. 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
```

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 actually they 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,

$$(\text{abs } x) \Leftrightarrow (\text{if } (\text{minusp } x) (- x) x)$$

For a complex number z , the absolute value may be computed as

$$(\text{sqrt } (+ (\text{expt } (\text{realpart } z) 2) (\text{expt } (\text{imagpart } z) 2))))$$

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,

$$(\text{phase } x) \Leftrightarrow (\text{atan } (\text{realpart } x) (\text{imagpart } x))$$

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,

$$(\text{signum } x) \Leftrightarrow (\text{if } (\text{zerop } x) x (/ x (\text{abs } x)))$$

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, unless z is a complex zero, in which case the result is z .

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 $e^{i \theta}$ 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 `acos` 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.

<u>Condition</u>	<u>Cartesian locus</u>	<u>Range of result</u>
$y = 0 \quad x > 0$	Positive x -axis	0
$y > 0 \quad x > 0$	Quadrant I	$0 < \text{result} < \pi/2$
$y > 0 \quad x = 0$	Positive y -axis	$\pi/2$
$y > 0 \quad x < 0$	Quadrant II	$\pi/2 < \text{result} < \pi$
$y = 0 \quad x < 0$	Negative x -axis	π
$y < 0 \quad x < 0$	Quadrant III	$-\pi < \text{result} < -\pi/2$
$y < 0 \quad x = 0$	Negative y -axis	$-\pi/2$
$y < 0 \quad x > 0$	Quadrant IV	$-\pi/2 < \text{result} < 0$
$y = 0 \quad x = 0$	Origin	error

Actually, the $<$ signs in the above table ought to be \leq 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` whose range is 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

[Constant]

This global variable has as its value the best possible approximation to π in *long* floating-point format.

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 134).

`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 [11]. The contents of this section are borrowed largely from that proposal.

Compatibility note: The branch cuts defined here differ in a few very minor respects from those advanced by W. Kahan, who considers not only the "usual" definitions but also the special modifications necessary for IEEE proposed floating-point arithmetic, which has infinities and minus zero as explicit computational objects. For example, he proposes that $\sqrt{-4+0i} = 2i$, but $\sqrt{-4-0i} = -2i$.

It is likely that the differences between the APL proposal and Kahan's proposal will be ironed out, perhaps in 1983. If so, COMMON LISP will be changed as necessary to be compatible with these other groups. Any changes from the specification below are likely to be quite minor.

An implementation of COMMON LISP is not required to support complex numbers before January 1, 1984. It is expected that the compatibility problems will have been resolved by then. As of that date, every COMMON LISP implementation will be expected to support complex numbers as then defined.

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 $z = x + yi$, $\log z$ is defined to be $(\log |z|) + i \text{phase}(z)$. 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:

$$\arctan z = -i \log ((1 + iz) \sqrt{1/(1 + z^2)})$$

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:

$$\operatorname{arsinh} z = \log (x + \sqrt{1 + x^2})$$

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:

$$\operatorname{arccosh} z = \log (x + (x+1)\sqrt{(x-1)/(x+1)})$$

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:

$$\operatorname{arctanh} z = \log ((1+x)\sqrt{1-1/x^2})$$

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 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:

$$\sin iz = i \sinh z$$

$$\cos iz = \cosh z$$

$$\tan iz = i \tanh z$$

$$\sinh iz = i \sin z$$

$$\cosh iz = \cos z$$

$$\operatorname{arcsin} iz = i \operatorname{arcsinh} z$$

$$\operatorname{arctan} iz = i \operatorname{arctanh} z$$

$$\operatorname{arcsinh} iz = i \operatorname{arcsin} z$$

$$\operatorname{arctanh} iz = i \operatorname{arctan} z$$

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 a `single-float` is produced. If the argument *other* is provided, then it must be a floating-point number, and *number* is converted to the same format as *other*. See also `coerce` (page 37).

`rational` *number* [Function]

`rationalize` *number* [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.

`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]
`ceiling number &optional divisor` [Function]
`truncate number &optional divisor` [Function]
`round number &optional divisor` [Function]

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 that is not larger than the argument.

`ceiling` converts its argument by truncating towards positive infinity; that is, the result is the smallest integer that is not smaller than the argument.

`truncate` 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 that is even (divisible by two).

Here is a table showing what the four functions produce when given various arguments.

<u>Argument</u>	<u>floor</u>	<u>ceiling</u>	<u>truncate</u>	<u>round</u>
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 90) 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`, `ceiling`, `truncate`, and `round` are more accurate than names like `fix` that 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 `truncate`. `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 `truncate`. 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 divisor` [Function]

`remainder number divisor` [Function]

`mod` performs the operation `floor` (page 135) on its two arguments, and returns the *second* result of `floor` as its only result. Similarly, `rem` performs the operation `truncate` (page 135) on its arguments, and returns the *second* result of `truncate` 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.

<code>(mod 13 4) => 1</code>	<code>(rem 13 4) => 1</code>
<code>(mod -13 4) => 3</code>	<code>(rem -13 4) => -1</code>
<code>(mod 13 -4) => -3</code>	<code>(rem 13 -4) => 1</code>
<code>(mod -13 -4) => -1</code>	<code>(rem -13 -4) => -1</code>
<code>(mod 13.4 1) => 0.4</code>	<code>(rem 13.4 1) => 0.4</code>
<code>(mod -13.4 1) => 0.6</code>	<code>(rem -13.4 1) => -0.4</code>

`ffloor number &optional divisor` [Function]

`fceiling number &optional divisor` [Function]

`ftruncate number &optional divisor` [Function]

`fround number &optional divisor` [Function]

These functions are just like `floor`, `ceiling`, `truncate`, 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 argument 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
```

<code>float-significand</code>	<i>float</i>	[Function]
<code>float-exponent</code>	<i>float</i>	[Function]
<code>scale-float</code>	<i>float integer</i>	[Function]
<code>float-radix</code>	<i>float</i>	[Function]
<code>float-sign</code>	<i>float1</i> &optional <i>float2</i>	[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; then `float-significand` divides the argument by an integral power of b so as to bring its value between $1/b$ (inclusive) and 1 (exclusive), and returns the quotient. If the argument is zero, however, the result equals the argument.

The function `float-exponent` performs a similar operation, but then returns the integer exponent e to which b must be raised to produce the appropriate power for the division. If the argument is zero, any integer value may be returned, provided that the identity shown below for `scale-float` holds.

The function `scale-float` takes a floating-point number f and an integer k , and returns $(* f (\text{expt } (\text{float } b) k))$. (The use of `scale-float` may be much more efficient than using exponentiation and multiplication, and avoids intermediate overflow and underflow if the final result is representable.)

Note that $(\text{scale-float } (\text{float-fraction } f) (\text{float-exponent } f)) \leq f$.

The function `float-radix` returns (as an integer) the radix b of the floating-point argument.

The function `float-sign` returns a floating-point number z such that z and *float1* have the same sign and also such that z and *float2* have the same absolute value. The argument *float2* defaults to the value of $(\text{float } 1 \text{ } \text{float1})$; $(\text{float-sign } x)$ therefore always produces a 1.0 or -1.0 according to the sign of x .

Rationale: These functions allow the writing of machine-independent, or at least machine-parameterized, floating-point software of reasonable efficiency.

<code>complex</code>	<i>realpart</i> &optional <i>imagpart</i>	[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 $(* \text{realpart } 0)$ 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).

<code>realpart</code>	<i>number</i>	[Function]
<code>imagpart</code>	<i>number</i>	[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 $(* \text{number } 0)$ (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).

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.

<code>lognand integer1 integer2</code>	[Function]
<code>lognor integer1 integer2</code>	[Function]
<code>logandc1 integer1 integer2</code>	[Function]
<code>logandc2 integer1 integer2</code>	[Function]
<code>logorc1 integer1 integer2</code>	[Function]
<code>logorc2 integer1 integer2</code>	[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
<code>logand</code>		0	0	0	1	and
<code>logior</code>		0	1	1	1	inclusive or
<code>logxor</code>		0	1	1	0	exclusive or
<code>logeqv</code>		1	0	0	1	equivalence (exclusive nor)
<code>lognand</code>		1	1	1	0	not-and
<code>lognor</code>		1	0	0	0	not-or
<code>logandc1</code>		0	1	0	0	and-complement of arg1 with arg2
<code>logandc2</code>		0	0	1	0	and arg1 with complement of arg2
<code>logorc1</code>		1	1	0	1	or complement of arg1 with arg2
<code>logorc2</code>		1	0	1	1	or arg1 with complement of arg2

<code>boole op integer1 integer2</code>	[Function]
<code>boole-clr</code>	[Constant]
<code>boole-set</code>	[Constant]
<code>boole-1</code>	[Constant]
<code>boole-2</code>	[Constant]
<code>boole-c1</code>	[Constant]
<code>boole-c2</code>	[Constant]
<code>boole-and</code>	[Constant]
<code>boole-ior</code>	[Constant]
<code>boole-xor</code>	[Constant]
<code>boole-eqv</code>	[Constant]
<code>boole-nand</code>	[Constant]
<code>boole-nor</code>	[Constant]

`boole-andc1` [Constant]
`boole-andc2` [Constant]
`boole-orc1` [Constant]
`boole-orc2` [Constant]

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`:

	<i>integer1</i>	0	0	1	1	
	<i>integer2</i>	0	1	0	1	<i>Operation performed</i>
<code>boole-clr</code>		0	0	0	0	always 0
<code>boole-set</code>		1	1	1	1	always 1
<code>boole-1</code>		0	0	1	1	<i>integer1</i>
<code>boole-2</code>		0	1	0	1	<i>integer2</i>
<code>boole-c1</code>		1	1	0	0	complement of <i>integer1</i>
<code>boole-c2</code>		1	0	1	0	complement of <i>integer2</i>
<code>boole-and</code>		0	0	0	1	and
<code>boole-ior</code>		0	1	1	1	inclusive or
<code>boole-xor</code>		0	1	1	0	exclusive or
<code>boole-eqv</code>		1	0	0	1	equivalence (exclusive nor)
<code>boole-nand</code>		1	1	1	0	not-and
<code>boole-nor</code>		1	0	0	0	not-or
<code>boole-andc1</code>		0	1	0	0	and complement of <i>integer1</i> with <i>integer2</i>
<code>boole-andc2</code>		0	0	1	0	and <i>integer1</i> with complement of <i>integer2</i>
<code>boole-orc1</code>		1	1	0	1	or complement of <i>integer1</i> with <i>integer2</i>
<code>boole-orc2</code>		1	0	1	1	or <i>integer1</i> with complement of <i>integer2</i>

`boole` can therefore compute all sixteen logical functions on two arguments. In general,

$$(\text{boole } \text{boole-and } x \ y) \iff (\text{logand } x \ y)$$

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.

$$(\text{logbitp } j \ (\text{lognot } x)) \iff (\text{not } (\text{logbitp } j \ x))$$

`logtest` *integer1 integer2* [Function]

`logtest` is a predicate that is true if any of the bits designated by the 1's in *integer1* are 1's in *integer2*.

$$(\text{logtest } x \ y) \iff (\text{not } (\text{zerop } (\text{logand } x \ y)))$$

`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 $\text{floor}(\text{integer} * 2^{\text{count}})$.

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
```

The following identity always holds:

```
(logcount x) <=> (logcount (- (+ x 1)))
```

`integer-length integer`

[Function]

This function performs the computation

```
[ceiling](log2(if integer < 0 then - integer else integer+1))
```

This is useful in two different ways. First, if *integer* is non-negative, then its value can be represented in unsigned binary form in a field whose width in bits is at least (`integer-length integer`). Second, regardless of the sign of *integer*, its value can be represented in signed binary

two's-complement form in a field whose width in bits is at least $(+ (\text{integer-length } \textit{integer}) 1)$:

For example:

```
(integer-length 0) => 0
(integer-length 1) => 1
(integer-length 3) => 2
(integer-length 4) => 3
(integer-length 7) => 3
(integer-length -1) => 0
(integer-length -4) => 2
(integer-length -7) => 3
(integer-length -8) => 3
```

Compatibility note: This function is similar to the MACLISP function `haulong`. One may define `haulong` as
 $(\text{haulong } x) \Leftrightarrow (\text{integer-length } (\text{abs } x))$

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 bits have weights $2^{\textit{position} + \textit{size} - 1}$ through $2^{\textit{position}}$.

`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”.

Compatibility note: The MACLISP function `hai part` can be implemented in terms of `ldb` as follows:

```
(defun hai part (integer count)
  (let ((x (abs integer)))
    (if (minusp count)
        (ldb (byte (- count) 0) x)
        (ldb (byte count (max (- (integer-length x) n) 0)) x))))
```

`setf` (page 66) may be used with `ldb`, provided that the argument *integer* is specified by a form that is a *place* form acceptable to `setf`, to modify a byte within the integer that is stored in that *place*. The effect is to perform a `dpb` (page 143) operation and then store the result back into the *place*.

`ldb-test` *bytespec integer* [Function]

`ldb-test` is a predicate that 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))
```

`setf` (page 66) may be used with `mask-field`, provided that the argument *integer* is specified by a form that is a *place* form acceptable to `setf`, to modify a byte within the integer that is stored in that *place*. The effect is to perform a `deposit-field` (page 144) operation and then store the result back into the *place*.

`dpb` *newbyte bytespec integer* [Function]

Returns a number that 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 that 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` *number &optional state*

[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 $1/n$. (The qualifier “approximate” is used because of implementation considerations; in practice the deviation from uniformity should be quite small.)

The argument *state* must be an object of type `random-state`; it defaults to the value of the variable `*random-state*`. This object is used to maintain the state of the pseudo-random-number generator, and is altered as a side effect of the `random` operation.

Compatibility note: `random` of zero arguments as defined in MACLISP 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 `(integer-length n)` (see `integer-length` (page 141)), 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 $X=N/M$. See also [7]. If one takes $M = 2^f$, where f is the length of the significand 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 significand 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 significand 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, an object of type `random-state`, that encodes the internal state of the random-number generator that `random` uses by default. 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.

make-random-state &optional state

[Function]

This function returns a new object of type `random-state`, 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

that has been “randomly” initialized by some means (such as by a time-of-day clock).

`random-state-p` *object* [Function]

`random-state-p` is true if its argument is a random-state object, and otherwise is false.

`(random-state-p x) <=> (typep x 'random-state)`

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-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

$$(\text{not } (= (\text{float } 1 e) (+ e (\text{float } 1 e))))$$

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

$$(\text{not } (= (\text{float } 1 e) (- e (\text{float } 1 e))))$$

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 that is the upper exclusive bound on values produced by the function `char-code` (page 154), 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`.

`char-font-limit`

[*Constant*]

The value of `char-font-limit` is a non-negative integer that is the upper exclusive bound on values produced by the function `char-font` (page 154), 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.

`char-bits-limit`

[*Constant*]

The value of `char-bits-limit` is a non-negative integer that is the upper exclusive bound on values produced by the function `char-bits` (page 154), 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.

13.1. Predicates on Characters

The predicate `characterp` (page 54) 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.

`graphic-charp` *char* [Function]

The argument *char* must be a character object. `graphic-charp` 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 <page>, 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, and otherwise is false. Any character that satisfies `standard-charp` also satisfies `string-charp`; others may also.

`alpha-charp` *char* [Function]

The argument *char* must be a character object. `alpha-charp` 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 155)), or if the argument is lower-case and there is a corresponding upper-case character (which can be obtained using char-upcase (page 155)).

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.

digit-charp *char* &optional (*radix* 10.) [Function]

The argument *char* must be a character object, and *radix* must be a non-negative integer.

digit-charp 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 that 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. digit-charp 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 (digit-charp (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 (alpha-charp x) (digit-charp x))
```

Alphanumeric characters are therefore necessarily graphic (as defined by graphic-charp (page

150)).

Of the standard characters (as defined by `standard-charp`), the characters "0" through "9", "A" through "Z", and "a" through "z" are alphanumeric.

<code>char=</code>	<i>character &rest more-characters</i>	[Function]
<code>char/=</code>	<i>character &rest more-characters</i>	[Function]
<code>char<</code>	<i>character &rest more-characters</i>	[Function]
<code>char></code>	<i>character &rest more-characters</i>	[Function]
<code>char<=</code>	<i>character &rest more-characters</i>	[Function]
<code>char>=</code>	<i>character &rest more-characters</i>	[Function]

The arguments must all be character objects. These functions compare the objects using the implementation-dependent total ordering on characters, in a manner analogous to numeric comparisons by `=` (page 122) and related function.

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 122) on their code attributes.

Notice that the total ordering is not necessarily the same as the total ordering on the integers produce by applying `char-int` (page 155) to the characters. Also, while alphabetic characters of a given case must be properly ordered, they need not be contiguous; therefore `(char<= #\a x #\z)` is *not* a valid way of determining whether or not `x` is a lower-case letter, for example; that is why a separate `lower-casep` (page 151) predicate is provided.

For example:

(char= #\d #\d) is true	(char/= #\d #\d) is false
(char= #\d #\x) is false	(char/= #\d #\x) is true
(char= #\d #\D) is false	(char/= #\d #\D) is true
(char= #\d #\d #\d #\d) is true	(char/= #\d #\d #\d #\d) is false
(char= #\d #\d.#\x #\d) is false	(char/= #\d #\d #\x #\d) is false
(char= #\d #\y #\x #\c) is false	(char/= #\d #\y #\x #\c) is true
(char= #\d #\c #\d) is false	(char/= #\d #\c #\d) is false
(char< #\d #\x) is true	(char<= #\d #\x) is true
(char< #\d #\d) is false	(char<= #\d #\d) is true
(char< #\a #\e #\y #\z) is true	(char<= #\a #\e #\y #\z) is true
(char< #\a #\e #\e #\y) is false	(char<= #\a #\e #\e #\y) is true
(char> #\e #\d) is true	(char>= #\e #\d) is true
(char> #\d #\c #\b #\a) is true	(char>= #\d #\c #\b #\a) is true
(char> #\d #\d #\c #\a) is false	(char>= #\d #\d #\c #\a) is true
(> #\e #\d #\b #\c #\a) is false	(>= #\e #\d #\b #\c #\a) is false
(> #\z #\A) <i>may be true or false</i>	(> #\Z #\a) <i>may be true or false</i>

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 56) and equal (page 56) compare character objects in the same way that char= does.

char-equal <i>character &rest more-characters</i>	[Function]
char-not-equal <i>character &rest more-characters</i>	[Function]
char-lessp <i>character &rest more-characters</i>	[Function]
char-greaterp <i>character &rest more-characters</i>	[Function]
char-not-greaterp <i>character &rest more-characters</i>	[Function]
char-not-lessp <i>character &rest more-characters</i>	[Function]

The predicate char-equal is like char=, and similarly for the others, except according to a different ordering such that differences of bits attributes and case are ignored, and font information is taken into account in an implementation-dependent manner. For the standard characters, the ordering is such that A=a, B=b, and so on, up to Z=z, and furthermore either 9<A or Z<0.

For example:

```
(char-equal #\A #\a) is true
(char= #\A #\a) is false
(char-equal #\A #\Control-A) is true
```

The ordering may depend on the font information. For example, an implementation might decree that (char-equal #\p #\p) be true, but that (char-equal #\p #\pi) be false (where #\pi is a lower-case “p” in some font).

13.2. Character Construction and Selection

`character object` [Function]

The function `character` coerces its argument to be a character if possible; see `coerce` (page 37).

`(character x) <=> (coerce x 'character)`

`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 149).

`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 149).

`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 149).

`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 151) is true of *char*, and `uppercasep` (page 151) 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 151) is true of *char*, and `lowercasep` (page 151) 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-weight weight &optional (radix 10.) (bits 0) (font 0)` [Function]

All arguments must be integers. `digit-weight` 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 `digit-charp` (page 151)). It returns such a character if that is possible, and otherwise returns `nil`.

`digit-weight` 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*.

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.

`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 that have zero font and bits attributes and that are non-graphic (do not satisfy the predicate `graphic-charp` (page 150)) 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>`, `<page>`, `<rubout>`, `<linefeed>`, and `<backspace>` have the respective names `tab`, `page`, `rubout`, `linefeed`, and `backspace`.

Characters that 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
```

`setf` (page 66) may be used with `char-bit`, provided that the argument *char* is specified by a form that is a *place* form acceptable to `setf`, to modify a bit of the character stored in that *place*. The effect is to perform a `set-char-bit` (page 157) operation and then store the result back into the *place*.

`set-char-bit` *char name newvalue*

[Function]

`set-char-bit` takes a character object *char*, the name of a bit, and a flag. A character is returned that 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 vectors (one-dimensional arrays). While these are different data structures with different structural properties leading to different algorithmic uses, they do have a common property: each contains an ordered set of elements.

There are some operations that 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:

<code>elt</code>	<code>reverse</code>	<code>map</code>	<code>remove</code>	<code>remove-duplicates</code>
<code>subseq</code>	<code>nreverse</code>	<code>some</code>	<code>delete</code>	<code>delete-duplicates</code>
<code>copy-seq</code>	<code>concatenate</code>	<code>every</code>	<code>position</code>	<code>find</code>
<code>fill</code>	<code>length</code>	<code>notany</code>	<code>mismatch</code>	<code>search</code>
<code>replace</code>	<code>sort</code>	<code>notevery</code>	<code>maxprefix</code>	<code>substitute</code>
<code>count</code>	<code>merge</code>	<code>reduce</code>	<code>maxsuffix</code>	<code>nsubstitute</code>

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 `eq1` to that item. (A test other than `eq1` can be specified by the `:test` or `:test-not` keyword.) 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 `eq1` 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.

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 $start \leq end$; 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 that 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 keyword arguments `:start1`, `:end1`, `:start2`, and `:end2` are 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 165).

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 169).

As a rule, whenever a sequence function must construct and return a new vector, it is always a simple vector.

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.

`setf` (page 66) may be used with `elt` to destructively replace a sequence element with a new value.

`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*.

`setf` (page 66) may be used with `subseq` to destructively replace a subsequence with a sequence of new values; see also `replace` (page 165).

`copy-seq` *sequence* [Function]

A copy is made of the argument *sequence*; the result is equal to the argument but not eq to it.

`(copy-seq x) <=> (subseq x 0)`

but the name `copy-seq` is more perspicuous when applicable.

`length` *sequence* [Function]

The number of elements in *sequence* is returned as a non-negative integer. If the sequence is a vector with a fill pointer, the “active length” as specified by the fill pointer is returned. See section 17.6 (page 199).

`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*.

`make-sequence` *type size* &key *:initial-element* [Function]

This returns a sequence of type *type* and of length *size*, each of whose elements has been initialized to the *:initial-element* argument. If specified, the *:initial-element* argument must be an object that can be an element of a sequence of type *type*.

For example:

```
(make-sequence '(vector double-float) 100 :initial-element 1d0)
```

If an *:initial-element* argument is not specified, then the sequence will be initialized in an implementation-dependent way.

14.2. Catenating, Mapping, and Reducing Sequences

`concatenate` *result-type* &rest *sequences* [Function]

The result is a new sequence that 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 `concatenate` differs from `append`). The type of the result is specified by *result-type*, which must be a subtype of `sequence`, as for the function `coerce` (page 37). 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 `concatenate` 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, `concatenate` returns a new empty sequence of type *result-type*.

If one argument is provided, and it has the type specified by *result-type*, `concatenate` is required to copy the argument rather than simply returning it. If a copy is not required, but only possible type-conversion, then the `coerce` (page 37) function may be appropriate.

`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 `coerce` (page 37). In addition, one may specify `nil` for the result type, meaning that no result sequence is to be produced; in this case the *function* is invoked only for effect, and `map` returns `nil`. This gives an effect similar to that of `mapc` (page 85).

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. However, standard computer science literature, and in particular the recent wave of papers on "functional programming", have come to use `map` to mean what in the past LISP people have called `mapcar`. To simplify things henceforth, COMMON LISP follows current usage, and what was formerly called `map` is named `mapl` (page 85) in COMMON LISP.

For example:

```
(map 'list #'- '(1 2 3 4)) => (-1 -2 -3 -4)
(map 'string #'(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.

reduce *function sequence* &key :from-end :start :end :initial-value [Function]

The specified subsequence of the *sequence* is "reduced" using the *function*, which must accept two arguments. The reduction is left-associative, unless the :from-end argument is true (it defaults to nil), in which case it is right-associative. If an :initial-value argument is given, it is logically placed before the subsequence (after it if :from-end is true) and included in the reduction operation. If no :initial-value is given, and the specified subsequence is empty, then the *function* is called with zero arguments, and reduce returns whatever the function does. (This is the only case where the *function* is called with other than two arguments.)

For example:

```
(reduce #' + '(1 2 3 4)) => 10
(reduce #' - '(1 2 3 4)) <=> (- (- (- 1 2) 3) 4) => -8
(reduce #' - '(1 2 3 4) :from-end t) ; Alternating sum.
  <=> (- 1 (- 2 (- 3 4))) => -2
(reduce #' list '(1 2 3 4)) => (((1 2) 3) 4)
(reduce #' list '(1 2 3 4) :from-end t) => (1 (2 (3 4)))
(reduce #' list '(1 2 3 4) :initial-value 'foo)
  => (((foo 1) 2) 3) 4)
(reduce #' list '(1 2 3 4) :from-end t :initial-value 'foo)
  => (1 (2 (3 (4 foo))))
```

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 :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:

$$(\min (- \text{end1 } \text{start1}) (- \text{end2 } \text{start2}))$$

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 [Function]
:count :key

`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* that 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 [Function]
:count :key

`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* that 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)
```

Compatibility note: In MACLISP, the `delete` function uses an equal comparison rather than `eq`, which is the default test for `delete` in COMMON LISP. Where in MACLISP one would write `(delete x y)` one must in COMMON LISP write `(delete x y :test #'equal)`.

