


XEROX

Xerox Development Environment



Mesa Language Manual

**XDE3.0-3001
Version 3.0
November 1984**

PRELIMINARY

**Office Systems Division
Xerox Corporation
3450 Hillview Avenue
Palo Alto, California 94304**

Notice

This manual is the current release of the Xerox Development Environment (XDE) and may be revised by Xerox without notice. No representations or warranties of any kind are made relative to this manual and use thereof, including implied warranties of merchantability and fitness for a particular purpose or that