`remove-duplicates` *sequence* &key :test :test-not :start :end [Function]

`delete-duplicates` *sequence* &key :test :test-not :start :end [Function]

The elements of *sequence* are examined, and if any two match then one is discarded. The result is a sequence of the same kind as the argument sequence with enough elements removed so that no two of the remaining elements match. `remove-duplicates` is the non-destructive version of this operation, whereas `delete-duplicates` may destroy the argument *sequence*.

`remove-duplicates` is useful for converting a sequence into a canonical form suitable for representing a set.

```
(remove-duplicates '(a b c b d d e))
=> (a c b d e) or (a b c d e)
```

`substitute` *newitem olditem sequence* &key :from-end :test :test-not [Function]
:start :end :count :key

`substitute-if` *newitem test sequence* &key :from-end :start :end [Function]
:count :key

`substitute-if-not` *newitem test sequence* &key :from-end :start :end [Function]

:count :key

The result is a sequence of the same kind as the argument *sequence* that 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` [Function]

`:start` `:end` `:count` `:key`

`nsubstitute-if` *newitem test sequence* &key `:from-end` `:start` `:end` [Function]

`:count` `:key`

`nsubstitute-if-not` *newitem test sequence* &key `:from-end` `:start` `:end` [Function]

`:count` `:key`

This is the destructive counterpart to `substitute`. The result is a sequence of the same kind as the argument *sequence* that 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 :key [Function]
          :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 :key [Function]
          :start1 :start2 :end1 :end2
maxsuffix sequence1 sequence2 &key :from-end :test :test-not :key [Function]
          :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 :key           [Function]
      :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.

14.5. Sorting and Merging

```
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 `: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 211) 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 expressions 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 argument, this is accomplished by permuting the elements in place. In the case of a list, the list is destructively reordered in the same manner as for `nreverse` (page 162). 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 foovector #'string-lessp :key #'mostcar)
```

If `foovector` 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 foovector 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 211) 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 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 175).

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 53). 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
```

See `first` (page 176). The *car* of a cons may be altered by using `replace` (page 181) or `setf` (page 66).

`cdr x`

[Function]

Returns the *cdr* of *x*, which must be a cons or `()`; that is, *x* must satisfy the predicate `listp` (page 53). 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)
```

See `rest` (page 176). The *cdr* of a cons may be altered by using `replace` (page 181) or `setf` (page 66).

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 *caddr* 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 211), which will automatically define new record data types and access functions for instances of them.

Any of these functions may be used to specify a *place* for *setf* (page 66).

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 &key :test :test-not

[Function]

This is a predicate that is true if *x* and *y* are isomorphic trees with identical leaves; that is, if *x* and *y* are atoms that satisfy the test (by default *eq1*), or if they are both *conses* and their *cars* are *tree-equal* and their *cdrs* are *tree-equal*. Thus *tree-equal* recursively compares *conses* (but not any other objects that have components). See *equal* (page 56), 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 161), 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.

`nth` may be used to specify a *place* to `setf` (page 66); when `nth` is used in this way, the argument *n* must be less than the length of the *list*.

<i>first list</i>	[Function]
<i>second list</i>	[Function]
<i>third list</i>	[Function]
<i>fourth list</i>	[Function]
<i>fifth list</i>	[Function]
<i>sixth list</i>	[Function]
<i>seventh list</i>	[Function]
<i>eighth list</i>	[Function]
<i>ninth list</i>	[Function]
<i>tenth list</i>	[Function]

These functions are sometimes convenient for accessing particular elements of a list. *first* is the same as *car* (page 173); *second* is the same as *cadr*; and so on. Note that the ordinal numbering used here is one-origin, as opposed to the zero-origin numbering used by *nth* (page 175):

```
(fifth x) <=> (nth 4 x)
```

??? Query: Should these be general sequence functions?

rest list [Function]

rest means the same as *cdr*, but mnemonically complements *first*.

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* &key :initial-element

[Function]

This creates and returns a list containing *size* elements, each of which is initialized to the *:initial-element* argument (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 that 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 `concatenate` (page 162) can perform a similar operation, but always copies all its arguments. See also `nconc` (page 178), which is like `append` but destroys all arguments but the last.

The last argument actually need not be a list, but may be any LISP object, which becomes the tail end of the constructed list. For example, `(append '(a b c) 'd) => (a b c . d)`.

`(append x '())` is an idiom once frequently used to copy the list *x*, but the `copy-list` function is more appropriate to this task.

`copy-list list` [Function]

Returns a list that is equal to *list*, but not eq. Only the top level of list-structure is copied; that is, `copy-list` 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 `copy-seq` (page 161).

`copy-alist list` [Function]

`copy-alist` is for copying association lists. The top level of list structure of *list* is copied, just as `copy-list` does. In addition, each element of *list* that is a cons is replaced in the copy by a new cons with the same *car* and *cdr*.

`copy-tree object` [Function]

`copy-tree` 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 `copy-tree` 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 potentially more efficient. Both *x* and *y* should be lists. The argument *x* is copied, not destroyed. Compare this with `nreconc` (page 179), which destroys its first argument.

`nconc &rest lists` [Function]

`nconc` takes lists as arguments. It returns a list that is the arguments concatenated together. The arguments are changed, rather than copied. (Compare this with `append` (page 177), 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 replaced 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; if the `*princircle*` (page 248) switch were non-`nil`, it would be printed as (a b c . #1=(d e f . #1#)).

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 178).

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 66). 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 comparisons using the **:test** predicate, which defaults to **eq1**), then the *item* is consed onto the front of the list, and the augmented list is stored back into *place* and returned; otherwise the unaugmented list is returned. The form *place* may be any form acceptable as a generalized variable to **setf** (page 66). If the list held in *place* is viewed as a set, then **pushnew** adjoins an element to the set; see **adjoin** (page 184). **pushnew** returns **nil**.

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)) => (5 b c) and x is unchanged
```

The effect of (**pushnew** *item place* **:test** *p*) is roughly equivalent to

```
(setf place (adjoin item place :test p))
```

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

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 66). 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) => ()
```

ldiff *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*. **ldiff** (meaning "list difference") 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 the entire *list* is returned. The argument *list* is not destroyed.

For example:

```
(setq x '(a b c d e))
(setq y (caddr x)) => (d e)
(ldiff x y) => (a b c)
but
(ldiff '(a b c d) '(c d)) => (a b c d)
since the sublist was not eq to any part of the list.
```

15.3. Alteration of List Structure

The functions **rplaca** and **rplacd** may be used to make alterations in already-existing list structure; that is, to change the cars and cdrs of existing conses. One may also use **setf** (page 66) in conjunction with **car** and **cdr** (page 173).

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 178), **nreverse** (page 162), **nreconc** (page 179), and **nbutlast** (page 180) functions already described, and the **delete** (page 166) 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* must 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* must 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)
```

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. They come in non-destructive and destructive varieties, and specify substitution either by two arguments or by an association list.

`subst new old tree &key :test :test-not :key` [Function]

`subst-if predicate new tree &key :key` [Function]

`subst-if-not predicate new tree &key :key` [Function]

`(subst new old tree)` substitutes *new* for every leaf of *tree* (whether a *car* or a *cdr*) such that *old* and the leaf satisfy the test, and returns the modified copy of *tree*. The original *tree* is unchanged, but the result tree may share with parts of the argument *tree*.

Compatibility note: In MACLISP, `subst` is guaranteed *not* to share with the *tree* argument, and `(subst nil nil x)` was used as an idiom for copying the tree *x*. In COMMON LISP, the function `copy-tree` (page 178) should be used to copy a tree, as the `subst` idiom will not work.

For example:

```
(subst 'tempest 'hurricane
      '(shakespeare wrote (the hurricane)))
=> (shakespeare wrote (the tempest))
(subst 'foo 'nil '(shakespeare wrote (twelfth night)))
=> (shakespeare wrote (twelfth night . foo) . foo)
```

This function is not destructive; that is, it does not change the *car* or *cdr* of any already-existing list structure. One possible definition of `subst`:

```
(defun subst (old new tree &rest x &key test test-not key)
  (cond ((atom tree)
        (if (satisfies-the-test old tree :test test
                                :test-not test-not :key key)
            new tree))
        (t (let ((a (apply #'subst old new (car tree) x))
                  (d (apply #'subst old new (cdr tree) x)))
              (if (and (eq a (car tree)) (eq d (cdr tree)))
                  tree
                  (cons a d)))))))
```

For example:

See also `substitute` (page 166), which substitutes for top-level elements of sequence.

`nsubst new old tree &key :test :test-not :key` [Function]

`nsubst-if predicate new tree &key :key` [Function]

`nsubst-if-not predicate new tree &key :key` [Function]

`nsubst` is a destructive version of `subst`. The list structure of *tree* is altered by destructively replacing with *new* each leaf of the *tree* such that *old* and the leaf satisfy the test.

`sublis alist tree &key :test :test-not :key` [Function]

`sublis` makes substitutions for symbols in a tree (a structure of conses). The first argument to `sublis` is an association list. The second argument is the tree in which substitutions are to be made, as for `subst` (page 182). `sublis` looks at all leaves in the tree; if a leaf appears as a key in the association list (that is, the key and the leaf satisfy the test), it is replaced by the object it is associated with. This operation is non-destructive. In effect, `sublis` can perform several `subst` operations simultaneously.

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 &key :test :test-not :key` [Function]

`nsublis` is like `sublis` but destructively modifies the relevant leaves of the *tree*.

15.5. Using Lists as Sets

COMMON LISP includes functions that 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 &key :key` [Function]

The *list* is searched for an element that satisfies the test. If none is found, `nil` is returned; otherwise, the tail of *list* beginning with the first element that satisfied the test is returned. The *list* is searched on the top level only. These functions are suitable for use as predicates.

For example:

```
(member 'snerd '(a b c d)) => nil
(member-if #'numberp '(a #\Space 5/3 foo)) => (5/3 foo)
(member 'a '(g (a y) c a d e a f)) => (a d e a f)
```

Note, in the last example, 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.

See also `find` (page 167) and `position` (page 168).

Compatibility note: In MACLISP, the `member` function uses an `equal` comparison rather than `eq`, which is

the default test for `member` in COMMON LISP. Where in MACLISP one would write `(member x y)` one must in COMMON LISP write `(member x y :test #'equal)`.

`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 `ldiff` (page 181).

`adjoin` *item list* &key :test :test-not :key [Function]

`adjoin` is used to add an element to a set, provided that it is not already a member. The equality test defaults to `eq`.

```
(adjoin item list) <=> (if (member item list) list (cons item list))
```

See `pushnew` (page 179).

`union` *list1 list2* &key :test :test-not :key [Function]

`nunion` *list1 list2* &key :test :test-not :key [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 :key [Function]

`nintersection` *list1 list2* &key :test :test-not :key [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.)

`set-difference list1 list2 &key :test :test-not :key` [Function]

`nset-difference list1 list2 &key :test :test-not :key` [Function]

`set-difference` returns a list of elements of `list1` that do not appear in `list2`. This operation is not destructive.

`nset-difference` is the destructive version of `setdifference`. This operation may destroy `list1`.

`set-exclusive-or list1 list2 &key :test :test-not :key` [Function]

`nset-exclusive-or list1 list2 &key :test :test-not :key` [Function]

`set-exclusive-or` returns a list of elements that 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 :key` [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 186) 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 187) is provided. Other variants of a-list searches can be constructed using the function `find` (page 167) or `member` (page 183).

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 that 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-if` *predicate a-list* [Function]

`assoc-if-not` *predicate a-list* [Function]

Each of these searches the association list *a-list*. The value is the first pair in the *a-list* such that the *car* of the pair satisfies the test, or `nil` if there is none such.

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 `eq`, as the default comparison test.

```
(assoc item list :test fn)
<=> (find item list :test fn :key #'car)
```

See `find` (page 167) and `position` (page 168).

Compatibility note: In `MACLISP`, the `assoc` function uses an `equal` comparison rather than `eq`, which is the

default test for `assoc` in COMMON LISP. Where in MACLISP one would write `(assoc x y)` one must in COMMON LISP write `(assoc x y :test #'equal)`.

`rassoc item a-list &key :test :test-not` [Function]

`rassoc-if predicate a-list` [Function]

`rassoc-if-not predicate a-list` [Function]

`rassoc` is the reverse form of `assoc`; it searches for a pair whose *cdr* satisfies the test, rather than the *car*. If the *a-list* is considered to be a mapping, then `rassoc` treats the *a-list* as representing the inverse mapping.

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)
```


Chapter 16

Hash Tables

A hash table is a LISP object that can efficiently map a given LISP object to another LISP object. 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`, `eq1`, or `equal`. In other words, there are hash tables that hash on Lisp *objects* (using `eq` or `eq1`) and there are hash tables that 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`.

16.1. Hash Table Functions

This section documents the functions for hash tables, which use *objects* as keys and associate other objects with them.

`make-hash-table &key :test :size :rehash-size :rehash-threshold` [Function]

This function creates and returns a new hash table. The `:test` argument determines how keys are compared; it must be one of the three values `#'eq`, `#'eql`, or `#'equal`, or one of the three symbols `eq`, `eql`, or `equal`.

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 value for this argument is implementation-dependent.

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 value for this argument is implementation-dependent.

For example:

```
(make-hash-table :rehash-size 1.5
                 :size (* number-of-widgets 43))
```

hash-table-p *object*

[Function]

hash-table-p is true if its argument is a hash table, and otherwise is false.

```
(hash-table-p x) <=> (typep x 'hash-table)
```

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.

setf (page 66) may be used with **gethash** to make new entries in a hash table. In this context, the *default* argument should not be specified to **gethash**. If an entry with the specified *key* already exists, it is removed before the new entry is added.

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.

hash-table-count *hash-table*

[Function]

This returns the number of entries in the *hash-table*. When a hash table is first created or has been cleared, the number of entries is zero.

16.2. 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 16a

Arrays

An array is an object with components arranged according to a rectilinear coordinate system. Arrays in COMMON LISP may have any number of dimensions, including zero. (A zero-dimensional array has exactly one element.) Every COMMON LISP implementation must support arrays with up to 63 dimensions. Each dimension is a non-negative integer; if any dimension of an array is zero, the array has no elements.

An array may be a *general array*, meaning each element may be any LISP object, or it may be a *specialized array*, meaning that each element must be of a given restricted type.

One-dimensional arrays are called vectors. General vectors may contain any LISP object. Vectors whose elements are restricted to type `string-char` are called *strings*. Vectors whose elements are restricted to type `bit` are called *bit-vectors*.

17.1. Array Creation

```
make-array dimensions &key :element-type :initial-element           [Function]
          :initial-contents :adjustable :fill-pointer
          :displaced-to :displaced-index-offset
```

This is the primitive function for making arrays. The *dimensions* argument 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. Note that if *dimensions* is `nil` then a zero-dimensional array is created. For convenience when making a one-dimensional array, the single dimension may be provided as an integer rather than a list of one integer.

An implementation of COMMON LISP may impose a limit on the rank of an array, but this limit may not be smaller than 63. Therefore, any COMMON LISP program may assume the use of arrays of rank 63 or less. In any case, the rank of an array must be a fixnum. The implementation-dependent limit on array rank is reflected in `array-rank-limit` (page 195).

The `:element-type` argument should be the name of the type of the elements of the array; an array is constructed of the most specialized type that can nevertheless accommodate elements 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-element` argument may be used to initialize each element of the array. The value must be of the type specified by the `:element-type` argument. If the `:initial-element` option is omitted, the initial values of the array elements are undefined (unless the `:initial-contents` or `:displaced-to` option is used). The `:initial-element` 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) (1 2 3))
              ((d e f) (3 1 2))
              ((g h i) (2 3 1))
              ((j k l) (0 0 0))))
```

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-element` or `:displaced-to` option is used). The `:initial-contents` option may not be used with the `:initial-element` or `:displaced-to` option.

The `:adjustable` argument, if specified and not `nil`, indicates that it must be possible to alter the array's size dynamically after it is created.

The `:fill-pointer` argument specifies that the array should have a fill pointer. If this option is specified and not `nil`, the array must be one-dimensional. The value is used to initialize the fill pointer for the array. If the value `t` 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; `make-array` will create an *indirect* or *shared* array that 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-element` 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.

array-rank-limit

[Constant]

The value of `array-rank-limit` is a positive integer that is the upper exclusive bound on the rank of an array. This value will not be smaller than 64; therefore every COMMON LISP implementation supports arrays whose rank is between 0 and 63 (inclusive).

`make-simple-vector` *length* &key *:element-type* *:initial-element*
:initial-contents

[Function]

`make-vector` is like `make-array` (page 193), but guarantees to return a simple 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-element`, and `:initial-contents` keyword arguments are as for `make-array`.

The function `make-array` in fact always returns a simple vector if requested to make a one-dimensional array that is not displaced, has no fill pointer, and does not have adjustable size. However, `make-simple-vector` may be easier for clearer to use in some situations.

vector &rest *objects*

[Function]

The function `vector` is a convenient means for creating a simple general vector with specified initial contents. It is analogous to the function `list`.

```
(vector  $a_1$   $a_2$  ...  $a_n$ )
=> (make-simple-vector  $n$  :type  $t$ 
      :initial-contents '( $a_1$   $a_2$  ...  $a_n$ ))
```

17.2. 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.

`aref` is unusual among the functions that operate on arrays in that it completely ignores fill pointers. `aref` can access without error any array element, whether active or not.

`setf` (page 66) may be used with `aref` to destructively replace an array element with a new value.

17.3. Array Information

array-element-type *array*

[Function]

`array-element-type` returns a type specifier for the set of objects that can be stored in the *array*. This set may be larger than the set requested when the array was created; for example, the result of

```
(array-element-type (make-array 5 :element-type '(mod 5)))
```

could be `(mod 5)`, `(mod 8)`, `fixnum`, `t`, or any other supertype of `(mod 5)`.

array-rank *array*

[Function]

Returns the number of dimensions (axes) of *array*. This will be a non-negative integer. See `array-rank-limit` (page 195).

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 *array* *axis-number*

[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*. If the *array* is a vector with a fill pointer, `array-dimension` returns the total size of the vector, including inactive elements, not the size indicated by the fill pointer.

Compatibility note: This is similar to the Lisp Machine LISP function `array-dimension-n`, but takes its

arguments in the other order, and is zero-origin for consistency instead of one-origin. 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.

17.4. Access Functions for Simple Vectors

The functions in this section are equivalent in operation to `aref` (page 196) corresponding more general functions, but require arguments to be simple vectors (of general or specialized type). These functions are provided primarily for reasons of efficiency and convenience.

`svref` *simple-vector* *index* [Function]

The element of the *simple-vector* specified by the integer *index* is returned. The *index* must be non-negative and less than the length of the vector.

`setf` (page 66) may be used with `svref` to destructively replace a simple-vector element with a new value.

`sgvref` *general-vector* *index* [Function]

The element of the *general-vector* specified by the integer *index* is returned. The vector must be a simple general vector, not a non-simple or specialized one. The *index* must be non-negative and less than the length of the vector.

`setf` (page 66) may be used with `sgvref` to destructively replace an element with a new value.

17.5. Functions on Arrays of Bits

`bit` *bit-array* &rest *subscripts* [Function]

`bit` is exactly like `aref` (page 196) but requires an array of bits, that is, one of type (`array bit`). The result will always be 0 or 1.

`setf` (page 66) may be used with `bit` to destructively replace a bit-array element with a new value.

bit-and *bit-array bit-array-2* &optional *result-bit-array* [Function]
 bit-ior *bit-array bit-array-2* &optional *result-bit-array* [Function]
 bit-xor *bit-array bit-array-2* &optional *result-bit-array* [Function]
 bit-eqv *bit-array bit-array-2* &optional *result-bit-array* [Function]
 bit-nand *bit-array1 bit-array2* &optional *result-bit-array* [Function]
 bit-nor *bit-array1 bit-array2* &optional *result-bit-array* [Function]
 bit-andc1 *bit-array1 bit-array2* &optional *result-bit-array* [Function]
 bit-andc2 *bit-array1 bit-array2* &optional *result-bit-array* [Function]
 bit-orc1 *bit-array1 bit-array2* &optional *result-bit-array* [Function]
 bit-orc2 *bit-array1 bit-array2* &optional *result-bit-array* [Function]

These functions perform bit-wise logical operations on bit-arrays. All of the arguments to any of these functions must be bit-arrays of the same rank and dimensions. The result is a bit-array of matching rank and dimensions, such that any given bit of the result is produced by operating on corresponding bits from each of the arguments.

If the third argument is *nil* or omitted, a new array is created to contain the result. If the third argument is a bit-array, the result is destructively placed into that array. If the third argument is *t*, then the first argument is also used as the third argument; that is, the result is placed back in the first array.

The following table indicates what the result bit is for each operation when two arguments are given. (Those operations that accept an indefinite number of arguments are commutative and associative, and require at least one argument.)

	<i>argument1</i>	0	0	1	1	
	<i>argument2</i>	0	1	0	1	<i>Operation name</i>
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 <i>argument1</i> with <i>argument2</i>
bit-andc2		0	0	1	0	and <i>argument1</i> with complement of <i>argument2</i>
bit-orc1		1	1	0	1	or complement of <i>argument1</i> with <i>argument2</i>
bit-orc2		1	0	1	1	or <i>argument1</i> with complement of <i>argument2</i>

For example:

```

(bit-and #*1100 #*1010) => #*1000
(bit-xor #*1100 #*1010) => #*0110
(bit-andc1 #*1100 #*1010) => #*0100
  
```

See `logand` (page 138) and related functions.

`bit-not` *bit-array* &optional *result-bit-array* [Function]

The first argument must be an array of bits. A bit-array of matching rank and dimensions is returned that contains a copy of the argument with all the bits inverted. See `lognot` (page 140).

If the second argument is `nil` or omitted, a new array is created to contain the result. If the second argument is a bit-array, the result is destructively placed into that array. If the second argument is `t`, then the first argument is also used as the second argument; that is, the result is placed back in the first array.

17.6. Fill Pointers

A set of functions for manipulating a *fill pointer* are provided in COMMON LISP to make it easy to incrementally fill in the contents of a vector, and more generally to allow efficient varying of the length of a vector. For example, a string with a fill pointer has most of the characteristics of a PL/I varying string.

The fill pointer is a non-negative integer no larger than the total number of elements in the vector (as returned by `array-dimension` (page 196)); it is the number of “active” or “filled-in” elements in the vector. The fill pointer constitutes the “active length” of the vector; all vector elements whose index is less than the fill pointer are active, and the others are inactive. Nearly all functions that operate on the contents of a vector will operate only on the active elements. An important exception is `aref` (page 196), which can be used to access any vector element whether in the active region of the vector or not. It is important to note that vector elements not in the active region are still considered part of the vector. (An implication of this for implementors is that vector elements outside the active region may not be garbage-collected.)

Only vectors (one-dimensional arrays) may have fill pointers; multi-dimensional arrays may not. (Note, however, that one can create a multi-dimensional arrays that is *displaced* to a vector that has a fill pointer.)

`reset-fill-pointer` *vector* &optional *index* [Function]

The fill pointer of *vector* is reset to *index*, which defaults to zero. The *index* must be a non-negative integer not greater than the dimension of the vector; see `array-dimension` (page 196).

`vector-push` *vector* *new-element* [Function]

vector must be a one-dimensional array that has a fill pointer, and *new-element* may be any object. `vector-push` attempts to store *new-element* in the element of the vector designated by the fill pointer, and increase the fill pointer by one. If the fill pointer does not designate an element of the vector (specifically, when it gets too big), it is unaffected and `vector-push` returns `nil`. Otherwise, the store and increment take place and `vector-push` returns the *former* value of the fill pointer (one less than the one it leaves in the vector); thus the value of `vector-push` is the index of the new element pushed.

`vector-push-extend` *vector* *x* &optional *extension* [Function]

`vector-push-extend` is just like `vector-push` except that if the fill pointer gets too large, the vector is extended (using `adjust-array` (page 200)) so that it can contain more elements; it never "fails" the way `vector-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 vector if it must be extended.

`vector-pop` *vector* [Function]

vector must be a one-dimensional array that has a fill pointer. It is an error if the fill pointer is zero. The fill pointer is decreased by one, and the vector element designated by the new value of the fill pointer is returned.

17.7. Changing the Dimensions of an Array

`adjust-array` *array* *new-dimensions* &key *:type* *:initial-element* [Function]
 :initial-contents *:fill-pointer*
 :displaced-to *:displaced-index-offset*

`adjust-array` takes an array and a number of other arguments as for `make-array` (page 193). The number of dimensions specified by *new-dimensions* must equal the rank of *array*.

`adjust-array` returns an array of the same type and rank as *array*, with the specified *new-dimensions*. In effect, the *array* argument itself is modified to conform to the new specifications, but this may be achieved either by modifying the *array* or by creating a new array and modifying the *array* argument to be *displaced* to the new array.

In the simplest case, one specifies only the *new-dimensions* and possibly an *:initial-element* argument. 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 the *:initial-element*; if this argument is not provided, then the initial contents of any new elements are undefined.

If *:type* is specified, then *array* must be such that it could have been originally created with that type; otherwise an error is signalled. Specifying *:type* to `adjust-array` serves only to require such an error check.

If *:initial-contents* or *:displaced-to* is specified, then it is treated as for `make-array`. In this case none of the original contents of *array* appears in the new array.

If *:fill-pointer* is specified, the fill pointer of the *array* is reset as specified. An error is signalled if *array* had no fill pointer already.

`adjust-array` 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.

It is not permitted to call `adjust-array` on an array that was not created with the `:adjustable` option.

If `adjust-array` is applied to an *array* that is displaced to another array *x*, then afterwards neither *array* nor the returned result is displaced to *x* unless such displacement is explicitly re-specified in the call to `adjust-array`.

Example: suppose that the 4-by-4 array *m* has the following contents:

alpha	beta	gamma	delta
epsilon	zeta	eta	theta
iota	kappa	lambda	mu
nu	xi	omicron	pi

Then the result of

```
(adjust-array m '(3 5) :initial-element 'baz)
```

is a 3-by-5 array with contents

alpha	beta	gamma	delta	baz
epsilon	zeta	eta	theta	baz
iota	kappa	lambda	mu	baz

Note that if array *a* is created displaced to array *b* and subsequently array *b* is given to `adjust-array`, array *a* will still be displaced to array *b*; the effects of this displacement and the rule of row-major storage order must be taken into account.

192j

Chapter 17

Strings

A string is a specialized kind of vector (one-dimensional array) whose elements are characters. Specifically, the type `string` is identical to the type `(vector string-char)`, which in turn is the same as `(array string-char (*))`.

As a rule, any string-specific function whose name begins with the prefix “string” will accept a symbol instead of a string as an argument *provided* that 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 `eq1` 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 54).

Note that strings, like all vectors, may have fill pointers (though such strings are not *simple*). String operations generally operate only on the active portion of the string (below the fill pointer). See `reset-fill-pointer` (page 199) and related functions.

17.1. String Access

`char string index`

[Function]

The given *index* must be a non-negative integer less than the length of *string*, which must be a (not necessarily simple) 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 150).) 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) => #\1
```

See `aref` (page 196) and `elt` (page 161). In effect,

```
(char s j) <=> (aref (the string s) j)
```

`setf` (page 66) may be used with `char` to destructively replace a character within a string.

17.2. String Comparison

```
string= string1 string2 &key :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 56) 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" "frog" :start1 1 :end1 3 :start2 2)
is true
```

```
string-equal string1 string2 &key :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 153) is true of them.

For example:

```
(string-equal "foo" "Foo") is true
```

```
string< string1 string2 &key :start1 :end1 :start2 :end2 [Function]
```

```
string> string1 string2 &key :start1 :end1 :start2 :end2 [Function]
```

```
string<= string1 string2 &key :start1 :end1 :start2 :end2 [Function]
```

```
string>= string1 string2 &key :start1 :end1 :start2 :end2 [Function]
```

```
string/= string1 string2 &key :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 152), 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 :start1 :end1 :start2 :end2      [Function]
string-greaterp string1 string2 &key :start1 :end1 :start2 :end2  [Function]
string-not-lessp string1 string2 &key :start1 :end1 :start2 :end2  [Function]
string-not-greaterp string1 string2 &key :start1 :end1 :start2 :end2 [Function]
string-not-equal string1 string2 &key :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 as if `char-lessp` (page 153) were used instead of `char<` (page 152) for comparing characters.

17.3. String Construction and Manipulation

```
make-string size &key :initial-element      [Function]
```

This returns a string of length *size*, each of whose characters has been initialized to the `:initial-element` argument. If an `:initial-element` argument 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 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 any sequence containing characters.

??? Query: Should this be generalized to any sequence, calling them `trim`, `left-trim`, and `right-trim`? Or just one function `trim`, also taking a `:from-end` keyword?

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 155) 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 155)).

The keyword arguments `:start` and `:end` delimit the portion of the string to be affected. The result is always of the same length as *string*, however.

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-upcase "Dr. Livingston, I presume?" 6 10)
=> "Dr. LiVINGston, I presume?"
```

`string-capitalize` produces a copy of *string* such that every word (subsequence of case-modifiable characters or digits delimited by non-case-modifiable non-digits) has its first character, if case-modifiable, in upper-case and any other case-modifiable characters in lower-case.

For example:

```
(string-capitalize " hello ") => " Hello "
(string-capitalize
 "occludEd cASEmenTs F0reSTAl1 iNADVertent DEFenestration")
=> "Occluded Casements Forestall Inadvertent Defenestration"
(string-capitalize 'kludgy-hash-search) => "Kludgy-Hash-Search"
(string-capitalize "DON'T!") => "Don'T!" ;not "Don't!"
(string-capitalize "pipe 13a, foo16c") => "Pipe 13a, Foo16c"
```

`nstring-upcase` *string* &key :start :end [Function]
`nstring-downcase` *string* &key :start :end [Function]
`nstring-capitalize` *string* &key :start :end [Function]

These functions are just like `string-upcase`, `string-downcase`, and `string-capitalize` (page 196), but destructively modify the argument *string* by altering case-modifiable characters as necessary.

The keyword arguments `:start` and `:end` delimit the portion of the string to be affected. The argument *string* is returned as the result.

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 `prin1-to-string` (page 259), `princ-to-string` (page 259), or `format` (page 261).

Chapter 18

Structures

COMMON LISP provides a facility for creating named record structures with named components. In effect, the user can define 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 defined.

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 that 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 that are quite arbitrary and hard to remember. Someone looking at `(caddr 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 to find all the places in the code to be changed to match (not all occurrences of `caddr` 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 `(caddr 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 52), for example. The `defstruct` facility provides all of this.

`defstruct` itself is a macro that 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`, that returns the *x*-position of the ship; `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 52), 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 that is true if its argument is a ship, and is false otherwise.
- A function called `make-ship` is defined that, 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.

- One may use `setf` to alter the components of a `ship`:

```
(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 70) form for each access function.

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* [*doc-string*] {*slot-description*}⁺ [Macro]

Defines a record-structure data type. A general call to `defstruct` looks like this:

```
(defstruct (name option-1 option-2 ...)
  doc-string
  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 52) 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 19.5 (page 214).

If the optional documentation string *doc-string* is present, then it is attached to the *name* as a documentation string of type `structure`; see `documentation` (page 301).

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 that 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 19.4.

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 a constructor function named *make-name*. All names of automatically created functions are symbols of the same package (if any) to which the structure *name* itself belongs. Also, all such functions may be declared *inline* at the discretion of the implementation to improve efficiency; if you do not want some function declared *inline*, follow the *defstruct* form with a *notinline* declaration to override any automatic *inline* declaration.

Because evaluation of a *defstruct* form causes many functions 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 that identifies the structure. In the example above, all the slot names start with "ship-". The *:conc-name* (page 214) option can be used to provide such prefixes automatically.

18.3. Using the Automatically Defined Constructor Function

After you have defined a new structure with *defstruct*, you can create instances of this structure by using the constructor function. By default, *defstruct* defines this function automatically. For a structure named *foo*, the constructor function is normally named *make-foo*; you can specify a different name by giving it as the argument to the *:constructor* (page 215) option, or specify that you don't want a normal constructor function at all by using *nil* as the argument.

A call to a constructor function, in general, has the form

```
(name-of-constructor-function
  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 function.) 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-function call, so that if, for example, the form (`gensym`) were used as an initialization form, either in the constructor-function call or as the default form in the `defstruct` declaration, then every call to the constructor function would call `gensym` once to generate a new symbol.

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. `setf` (page 66) 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.

: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 a constructor function 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.

Specifying this option has the effect of forcing a specific representation, and of forcing the components to be stored in successive elements of the specified representation.

Normally this option is not specified, in which case the structure is represented in an implementation-dependent manner, and the `:named` option is assumed unless `:unnamed` is explicitly specified.

vector Use a general vector, storing components as vector elements. This is normally `:named`. The first component is vector element 1 if the structure is `:named`, and element 0 if it is `:unnamed`.

(vector *element-type*)

A specialized vector may be used, in which case every component must

be of a type that can be stored in such a vector. The first component is vector element 1 if the structure is `:named`, and element 0 if it is `:unnamed`.

`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`. The first component the *cdr* if the structure is `:named`, and the *car* if it is `:unnamed`.

`: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 52). 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 that is explained in section 19.6 (page 217).

`: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 215).

`: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)
                    (:conc-name 'astro))
  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 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. Moreover, `astronaut` will have its own access functions for components defined by the `person` structure. 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
(astro-name x) => buzz
(astro-favorite-beverage x) => tang
```

The difference between the access functions `person-name` and `astro-name` is that `person-name` may be correctly applied to any `person`, including an `astronaut`, while `astro-name` may be correctly applied only to an `astronaut`. (An implementation may or may not check for incorrect use of access functions.)

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* that 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 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))
```

`:print-function`

This option may be used only with `:named` structures. The argument to this option should be a function of three arguments 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, and an integer indicating the current depth (to be compared against `*prinlevel*` (page 252)). The printing function should observe the values of such printer-control variables as `*prinescape*` (page 248) and `*prinpretty*` (page 248).

`: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, 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 functions 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 49) special form. For example,

```
(:eval-when (:eval :compile))
```

will cause the functions to be defined only when the code is running interpreted or inside the compiler.

18.6. By-position Constructor Functions

If the `:constructor` (page 205) 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.

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 form`

[Function]

The *form* is evaluated in the current dynamic environment and a null lexical environment. Whatever results from the evaluation is returned from the call to `eval`.

Note that when you write a call to `eval` *two* levels of evaluation occur on the argument form you write. First the argument form is evaluated, as for arguments to any function, by the usual argument evaluation mechanism (which involves an implicit use of `eval`). Then the argument is passed to the `eval` function, where another evaluation occurs.

For example:

```
(eval (list 'cdr (car '((quote (a . b)) c)))) => b
```

The argument form `(list 'cdr (car '((quote (a . b)) c)))` is evaluated in the usual way to produce the argument `(cdr (quote (a . b)))`; this is then given to `eval` because `eval` is being called explicitly, and `eval` evaluates its argument `(cdr (quote (a . b)))` to produce `b`.

If all that is required for some application is to obtain the current dynamic value of a given symbol, the function `symbol-value` (page 62) may be more efficient than `eval`.

`*evalhook*`

[Variable]

If the value of `*evalhook*` is not `nil`, then `eval` behaves in a special way. The non-`nil` value of `*evalhook*` should be a function that takes arguments according to a lambda-list that looks like `(form &rest env)`; this is called the *hook function*. When a form is to be evaluated (any form at all, even a number or a symbol), whether implicitly or via an explicit call to `eval`, no attempt is made to evaluate the form. Instead, the hook function is invoked, and passed the form to be evaluated as its first argument. The hook function is then responsible for evaluating the form; whatever is returned by the hook function is assumed to be the result of evaluating the form.

The other arguments passed to the hook function contain information about the lexical environment in an implementation-dependent format. These arguments are suitable for the

function `*eval` (page 220) and `evalhook` (page 220).

The `*evalhook*` feature is provided as an aid to debugging. The `step` (page 303) facility is implemented around this hook.

If a non-local exit causes a throw back to the top level of LISP, perhaps because an error could not be corrected, then `*evalhook*` is automatically reset to `nil`.

`*eval form &rest env` [Function]

This function is just like `eval`, but treats `env` as a specification of the lexical environment in which to evaluate the `form`. The format of `env` is implementation-dependent, and may be required to consist of a certain number of arguments, but anything that is passed to a hook function because of the `*evalhook*` feature will be acceptable.

Note that if a hook function simply calls `*eval` to evaluate the form, an endless loop may occur, because `*eval` will invoke the hook function on its argument if `*evalhook*` is not `nil`. See `evalhook` (page 220).

`evalhook form hookfn &rest env` [Function]

The `evalhook` function is provided to make it easier to exploit the `*evalhook*` feature. The `form` is evaluated with `*evalhook*` bound to `hookfn`, which should be a hook function or `nil`. The `env` arguments are used as the lexical environment, as for `*eval` (page 220). The check for a hook function is *bypassed* for the evaluation of the `form` itself, but not for subsidiary evaluations, such as of subforms. It is this one-shot bypass that makes `evalhook` so useful.

Here is an example of a very simple tracing routine that uses the `*evalhook*` feature:

```
(defvar *hooklevel* 0)

(defun hook (x)
  (let ((*evalhook* 'hook-function))
    (eval x)))

(defun hook-function (form &rest env)
  (let ((*hooklevel* (+ *hooklevel* 1)))
    (format trace-output "~%~V@TForm: ~S"
            (* *hooklevel* 2) form)
    (let ((values (multiple-value-list
                    (apply #'evalhook form env))))
      (format trace-output "~%~V@TValue:~{~S ~}"
              (* *hooklevel* 2) values))))
```

Using these routines, one might see the following interaction:

```

(hook '(cons (floor *base* 2) 'b))
Form: (CONS (FLOOR *BASE* 2) (QUOTE B))
Form: (FLOOR *BASE* 3)
Form: *BASE*
Value: 10
Form: 3
Value: 3
Value: 3 1
Form: (QUOTE B)
Value: B
Value: (3 . B)
(3 . B)

```

`constantp` *object*

[Function]

The predicate `constantp` is true of any LISP object that may be considered to have a constant value. This includes all self-evaluating objects (numbers, characters, strings, bit-vectors, and keywords) as well as constant symbols declared by `defconstant` (page 48) such as `nil` (page 51), `t` (page 51), and `pi` (page 130).

19.2. The Top-Level Loop

Normally one interacts with LISP through a “top level read-eval-print loop”, so called because it is the highest level of control and consists of an endless loop that reads an expression, evaluates it, and prints the results. One has an effect on the state of the LISP system only by invoking actions that have side effects.

The precise nature of the top-level loop for COMMON LISP is purposely not specified rigorously here, so that implementors can experiment to improve the user interface. For example, an implementor may choose to require line-at-a-time input, or may provide a fancy editor or complex graphics-display interface. An implementor may choose to prompt explicitly for input, or may choose (as MACLISP does) not to clutter up the transcript with prompts.

The top-level loop is required to trap all throws and recover gracefully. It is also required to print all values resulting from evaluation of a form, perhaps on separate lines. If a form returns zero values, as little as possible should be printed.

The following variables are maintained by the top-level loop as a limited safety net, in case the user forgets to save an interesting input expression or output value. (Note that the names of some these variables violate the convention that names of global variables begin and end with an asterisk.) These are intended primarily for user interaction, which is why they have short names. Use of these variables should be avoided in programs.

+	[Variable]
++	[Variable]
+++	[Variable]

While a form is being evaluated by the top-level loop, the variable + is bound to the previous form read by the loop. The variable ++ holds the previous value of + (that is, the form evaluated two interactions ago), and +++ holds the previous value of ++.

-	[Variable]
---	------------

While a form is being evaluated by the top-level loop, the variable - is bound to the form itself; that is, it is the value about to be given to + once this interaction is done.

*	[Variable]
**	[Variable]
***	[Variable]

While a form is being evaluated by the top-level loop, the variable * is bound to the result printed at the end of the last time through the loop; that is, it is the value produced by evaluating the form in +. If several values were produced, * contains the first value only (or nil if zero values were produced). The variable ** holds the previous value of * (that is, the result printed two interactions ago), and *** holds the previous value of **.

If the evaluation of + was aborted for some reason, * will have the value nil; this is so that + and *, ++ and **, and +++ and *** will be correspond properly.

/	[Variable]
//	[Variable]
///	[Variable]

While a form is being evaluated by the top-level loop, the variable / is bound to a list of the results printed at the end of the last time through the loop; that is, it is a list of all values produced by evaluating the form in +. The value of * should always be equal to the car of the value of /. The variable // holds the previous value of / (that is, the results printed two interactions ago), and /// holds the previous value of //.

If the evaluation of + was aborted for some reason, / will have the value nil; this is so that + and /, ++ and //, and +++ and /// will be correspond properly.

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 that 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 253) and `read-char` (page 255), take a stream argument that 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 258) and `write-char` (page 259), take a stream argument that 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 271).

terminal-io [Variable]

The value of ***terminal-io*** is ordinarily the stream that connects to the user's console.

trace-output [Variable]

The value of ***trace-output*** is the stream on which the **trace** (page 302) function prints its output.

standard-input, ***standard-output***, ***error-output***, ***trace-output***, and ***query-io*** are initially bound to synonym streams that pass all operations on to the stream that is the value of ***terminal-io***. (See **make-synonym-stream** (page 224).) 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 that 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 286) and **open** (page 283). 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 that 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.

`make-broadcast-stream &rest streams` [Function]

Returns a stream that 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 that 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 that 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-two-way-stream input-stream output-stream` [Function]

Returns a bidirectional stream that 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 that 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 that 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 that 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).

`with-input-from-string` (*var string* {*keyword value*}*) {*declaration*}* {*form*}* [Macro]

The body is executed as an implicit `progn` with the variable *var* bound to a character input stream that supplies successive characters from the value of the form *string*. `with-input-from-string` returns the results from the last *form* of the body.

The input stream is automatically closed on exit from the `with-input-from-string` form. It is best to regard the stream as having dynamic extent.

The following keyword options may be used:

- `:index` The form after the `:index` keyword should be a *place* acceptable to `setf`. If the `with-input-from-string` form is exited normally, then the *place* will have stored into it the index into the *string* indicating the first character not read (the length of the string if all characters were used). The *place* is not updated as reading progresses, but only at the end of the operation.
- `:start` The `:start` keyword takes an argument indicating, in the manner usual for sequence functions, the beginning of a substring of *string* to be used.
- `:end` The `:end` keyword takes an argument indicating, in the manner usual for sequence functions, the end of a substring of *string* to be used.

For example:

```
(with-input-from-string (s "Why a Duck?" :index j :start 6)
  (read s)) => duck?
```

As a side effect, the variable *j* is set to 10.

`with-output-to-string` (*var [string]*) {*declaration*}* {*form*}* [Macro]

The body is executed as an implicit `progn` with the variable *var* bound to a character output stream. All output to that stream is saved in a string. If no *string* argument is provided, then the value of `with-output-to-string` is a string containing all the collected output. If *string* is specified, it must be a string with a fill pointer, the output is incrementally appended to the string (see `vector-push` (page 199)); in this case `with-output-to-string` returns the results from the last *form* of the body.

The output stream is automatically closed on exit from the `with-output-to-string` form. It is best to regard the stream as having dynamic extent.

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.

`stream-element-type` *stream*

[Function]

A type specifier is returned to indicate what objects may be read from or written to the *stream*. Streams created by `open` (page 283) will have an element type restricted to a subset of character or integer, but in principle a stream may conduct transactions using any LISP objects.

`close` *stream* &key `:abort`

[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 the `:abort` parameter 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.

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 `pr int` (page 258) 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 56) 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 `pr int` 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 `pr int`, 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 89)). 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: What MACLISP calls a "single character object" (tokens of type *single*) are not provided for explicitly in COMMON LISP. They can be viewed as simply a kind of macro character. That is, the effect of `(setsyntax '$ 'single`

```

nil) in MACLISP can be achieved in COMMON LISP by
      (set-macro-character '$ #'(lambda (stream char)
                                (declare (ignore stream char))
                                '$))

```

The characters of the standard character set initially have the attributes shown in Table 22-1. Note that the square brackets, braces, question mark, and exclamation point (that is, “[”, “]”, “{”, “}”, “?”, and “!”) are normally defined to be constituents, but are not used for any purpose in standard COMMON LISP syntax and do not occur in the names of built-in COMMON LISP functions or variables. These characters are explicitly reserved to the user, primarily for use as macro characters if desired.

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 22-2, 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 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 252) 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 cases where there are two or more colons, or where a colon appears at the end of the token, presently do not mean anything in COMMON LISP and are reserved for future use; see chapter 11 (page 115). If there is a single non-final colon, it divides the token into two parts. The first part specifies a package. A null first part indicates the `keyword` package; otherwise it is interpreted as the name of a symbol in the current package, and that symbol must name a package. The second part is the name of the symbol.

If a symbol token contains no package markers, then the entire token is the name of the symbol. The symbol is looked up in the default package; see `*package*` (page 117).

<tab> <i>whitespace</i>	<page> <i>whitespace</i>	<return> <i>whitespace</i>
<space> <i>whitespace</i>	@ <i>constituent</i>	' <i>terminating macro character</i>
! <i>constituent</i>	A <i>constituent</i>	a <i>constituent</i>
" <i>terminating macro character</i>	B <i>constituent</i>	b <i>constituent</i>
# <i>terminating macro character</i>	C <i>constituent</i>	c <i>constituent</i>
\$ <i>constituent</i>	D <i>constituent</i>	d <i>constituent</i>
% <i>constituent</i>	E <i>constituent</i>	e <i>constituent</i>
& <i>constituent</i>	F <i>constituent</i>	f <i>constituent</i>
' <i>terminating macro character</i>	G <i>constituent</i>	g <i>constituent</i>
(<i>terminating macro character</i>	H <i>constituent</i>	h <i>constituent</i>
) <i>terminating macro character</i>	I <i>constituent</i>	i <i>constituent</i>
* <i>constituent</i>	J <i>constituent</i>	j <i>constituent</i>
+ <i>constituent</i>	K <i>constituent</i>	k <i>constituent</i>
, <i>terminating macro character</i>	L <i>constituent</i>	l <i>constituent</i>
- <i>constituent</i>	M <i>constituent</i>	m <i>constituent</i>
. <i>constituent</i>	N <i>constituent</i>	n <i>constituent</i>
/ <i>constituent</i>	O <i>constituent</i>	o <i>constituent</i>
0 <i>constituent</i>	P <i>constituent</i>	p <i>constituent</i>
1 <i>constituent</i>	Q <i>constituent</i>	q <i>constituent</i>
2 <i>constituent</i>	R <i>constituent</i>	r <i>constituent</i>
3 <i>constituent</i>	S <i>constituent</i>	s <i>constituent</i>
4 <i>constituent</i>	T <i>constituent</i>	t <i>constituent</i>
5 <i>constituent</i>	U <i>constituent</i>	u <i>constituent</i>
6 <i>constituent</i>	V <i>constituent</i>	v <i>constituent</i>
7 <i>constituent</i>	W <i>constituent</i>	w <i>constituent</i>
8 <i>constituent</i>	X <i>constituent</i>	x <i>constituent</i>
9 <i>constituent</i>	Y <i>constituent</i>	y <i>constituent</i>
: <i>constituent</i>	Z <i>constituent</i>	z <i>constituent</i>
; <i>terminating macro character</i>	[<i>constituent</i>	{ <i>constituent</i>
< <i>constituent</i>	\ <i>escape character</i>	<i>terminating macro character</i>
= <i>constituent</i>] <i>constituent</i>	} <i>constituent</i>
> <i>constituent</i>	^ <i>constituent</i>	~ <i>constituent</i>
? <i>constituent</i>	_ <i>constituent</i>	<rubout> <i>constituent</i>
<backspace> <i>constituent</i>	<linefeed> <i>whitespace</i>	

Table 21-1: Standard Character Syntax Attributes

```

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

The interpretation of standard characters within extended tokens is shown in Table 22-3. 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 22-3 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 253) 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:

<tab>	alphabetic *	{	alphabetic
<linefeed>	alphabetic *		alphabetic *
<page>	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

```
(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 simple 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.
- ',*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 x_j as above.
- `'(x1 x2 x3 ... xn . ,@form)` is an error.
- `'#(x1 x2 x3 ... xn)` may be interpreted to mean `(make-simple-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 178) instead of `append` (page 177), 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 (copy-list d))
```

(There is no good reason why `copy-list` 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.

The currently-defined sharp-sign constructs are described below and summarized in Table 22-4; 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.
page	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 #\

#<tab> signals error	#<page> signals error	#<return> signals error
#<space> signals error	#@ undefined	#' undefined
#! undefined	#A array	#a array
#" undefined	#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
##* bit-vector	#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- <i>n</i> rational	#r radix- <i>n</i> 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

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 154) 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 code that 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 simple 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) preceded by `"#*"` is read as a simple bit-vector containing those bits, the leftmost bit in the series being bit 0 of the bit-vector.

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 bits are specified, and if too few are specified the last one is used to fill all remaining elements of the bit-vector.

For example:

```
#*101111
#6*101111
#6*101
#6*1011
```

all mean the same thing: a vector of length 6 with elements a, b, and four instances of c.

#: *#:foo* requires *foo* to have the syntax of an unqualified symbol name (no embedded colons). It denotes an *uninterned* symbol whose name is *foo*. Every time this syntax is encountered a different uninterned symbol is created.

#. *#.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 49), 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 *#brational* reads *rational* in binary (radix 2). For example, *#B1101* \Leftrightarrow 13, and *#b101/11* \Leftrightarrow 5/3.

#O *#orational* reads *rational* in octal (radix 8). For example, *#o37/15* \Leftrightarrow 31/13, and *#o777* \Leftrightarrow 511.

#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). For example, *#xF00* \Leftrightarrow 3840.

#nR *#radixr**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.

#nA The syntax *#nA**form* constructs an *n*-dimensional array, using *form* as the value of the `:initial-contents` argument to `make-array` (page 193).

#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 201), 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.

#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.

The syntax *##*, where *n* is a required unsigned decimal integer, serves as a reference to some object labelled by *#n=*; that is, *##* 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 252) 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) ...)
```

A reference *##* may only occur after a label *#n=*; forward references are not permitted.

#+ 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 308).

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 `lisp` is true. Then the expressions on the left below are read the same as those on the right in implementation A:

```
(cons #+spice "Spice" #+lisp "Lisp" x)  (cons "Spice" x)
(setq a '(1 2 #+perq 43 #+(not perq) 27)) (setq a '(1 2 43))
(let ((a 3) #+(or spice lisp) (b 3))      (let ((a 3) (b 3))
  (foo a))                                (foo a))
```

In implementation B, however, they are read in this way:

```
(cons #+spice "Spice" #+lisp "Lisp" x)  (cons "Lisp" x)
(setq a '(1 2 #+perq 43 #+(not perq) 27)) (setq a '(1 2 27))
(let ((a 3) #+(or spice lisp) (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.

#- `#-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>**, **#<page>**

A `#` followed by a standard whitespace character is not legal reader syntax. This is so that abbreviated forms produced via `*prinlevel*` (page 252) 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 252) 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 variable **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.

readtablep *object*

[Function]

readtablep is true if its argument is a readtable, and otherwise is false.

```
(readtablep x) <=> (typep x 'readtable)
```

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 245)), 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 254), 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 char)
  (declare (ignore char))
  (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 char)
  (declare (ignore char))
  ;; First swallow the rest of the current input line.
  (do () ((char= (read-char stream) #\Return)))
  ;; Return zero values.
  (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 236), 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 subchar arg)
  (declare (ignore subchar arg))
  (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; `*prinescape*` is one of the most important.

`*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 258) 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 258) effectively binds `*prinescape*` to `t`.

Compatibility note: This flag controls 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 calls this variable *base*, and its default value is 8, not 10.

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. To prevent confusion of the letter "O" and the digit "0", and of the letter "B" with the digit "b", the radix specifier is always printed using lower-case letters. 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**; they are obtained as if by the *numerator* (page 134) and *denominator* (page 134) functions, and so ratios are always printed in lowest form.

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 253), however, then the exponent marker "E" is used.

Characters. When **prinescape** (page 238) 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 238) 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*.

The rules for package qualifiers are as follows. When the symbol is printed, if it is in the keyword package

then it is printed with a preceding colon; otherwise, if it is present in the current package, it is printed without any qualification; otherwise, it is printed with qualification. See `*package*` (page 117).

Implementation note: The syntax `"foo:|string|"` presently has no defined meaning in COMMON LISP. If a package qualifier must be printed, then vertical-bar syntax may not be used for either the name of the package or the name of the symbol; instead, individual escape characters must be used. This means that one must print `"si:\(\ \. \ \)"` instead of `"si:|(.)|"`, for example. This unpleasant form of output should not occur very often in practice.

A symbol that is uninterned (has no home package) is printed preceded by `"#:"`.

Implementation note: Because the `"#:"` syntax does not intern the following symbol, it is necessary to use circular-list syntax if `*princircle*` (page 238) is not `nil` and the same uninterned symbol appears several times in an expression to be printed. For example, the result of

```
(let ((x (make-symbol "FOO"))) (list x x))
```

would be printed as `"(#:foo #:foo)"` if `*princircle*` were `nil`, but as `"(#1=#:foo #1#)"` if `*princircle*` were not `nil`.

The case in which symbols are printed is controlled by the variable `*princase*` (page 250).

`*princase*`

[Variable]

The `read` (page 253) function normally converts lower-case letters appearing in symbols to upper case, so that internally print names normally contain only upper-case characters. However, users may prefer to see output in lower case or mixed case. This variable controls the case (upper or lower) in which to print any upper-case characters in the names of symbols when vertical-bar syntax is not used. The value of `*princase*` should be one of the keywords `:upcase`, `:downcase`, or `:capitalize`.

Lower-case characters in the internal print name are always printed in lower case, and are preceded by an escape character. Upper-case characters in the internal print name are printed in upper case, lower case, or in mixed case so as to capitalize words, according to the value of `*princase*`. The convention for what constitutes a "word" is the same as for the function `string-capitalize` (page 196).

Strings. The characters of the string are output in order. If `*prinscape*` (page 238) is not `nil`, a double quote `"` is output beforehand and afterward, and all and double quotes and escape characters are preceded by `\`. The printing of strings is not affected by `*prinarray*` (page 252).

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 it 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 252) and `*prinlength*` (page 252).

Bit-vectors. A bit-vector is printed as “#*” followed by the bits of the bit-vector in order. If `*prinarray*` (page 252) is `nil`, however, then the bit-vector is printed in a format (using “#<”) that is concise but not readable.

Vectors. Any vector other than a string or bit-vector is printed using general-vector syntax; this means that information about specialized vector representations will be lost. 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 252) and `*prinlength*` (page 252). If `*prinarray*` (page 252) is `nil`, however, then the vector is printed in a format (using “#<”) that is concise but not readable.

Arrays. Normally any array other than a vector is printed using “#nA” format. Let n be the rank of the array. Then “#” is printed, then n as a decimal integer, then “A”, then n left parentheses. Next the elements are scanned in row-major order. Imagine the array indices being enumerated in odometer fashion. Every time the index for dimension j is incremented, first a right parenthesis is printed; then if dimension j has overflowed, dimension $j-1$ is incremented; then a left parenthesis is printed. Finally, n right parentheses are printed. This causes the contents to be printed in a format suitable for the `:initial-contents` argument to `make-array` (page 193). The lists effectively printed by this procedure are subject to `*prinlevel*` (page 252) and `*prinlength*` (page 252). If `*prinarray*` (page 252) is `nil`, however, then the array is printed in a format (using “#<”) that is concise but not readable.

Structures defined by `defstruct` (page 201) 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`.

<i>v</i>	<i>n</i>	<i>Output</i>
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 ...)))

Another way to cut down on the volume of printing is to disable the printing of array contents.

`*prinarray*` [Variable]

If `prinarray` is `nil`, then the contents of arrays other than strings are never printed. Instead, arrays are printed in a concise form using "#<" that gives enough information for the user to be able to identify the array, but does not include the entire array contents. If `prinarray` is not `nil`, non-string arrays are printed using "#(", "#*", or "#nA" syntax.

21.2. Input Functions

21.2.1. Input from ASCII Streams

Many input functions take optional arguments called *input-stream*, *eof-errorp*, 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 213). One may also specify *t* as a stream, meaning the value of the special variable **terminal-io** (page 214).

Rationale: Allowing the use of *t* provides some semblance of MACLISP compatibility.

The *eof-errorp* 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 *eof-errorp* is true (the default), an error will be signalled at end of file. If it is false, then no error is signalled, and instead the function returns *eof-value*.

Functions such as *read* (page 253) that read an “object” rather than a single character will always signal an error, regardless of *eof-errorp*, 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 act according to *eof-errorp*. Similarly, the function *read-line* (page 255) 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.

read &optional *input-stream eof-errorp 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 a type specifier symbol for a specific floating-point format; these include *short-float*, *single-float*, *double-float*, *long-float*, and may include implementation-specific types as well. The default value is *single-float*.

read-default-float-format 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-errorp 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 `unread-char` (page 255)) 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 char)
  (declare (ignore char))
  (do ((path (list (read-preserving-whitespace stream))
                  (cons (progn (read-char stream)
                            (read-preserving-whitespace
                              stream))
                        path))))
      ((not (char= (peek-char 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 `peek-char` 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 readable.) 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 char arg)
  (declare (ignore char arg))
  (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`.

`read-line` &optional *input-stream eof-errorp eof-value* [Function]
`read-line` 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. See `write-line` (page 259).

`read-char` &optional *input-stream eof-errorp eof-value* [Function]
`read-char` inputs one character from *input-stream* and returns it as a character object.

`unread-char` *character* &optional *input-stream* [Function]
`unread-char` 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*. `unread-char` returns `nil`.

One may only apply `unread-char` to the character most recently read from *input-stream*; moreover, one may not invoke `unread-char` twice consecutively without an intervening `read-char` 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 `unread-char` 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.

`peek-char` &optional *peek-type input-stream eof-errorp eof-value* [Function]
 What `peek-char` does depends on the *peek-type*, which defaults to `nil`. With a *peek-type* of `nil`, `peek-char` 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 `read-char` and then `unread-char` in succession.

If *peek-type* is `t`, then `peek-char` 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 `peek-char` skips over input characters until a character that is `char=` (page 152) to that object is found; that character is left in the input stream.

Characters passed over by `peek-char` are echoed if *input-stream* is interactive.

`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 `read-char` (page 245) 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.

`read-char-no-hang` &optional *input-stream eof-errorp eof-value* [Function]

This function is exactly like `read-char` (page 245), 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 256) 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-p eof-errorp 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-p*, if provided and not `nil`, indicates that the operation should preserve

whitespace as for `read-preserving-whitespace` (page 243).

The arguments *eof-errorp* and *eof-value* control the action if the end of the (sub)string is reached before the operation is completed, as with other reading functions; reaching the end of the string is treated as any other end-of-file event.

`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-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

`read-byte` *binary-input-stream* &optional *eof-errorp eof-value* [Function]
`read-byte` reads one byte from the *binary-input-stream* and returns it in the form of an integer.

`read-binary-object` *type binary-input-stream* &optional *eof-errorp eof-value* [Function]
`read-binary-object` reads an object of the specified *type* from the *binary-input-stream*. The object is assumed to be encoded in the manner used by `write-binary-object` (page 260); the object is guaranteed to be read properly only if the exact same *type* is specified to `read-binary-object` as was specified to `write-binary-object` to originally encode the object, and if the *:type* (page 283) option for the input stream matches that for the output stream given to `write-binary-object`.

The *eof-errorp* and *eof-value* options apply only if the *binary-input-stream* is at the end of file before the operation is begun. If the *type* requires more than one byte to be read and end-of-file is encountered before enough bytes have been read, an error is signalled.

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 213). If it is *t*, the value of the variable **terminal-io** (page 214) is used.

```
write object &key :stream :prinescape :prinradix :base [Function]
      :princircle :prinpretty :prinlevel :prinlength
      :princase :prinarray
```

The printed representation of *object* is written to the output stream specified by *:stream*, which defaults to the value of **standard-output** (page 213).

The other keyword arguments specify values used to control the generation of the printed representation. Each defaults to the global variable of the same name; see **prinescape** (page 238), **prinradix** (page 239), **base** (page 239), **princircle** (page 238), **prinpretty** (page 238), **prinlevel** (page 242), **prinlength** (page 242), **princase** (page 240), and **prinarray** (page 242). (This is the means by which these variables affect printing operations: supplying default values for the *write* function.) Note that the printing of symbols is also affected by the value of the variable **package** (page 117).

```
prin1 object &optional output-stream [Function]
print object &optional output-stream [Function]
pprint 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 243). *prin1* returns *object*.

```
(prin1 object output-stream)
<=> (write object :stream output-stream :prinescape t)
```

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*.

pprint is just like *print* except that the trailing space is omitted, and the *object* is printed with the **prinpretty** (page 238) flag non-*nil* to produce “pretty” output. *pprint* 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 243). `princ` returns *object*.

```
(prin1 object output-stream)
<=> (write object :stream output-stream :prinscape nil)
```

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.

```
write-to-string object &key :prinscape :prinradix :base [Function]
                        :princircle :prinpretty :prinlevel :prinlength
                        :princase :prinarray
```

```
prin1-to-string object [Function]
```

```
princ-to-string object [Function]
```

The object is effectively printed, as if by `write` (page 248), `prin1` (page 248), or `princ` (page 248), and the characters that would be output are made into a string and returned.

```
write-char character &optional output-stream [Function]
```

`write-char` outputs the *character* to *output-stream*, and returns `nil`.

```
write-string string &optional output-stream [Function]
```

```
write-line string &optional output-stream [Function]
```

`write-string` writes the characters of the *string* to the *output-stream*. `write-line` does the same thing, but then outputs a newline afterwards. (See `read-line` (page 245).) In some implementations these may be significantly more efficient than an explicit loop using `write-char`.

```
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.

`finish-output` &optional *output-stream* [Function]
`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 `finish-output` attempts to ensure that all output sent to *output-stream* has reached its destination, and only then returns `nil`. `force-output` initiates the emptying of any internal buffers, but returns `nil` without waiting for completion or acknowledgement.

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 precise actions of all three of these operations are implementation-dependent.

The function `format` (page 261) is very useful for producing nicely formatted text, producing good-looking messages, and so on. `format` can generate a string or output to a stream.

21.3.2. Output to Binary Streams

`write-byte` *integer binary-output-stream* [Function]
`write-byte` writes one byte, the value of *integer*. It is an error if *integer* is not of the type specified as the `:type` argument to `open` (page 283) when the stream was created.

`write-binary-object` *object type binary-output-stream* [Function]
The *object* is encoded as a stream of bytes and written to the *binary-output-stream*. The *object* must be of the type specified by *type*. The encoding used may depend on the `:type` (page 283) of the stream and on the specified *type*. For example, the integer 126 may be encoded in different ways depending on whether the *type* specified is `integer` or `(byte 8)`.

The *type* specified must be one of the following types or a subtype of one: `number`, `character`, or `(array x)` where *x* is a subtype of `integer` or `character`.

The encoding is implementation-dependent. However, the function `read-binary-object` (page 247) may be used in the same implementation to read back an object encoded by `write-binary-object`. (These functions are intended to provide efficient storage of data in an implementation-dependent format.)

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 213).

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 248)). 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 "()".

`~mincolA` 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, padcharA` 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 248) rather than `princ`). The output is therefore suitable for input to `read` (page 243). `~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.

`~mincolD` 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, padcharD` 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, commacharD`.

`~B` *Binary*. This is just like `~D` but prints in binary radix (radix 2) instead of decimal. The full form is therefore `~mincol, padchar, commacharB`.

`~O` *Octal*. This is just like `~D` but prints in octal radix (radix 8) instead of decimal. The full form is therefore `~mincol, padchar, commacharO`.

`~X` *Hexadecimal*. This is just like `~D` but prints in hexadecimal radix (radix 16) instead of decimal. The full form is therefore `~mincol, padchar, commacharX`.

`~R` *Radix*. `~nR` 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, commacharR`.

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 eq1 to the integer 1, a lower-case "s" is printed; if *arg* is eq1 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. $\sim nF$ 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 $\sim A$ 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"
```

Compatibility note: This is not the same as FORTRAN "F" format.

??? Query: Sigh. If I had my druthers, $\sim E$, $\sim F$, and $\sim G$ would be the same as FORTRAN E, F, and G formats; they are widely known and understood.

$\sim E$ *Exponential.* *arg* is printed in exponential notation. This is identical to $\sim F$, 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.

Compatibility note: This is not the same as FORTRAN "E" format.

$\sim C$ *Character.* The next *arg* should be a character; it is printed according to the modifier flags.

$\sim C$ prints the character in an implementation-dependent abbreviated format. This format should be culturally compatible with the host environment. $\text{fi } \sim :C$ 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.

$\sim :@C$ prints what $\sim :C$ 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.

$\sim @C$ prints the character in a way that the LISP reader can understand, using "#\" syntax.

Rationale: In some implementations the $\sim S$ directive would accomplish this also, but the $\sim C$ directive is compatible with LISP dialects that do not have a character data type.

$\sim \%$ Outputs a newline (see *terpri* (page 249)). $\sim n\%$ outputs *n* newlines. No *arg* is used. Simply putting a newline in the control string would work, but $\sim \%$ is often used because it makes the control string look nicer in the middle of a LISP program.

$\sim \&$ Unless the stream knows that it is already at the beginning of a line, this outputs a newline (see *fresh-line* (page 249)). $\sim n\&$ does a *fresh-line* operation and then outputs *n-1* newlines.

$\sim |$ Outputs a page separator character, if possible. $\sim n|$ does this *n* times. | is vertical bar, not capital I.

$\sim \sim$ *Tilde.* Outputs a tilde. $\sim n\sim$ 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, colincT` 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 that 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`.

`~G` *Goto*. Goes to the *n*th *arg*, where 0 means the first one; *n* defaults to 0, so `~G` goes back to the first *arg*. Directives after a `~nG` 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 `~*`.

`~?` *Indirection.* The next *arg* must be a string; it is processed as part of the control string as if it had appeared in place of the `~?` construct.

The format directives after this point are much more complicated than the foregoing; they constitute "control structures" that can perform case conversion, conditional selection, iteration, justification, and non-local exits. Used with restraint, they can perform powerful tasks. Used with wild abandon, they can produce completely unreadable and unmaintainable code.

`~(str~)` *Conditional Case conversion.* The contained control string *str* is processed, and what it produces is subject to case conversion. With no flags, all case-modifiable characters are forced to lower case. `~;:(` capitalizes all words, as if by `string-capitalize` (page 196). `~;@` capitalizes just the first word, and forces the rest to lower case. `~;:@` forces all case-modifiable characters to upper case.

For example:

```
(format nil "~@R ~(~@R~)" 14 14) => "XIV xiv"
(defun f (n) (format nil "~@(~R~) error~:P detected." n))
(f 0) => "Zero errors detected."
(f 1) => "One error detected."
(f 23) => "Twenty-three errors detected."
```

`~[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~] 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~;;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 that follows it. That clause is processed that 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 that 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. This case may be handled more clearly by the `~?` directive, but this general feature of `~{` is more powerful than `~?`.

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  "
```

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 ~^ 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 that 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*` (page 214), which normally is synonymous with `*terminal-io*` (page 214) but can be rebound to another stream for special applications.

`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 249)). 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 214).

For example:

```
(y-or-n-p "Cannot establish connection.  Retry? ")
```

`y-or-n-p` should only be used for questions that 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

249)). 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 214).

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.

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:

<u>System</u>	<u>File name</u>
TOPS-20	<LISP10>FORMAT.FASL.13
TOPS-10	FORMAT.FAS[1,4]
ITS	LISP10;FORMAT FASL
MULTICS	>udd>Lisp10>format.fasl
TENEX	<LISP10>FORMAT.FASL;13
VAX VMS	[LISP10]FORMAT.FAS;13
UNIX	/usr/lisp10/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 that 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 this is 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 283).

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 be used internally for various purposes; in particular, parsing a namestring that does not specify certain components will result in a pathname with missing components.

A component of a pathname can also be the keyword `: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 that just needs to be remembered and copied from place to place.

The type is always a string or `nil` 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` 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]

The *pathname* function converts its argument to be a pathname. The argument may be a pathname, a string or symbol, or a stream.

truename thing [Function]

The *truename* function converts *thing* to be a pathname, and then endeavors to discover the "true name" of the file associated with that pathname within the file system. 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 282)).

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"
```

??? Query: If the file is not found, should truename signal an error, return nil, or just quietly return an untranslated pathname?

`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 of pathname components; 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 *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 *convention* is nil then the host name is extracted from the default pathname in *defaults* and used to determine the syntax convention. The *defaults* argument defaults to the value of `*default-pathname-defaults*` (page 281).

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 *convention* is specified and not nil, and *thing* contains a manifest host name, an error is signalled if the conventions do not match.

`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 281). *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 23.1.3 (page 280). 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 that 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 281). 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 `symbol-plist` (page 109)).

`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 that 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 `true-name` (page 266)).

`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 that 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 281)). That is,

```
(merge-pathname-defaults (enough-namestring pathname defaults)
                          defaults)
=> (parse-pathname pathname)
```

`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 of 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. Programs that have a default type for the files they manipulate usually 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.

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.

See also ***load-pathname-defaults*** (page 289).

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.

`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 282)) is *actual-pathname*.

An example of how this would be used is in connection with `truename`s. 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 string always returns a pathname whose device component is `nil`.

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 217) may be used to close a file; the functions described below may be used to open them. The basic operation is `open`, but `with-open-file` is usually more convenient for most applications.

`open filename &key :direction :type :if-exists :if-does-not-exist` [Function]

Returns a stream that is connected to the file specified by *filename*. The keyword arguments specify what kind of stream to produce and how to handle errors:

:direction This argument specifies whether the stream should handle input, output, or both.

:input The result will be an input stream. This is the default.

:output The result will be an output stream.

:io The result will be a bidirectional stream.

:probe The result will be a no-directional stream (in effect, the stream is created and then closed). This is useful for determining whether a file exists without actually setting up a complete stream.

:type This argument specifies the type of the unit of transaction for the stream. As a rule, anything that can be recognized as being a finite subtype of `character` or `integer` is acceptable. In particular, the following types are recognized:

string-char The unit of transaction is a string-character. The functions `read-char` (page 245) and/or `write-char` (page 249) may be used on the stream. This is the default.

character The unit of transaction is any character, not just a string-character. The functions `read-char` (page 245) and/or `write-char` (page 249) may be used on the stream.

standard-char

The unit of transaction is a standard character. The functions `read-char` (page 245) and/or `write-char` (page 249) may be used on the stream. This option may be used to guarantee that no non-standard character will be read from an input source.

(unsigned-byte *n*)

The unit of transaction is an unsigned byte (a non-negative integer) of size *n*. The functions `read-byte` (page 247) and/or `write-byte` (page 250) may be used on the stream.

- unsigned-byte** The unit of transaction is an unsigned byte (a non-negative integer); the size of the byte is determined by the file system. The functions `read-byte` (page 247) and/or `write-byte` (page 250) may be used on the stream.
- (signed-byte *n*)** The unit of transaction is a signed byte of size *n*. The functions `read-byte` (page 247) and/or `write-byte` (page 250) may be used on the stream.
- signed-byte** The unit of transaction is a signed byte of size *n*. the size of the byte is determined by the file system. The functions `read-byte` (page 247) and/or `write-byte` (page 250) may be used on the stream.
- bit** The unit of transaction is a bit (values 0 and 1). The functions `read-byte` (page 247) and/or `write-byte` (page 250) may be used on the stream.
- (mod *n*)** The unit of transaction is a non-negative integer less than *n*. The functions `read-byte` (page 247) and/or `write-byte` (page 250) may be used on the stream.
- :default** The unit of transaction is to be determined by the file system, based on the file it finds. The type can be determined by using the function `stream-element-type` (page 217).
- :if-exists** This argument specifies the action to be taken if the `:direction` is `:output` or `:io` and a file of the specified name already exists. If the direction is `:input` or `:probe`, this argument is ignored.
- :error** Signal an error. This is the default when the version component of the filename is not `:newest`.
- :new-version** Create a new file with the same file name, but with a larger version number. This is the default when the version component of the filename is `:newest`.
- :rename** Rename the existing file to some other name, and then create a new file with the specified name.
- :rename-and-delete** Rename the existing file to some other name and then delete it (but don't expunge it, on those systems that distinguish deletion from expunging). Then create a new file with the specified name.

<code>:overwrite</code>	The existing file is used, and output operations on the stream will destructively modify the file. If the <code>:direction</code> is <code>:io</code> , the file is opened in a bidirectional mode that allows both reading and writing. The file pointer is initially positioned at the beginning of the file. This mode is most useful when the <code>file-position</code> (page 287) function can be used on the stream.
<code>:append</code>	The existing file is used, and output operations on the stream will destructively modify the file. The file pointer is initially positioned at the end of the file. If the <code>:direction</code> is <code>:io</code> , the file is opened in a bidirectional mode that allows both reading and writing.
<code>:supersede</code>	Supersede the existing file. If possible, the implementation should arrange not to destroy the old file until the new stream is closed, against the possibility that the stream will be closed in "abort" mode.
<code>nil</code>	Do not create a file or even a stream. Instead, simply return <code>nil</code> to indicate failure.
<code>:if-does-not-exist</code>	This argument specifies the action to be taken if a file of the specified name does not already exist.
<code>:error</code>	Signal an error. This is the default if the <code>:direction</code> is <code>:input</code> , or if the <code>:if-exists</code> argument is <code>:overwrite</code> or <code>:append</code> .
<code>:create</code>	Create an empty file with the specified name, and then proceed as if it had already existed. This is the default if the <code>:direction</code> is <code>:output</code> or <code>:io</code> , and the <code>:if-exists</code> argument is anything but <code>:overwrite</code> or <code>:append</code> .
<code>nil</code>	Do not create a file or even a stream. Instead, simply return <code>nil</code> to indicate failure. This is the default if the <code>:direction</code> is <code>:probe</code> .

When the caller is finished with the stream, it should close the file by using the `close` (page 217) function. The `with-open-file` (page 286) 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 that uses `open` directly should use the special form `unwind-protect` (page 94) to close the file if an abnormal exit occurs.

`with-open-file` *bindspec* {*form*}* [Macro]
 (`with-open-file` (*stream filename . options*) . *body*) evaluates the forms of *body* (an implicit progn) with the variable *stream* bound to a stream that reads or writes the file named by the value of *filename*. The *options* are evaluated, and are used as keyword arguments to the function `open` (page 273).

When control leaves the body, either normally or abnormally (such as by use of `throw` (page 95)), the file is automatically closed. If a new output file is being written, and control leaves abnormally, the file is aborted and the file system is left, so far as possible, as if the file 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 may be a string, a pathname, or a stream.

For example:

```
(with-open-file (ifile name :direction :input)
  (with-open-file (ofile (merge-pathname-defaults ifile
                                                    nil
                                                    "out"))
    :direction :output
    :if-exists :supersede)
  (transduce-file ifile ofile)))
```

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 66).

`rename-file` *file new-name* &optional *error-p* [Function]
file can be a filename or a stream that 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 that 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 that 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 266)).

`file-creation-date` *file* [Function]

file can be a filename or a stream that 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 that 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.

`file-position` *file-stream* &optional *position* [Function]

`file-position` returns or sets the current position within a random-access file.

`(file-position 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 stream, the position is in units of characters; for a binary file, the position is in bytes.

`(file-position 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 287) function returns the length beyond which `file-position` may not access). With two arguments, `file-position` 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 that 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 stream, the position is in units of characters; for a binary stream, 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 47), `defmacro` (page 99), and `defvar` (page 48), 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 that can be loaded more quickly.

`load` &optional *filename* &key `:verbose` `:print` `:if-does-not-exist` [Function]
`:set-default-pathname`

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.

If the first argument is a stream rather than a pathname, then `load` determines what kind of stream it is and loads directly from the stream.

The `:verbose` argument (which defaults to the value of `*load-verbose*` (page 288)), if true, permits `load` to print a message in the form of a comment to `*standard-output*` (page 213) indicating what file is being loaded and other useful information.

The `:print` argument (default `nil`), if true, causes the value of each expression loaded to be printed to `*standard-output*` (page 213). If a binary file is being loaded, then what is printed may not reflect precisely the contents of the source file, but nevertheless some information will be printed, including the name of each function loaded.

If a file is successfully loaded, `load` always returns a non-`nil` value. If `:if-does-not-exist` is specified and is `nil`, `load` just returns `nil` rather than signalling an error if the file does not exist.

`load` maintains a default filename in the variable `*load-pathname-defaults*` (page 289), used to default missing components of the *filename* argument; thus (`load`) will load the same file previously loaded. (The function `compile-file` (page 300) also uses and sets these pathname defaults.) The `:set-pathname-defaults` argument (which defaults to the value of `*load-set-pathname-defaults*`), if true, causes `load` to update `load-pathname-defaults` from its first argument.

`*load-verbose*` [Variable]

This variable provides the default for the `:verbose` argument to `load` (page 288). Its initial value is implementation-dependent.

`*load-set-default-pathname*` [Variable]

This variable provides the default for the `:set-default-pathname` argument to `load` (page 288). Its initial value is implementation-dependent.

See also `*compile-file-set-default-pathname*` (page 300).

load-pathname-defaults**[Variable]**

This is the `pathname-defaults` pathname for the `load` (page 278) and `compile-file` (page 300) functions. Other functions may share these defaults if they deem that to be an appropriate user interface.

22.5. Accessing Directories**directory *pathname*****[Function]**

A list of pathnames is returned, one for each file in the file system that matches the given *pathname*. For each such file, the `true-name` (page 266) for that file appears in the result list. Keywords such as `:wild` and `:newest` may be used in `:pathname` to indicate the search space.

Chapter 23

Errors

COMMON LISP handles errors through a system of *conditions*. One may establish *handlers* that gain control when 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 may deal with it.

The condition mechanism is completely general and can be used for purposes other than “error” handling.

Every condition is named by a symbol, typically a keyword. When an unusual situation occurs, such as an error, a condition is signalled. Handlers are established with dynamic scope, and so the most recently established handler for the condition will be invoked.

23.1. Signalling Conditions

`signal` *condition-name* &rest *args* [Function]

This searches through all currently established condition handlers, starting with the most recent. If it finds one established to handle *condition-name* or to handle any condition, then it calls that handler, giving it precisely the arguments that were given to `signal`, including the *condition-name* as the first argument. The dynamic environment (such as catchers and special variable bindings) is *not* unwound; the handler is invoked in the dynamic environment of the call to `signal`.

If a handler returns values, and the first value returned by the handler is not `nil`, the handler is said to be *willing*; all the values it returns are returned from the call to `signal`. Otherwise, `signal` will continue searching for another matching handler. If no matching and willing handler is found, then `signal` returns `nil`.

It is possible for a handler to effectively handle the error other than by returning values; it may, for example, call `throw` (page 95).

23.2. Establishing Handlers

`condition-bind` *bindings* {*form*}* [Macro]

This is used to establish handlers for conditions, then perform the *body* in that established handler environment. The handlers established have dynamic scope. The format is:

```
(condition-bind ((condition-name-1 handler-1)
                (condition-name-2 handler-2)
                ...
                (condition-name-m handler-m))
  form1
  form2
  ...
  formn)
```

Each *condition-name-j* is either the name of a condition or a list of names of conditions. Each *handler-j* is a form that is evaluated to produce a handler function; they are evaluated in order from *handler-1* to *handler-m*; only after all the *handler-j* forms are evaluated is any handler established. The condition handlers are established in the order shown, such that if a condition is signalled *handler-m*, as the most recently established, will be the first one examined.

??? Query: This differs from Lisp Machine LISP. However, if shadowing within a single `condition-bind` is to be permitted, this is the more logical definition. Perhaps it would be better not to allow such shadowing?

The expressions *formj* are then evaluated as an implicit progn. The `condition-bind` form returns whatever *formn* returns (`nil` if there are no forms in the body). The established conditions become disestablished when the `condition-bind` form is exited.

`condition-bind` also establishes a context limiting the extent of effectiveness of `condition-psetq`.

`condition-psetq` {*spec*}* [Macro]

The `condition-psetq` form is used to establish condition handlers as a side effect. It takes the form:

For example:

```
(condition-psetq condition-name-1 handler-1
                condition-name-2 handler-2
                ...
                condition-name-n handler-n)
```

Each *condition-name-j* is either the name of a condition or a list of names of conditions. Each *handler-j* is a form that is evaluated to produce a handler function. The *handler-j* forms are all evaluated in order, and only then are the results established as condition handlers, in such an order that *handler-1* is examined first when a condition is signalled.

The conditions established by `condition-psetq` remain established until execution is unwound (either normally or by being thrown) past the most recent `condition-bind`. (If no `condition-bind` is in effect, `condition-psetq` effectively establishes globally defined handlers. Multiple uses of `condition-psetq` cause the most recently established handler to be

tried first when a condition is signalled. For example, consider:

For example:

```
(condition-psetq :wrong-type-argument 'default-wta-handler)
(+ 23 nil)
(condition-psetq :wrong-type-argument 'hairly-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 `hairly-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.

??? Query: Need to have a way to disestablish a handler established by `condition-psetq`?

`condition-case form` `{{(condition-names {form}*)}`* [Macro]

This is a form sometimes more convenient than `condition-bind` (page 282) for executing a form with certain condition handlers established. The handlers established have dynamic scope. The format is:

```
(condition-case form
  clause1
  clause2
  ...
  clausen)
```

Each *clause* is similar to a `case` (page 78) clause; it is a list whose first element specifies keys, in this case a list of condition names or a single condition name. The remainder of each clause is a body, a list of forms constituting an implicit `progn`. For each condition mentioned a handler is established, such that if the condition occurs a `throw` (page 95) is performed to unwind the dynamic environment back to the point of the `condition-case`; the body of the corresponding clause is then executed, and whatever is produced by the last form in the body is returned as the value of the `condition-case` form. Note that when the body of any `condition-case` clause is executed, all the handlers established by that `condition-case` have already been disestablished.

Once these handlers are established, the *form* is evaluated. If evaluation of the *form* does not cause a condition to invoke one of the handlers, the established handlers become disestablished, and whatever the *form* produced is returned from the `condition-case` form.

23.3. Error Handlers

Certain conditions names are used by the COMMON LISP system to signal error conditions. Like all condition handlers, an error handler will receive as arguments all the arguments given to `signal`. By convention, however, the arguments for a signalled error have the following interpretation:

- *condition-name*. As for any condition handler, the first argument is the name of the condition.

- *proceed-flag*. If this is not `nil`, then the handler can expect to correct the error by returning suitable values (see below); the signaller will be prepared to retry the failed operation or otherwise recover. If this is `nil`, then if the handler returns a value (other than `nil` to decline handling of the condition) a `:failed-handler` error will be signalled.

- *function-name*. If this is not `nil`, it is the name of the function that signalled the error.

??? Query: Here is an odd idea: let every `defun` implicitly bind a lexical variable named, say, `name-of-this-function` to the name of the function, in much the same way that it establishes an implicit block for use by `return`. Then `ferror` could be a macro such that `(ferror ...)` expanded into

```
(*ferror name-of-this-function ...)
```

and thereby capture the function name automatically.

- *control-string*. A string suitable for use with the remaining arguments for `format` (page 251).
- *other-arguments*. Other arguments; these vary with the condition involved.

An error handler can do some processing and then make one of three responses to the error. It can return `nil` to decline handling the error, in which case some other handler will be given the opportunity. It can call `throw` or `signal` in an attempt to make a non-local exit. Finally, it can return several values of which the first is `:return`, in an attempt to correct the error. If the handler returns values and the first value is not `nil` or `:return`, or if the first value is `:return` and the signaller is not prepared to correct the error, then (by convention) a `:failed-handler` error will be signalled.

The function `error` (page 295) is the primary means for signalling a correctable error. If a handler returns several values of which the first is `:return`, then `error` will return all the values except the first (the `:return` keyword) as the values from the call to `error`.

23.4. Signalling Errors

LISP programs can signal errors by using one of the functions `error` (for “trivial errors”), `ferror` (for fatal error) or `cerror` (for correctable error). `error` is the easiest way to signal an error, but it provides for neither distinguishing types of errors nor recovering from errors. `ferror` distinguishes among various types of errors, and `cerror` further allows recovery from the error. High-quality software packages should endeavor to use `cerror` or `ferror` whenever appropriate.

`error` *control-string* &rest *args*

[Function]

`error` takes a *control-string* and other arguments suitable for `format` (page 251). It signals an error using the condition name `:error`.

```
(error s x y ...) <=> (ferror :error s x y ...)
```

ferror *condition-name control-string &rest args* [Function]

cerror *condition-name control-string &rest args* [Function]

ferror signals the error condition *condition-name*. The remaining arguments to **ferror** should be suitable for **format** (page 251). The error condition signalled is not correctable; function **ferror** never returns. (A call to **ferror** may be terminated by a non-local exit, such as a **throw**, however.)

If no handler can be found for *condition-name*, then a **:missing-handler** error is recursively signalled.

If a handler attempts to correct the error (by returning as first value the keyword **:return**), then **ferror** signals an uncorrectable **:failed-handler** error.

cerror is similar to **ferror**, but signals a correctable error (the second argument given to **signal** is **t**). If a handler attempts to correct the error (by returning as first value the keyword **:return**), then **cerror** returns as its values all values returned by the handler but the first (that is, the keyword **:return** is not returned by **cerror**).

assert *test &optional control-string &rest args* [Macro]

The *test* is evaluated. If the result is not **nil**, **assert** returns **nil**.

If the result is **nil**, then a correctable **:failed-assertion** error is signalled; see **cerror** (page 295). If a *control-string* is supplied, it and the *args* are used in signalling the error; otherwise a default message is provided. The *control-string* and *args* are not evaluated unless the *test* is **nil**. If the handler “corrects” the error, then the *test* is re-evaluated; thus **assert** iterates until the *test* is satisfied. (If the *test* is evaluated several times, then *control-string* and the *args* may also be evaluated multiple times, so it is best if they are free of side effects.)

For example:

```
(assert (<= 2 base 36))
(assert (apply #'= (array-dimensions a))
        "The array ~S is not equidimensional." a)
```

check-type *place typespec &optional string* [Macro]

The *place* must be a form acceptable to **setf** (page 66). If the value of *place* is of the type specified by *typespec*, then **check-type** simply returns **t**. Otherwise, it signals a **:wrong-type-argument** correctable error. The message will mention the *string*, which should be an English name or phrase for the type; if *string* is omitted (as it frequently is), a name is derived automatically from the *typespec*. If a correction value is returned, then **setf** is used to install the correction value in *place*, and the test is then repeated. Thus **check-type** will not terminate until *place* contains a value of the specified type.

The *typespec* is not evaluated, but the *string* (if supplied) is.

For example:

```
(check-type *readtable* readtable)

(defun primep-the-hard-way (x)
  (check-type x integer)
  (do ((j 1 (+ j 1)))
      ((> j x) t)
    (when (> (gcd j x) 1) (return nil))))
```

23.5. 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. The arguments listed are those that are to follow the *control-string* as arguments to the handler.

If an error is signalled correctly, the term “correction values” refers to the values returned by the handler along with the keyword `:return`. The signaller may use these values in any manner, but the conventional use is described below. A handler may not return valid correction values, of course; a prudent signaller will re-check correction values using the same test that led to the signalling of the original error.

:error This catch-all condition name is used by the `error` (page 284) function. There is no particular convention regarding arguments or correction values.

:wrong-type-argument Requires *type* and *value*, where the first is a type specifier indicating what type of value is required, and the second is the value being complained about. The correction value should be a new object to be used in place of the one that was of the wrong type.

:contradictory-arguments Requires *function* and a list of all the arguments given to that function. 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. The correction value should be a list containing a new set of arguments to the same function.

??? Query: This differs from the Lisp Machine LISP `:inconsistent-arguments` error, where the correction value is simply used as the return value from the function that got the bad arguments.

:too-few-arguments Requires *function* and a list of all the arguments given to that function. This condition indicates that not enough arguments were passed to satisfy the required parameters. The correction value should be a list containing a new set of arguments to the same function.

:too-many-arguments Requires *function* and a list of all the arguments given to that function. This condition indicates that too many arguments were passed to a function. The correction value should be a list containing a new set of arguments to the same function.

:unexpected-keyword

Requires *function* and *keyword*, the latter being the purported name of a keyword passed as a keyword argument. The *keyword* is not a valid keyword for *function*. Correction values are ignored; correcting the error causes the keyword and its associated value to be ignored.

:invalid-form

Requires one argument *form*. The so-called *form* was not a meaningful form for `eval`. Probably it was of a bad data type. If the error proceeded, the value returned should be a new form to be evaluated in place of the bad form.

:unbound-variable

Requires a *symbol*. The *symbol* has no dynamic value associated with it (see `symbol-value` (page 62) and `set` (page 64)). The correction value is used to satisfy the request for the symbol's value as a dynamic variable.

:invalid-variable

Requires an *object*. An attempt was made to bind or assign to *object* as a variable, but it is not a symbol, or is a symbol but is a constant such as `nil` (page 51) or `pi` (page 130). The correction value should be a symbol to be used in place of *object* as the name of the variable.

:undefined-function

Requires *function-name*. The symbol *function-name* had no function definition. The first correction value is used as a function instead. (If this correction value is not a function after all, presumably an `invalid-function` error should ensue, but this check is the responsibility of the signaller.)

:invalid-function

Requires an *object*. An attempt was made to invoke *object* as a function, but it is not suitable for calling. If the *object* is a symbol, then perhaps it has no function definition, or is the name of a macro or special form rather than of a function. The correction value is used in place of *object* as a function to be invoked.

:failed-assertion

There is no particular convention on arguments. This is signalled correctly when the test for an `assert` (page 285) form fails. The correction values are ignored.

:failed-handler

Requires a *condition-name* and all the other arguments given to the handler. The handler that handled the condition returned the keyword `:return` for an uncorrectable error, or returned a first value other than `:return` or `nil`. This error is normally signalled uncorrectably.

:missing-handler

Requires a *condition-name* and all the other arguments given to the handler. One of the error-signalling functions `error` (page 284), `cerror` (page 285), or `ferror` (page 285) could not locate a handler for the specified *condition-name*. This error is normally signalled correctly; the correction value is the handler to use.

Chapter 24

Miscellaneous Features

24.1. The Compiler

The compiler is a program that may make code run faster, by translating programs into an implementation-dependent form that 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 49) 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 compiled-function object produced. If *name* is a symbol, then the compiled-function object is installed as the global function definition of the symbol and the symbol is returned. If *name* is `nil`, then the compiled-function 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 that computes  $b^2 - 4ac$ 
```

`compile-file` &optional *input-pathname* &key :output-file [Function]
 :set-default-pathname

The *input-pathname* must be a valid file specifier, such as a pathname. The defaults for *input-filename* are taken from the variable `*load-pathname-defaults*` (page 279). The file should be a LISP source file; its contents are compiled and written as a binary object (“FASL”) file.

The `:output` argument may be used to specify an output pathname; it defaults in a manner appropriate to the implementation’s file system conventions.

If the `:set-default-pathname` argument is true, then `compile-file` will set `*load-pathname-defaults*` (page 279) in such a way that `(load)` will load the newly compiled file and `(compile-file)` will recompile the source for that file.

`*compile-file-set-default-pathname*` [Variable]

This variable provides the default for the `:set-default-pathname` argument to `compile-file` (page 300). Its initial value is implementation-dependent.

`disassemble` *name-or-compiled-function* [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.

24.2. Documentation

A simple facility is provided for attaching strings to symbols for the purpose of on-line documentation. Rather than using the property list of the symbol, a separate function `documentation` is provided so that implementations can optimize the storage of documentation strings.

`documentation` *symbol doc-type*

[Function]

This function returns the documentation string of type *doc-type* for the *symbol*, or `nil` if none exists. Both arguments must be symbols. Some kinds of documentation are provided automatically by certain COMMON LISP constructs if the user writes an optional documentation string within them:

<u>Construct</u>	<u>Documentation Type</u>
<code>defvar</code> (page 48)	variable
<code>defparameter</code> (page 48)	variable
<code>defconstant</code> (page 48)	variable
<code>defun</code> (page 47)	function
<code>defmacro</code> (page 99)	function
<code>defstruct</code> (page 201)	structure
<code>deftype</code> (page 36)	type
<code>defsetf</code> (page 70)	setf

In addition, names of special forms may also have `function` documentation. (Macros and special forms are not really functions, of course, but it is convenient to group them with functions for documentation purposes.)

`setf` (page 66) may be used with `documentation` to update documentation information.

24.3. Modules

`provide` *module-name*

[Function]

`require` *module-name* &optional *pathname*

[Function]

A *module-name* should be a string or a symbol. Calling `provide` notes the fact that a program module of the specified name 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 278) in an attempt to obtain the necessary module from the file system. After the loading process is done, if the module still has not been provided, then an error is signalled. The *pathname* defaults in an implementation-dependent way that may depend on the *module-name*. (Typically, the name of the module 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 provide facilities 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 module is whimsical, but based on one actually provided in the MACLISP library.) 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 127).

```
(setq *package* (make-package 'timestamp))
(provide 'timestamp)
(export '(timestamp moonprinc sunprinc))
(require 'moonphase)
(require 'suncalc "/usr/gls/chutzpah/suncalc")
(shadow 'sqrt)
```

```
;;; Location of University of Southern North Dakota at Hoople.
(defconstant latitude 48.503 "Latitude of U. of S.N.D. at H.")
(defconstant longitude 97.61 "Longitude of U. of S.N.D. at H.")
```

```
(defun timestamp ...)
```

```
(defun moonprinc ...)
```

```
(defun sunprinc ...)
```

```
(defun stamp-utility ...)
```

It is important that the call to `provide` precede the calls to `require`. For suppose that the `moonphase` module needs to use `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.

24.4. Debugging Tools

The utilities described in this section are sufficiently complex and sufficiently dependent on the host environment that their complete definition necessarily belongs to either the yellow pages or the red pages. However, they are also sufficiently useful as to warrant mention here, to ensure that every implementation provides some version of them, however clever or however simple.

`trace {function-name}*` [Macro]
`untrace {function-name}*` [Macro]

Invoking `trace` with one or more function names (symbols) causes the functions named to be "traced". Henceforth, whenever such a function is invoked, information about the call, the arguments passed, and the eventually returned values, if any, will be printed to the stream that is the value of `*trace-output*` (page 214).

For example:

```
(trace fft gcd chase-pacman)
```

If a function call is open-coded (possibly as a result of an `inline` declaration), then such a call may not produce trace output.

Invoking `untrace` with one or more function names will cause those functions not to be traced

any more.

Tracing an already-traced function, or untracing a function not currently being traced, should produce no harmful effects, but may produce a warning message.

Calling `trace` with no argument forms will return a list of functions currently being traced.

Calling `untrace` with no argument forms will cause all currently traced functions to be no longer traced.

`trace` and `untrace` may also accept additional implementation-dependent argument formats. The format of the trace output is implementation-dependent.

`step form`

[Macro]

This evaluates *form*, and returns what *form* returns. However, the user is allowed to interactively "single-step" through the evaluation of *form*, at least through those evaluation steps that are performed interpretively. The nature of the interaction is implementation-dependent. However, implementations are encouraged to respond to the typing of the character "?" by providing help including a list of commands.

`time form`

[Macro]

This evaluates *form*, and returns what *form* returns. However, as a side effect, various timing data and other information is printed to the stream that is the value of `*trace-output*` (page 214). The nature and format of the printed information is implementation-dependent. However, implementations are encouraged to provide such information as elapsed real time, machine run time, storage management statistics, and so on.

Compatibility note: This facility is inspired by the INTERLISP facility of the same name. Note that the MACLISP/Lisp Machine LISP function `time` does something else entirely, namely return a quantity indicating relative elapsed real time.

`describe object`

[Function]

`describe` prints, to the stream in the variable `*standard-output*` (page 213), information about the *object*. Sometimes it will describe something that it finds inside something else; such recursive descriptions are indented appropriately. For instance, `describe` of a symbol will exhibit the symbol's value, its definition, and each of its properties. `describe` of a floating-point number will exhibit its internal representation in a way that is useful for tracking down roundoff errors and the like. The nature and format of the output is implementation-dependent.

`describe` always returns its argument.

`inspect` *object* [Function]

`inspect` is an interactive version of `describe`. The nature of the interaction is implementation-dependent, but the purpose of `inspect` is to make it easy to wander through a data structure, examining and modifying parts of it. Implementations are encouraged to respond to the typing of the character “?” by providing help, including a list of commands.

`room` &optional *x* [Function]

`room` prints, to the stream in the variable `*standard-output*` (page 213), information about the state of internal storage and its management. This might include descriptions of the amount of memory in use and the degree of memory compaction, possibly broken down by internal data type if that is appropriate. The nature and format of the printed information is implementation-dependent. The intent is to provide information that may help a user to tune his program to a particular implementation.

`(room nil)` prints out a minimal amount of information. `(room t)` prints out a maximal amount of information. Simply `(room)` prints out an intermediate amount of information that is likely to be useful.

`ed` &optional *x* [Function]

If the implementation provides a resident editor, this function should invoke it.

`(ed)` or `(ed nil)` simply enters the editor, leaving you in the same state as the last time you were in the editor.

`(ed pathname)` edits the contents of the file specified by *pathname*. The *pathname* may be an actual pathname or a string.

`(ed symbol)` tries to let you edit the text for the function named *symbol*. The means by which the function text is obtained is implementation-dependent; it might involve searching the file system, or pretty-printing resident interpreted code, for example.

`dribble` &optional *pathname* [Function]

`(dribble pathname)` rebinds `*standard-input*` (page 213) and `*standard-output*` (page 213) so as to send a record of the input/output interaction to a file named by *pathname*. The primary purpose of this is to create a readable record of an interactive session.

`(dribble)` terminates the recording of input and output and closes the dribble file.

`apropos` *string* &optional *package* [Function]

`(apropos string)` tries to find all symbols whose print-names contain *string* as a substring. Whenever it finds a symbol, it prints out the symbol's name; in addition, information about the function definition and dynamic value of the symbol, if any, is printed. If *package* is specified and not `nil`, then only that package is searched; otherwise “all” packages are searched, as if by

`do-all-symbols` (page 119). The information is printed to the stream that is the value of `*standard-output*` (page 213). `apropos` returns `t`.

24.5. Environment Inquiries

24.5.1. Time Functions

Time is represented in three different ways in COMMON LISP: Decoded Time, Universal Time, and Internal Time. The first two representations are used primarily to represent “real” (calendar) time, and are precise only to the second. Internal Time is used primarily to represent measurements of “computer” time (such as run time), and is precise to some implementation-dependent fraction of a second, as specified by `internal-time-units-per-second` (page 306). Decoded Time format is used only for absolute time indications. Universal Time and Internal Time formats are used for both absolute and relative times.

Decoded Time format represents time of day as a number of components:

- *Second*: an integer between 0 and 59, inclusive.
- *Minute*: an integer between 0 and 59, inclusive.
- *Hour*: an integer between 0 and 23, inclusive.
- *Date*: an integer between 1 and 31, inclusive (the upper limit actually depends on the month and year, of course).
- *Month*: an integer between 1 and 12, inclusive; 1 means January, 12 means December.
- *Year*: an integer indicating the year A.D. However, if this integer is between 0 and 99, the “obvious” year is used; more precisely, that year is assumed that is equal to the integer modulo 100 and within fifty years of the current year (inclusive backwards and exclusive forwards). Thus, in the year 1978, year 28 is 1928 but year 27 is 2027. (Functions that return time in this format always return a full year number.)

Compatibility note: This is incompatible with the Lisp Machine LISP definition in two ways. First, in Lisp Machine LISP a year between 0 and 99 always has 1900 added to it. Second, in Lisp Machine LISP time functions return the abbreviated year number between 0 and 99, rather than the full year number. The incompatibility is prompted by the imminent arrival of the twenty-first century. Note that `(mod year 100)` always reliably converts a year number to the abbreviated form, while the inverse conversion can be very difficult.

- *Day-of-week*: an integer between 0 and 6, inclusive; 0 means Monday, 1 means Tuesday, and so on, and 6 means Sunday.

??? Query: How did this happen? One would expect Sunday to be either 0 or 1.

- *Daylight-savings-time-p*: a flag that, if not `nil`, indicates that daylight savings time is in effect.

- *Time-zone*: an integer specified as the number of hours west of GMT (Greenwich Mean Time). For example, in Massachusetts the time-zone is 5, and in California it is 8. Any adjustment for daylight savings time is separate from this.

Universal Time represents time as a single integer. For relative time purposes, this is a number of seconds. For absolute time, this is the number of seconds since midnight, January 1, 1900 GMT. Thus the time 1 is 00:00:01 (that is, 12:00:01 AM) on January 1, 1900 GMT. Similarly, the time 2398291201 corresponds to time 00:00:01 on January 1, 1976 GMT. Recall that the year 1900 was *not* a leap year; for the purposes of COMMON LISP, a year is a leap year iff its number is divisible by 4, except that years divisible by 100 are *not* leap years, except that years divisible by 400 *are* leap years. Universal Time format is used as a standard time representation within the ARPANET; see [5].

Internal Time also represents time as a single integer, in terms of an implementation-dependent unit. Relative time is measured as a number of these units. Absolute time is relative to an arbitrary time base, typically the time at which the system began running.

`internal-time-units-per-second`

[*Constant*]

This value is an integer, the implementation-dependent number of internal time units in a second. (The internal time unit must be chosen so that one second is an integral multiple of it.)

Rationale: The reason for allowing the internal time units to be implementation-dependent is so that `get-internal-time` (page 306) can execute with minimum overhead. The idea is that it should be very likely that a fixnum will suffice as the returned value from `get-internal-time`. This probability can be tuned to the implementation by trading off the speed of the machine against the word size. Any particular unit will be inappropriate for some implementations: a microsecond is too long for a very fast machine such as an S-1, while a much smaller unit would force many implementations to return bignums for most calls to `get-internal-time`, rendering that function less useful for accurate timing measurements.

`get-decoded-time`

[*Function*]

The current time is returned in Decoded Time format. Nine values are returned: *second*, *minute*, *hour*, *date*, *month*, *year*, *day-of-week*, *daylight-savings-time-p*, and *time-zone*.

Compatibility note: In Lisp Machine LISP the *time-zone* is not currently returned. Consider, however, the use of COMMON LISP in some mobile vehicle. It is entirely plausible that the time-zone might change from time to time.

`get-universal-time`

[*Function*]

The current time of day is returned as a single integer in Universal Time format.

`get-internal-time`

[*Function*]

The current time is returned as a single integer in Internal Time format. The precise meaning of this quantity is implementation-dependent; it may measure real time, run time, CPU cycles, or some other quantity.

??? Query: How can this notion be made meaningful and portable?

`sleep` *seconds* [Function]
 (`sleep` *n*) causes execution to cease and become dormant for approximately *n* seconds of real time, whereupon execution is resumed. The argument may be any non-negative non-complex number. `sleep` returns `nil`.

`decode-universal-time` *universal-time* &optional *time-zone* [Function]
 The time specified by *universal-time* in Universal Time format is converted to Decoded Time format. Nine values are returned: *second*, *minute*, *hour*, *date*, *month*, *year*, *day-of-week*, *daylight-savings-time-p*, and *time-zone*.

Compatibility note: In Lisp Machine LISP the *time-zone* is not currently returned. Consider, however, the use of COMMON LISP in some mobile vehicle. It is entirely plausible that the time-zone might change from time to time.

The *time-zone* argument defaults to the current time-zone.

`encode-universal-time` *second* *minute* *hour* *date* *month* *year* &optional *time-zone* [Function]
 The time specified by the given components of Decoded Time format is encoded into Universal Time format and returned. If you don't specify *time-zone*, it defaults to the current time-zone adjusted for daylight savings time. If you provide *time-zone* explicitly, no adjustment for daylight savings time is performed.

24.5.2. Other Environment Inquiries

For any of the following functions, if no appropriate and relevant result can be produced, `nil` is returned instead of a string.

Rationale: These inquiry facilities are functions rather than variables against the possibility that a COMMON LISP process might migrate from machine to machine. This need not happen in a distributed environment; consider, for example, dumping a core image file containing a compiler and then shipping it to another site.

`lisp-implementation-type` [Function]
 A string is returned that identifies the generic name of the particular COMMON LISP implementation. Examples: "Spice LISP", "Zetalisp".

`lisp-implementation-version` [Function]
 A string is returned that identifies the version of the particular COMMON LISP implementation; this information should be of use to maintainers of the implementation. Examples: "1192", "53.7 with complex numbers", "1746.9A, NEWIO 53, ETHER 5.3".

`machine-type` [Function]

A string is returned that identifies the generic name of the computer hardware on which COMMON LISP is running. Examples: "DEC PDP-10", "DEC VAX-11/780".

`machine-version` [Function]

A string is returned that identifies the version of the computer hardware on which COMMON LISP is running. Example: "KL10, microcode 9".

`machine-instance` [Function]

A string is returned that identifies the particular instance of the computer hardware on which COMMON LISP is running; this might be a local nickname, for example, and/or a serial number. Examples: "MIT-MC", "CMU GP-VAX".

`host-software-type` [Function]

A string is returned that identifies the generic name of any relevant host software. Examples: "Spice", "TOPS-20", "ITS".

`host-software-version` [Function]

A string is returned that identifies the version of the particular COMMON LISP implementation; this information should be of use to maintainers of the implementation.

`short-site-name` [Function]

`long-site-name` [Function]

A string is returned that identifies the physical location of the computer hardware. Examples of short names: "MIT AI Lab", "CMU-CSD". Examples of long names:

```
"MIT Artificial Intelligence Laboratory"
"Massachusetts Institute of Technology
Artificial Intelligence Laboratory"
"Carnegie-Mellon University Computer Science Department"
```

See also `user-homedir-pathname` (page 269) and `init-file-pathname` (page 269).

`*features*` [Variable]

The value of the variable `*features*` should be a list of symbols that name "features" provided by the implementation. Most such symbols will be implementation-specific; typically a name for the implementation will be included. One standard feature name is `ieee-floating-point`, which should be present if and only if full IEEE proposed floating-point arithmetic [6] is supported.

The value of this variable is used by the `#+` and `#-` reader syntax; see page 233.

24.6. Identity Function

`identity` *object*

[Function]

The *object* is returned as the value of `identity`. This function is useful primarily as an argument to other functions.

References

1. Brooks, Rodney A.; Gabriel, Richard P.; and Steele, Guy L. Jr. "An Optimizing Compiler for Lexically Scoped LISP." *Proceedings of the 1982 Symposium on Compiler Construction*. ACM SIGPLAN (Boston, June 1982), 261-275. 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 (Jan. 1980), 68-79. 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 (March 1981), 62. These are errata for [2].
4. Fateman, Richard J. "Reply to an Editorial." *ACM SIGSAM Bulletin* 25 (March 1973), 9-11.
5. Harrenstien, Kenneth L. *Time Server*. Request for Comments (RFC) 738 (NIC 42218), ARPANET Network Working Group (Oct. 1977). Available from the ARPANET Network Information Center.
6. IEEE Computer Society Standard Committee, Microprocessor Standards Subcommittee, Floating-Point Working Group. "A Proposed Standard for Binary Floating-Point Arithmetic." *Computer* 14, 3 (March 1981), 51-62.
7. Knuth, Donald E.. *The Art of Computer Programming*. Volume 2: *Seminumerical Algorithms*. Addison-Wesley (Reading, Massachusetts, 1969).
8. Marti, J.; Hearn, A.C.; Griss, M.L.; and Griss, C. "Standard LISP Report." *SIGPLAN Notices* 14, 10 (Oct. 1979), 48-68.
9. Moon, David. *MacLISP Reference Manual, Revision 0*. M.I.T. Project MAC (Cambridge, Massachusetts, April 1974).
10. Moore, J. Strother II. *The InterLISP Virtual Machine Specification*. Tech. Rept. CSL 76-5, Xerox Palo Alto Research Center (Palo Alto, California, Sept. 1976).
11. Penfield, Paul, Jr. "Principal Values and Branch Cuts in Complex APL." *APL 81 Conference Proceedings*. ACM SIGAPL (San Francisco, Sept. 1981), 248-256. Proceedings published as *APL Quote Quad* 12, 1 (September 1981).
12. Steele, Guy Lewis Jr., and Sussman, Gerald Jay. *The Revised Report on SCHEME: A Dialect of LISP*. AI Memo 452, MIT Artificial Intelligence Lab. (Cambridge, Massachusetts, Jan. 1978).
13. Suzuki, Norihisa. "Analysis of Pointer 'Rotation'." *Comm. ACM* 25, 5 (May 1982), 330-335.
14. Teitelman, Warren, et al. *InterLISP Reference Manual*. Xerox Palo Alto Research Center (Palo Alto, California, 1978). Third revision.
15. Weinreb, Daniel, and Moon, David. *LISP Machine Manual, Fourth Edition*. MIT Artificial Intelligence Lab. (Cambridge, Massachusetts, July 1981).

COMMON LISP Summary

COMMON LISP Summary

sample-function <i>arg1 arg2</i> &optional <i>arg3 arg4</i>	[Function]
sample-variable	[Variable]
sample-constant	[Constant]
sample-special-form [<i>name</i>] (<i>{var}</i> *) <i>{form}</i> ⁺	[Special form]
sample-macro <i>var</i> <i>{tag statement}</i> *	[Macro]
deftype <i>name lambda-list</i> <i>{declaration doc-string}</i> * <i>{form}</i> *	[Macro]
coerce <i>object result-type</i>	[Function]
type-of <i>object</i>	[Function]
lambda-list-keywords	[Constant]
defun <i>name lambda-list</i> <i>{declaration doc-string}</i> * <i>{form}</i> *	[Macro]
defvar <i>name</i> [<i>initial-value</i> [<i>documentation</i>]]	[Macro]
defparameter <i>name initial-value</i> [<i>documentation</i>]	[Macro]
defconstant <i>name initial-value</i> [<i>documentation</i>]	[Macro]
eval-when (<i>{situation}</i> *) <i>{form}</i> *	[Function]
nil	[Constant]
t	[Constant]
typep <i>object type</i>	[Function]
subtypep <i>type1 type2</i>	[Function]
null <i>object</i>	[Function]
symbolp <i>object</i>	[Function]
atom <i>object</i>	[Function]
consp <i>object</i>	[Function]
listp <i>object</i>	[Function]
numberp <i>object</i>	[Function]
integerp <i>object</i>	[Function]
rationalp <i>object</i>	[Function]
floatp <i>object</i>	[Function]
complexp <i>object</i>	[Function]
characterp <i>object</i>	[Function]
stringp <i>object</i>	[Function]
bit-vector-p <i>object</i>	[Function]
vectorp <i>object</i>	[Function]
simple-string-p <i>object</i>	[Function]
simple-bit-vector-p <i>object</i>	[Function]
simple-vector-p <i>object</i>	[Function]
arrayp <i>object</i>	[Function]
functionp <i>object</i>	[Function]
compiled-function-p <i>object</i>	[Function]
commonp <i>object</i>	[Function]

<code>eq x y</code>	[Function]
<code>eq1 x y</code>	[Function]
<code>equal x y</code>	[Function]
<code>equalp x y</code>	[Function]
<code>not x</code>	[Function]
<code>and {form}*</code>	[Special form]
<code>or {form}*</code>	[Special form]
<code>quote object</code>	[Special form]
<code>function fn</code>	[Special form]
<code>symbol-value symbol</code>	[Function]
<code>symbol-function symbol</code>	[Function]
<code>boundp symbol</code>	[Function]
<code>fboundp symbol</code>	[Function]
<code>macro-p symbol</code>	[Function]
<code>special-form-p symbol</code>	[Function]
<code>setq {var form}*</code>	[Special form]
<code>psetq {var form}*</code>	[Macro]
<code>set symbol value</code>	[Function]
<code>fset symbol value</code>	[Function]
<code>makunbound symbol</code>	[Function]
<code>fmakunbound symbol</code>	[Function]
<code>setf {place newvalue}*</code>	[Macro]
<code>psetf {place newvalue}*</code>	[Macro]
<code>shiftf place {place}* newvalue</code>	[Macro]
<code>rotatef {place}*</code>	[Macro]
<code>defsetf access-fn {update-fn [doc-string] lambda-list lambda-list {declaration doc-string}* {form}*}</code>	[Macro]
<code>apply function arg &rest more-args</code>	[Function]
<code>funcall fn &rest arguments</code>	[Function]
<code>progn {form}*</code>	[Special form]
<code>prog1 first {form}*</code>	[Macro]
<code>prog2 first second {form}*</code>	[Macro]
<code>let ({var (var value)}*) {form}*</code>	[Macro]
<code>let* ({var (var value)}*) {form}*</code>	[Special form]
<code>compiler-let ({var (var value)}*) {form}*</code>	[Macro]
<code>progv symbols values {form}*</code>	[Special form]
<code>flet ({(name lambda-list {declaration doc-string}* {form}*)}*) {form}*</code>	[Special form]
<code>labels ({(name lambda-list {declaration doc-string}* {form}*)}*) {form}*</code>	[Special form]
<code>macrolet ({(name varlist {declaration doc-string}* {form}*)}*) {form}*</code>	[Special form]
<code>cond {{test {form}*}}*</code>	[Macro]
<code>if pred then [else]</code>	[Special form]
<code>when pred {form}*</code>	[Macro]
<code>unless pred {form}*</code>	[Macro]

<code>case keyform {{{key}*} {form}*}</code> *	[Macro]
<code>typecase keyform {(type {form}*)}</code> *	[Macro]
<code>block name {form}</code> *	[Special form]
<code>return-from name [result]</code>	[Special form]
<code>return [result]</code>	[Macro]
<code>loop {form}</code> *	[Macro]
<code>do ({(var [init [step]])}*} (end-test {form})* {declaration}* {tag statement}*</code>	[Macro]
<code>do* ({(var [init [step]])}*} (end-test {form})* {declaration}* {tag statement}*</code>	[Macro]
<code>dolist (var listform [resultform]) {declaration}* {tag statement}*</code>	[Macro]
<code>dotimes (var countform [resultform]) {declaration}* {tag statement}*</code>	[Macro]
<code>mapcar function list &rest more-lists</code>	[Function]
<code>maplist function list &rest more-lists</code>	[Function]
<code>mapc function list &rest more-lists</code>	[Function]
<code>mapl function list &rest more-lists</code>	[Function]
<code>mapcan function list &rest more-lists</code>	[Function]
<code>mapcon function list &rest more-lists</code>	[Function]
<code>tagbody {tag statement}*</code>	[Special form]
<code>prog ({var (var [init])}*} {declaration}* {tag statement}*</code>	[Macro]
<code>prog* ({var (var [init])}*} {declaration}* {tag statement}*</code>	[Macro]
<code>go tag</code>	[Special form]
<code>values &rest args</code>	[Function]
<code>values-list list</code>	[Function]
<code>multiple-value-list form</code>	[Macro]
<code>multiple-value-call function {form}</code> *	[Special form]
<code>multiple-value-prog1 form {form}</code> *	[Special form]
<code>multiple-value-bind ({var}*} values-form {declaration}* {form}</code> *	[Macro]
<code>multiple-value variables form</code>	[Macro]
<code>catch tag {form}</code> *	[Special form]
<code>catch-all catch-function {form}</code> *	[Special form]
<code>unwind-all catch-function {form}</code> *	[Special form]
<code>unwind-protect protected-form {cleanup-form}</code> *	[Special form]
<code>throw tag result</code>	[Special form]
<code>macro name (var) {declaration}* {form}</code> *	[Macro]
<code>defmacro name lambda-list {declaration doc-string}* {form}</code> *	[Macro]
<code>macroexpand form &rest environment</code>	[Function]
<code>macroexpand-1 form &rest environment</code>	[Function]
<code>*macroexpand-hook*</code>	[Variable]
<code>declare {declaration-form}</code> *	[Special form]
<code>locally {declaration}* {form}</code> *	[Macro]
<code>the value-type form</code>	[Special form]
<code>get symbol indicator &optional default</code>	[Function]
<code>remprop symbol indicator</code>	[Function]
<code>symbol-plist symbol</code>	[Function]

<code>getf</code> <i>place indicator</i> &optional <i>default</i>	[Function]
<code>putf</code> <i>place indicator newvalue</i>	[Macro]
<code>remf</code> <i>place indicator</i>	[Macro]
<code>get-properties</code> <i>place indicator-list</i>	[Function]
<code>symbol-print-name</code> <i>sym</i>	[Function]
<code>samepnamep</code> <i>sym1 sym2</i>	[Function]
<code>make-symbol</code> <i>print-name</i>	[Function]
<code>copy-symbol</code> <i>sym</i> &optional <i>copy-props</i>	[Function]
<code>gensym</code> &optional <i>x</i>	[Function]
<code>gentemp</code> <i>prefix</i> &optional <i>package</i>	[Function]
<code>symbol-package</code> <i>sym</i>	[Function]
<code>keywordp</code> <i>symbol</i>	[Function]
<code>make-package</code> <i>package-name</i> &optional <i>copy-from</i>	[Function]
<code>*package*</code>	[Variable]
<code>packagep</code> <i>object</i>	[Function]
<code>package</code> <i>package</i>	[Function]
<code>package-name</code> <i>package</i>	[Function]
<code>intern</code> <i>string-or-symbol</i> &optional <i>package</i>	[Function]
<code>unintern</code> <i>string-or-symbol</i> &optional <i>package</i>	[Function]
<code>internedp</code> <i>string-or-symbol</i> &optional <i>package</i>	[Function]
<code>export</code> <i>symbols</i>	[Function]
<code>shadow</code> <i>symbols</i>	[Function]
<code>do-symbols</code> (<i>var</i> [<i>package</i>] [<i>result-form</i>]) { <i>declaration</i> }* { <i>tag</i> <i>statement</i> }*	[Macro]
<code>do-all-symbols</code> (<i>var</i> [<i>result-form</i>]) { <i>declaration</i> }* { <i>tag</i> <i>statement</i> }*	[Macro]
<code>zerop</code> <i>number</i>	[Function]
<code>plusp</code> <i>number</i>	[Function]
<code>minusp</code> <i>number</i>	[Function]
<code>oddp</code> <i>integer</i>	[Function]
<code>evenp</code> <i>integer</i>	[Function]
<code>=</code> <i>number</i> &rest <i>more-numbers</i>	[Function]
<code>/=</code> <i>number</i> &rest <i>more-numbers</i>	[Function]
<code><</code> <i>number</i> &rest <i>more-numbers</i>	[Function]
<code>></code> <i>number</i> &rest <i>more-numbers</i>	[Function]
<code><=</code> <i>number</i> &rest <i>more-numbers</i>	[Function]
<code>>=</code> <i>number</i> &rest <i>more-numbers</i>	[Function]
<code>max</code> <i>number</i> &rest <i>more-numbers</i>	[Function]
<code>min</code> <i>number</i> &rest <i>more-numbers</i>	[Function]
<code>+</code> &rest <i>numbers</i>	[Function]
<code>-</code> <i>number</i> &rest <i>more-numbers</i>	[Function]
<code>*</code> &rest <i>numbers</i>	[Function]
<code>/</code> <i>number</i> &rest <i>more-numbers</i>	[Function]
<code>1+</code> <i>number</i>	[Function]
<code>1-</code> <i>number</i>	[Function]

<i>incf place [delta]</i>	[Macro]
<i>decf place [delta]</i>	[Macro]
<i>conjugate number</i>	[Function]
<i>gcd &rest integers</i>	[Function]
<i>lcm integer &rest more-integers</i>	[Function]
<i>exp number</i>	[Function]
<i>expt base-number power-number</i>	[Function]
<i>log number &optional base</i>	[Function]
<i>sqrt number</i>	[Function]
<i>isqrt integer</i>	[Function]
<i>abs number</i>	[Function]
<i>phase number</i>	[Function]
<i>signum number</i>	[Function]
<i>sin radians</i>	[Function]
<i>cos radians</i>	[Function]
<i>tan radians</i>	[Function]
<i>cis radians</i>	[Function]
<i>asin number</i>	[Function]
<i>acos number</i>	[Function]
<i>atan y &optional x</i>	[Function]
<i>pi</i>	[Constant]
<i>sinh number</i>	[Function]
<i>cosh number</i>	[Function]
<i>tanh number</i>	[Function]
<i>asinh number</i>	[Function]
<i>acosh number</i>	[Function]
<i>atanh number</i>	[Function]
<i>float number &optional other</i>	[Function]
<i>rational number</i>	[Function]
<i>rationalize number</i>	[Function]
<i>numerator rational</i>	[Function]
<i>denominator rational</i>	[Function]
<i>floor number &optional divisor</i>	[Function]
<i>ceiling number &optional divisor</i>	[Function]
<i>truncate number &optional divisor</i>	[Function]
<i>round number &optional divisor</i>	[Function]
<i>mod number divisor</i>	[Function]
<i>remainder number divisor</i>	[Function]
<i>ffloor number &optional divisor</i>	[Function]
<i>fceiling number &optional divisor</i>	[Function]
<i>ftruncate number &optional divisor</i>	[Function]
<i>fround number &optional divisor</i>	[Function]
<i>float-significand float</i>	[Function]

<code>float-exponent</code> <i>float</i>	[Function]
<code>scale-float</code> <i>float integer</i>	[Function]
<code>float-radix</code> <i>float</i>	[Function]
<code>float-sign</code> <i>float1</i> &optional <i>float2</i>	[Function]
<code>complex</code> <i>realpart</i> &optional <i>imagpart</i>	[Function]
<code>realpart</code> <i>number</i>	[Function]
<code>imagpart</code> <i>number</i>	[Function]
<code>logior</code> &rest <i>integers</i>	[Function]
<code>logxor</code> &rest <i>integers</i>	[Function]
<code>logand</code> &rest <i>integers</i>	[Function]
<code>logeqv</code> &rest <i>integers</i>	[Function]
<code>lognand</code> <i>integer1 integer2</i>	[Function]
<code>lognor</code> <i>integer1 integer2</i>	[Function]
<code>logandc1</code> <i>integer1 integer2</i>	[Function]
<code>logandc2</code> <i>integer1 integer2</i>	[Function]
<code>logorc1</code> <i>integer1 integer2</i>	[Function]
<code>logorc2</code> <i>integer1 integer2</i>	[Function]
<code>boole</code> <i>op integer1 integer2</i>	[Function]
<code>boole-clr</code>	[Constant]
<code>boole-set</code>	[Constant]
<code>boole-1</code>	[Constant]
<code>boole-2</code>	[Constant]
<code>boole-c1</code>	[Constant]
<code>boole-c2</code>	[Constant]
<code>boole-and</code>	[Constant]
<code>boole-ior</code>	[Constant]
<code>boole-xor</code>	[Constant]
<code>boole-eqv</code>	[Constant]
<code>boole-nand</code>	[Constant]
<code>boole-nor</code>	[Constant]
<code>boole-andc1</code>	[Constant]
<code>boole-andc2</code>	[Constant]
<code>boole-orc1</code>	[Constant]
<code>boole-orc2</code>	[Constant]
<code>lognot</code> <i>integer</i>	[Function]
<code>logtest</code> <i>integer1 integer2</i>	[Function]
<code>logbitp</code> <i>index integer</i>	[Function]
<code>ash</code> <i>integer count</i>	[Function]
<code>logcount</code> <i>integer</i>	[Function]
<code>integer-length</code> <i>integer</i>	[Function]
<code>byte</code> <i>size position</i>	[Function]
<code>byte-size</code> <i>bytespec</i>	[Function]
<code>byte-position</code> <i>bytespec</i>	[Function]

<code>ldb</code> <i>bytespec integer</i>	[Function]
<code>ldb-test</code> <i>bytespec integer</i>	[Function]
<code>mask-field</code> <i>bytespec integer</i>	[Function]
<code>dpb</code> <i>newbyte bytespec integer</i>	[Function]
<code>deposit-field</code> <i>newbyte bytespec integer</i>	[Function]
<code>random</code> <i>number &optional state</i>	[Function]
<code>*random-state*</code>	[Variable]
<code>make-random-state</code> <i>&optional state</i>	[Function]
<code>random-state-p</code> <i>object</i>	[Function]
<code>most-positive-fixnum</code>	[Constant]
<code>most-negative-fixnum</code>	[Constant]
<code>most-positive-short-float</code>	[Constant]
<code>least-positive-short-float</code>	[Constant]
<code>least-negative-short-float</code>	[Constant]
<code>most-negative-short-float</code>	[Constant]
<code>most-positive-single-float</code>	[Constant]
<code>least-positive-single-float</code>	[Constant]
<code>least-negative-single-float</code>	[Constant]
<code>most-negative-single-float</code>	[Constant]
<code>most-positive-double-float</code>	[Constant]
<code>least-positive-double-float</code>	[Constant]
<code>least-negative-double-float</code>	[Constant]
<code>most-negative-double-float</code>	[Constant]
<code>most-positive-long-float</code>	[Constant]
<code>least-positive-long-float</code>	[Constant]
<code>least-negative-long-float</code>	[Constant]
<code>most-negative-long-float</code>	[Constant]
<code>short-float-epsilon</code>	[Constant]
<code>single-float-epsilon</code>	[Constant]
<code>double-float-epsilon</code>	[Constant]
<code>long-float-epsilon</code>	[Constant]
<code>short-float-negative-epsilon</code>	[Constant]
<code>single-float-negative-epsilon</code>	[Constant]
<code>double-float-negative-epsilon</code>	[Constant]
<code>long-float-negative-epsilon</code>	[Constant]
<code>char-code-limit</code>	[Constant]
<code>char-font-limit</code>	[Constant]
<code>char-bits-limit</code>	[Constant]
<code>standard-charp</code> <i>char</i>	[Function]
<code>graphic-charp</code> <i>char</i>	[Function]
<code>string-charp</code> <i>char</i>	[Function]
<code>alpha-charp</code> <i>char</i>	[Function]
<code>uppercasep</code> <i>char</i>	[Function]

<code>lowercasep</code> <i>char</i>	[Function]
<code>bothcasep</code> <i>char</i>	[Function]
<code>digit-charp</code> <i>char</i> &optional (<i>radix</i> 10.)	[Function]
<code>alphanumericp</code> <i>char</i>	[Function]
<code>char=</code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>char/=</code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>char<</code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>char></code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>char<=</code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>char>=</code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>char-equal</code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>char-not-equal</code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>char-lessp</code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>char-greaterp</code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>char-not-greaterp</code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>char-not-lessp</code> <i>character</i> &rest <i>more-characters</i>	[Function]
<code>character</code> <i>object</i>	[Function]
<code>char-code</code> <i>char</i>	[Function]
<code>char-bits</code> <i>char</i>	[Function]
<code>char-font</code> <i>char</i>	[Function]
<code>code-char</code> <i>code</i> &optional (<i>bits</i> 0) (<i>font</i> 0)	[Function]
<code>make-char</code> <i>char</i> &optional (<i>bits</i> 0) (<i>font</i> 0)	[Function]
<code>char-upcase</code> <i>char</i>	[Function]
<code>char-downcase</code> <i>char</i>	[Function]
<code>digit-weight</code> <i>weight</i> &optional (<i>radix</i> 10.) (<i>bits</i> 0) (<i>font</i> 0)	[Function]
<code>char-int</code> <i>char</i>	[Function]
<code>int-char</code> <i>integer</i>	[Function]
<code>char-name</code> <i>char</i>	[Function]
<code>name-char</code> <i>sym</i>	[Function]
<code>char-control-bit</code>	[Constant]
<code>char-meta-bit</code>	[Constant]
<code>char-super-bit</code>	[Constant]
<code>char-hyper-bit</code>	[Constant]
<code>char-bit</code> <i>char name</i>	[Function]
<code>set-char-bit</code> <i>char name newvalue</i>	[Function]
<code>elt</code> <i>sequence index</i>	[Function]
<code>subseq</code> <i>sequence start</i> &optional <i>end</i>	[Function]
<code>copy-seq</code> <i>sequence</i>	[Function]
<code>length</code> <i>sequence</i>	[Function]
<code>reverse</code> <i>sequence</i>	[Function]
<code>nreverse</code> <i>sequence</i>	[Function]
<code>make-sequence</code> <i>type size</i> &key <i>:initial-element</i>	[Function]
<code>concatenate</code> <i>result-type</i> &rest <i>sequences</i>	[Function]

<i>map result-type function sequence &rest more-sequences</i>	[Function]
<i>some predicate sequence &rest more-sequences</i>	[Function]
<i>every predicate sequence &rest more-sequences</i>	[Function]
<i>notany predicate sequence &rest more-sequences</i>	[Function]
<i>notevery predicate sequence &rest more-sequences</i>	[Function]
<i>reduce function sequence &key :from-end :start :end :initial-value</i>	[Function]
<i>fill sequence item &key :start :end</i>	[Function]
<i>replace sequence1 sequence2 &key :start1 :end1 :start2 :end2</i>	[Function]
<i>remove item sequence &key :from-end :test :test-not :start :end :count :key</i>	[Function]
<i>remove-if test sequence &key :from-end :start :end :count :key</i>	[Function]
<i>remove-if-not test sequence &key :from-end :start :end :count :key</i>	[Function]
<i>delete item sequence &key :from-end :test :test-not :start :end :count :key</i>	[Function]
<i>delete-if test sequence &key :from-end :start :end :count :key</i>	[Function]
<i>delete-if-not test sequence &key :from-end :start :end :count :key</i>	[Function]
<i>remove-duplicates sequence &key :test :test-not :start :end</i>	[Function]
<i>delete-duplicates sequence &key :test :test-not :start :end</i>	[Function]
<i>substitute newitem olditem sequence &key :from-end :test :test-not :start :end :count :key</i>	[Function]
<i>substitute-if newitem test sequence &key :from-end :start :end :count :key</i>	[Function]
<i>substitute-if-not newitem test sequence &key :from-end :start :end :count :key</i>	[Function]
<i>nsubstitute newitem olditem sequence &key :from-end :test :test-not :start :end :count :key</i>	[Function]
<i>nsubstitute-if newitem test sequence &key :from-end :start :end :count :key</i>	[Function]
<i>nsubstitute-if-not newitem test sequence &key :from-end :start :end :count :key</i>	[Function]
<i>find item sequence &key :from-end :test :test-not :start :end :key</i>	[Function]
<i>find-if test sequence &key :from-end :start :end :key</i>	[Function]
<i>find-if-not test sequence &key :from-end :start :end :key</i>	[Function]
<i>position item sequence &key :from-end :test :test-not :start :end :key</i>	[Function]
<i>position-if test sequence &key :from-end :start :end :key</i>	[Function]
<i>position-if-not test sequence &key :from-end :start :end :key</i>	[Function]
<i>count item sequence &key :from-end :test :test-not :start :end :key</i>	[Function]
<i>count-if test sequence &key :from-end :start :end :key</i>	[Function]
<i>count-if-not test sequence &key :from-end :start :end :key</i>	[Function]
<i>mismatch sequence1 sequence2 &key :from-end :test :test-not :key :start1 :start2 :end1 :end2</i>	[Function]
<i>maxprefix sequence1 sequence2 &key :from-end :test :test-not :key :start1 :start2 :end1 :end2</i>	[Function]

<code>ldiff list sublist</code>	[Function]
<code>rplaca x y</code>	[Function]
<code>rplacd x y</code>	[Function]
<code>subst new old tree &key :test :test-not :key</code>	[Function]
<code>subst-if predicate new tree &key :key</code>	[Function]
<code>subst-if-not predicate new tree &key :key</code>	[Function]
<code>nsubst new old tree &key :test :test-not :key</code>	[Function]
<code>nsubst-if predicate new tree &key :key</code>	[Function]
<code>nsubst-if-not predicate new tree &key :key</code>	[Function]
<code>sublis alist tree &key :test :test-not :key</code>	[Function]
<code>nsublis alist tree &key :test :test-not :key</code>	[Function]
<code>member item list &key :test :test-not :key</code>	[Function]
<code>member-if predicate list &key :key</code>	[Function]
<code>member-if-not predicate list &key :key</code>	[Function]
<code>tailp sublist list</code>	[Function]
<code>adjoin item list &key :test :test-not :key</code>	[Function]
<code>union list1 list2 &key :test :test-not :key</code>	[Function]
<code>nunion list1 list2 &key :test :test-not :key</code>	[Function]
<code>intersection list1 list2 &key :test :test-not :key</code>	[Function]
<code>nintersection list1 list2 &key :test :test-not :key</code>	[Function]
<code>set-difference list1 list2 &key :test :test-not :key</code>	[Function]
<code>nset-difference list1 list2 &key :test :test-not :key</code>	[Function]
<code>set-exclusive-or list1 list2 &key :test :test-not :key</code>	[Function]
<code>nset-exclusive-or list1 list2 &key :test :test-not :key</code>	[Function]
<code>subsetp list1 list2 &key :test :test-not :key</code>	[Function]
<code>acons key datum a-list</code>	[Function]
<code>pairlis keys data &optional a-list</code>	[Function]
<code>assoc item a-list &key :test :test-not</code>	[Function]
<code>assoc-if predicate a-list</code>	[Function]
<code>assoc-if-not predicate a-list</code>	[Function]
<code>rassoc item a-list &key :test :test-not</code>	[Function]
<code>rassoc-if predicate a-list</code>	[Function]
<code>rassoc-if-not predicate a-list</code>	[Function]
<code>make-hash-table &key :test :size :rehash-size :rehash-threshold</code>	[Function]
<code>hash-table-p object</code>	[Function]
<code>gethash key hash-table &optional default</code>	[Function]
<code>remhash key hash-table</code>	[Function]
<code>maphash function hash-table</code>	[Function]
<code>clrhash hash-table</code>	[Function]
<code>hash-table-count hash-table</code>	[Function]
<code>sxhash S-expression</code>	[Function]
<code>char string index</code>	[Function]
<code>string= string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]

<code>string-equal string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]
<code>string< string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]
<code>string> string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]
<code>string<= string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]
<code>string>= string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]
<code>string/= string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]
<code>string-lessp string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]
<code>string-greaterp string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]
<code>string-not-lessp string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]
<code>string-not-greaterp string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]
<code>string-not-equal string1 string2 &key :start1 :end1 :start2 :end2</code>	[Function]
<code>make-string size &key :initial-element</code>	[Function]
<code>string-trim character-bag string</code>	[Function]
<code>string-left-trim character-bag string</code>	[Function]
<code>string-right-trim character-bag string</code>	[Function]
<code>string-upcase string &key :start :end</code>	[Function]
<code>string-downcase string &key :start :end</code>	[Function]
<code>string-capitalize string &key :start :end</code>	[Function]
<code>nstring-upcase string &key :start :end</code>	[Function]
<code>nstring-downcase string &key :start :end</code>	[Function]
<code>nstring-capitalize string &key :start :end</code>	[Function]
<code>string x</code>	[Function]
<code>defstruct name-and-options [doc-string] {slot-description} ⁺</code>	[Macro]
<code>eval form</code>	[Function]
<code>*evalhook*</code>	[Variable]
<code>*eval form &rest env</code>	[Function]
<code>evalhook form hookfn &rest env</code>	[Function]
<code>constantp object</code>	[Function]
<code>+</code>	[Variable]
<code>++</code>	[Variable]
<code>+++</code>	[Variable]
<code>-</code>	[Variable]
<code>*</code>	[Variable]
<code>**</code>	[Variable]
<code>***</code>	[Variable]
<code>/</code>	[Variable]
<code>//</code>	[Variable]
<code>///</code>	[Variable]
<code>*standard-input*</code>	[Variable]
<code>*standard-output*</code>	[Variable]
<code>*error-output*</code>	[Variable]
<code>*query-io*</code>	[Variable]
<code>*terminal-io*</code>	[Variable]

<code>*trace-output*</code>	[Variable]
<code>make-synonym-stream</code> <i>symbol</i>	[Function]
<code>make-broadcast-stream</code> &rest <i>streams</i>	[Function]
<code>make-concatenated-stream</code> &rest <i>streams</i>	[Function]
<code>make-two-way-stream</code> <i>input-stream output-stream</i>	[Function]
<code>make-echo-stream</code> <i>input-stream output-stream</i>	[Function]
<code>make-string-input-stream</code> <i>string</i> &optional <i>start end</i>	[Function]
<code>make-string-output-stream</code> &optional <i>line-length</i>	[Function]
<code>get-output-stream-string</code> <i>string-output-stream</i>	[Function]
<code>with-input-from-string</code> (<i>var string</i> {keyword value}*) {declaration}* {form}*	[Macro]
<code>with-output-to-string</code> (<i>var [string]</i>) {declaration}* {form}*	[Macro]
<code>streamp</code> <i>object</i>	[Function]
<code>input-stream-p</code> <i>stream</i>	[Function]
<code>output-stream-p</code> <i>stream</i>	[Function]
<code>stream-element-type</code> <i>stream</i>	[Function]
<code>close</code> <i>stream</i> &key :abort	[Function]
<code>*readtable*</code>	[Variable]
<code>copy-readtable</code> &optional <i>from-readtable to-readtable</i>	[Function]
<code>readtablep</code> <i>object</i>	[Function]
<code>set-syntax-from-char</code> <i>to-char from-char</i> &optional <i>to-readtable from-readtable</i>	[Function]
<code>set-macro-character</code> <i>char function</i> &optional <i>non-terminating-p readtable</i>	[Function]
<code>get-macro-character</code> <i>char</i> &optional <i>readtable</i>	[Function]
<code>make-dispatch-macro-character</code> <i>char</i> &optional <i>non-terminating-p readtable</i>	[Function]
<code>set-dispatch-macro-character</code> <i>disp-char sub-char function</i> &optional <i>readtable</i>	[Function]
<code>get-dispatch-macro-character</code> <i>disp-char sub-char</i> &optional <i>readtable</i>	[Function]
<code>*prinescape*</code>	[Variable]
<code>*prinpretty*</code>	[Variable]
<code>*princircle*</code>	[Variable]
<code>*base*</code>	[Variable]
<code>*prinradix*</code>	[Variable]
<code>*princase*</code>	[Variable]
<code>*prinlevel*</code>	[Variable]
<code>*prinlength*</code>	[Variable]
<code>*prinarray*</code>	[Variable]
<code>read</code> &optional <i>input-stream eof-errorp eof-value</i>	[Function]
<code>*read-default-float-format*</code>	[Variable]
<code>read-preserving-whitespace</code> &optional <i>input-stream eof-errorp eof-value</i>	[Function]
<code>read-delimited-list</code> <i>char</i> &optional <i>input-stream</i>	[Function]
<code>read-line</code> &optional <i>input-stream eof-errorp eof-value</i>	[Function]
<code>read-char</code> &optional <i>input-stream eof-errorp eof-value</i>	[Function]
<code>unread-char</code> <i>character</i> &optional <i>input-stream</i>	[Function]
<code>peek-char</code> &optional <i>peek-type input-stream eof-errorp eof-value</i>	[Function]
<code>listen</code> &optional <i>input-stream</i>	[Function]

provide <i>module-name</i>	[Function]
require <i>module-name</i> &optional <i>pathname</i>	[Function]
trace { <i>function-name</i> }*	[Macro]
untrace { <i>function-name</i> }*	[Macro]
step <i>form</i>	[Macro]
time <i>form</i>	[Macro]
describe <i>object</i>	[Function]
inspect <i>object</i>	[Function]
room &optional <i>x</i>	[Function]
ed &optional <i>x</i>	[Function]
dribble &optional <i>pathname</i>	[Function]
apropos <i>string</i> &optional <i>package</i>	[Function]
internal-time-units-per-second	[Constant]
get-decoded-time	[Function]
get-universal-time	[Function]
get-internal-time	[Function]
sleep <i>seconds</i>	[Function]
decode-universal-time <i>universal-time</i> &optional <i>time-zone</i>	[Function]
encode-universal-time <i>second minute hour date month year</i> &optional <i>time-zone</i>	[Function]
lisp-implementation-type	[Function]
lisp-implementation-version	[Function]
machine-type	[Function]
machine-version	[Function]
machine-instance	[Function]
host-software-type	[Function]
host-software-version	[Function]
short-site-name	[Function]
long-site-name	[Function]
features	[Variable]
identity <i>object</i>	[Function]
(End of COMMON LISP summary.)	



Index

Index of Concepts

Index of Concepts

- Compatibility note 12, 19, 32, 38, 47, 53, 57, 72, 78, 82, 84, 86, 87, 90, 91, 94, 99, 102, 103, 106, 108, 110, 111, 118, 123, 124, 125, 129, 130, 135, 142, 143, 144, 163, 164, 166, 175, 176, 177, 182, 183, 186, 190, 201, 202, 220, 231, 234, 237, 238, 239, 249, 254, 261, 276, 293, 295, 296, 297
 Implementation note 12, 13, 14, 56, 95, 111, 125, 127, 128, 130, 138, 144, 145, 149, 175, 195, 240, 276, 290
 Query 32, 71, 108, 176, 195, 253, 254, 267, 282, 283, 284, 286, 295, 296
 Rationale 21, 40, 44, 69, 80, 121, 123, 125, 137, 199, 243, 245, 254, 296, 297
 ~% (new line) format directive 254
 ~& (fresh line) format directive 254
 ~ ((case conversion) format directive 256
 ~* (ignore argument) format directive 255
 ~< (justification) format directive 259
 ~<return> (ignore whitespace) format directive 254
 ~? (indirection) format directive 255
 ~~ (*Tilde*) format directive 254
 ~[(conditional) format directive 256
 ~^ (*loop escape*) format directive 260
 ~A (*Ascii*) format directive 252
 ~B (*Binary*) format directive 253
 ~C (*Character*) format directive 254
 ~D (*Decimal*) format directive 252
 ~E (*Exponential*) format directive 254
 ~F (*Floating-point*) format directive 253
 ~G (*Goto* argument) format directive 255
 ~O (*Octal*) format directive 253
 ~P (*Plural*) format directive 253
 ~R (*Radix*) format directive 253
 ~S (*S-expression*) format directive 252
 ~T (*Tabulate*) format directive 255
 ~X (*hexadecimal*) format directive 253
 ~{ (iteration) format directive 257
 ~| (new page) format directive 254
 " macro character 226
 # macro character 229
 ' macro character 225
 (macro character 223
) macro character 225
 , macro character 228
 ; macro character 225
 ' macro character 227
 | macro character 226
 A-list 185
 Access functions 200
 ADA 12, 58, 59
 ALGOL 28, 40, 86, 136
 APL 130
 Array 19
 predicate 55
 Association list 83, 185
 as a substitution table 183
 compared to hash table 189
 Atom
 predicate 53
 Bignum 11
 Bit string
 infinite 138
 integer representation 138
 Bit-vector
 predicate 54
 Byte 142
 Byte specifiers 142
 Car 18, 173
 Catch 93
 Cdr 18, 173
 Character
 predicate 54
 Character syntax 229
 Cleanup handler 94
 Comments 225
 Common data type
 predicate 55
 Compiled function
 predicate 55
 Complex number
 predicate 54
 Conditional
 and 58
 or 59
 during read 233
 Cons 18, 173
 predicate 53
 Constructor function 200
 Control structure 61
 Data type
 predicates 52
 Declaration
 declaration 105
 function 104
 function type 104
 ignore 105
 inline 105
 notinline 105
 optimize 105
 special 103
 type 104
 Declaration declaration 105
 Declarations 101
 Defstruct 199
 Denominator 12

- Device (pathname component) 264
- Directory (pathname component) 264
- Dotted list 173
- Dynamic exit 93
- Empty list
 - predicate 52
- Environment structure 61
- Extent 27
- False
 - when a predicate is 51
- Fixnum 11
- Floating-point number 13
 - predicate 54
- Flow of control 61
- Formatted output 251
- FORTRAN 2, 12, 15, 86, 136, 254
- Function
 - predicate 55
- Function declaration 104
- Function type declaration 104
- Hash table 189, 191
 - predicate 191
- Home directory 269
- Host (pathname component) 264
- Ignore declaration 105
- Implicit progn 61, 73, 74, 75, 76, 77, 81
- Indicator 107
- Init file 269
- Inline declaration 105
- Integer 11
 - predicate 53
- INTERLISP 1, 2, 3, 12, 32, 72, 87, 106, 108, 118, 129, 136, 163, 164, 175, 176, 190, 237, 293
- Iteration 80
- Keywords
 - for defstruct slot-descriptions 203
- LISP 1.5 86, 163
- Lisp Machine LISP 1, 2, 12, 19, 57, 78, 84, 87, 90, 91, 99, 108, 110, 111, 115, 124, 125, 129, 136, 163, 164, 175, 177, 190, 201, 202, 231, 234, 237, 276, 282, 286, 293, 295, 296, 297
- List 18, 173
 - predicate 53
 - See also dotted list
- List syntax 223
- Logical operators
 - on nil and non-nil values 58
- Logical pathnames 271
- MacLISP 1, 2, 12, 19, 21, 38, 47, 53, 82, 87, 94, 99, 102, 103, 104, 106, 108, 110, 118, 121, 123, 124, 125, 129, 136, 142, 143, 144, 159, 163, 166, 175, 182, 183, 186, 190, 211, 220, 231, 234, 237, 238, 239, 243, 249, 276, 291, 293
- Macro character 223
- Mapping 85
- Merging
 - of pathnames 265
 - sorted sequences 171
- Multiple values 89
 - returned by read-from-string 247
- Name (pathname component) 264
- Naming conventions
 - predicates 51
- NIL 1, 87, 108, 111, 136, 175, 177
- Non-local exit 93
- Notinline declaration 105
- Number 121
 - floating-point 13
 - predicate 53
- Numerator 12
- Optimize declaration 105
- Package
 - predicate 117
- Package cell 107
- Parsing 223
 - of pathnames 265
- PASCAL 23, 58, 59, 123
- PL/I 15, 136
- Plist 107
- Position
 - of a byte 142
- Predicate 51
- Predicates
 - true and false 51
- Print name 107, 110, 193
- Print-name
 - coercion to string 197
- Printed representation 219
- Printer 219, 238
- Property 107
- Property list 107
 - compared to association list 107
 - compared to hash table 189
- Pseudo-predicate 51, 151
- Querying the user 261
- Quote character 225
- Random-state
 - predicate 146
- Rank 19
- Rational
 - predicate 53
- Reader 219, 220
- Readable 234
 - predicate 235
- Record structure 199
- S-1 LISP 1, 2
- SCHEME 1
- Scope 27
- Set

- list representation 183
- Sets
 - bit-vector representation 138
 - infinite 138
 - integer representation 138
- Shadowing 28
- Sharp-sign macro characters 229
- Simple bit-vector
 - predicate 54
- Simple string
 - predicate 54
- Simple vector
 - predicate 54
- Size
 - of a byte 142
- Sorting 169
- Special declaration 103
- SPICE LISP 1, 108
- STANDARD LISP 2, 136
- String 193
 - predicate 54
- String syntax 226
- Structure 199
- Structured pathname components 265
- Substitution 182
- Symbol 9, 107
 - coercion to a string 193
 - coercion to string 197
 - predicate 53
- Symbol syntax 226

- Throw 93
- Tree 19
- True
 - when a predicate is 51
- Type (pathname component) 264
- Type declaration 104
- Type specifiers 31

- Unwind protection 94

- Vector
 - predicate 54
- Version (pathname component) 264

- Yes-or-no functions 261

Index of Variables and Constants

Index of Variables and Constants

- * 212
 - ** 212
 - *** 212
 - + 212
 - ++ 212
 - +++ 212
 - 212
 - / 212
 - // 212
 - /// 212
 - *base* 239, 248
 - *compile-file-set-default-pathname* 278, 290
 - *default-pathname-defaults* 267, 268, 269, 271
 - *error-output* 214
 - *evalhook* 209
 - *features* 234, 298
 - *load-pathname-defaults* 271, 278, 279, 290
 - *load-set-default-pathname* 278
 - *load-verbose* 278, 278
 - *macroexpand-hook* 100, 100
 - *package* 112, 115, 117, 221, 240, 248
 - *prinarray* 240, 241, 242, 248
 - *princase* 240, 240, 248
 - *princircle* 178, 238, 240, 248
 - *princescape* 207, 238, 239, 240, 248
 - *prinlength* 221, 233, 241, 242, 248
 - *prinlevel* 207, 234, 241, 242, 248
 - *prinpretty* 207, 238, 248
 - *prinradix* 239, 248
 - *query-io* 214, 261, 262
 - *random-state* 145
 - *read-default-float-format* 15, 239, 243
 - *readtable* 235, 235
 - *sample-variable* 5
 - *standard-input* 213, 243, 294
 - *standard-output* 213, 248, 251, 278, 293, 294, 295
 - *terminal-io* 214, 243, 248, 261
 - *trace-output* 214, 292, 293
- A
- B
- boole-1 139
 - boole-2 139
 - boole-and 139
 - boole-andc1 140
 - boole-andc2 140
 - boole-c1 139
 - boole-c2 139
 - boole-clr 139
 - boole-eqv 139
 - boole-ior 139
 - boole-nand 139
 - boole-nor 139
- boole-orc1 140
 - boole-orc2 140
 - boole-set 139
 - boole-xor 139
- C
- char-bits-limit 16, 149, 154
 - char-code-limit 149, 154
 - char-control-bit 156
 - char-font-limit 16, 149, 154
 - char-hyper-bit 156
 - char-meta-bit 156
 - char-super-bit 156
- D
- double-float-epsilon 147
 - double-float-negative-epsilon 147
- E
- F
- G
- H
- I
- internal-time-units-per-second 295, 296
- J
- K
- L
- lambda-list-keywords 47
 - least-negative-double-float 147
 - least-negative-long-float 147
 - least-negative-short-float 146
 - least-negative-single-float 146
 - least-positive-double-float 146
 - least-positive-long-float 147
 - least-positive-short-float 146
 - least-positive-single-float 146
 - long-float-epsilon 147
 - long-float-negative-epsilon 147
- M
- most-negative-double-float 147
 - most-negative-fixnum 11, 146
 - most-negative-long-float 147
 - most-negative-short-float 146
 - most-negative-single-float 146
 - most-positive-double-float 146
 - most-positive-fixnum 11, 49, 146

most-positive-long-float 147
most-positive-short-float 146
most-positive-single-float 146

N

nil 3, 29, 51, 211, 287

O

P

pi 29, 130, 211, 287

Q

R

S

sample-constant 5
short-float-epsilon 147
short-float-negative-epsilon 147
single-float-epsilon 147
single-float-negative-epsilon 147

T

t 49, 51, 211

U

V

W

X

Y

Z

Index of Keywords

Index of Keywords

A

:abort
 for close 217
 :append
 for if-exists option to open 275

B

:base
 for write 248
 for write-to-string 249

C

:conc-name
 for defstruct 202, 204
 :constructor
 for defstruct 202, 205, 207
 :contradictory-arguments
 for signal 286
 :count
 for delete 166
 for delete-if 166
 for delete-if-not 166
 for nsubstitute 167
 for nsubstitute-if 167
 for nsubstitute-if-not 167
 for remove 165
 for remove-if 165
 for remove-if-not 165
 for substitute 166
 for substitute-if 166
 for substitute-if-not 167
 :create
 for if-does-not-exist option to open 275

D

:default
 for type option to open 274
 :defaults
 for make-pathname 268
 :device
 for make-pathname 268
 :direction
 for open 273
 :directory
 for make-pathname 268

E

:end1
 for maxprefix 168
 for maxsuffix 168
 for mismatch 168
 for replace 165
 for search 169

for string-equal 194
 for string-greaterp 195
 for string-lessp 195
 for string-not-equal 195
 for string-not-greaterp 195
 for string-not-lessp 195
 for string/= 194
 for string< 194
 for string<= 194
 for string= 194
 for string> 194
 for string>= 194
 :end2
 for maxprefix 168
 for maxsuffix 168
 for mismatch 168
 for replace 165
 for search 169
 for string-equal 194
 for string-greaterp 195
 for string-lessp 195
 for string-not-equal 195
 for string-not-greaterp 195
 for string-not-lessp 195
 for string/= 194
 for string< 194
 for string<= 194
 for string= 194
 for string> 194
 for string>= 194
 :end
 for count 168
 for count-if 168
 for count-if-not 168
 for delete 166
 for delete-duplicates 166
 for delete-if 166
 for delete-if-not 166
 for fill 164
 for find 167
 for find-if 167
 for find-if-not 167
 for nstring-capitalize 197
 for nstring-downcase 197
 for nstring-upcase 197
 for nsubstitute 167
 for nsubstitute-if 167
 for nsubstitute-if-not 167
 for position 168
 for position-if 168
 for position-if-not 168
 for reduce 164
 for remove 165
 for remove-duplicates 166

for remove-if 165
 for remove-if-not 165
 for string-capitalize 196
 for string-downcase 196
 for string-upcase 196
 for substitute 166
 for substitute-if 166
 for substitute-if-not 167
 for with-input-from-string 216
 :error
 for if-does-not-exist option to open 275
 for if-exists option to open 274
 for signal 286

:eval-when
 for defstruct 207

F

:failed-assertion
 for signal 287
 :failed-handler
 for signal 287
 :from-end
 for count 168
 for count-if 168
 for count-if-not 168
 for delete 166
 for delete-if 166
 for delete-if-not 166
 for find 167
 for find-if 167
 for find-if-not 167
 for maxprefix 168
 for maxsuffix 168
 for mismatch 168
 for nsubstitute 167
 for nsubstitute-if 167
 for nsubstitute-if-not 167
 for position 168
 for position-if 168
 for position-if-not 168
 for reduce 164
 for remove 165
 for remove-if 165
 for remove-if-not 165
 for search 169
 for substitute 166
 for substitute-if 166
 for substitute-if-not 167

G

H

:host
 for make-pathname 268

I

:if-does-not-exist
 for load 278
 for open 275
 :if-exists

for open 274
 :include
 for defstruct 25,205
 :index
 for with-input-from-string 216
 :initial-element
 for make-list 177
 for make-sequence 162
 for make-string 195
 :initial-offset
 for defstruct 207
 :initial-value
 for reduce 164
 :input
 for direction option to open 273
 :invalid-form
 for signal 287
 :invalid-function
 for signal 287
 :invalid-variable
 for signal 287
 :invisible
 for defstruct slot-descriptions 203
 :io
 for direction option to open 273

J

K

:key
 for adjoin 184
 for count 168
 for count-if 168
 for count-if-not 168
 for delete 166
 for delete-if 166
 for delete-if-not 166
 for find 167
 for find-if 167
 for find-if-not 167
 for intersection 184
 for maxprefix 168
 for maxsuffix 168
 for member 183
 for member-if 183
 for member-if-not 183
 for merge 171
 for mismatch 168
 for nintersection 184
 for nset-difference 185
 for nset-exclusive-or 185
 for nsublis 183
 for nsubst 182
 for nsubst-if 182
 for nsubst-if-not 182
 for nsubstitute 167
 for nsubstitute-if 167
 for nsubstitute-if-not 167
 for union 184
 for position 168

- for position-if 168
- for position-if-not 168
- for remove 165
- for remove-if 165
- for remove-if-not 165
- for search 169
- for set-difference 185
- for set-exclusive-or 185
- for sort 169
- for stable-sort 169
- for sublis 183
- for subsetp 185
- for subst 182
- for subst-if 182
- for subst-if-not 182
- for substitute 166
- for substitute-if 166
- for substitute-if-not 167
- for union 184

L

M

- :missing-handler
- for signal 287

N

- :name
- for make-pathname 268
- :named
- for defstruct 205, 205
- :new-version
- for if-exists option to open 274

O

- :output-file
- for compile-file 290
- :output
- for direction option to open 273
- :overwrite
- for if-exists option to open 274

P

- :predicate
- for defstruct 205
- :prinarray
- for write 248
- for write-to-string 249
- :princase
- for write 248
- for write-to-string 249
- :princircle
- for write 248
- for write-to-string 249
- :prinescape
- for write 248
- for write-to-string 249
- :prinlength
- for write 248
- for write-to-string 249

- :prinlevel
- for write 248
- for write-to-string 249
- :prinpretty
- for write 248
- for write-to-string 249
- :prinradix
- for write 248
- for write-to-string 249
- :print-function
- for defstruct 23, 207
- :print
- for load 278
- :probe
- for direction option to open 273

Q

R

- :read-only
- for defstruct slot-descriptions 203
- :rehash-size
- for make-hash-table 190
- :rehash-threshold
- for make-hash-table 190
- :rename-and-delete
- for if-exists option to open 274
- :rename
- for if-exists option to open 274

S

- :set-default-pathname
- for compile-file 290
- for load 278
- :size
- for make-hash-table 190
- :start1
- for maxprefix 168
- for maxsuffix 168
- for mismatch 168
- for replace 165
- for search 169
- for string-equal 194
- for string-greaterp 195
- for string-lessp 195
- for string-not-equal 195
- for string-not-greaterp 195
- for string-not-lessp 195
- for string/= 194
- for string< 194
- for string<= 194
- for string= 194
- for string> 194
- for string>= 194
- :start2
- for maxprefix 168
- for maxsuffix 168
- for mismatch 168
- for replace 165
- for search 169

```

for string-equal 194
for string-greaterp 195
for string-lessp 195
for string-not-equal 195
for string-not-greaterp 195
for string-not-lessp 195
for string/= 194
for string< 194
for string<= 194
for string= 194
for string> 194
for string>= 194
:start
for count 168
for count-if 168
for count-if-not 168
for delete 166
for delete-duplicates 166
for delete-if 166
for delete-if-not 166
for fill 164
for find 167
for find-if 167
for find-if-not 167
for nstring-capitalize 197
for nstring-downcase 197
for nstring-upcase 197
for nsubstitute 167
for nsubstitute-if 167
for nsubstitute-if-not 167
for position 168
for position-if 168
for position-if-not 168
for reduce 164
for remove 165
for remove-duplicates 166
for remove-if 165
for remove-if-not 165
for string-capitalize 196
for string-downcase 196
for string-upcase 196
for substitute 166
for substitute-if 166
for substitute-if-not 167
for with-input-from-string 216
:stream
for write 248
:supersede
for if-exists option to open 275

T
:test-not
for adjoin 184
for assoc 186
for count 168
for delete 166
for delete-duplicates 166
for find 167
for intersection 184
for maxprefix 168
for maxsuffix 168
for member 183
for mismatch 168
for nintersection 184
for nset-difference 185
for nset-exclusive-or 185
for nsublis 183
for nsubst 182
for nsubstitute 167
for nunion 184
for position 168
for rassoc 187
for remove 165
for remove-duplicates 166
for search 169
for set-difference 185
for set-exclusive-or 185
for sublis 183
for subsetp 185
for subst 182
for substitute 166
for tree-equal 174
for union 184
:test
for adjoin 184
for assoc 186
for count 168
for delete 166
for delete-duplicates 166
for find 167
for intersection 184
for make-hash-table 190
for maxprefix 168
for maxsuffix 168
for member 183
for mismatch 168
for nintersection 184
for nset-difference 185
for nset-exclusive-or 185
for nsublis 183
for nsubst 182
for nsubstitute 167
for nunion 184
for position 168
for rassoc 187
for remove 165
for remove-duplicates 166
for search 169
for set-difference 185
for set-exclusive-or 185
for sublis 183
for subsetp 185
for subst 182
for substitute 166
for tree-equal 174
for union 184
:too-few-arguments
for signal 286
:too-many-arguments
for signal 286

```

:type

for make-pathname 268
for defstruct slot-descriptions 203
for defstruct 204
for open 247, 250, 273

U

:unbound-variable
for signal 287
:undefined-function
for signal 287
:unexpected-keyword
for signal 286
:unnamed
for defstruct 205

V

:verbose
for load 278
:version
for make-pathname 268

W

:wrong-type-argument
for signal 286

X

Y

Z

Index of Functions, Macros, and Special Forms

Index of Functions, Macros, and Special Forms

* 125
 *eval 209, 210, 210
 + 124
 - 124
 /= 122
 / 125
 1+ 125
 1- 125
 <= 122
 < 122, 152
 = 56, 121, 122, 152
 >= 122
 > 122

A

abs 128
 acons 107, 186
 acos 129
 acosh 130
 add-logical-pathname-host 272
 adjoin 179, 184
 alpha-charp 150
 alphanumericp 151
 and 40, 58, 77, 92
 append 177, 178, 228
 apply 23, 71, 91, 97, 99
 apropos 294
 aref 20, 67, 193
 arrayp 55
 ash 141
 asin 129
 asinh 130
 assert 285, 287
 assoc-if-not 186
 assoc-if 186
 assoc 185, 186
 atan 129
 atanh 130
 atom 53

B

back-translated-pathname 272
 bit-vector-p 54
 bit 67
 block 29, 30, 40, 48, 61, 79, 80, 82, 83, 87, 88, 92, 93
 boole 139
 bothcasep 151
 boundp 62, 63
 butlast 180
 byte-position 142
 byte-size 142
 byte 142

C

c...r 174
 caaaaar 67, 174
 caaadr 67, 174
 caaar 67, 174
 caadar 67, 174
 caaddr 67, 174
 caadr 67, 174
 caar 67, 174
 cadaar 67, 174
 cadadr 67, 174
 cadar 67, 174
 caddar 67, 174
 caddr 67, 174
 caddr 67, 174
 cadr 67, 174
 car 65, 67, 173, 176
 case 78, 78, 91, 283
 catch-all 29, 40, 91, 93
 catch 29, 40, 61, 91, 93
 cdaaar 67, 174
 cdaadr 67, 174
 cdaar 67, 174
 cdadar 67, 174
 cdaddr 67, 174
 cdadr 67, 174
 cdar 67, 174
 cddaar 67, 174
 cddadr 67, 174
 cddar 67, 174
 cdddar 67, 174
 cddddr 67, 174
 cddr 67, 174
 cdr 67, 173, 181
 ceiling 125, 135
 cerror 5, 284, 285, 285, 288
 char-bit 67, 157
 char-bits 149, 154
 char-code 149, 154
 char-downcase 151, 155, 196
 char-equal 57, 153, 194
 char-font 149, 154, 231
 char-greaterp 153
 char-int 152, 155
 char-lessp 153, 195
 char-name 156
 char-not-equal 153
 char-not-greaterp 153
 char-not-lessp 153
 char-upcase 151, 155, 196
 char/= 152
 char<= 152
 char< 152, 195
 char= 56, 152, 246

char>= 152
 char> 152
 char 67, 193
 character 38, 154
 characterp 54, 150
 check-type 285
 cis 128
 clear-input 246
 clear-output 250
 close 217, 272, 275
 clrhash 191
 code-char 154
 coerce 37, 38, 134, 154, 162, 163
 commonp 55
 compile-file 278, 279, 290, 290
 compile 289
 compiled-function-p 55
 compiler-let 74
 complex 15, 34, 137
 complexp 54, 122
 concatenate 162, 177
 cond 51, 59, 76, 78, 81, 92, 98
 condition-bind 282, 283
 condition-case 283
 condition-psetq 282
 conjugate 126
 cons 34, 174
 consp 53
 constantp 211
 copy-alist 178
 copy-list 178
 copy-readable 235
 copy-seq 161, 178
 copy-symbol 112
 copy-tree 178, 182
 cos 128
 cosh 130
 count-if-not 168
 count-if 168
 count 168

D

decf 70, 126
 declare 9, 40, 44, 48, 82, 101
 decode-universal-time 297
 defconstant 40, 48, 211, 291
 defmacro 36, 41, 47, 70, 76, 97, 98, 99, 100, 101, 102, 277, 291
 defparameter 48, 291
 defsetf 68, 70, 101, 102, 201, 291
 defstruct 11, 23, 25, 31, 36, 67, 170, 171, 174, 201, 233, 241, 291
 deftype 31, 36, 101, 102, 291
 defun 42, 46, 47, 75, 79, 99, 101, 102, 104, 277, 291
 defvar 48, 277, 291
 delete-duplicates 166
 delete-file 276
 delete-if-not 166
 delete-if 166
 delete 166, 181

denominator 134, 239
 deposit-field 67, 143, 144
 describe 293
 digit-charp 151, 155
 digit-weight 155
 directory-namestring 269
 directory 279
 disassemble 290
 do* 80, 80, 101
 do-all-symbols 101, 119, 294
 do-symbols 85, 101, 119
 do 29, 61, 64, 80, 80, 87, 92, 101
 documentation 37, 48, 49, 67, 99, 201, 291
 dolist 80, 84, 92, 101
 dotimes 80, 84, 92, 101
 dpb 67, 143, 143
 dribble 294

E

ed 294
 eighth 176
 elt 67, 161, 193
 encode-universal-time 297
 endp 83, 173, 175
 enough-namestring 269
 eq 55
 compared to equal 55
 eql 32, 56, 121, 123, 153
 equal 56, 153, 174, 194, 219
 equalp 49, 57
 error 5, 284, 286, 288
 eval-when 49, 98, 103, 207, 232, 289
 eval 91, 97, 209
 evalhook 100, 209, 210, 210
 evenp 122
 every 163
 exp 127
 export 118
 expt 127

F

fboundp 63, 63
 fceiling 136
 ferror 285, 288
 ffloor 136
 fifth 176
 file-author 277
 file-creation-date 277
 file-length 277, 277
 file-namestring 269
 file-position 275, 277
 fill 164
 find-if-not 167
 find-if 167
 find 167, 183, 185, 186
 finish-output 250
 first 173, 176
 flet 40, 63, 65, 75, 97, 101
 float-exponent 137
 float-radix 13, 137

float-sign 137
float-significand 137
float 130, 134
floatp 54, 122
floor 38, 89, 125, 135, 136
fmakunbound 63, 65
force-output 250
format 197, 250, 251, 284, 285
fourth 176
fresh-line 249, 254, 261
fround 136
fset 63, 65
ftruncate 136
funcall 23, 71, 91, 97, 100
function 30, 40, 42, 62
functionp 55

G

gcd 126
gensym 112, 112
gentemp 112, 112
get-decoded-time 296
get-dispatch-macro-character 237
get-internal-time 296, 296
get-macro-character 236
get-output-stream-string 215
get-properties 110
get-universal-time 296
get 66, 67, 108, 108, 109
getf 67, 70, 108, 109, 109, 110
gethash 67, 191
go 29, 40, 80, 81, 82, 84, 87, 89, 94
graphic-charp 150, 151, 156

H

hash-table-count 191
hash-table-p 55, 191
host-namestring 269
host-software-type 298
host-software-version 298

I

identity 299
if 40, 51, 59, 77, 77, 92, 98
imagpart 137
incf 70, 126
init-file-pathname 269, 298
input-stream-p 217
inspect 294
int-char 156
integer-length 141, 145
integerp 53, 122
intern 55, 111, 112, 117
internedp 118
intersection 184
isqrt 127

J

K

keywordp 113

L

labels 40, 42, 63, 65, 75, 97, 101
last 176
lcm 126
ldb-test 143
ldb 67, 143
ldiff 181, 184
length 161, 175
let* 40, 74, 83, 88, 91, 101
let 28, 39, 45, 73, 74, 75, 80, 83, 87, 88, 91, 101
lisp-implementation-type 297
lisp-implementation-version 297
list* 177
list-length 175
list 177
listen 246, 246
listp 53, 173
load 278, 278, 279, 291
locally 101, 103
log 127
logand 138
logandc1 139
logandc2 139
logbitp 141
logcount 141
logeqv 138
logior 138
lognand 139
lognor 139
lognot 140
logorc1 139
logorc2 139
logtest 140
logxor 138
long-site-name 298
loop 80, 80, 81, 83
lowercasep 151, 152, 155

M

machine-instance 298
machine-type 298
machine-version 298
macro-p 63, 63
macro 97, 101, 102
macroexpand-1 100, 100
macroexpand 63, 100
macrolet 40, 75, 97, 100, 101
make-array 33, 233, 241
make-broadcast-stream 215
make-char 154
make-concatenated-stream 215
make-dispatch-macro-character 237
make-echo-stream 215
make-hash-table 190
make-list 177
make-package 116
make-pathname 268
make-random-state 145

make-sequence 162
 make-string-input-stream 215
 make-string-output-stream 215
 make-string 195
 make-symbol 111
 make-synonym-stream 214, 214
 make-two-way-stream 215
 makunbound 62, 63, 65, 75
 map 38, 85, 97, 163
 mapc 85, 163
 mapcan 85
 mapcar 85
 mapcon 85
 maphash 191
 mapl 85, 163
 maplist 85
 mask-field 67, 143
 max 123
 maxprefix 168
 maxsuffix 168
 member-if-not 183
 member-if 183
 member 51, 183, 185
 merge-pathname-defaults 268
 merge 171
 min 124
 minusp 122
 mismatch 168
 mod 136
 multiple-value-bind 89, 90, 91, 101, 135
 multiple-value-call 35, 40, 70, 90, 91
 multiple-value-list 89, 90
 multiple-value-prog1 40, 72, 90, 91
 multiple-value 89, 91, 91, 92

N

name-char 156
 namestring 269
 nbutlast 180, 181
 nconc 86, 177, 178, 181, 228
 nintersection 184
 ninth 176
 not 52, 58
 notany 163
 notevery 163
 nreconc 178, 179, 181
 nreverse 82, 162, 170, 181
 nset-difference 185
 nset-exclusive-or 185
 nstring-capitalize 197
 nstring-downcase 197
 nstring-upcase 197
 nsublis 183
 nsubst-if-not 182
 nsubst-if 182
 nsubst 182
 nsubstitute-if-not 167
 nsubstitute-if 167
 nsubstitute 167
 nth 67, 118, 175, 176

nthcdr 176
 null 52, 58, 83
 numberp 53, 122
 numerator 134, 239
 nunion 184

O

oddp 122
 open ix, 22, 214, 217, 250, 265, 273, 276
 or 40, 59, 77, 92
 output-stream-p 217

P

package-name 117
 package 117
 packagep 55, 117
 pairlis 107, 186
 parse-namestring 267
 parse-number 247
 pathname-device 268
 pathname-directory 268
 pathname-host 268
 pathname-name 268
 pathname-plist 269
 pathname-type 268
 pathname-version 268
 pathname 266
 pathnamep 55, 268
 peek-char 245
 phase 128
 plusp 122
 pop 70, 180
 position-if-not 168
 position-if 168
 position 32, 168, 183, 186
 pprint 248
 prin1-to-string 197, 249
 prin1 12, 238, 248, 249, 252
 princ-to-string 197, 249
 princ 238, 248, 249, 252
 print 213, 219, 248
 probe-file 277
 prog* 87, 92, 101
 prog1 61, 72, 90, 91, 92
 prog2 61, 72
 prog 29, 80, 87, 92, 101
 progn 40, 47, 61, 72, 79, 80, 81, 91
 progv 40, 65, 75, 91
 provide 291
 psetf 68
 psetq 64, 81, 83
 push 70, 179
 pushnew 179, 184
 putf 70, 108, 109

Q

quote 40, 62, 63

R

random-state-p 55, 146

random 144
 rassoc-if-not 187
 rassoc-if 187
 rassoc 185, 187
 rational 38, 134
 rationalize 134
 rationalp 53, 122
 read-binary-object ix, 247, 250
 read-byte 247, 273, 274
 read-char-no-hang 246
 read-char 213, 245, 246, 273
 read-delimited-list 236, 244
 read-from-string 246
 read-line 243, 245, 249
 read-preserving-whitespace 243, 246
 read 7, 10, 21, 62, 110, 111, 213, 223, 240, 243, 243, 248,
 249, 252
 readtablep 55, 235
 realpart 137
 reduce 164
 remainder 136
 remf 70, 108, 110
 remhash 191
 remove-duplicates 166
 remove-if-not 86, 165
 remove-if 165
 remove 160, 165
 remprop 109, 110
 rename-file 276
 replace 161, 165
 require 291
 reset-fill-pointer 193
 rest 173, 176
 return-from 5, 30, 40, 48, 79, 80, 81, 89, 92, 94
 return 41, 48, 61, 79, 81, 82, 83, 84, 88, 92, 93, 119
 revappend 178, 179
 reverse 162
 room 294
 rotatef 69
 round 125, 135
 rplaca 65, 173, 181
 rplacd 173, 181

S

samepnamep 110
 sample-function 5
 sample-macro 5
 sample-special-form 5
 scale-float 137
 search 161, 169
 second 176
 set-char-bit 67, 157, 157
 set-difference 185
 set-dispatch-macro-character 237
 set-exclusive-or 185
 set-macro-character 236, 237
 set-syntax-from-char 235
 set 63, 64, 65, 287
 setf 64, 66, 68, 69, 70, 108, 109, 110, 126, 143, 157, 161,
 173, 174, 175, 179, 180, 181, 191, 194, 203, 276, 285, 291

setq 39, 40, 64, 65, 73, 74, 81, 83, 84, 92, 104
 seventh 176
 sgvref 67
 shadow 118
 shiftf 68
 short-site-name 298
 signal 281
 signum 128
 simple-bit-vector-p 54
 simple-string-p 54
 simple-vector-p 54
 sin 128
 sinh 130
 sixth 176
 sleep 297
 some 163
 sort 169
 special-form-p 63, 63
 sqrt 127, 292
 stable-sort 169
 standard-charp 55, 150
 step 210, 293
 stream-element-type ix, 217, 274
 streamp 55, 217
 string-capitalize 196, 197, 240, 256
 string-charp 55, 150, 193
 string-downcase 196
 string-equal 194
 string-greaterp 195
 string-left-trim 195
 string-lessp 195
 string-not-equal 195
 string-not-greaterp 195
 string-not-lessp 195
 string-right-trim 195
 string-trim 195
 string-upcase 196
 string/= 194
 string<= 194
 string< 194
 string= 194
 string>= 194
 string> 194
 string 197
 stringp 54, 193
 sublis 183
 subseq 67, 161
 subsetp 185
 subst-if-not 182
 subst-if 182
 subst 182, 183
 substitute-if-not 167
 substitute-if 166
 substitute 166, 182
 subtypep 52
 svref 67
 sxhash 191
 symbol-function 23, 63, 67
 symbol-package 113
 symbol-plist 67, 109, 269

symbol-print-name 110
 symbol-value 62, 65, 67, 117, 209, 287
 symbolp 53

Z
 zerop 122

T

tagbody 40, 80, 82, 83, 84, 87, 87, 88, 89
 tailp 184
 tan 128
 tanh 130
 tenth 176
 terpri 249, 254
 the 35, 40, 68, 106
 third 176
 throw 29, 40, 42, 61, 80, 81, 91, 95, 276, 281, 283
 time 293
 trace 214, 292
 translated-pathname 266, 272, 272
 tree-equal 57, 174
 truename 266, 269, 277, 279
 truncate 35, 125, 135, 136
 type-of 9, 38
 typecase 78, 91
 typep 9, 34, 37, 38, 52, 52, 200, 201, 205

U

unintern 118, 118, 119
 union 184
 unless 51, 59, 77, 91
 unread-char 243, 245
 untrace x, 292
 unwind-all ix, 29, 40, 93
 unwind-protect 29, 40, 91, 94, 275
 uppercasep 151, 155
 user-homedir-pathname 269, 298

V

values-list 90
 values 42, 61, 89, 89, 220
 vector-push 216
 vectorp 54

W

when 51, 59, 77, 77, 91
 with-input-from-string 216
 with-open-file 27, 214, 275, 276
 with-output-to-string 216
 write-binary-object ix, 247, 250
 write-byte 250, 273, 274
 write-char 213, 249, 273
 write-line 245, 249
 write-string 249
 write-to-string 249
 write 248, 249

X

Y

y-or-n-p 261
 yes-or-no-p 214, 261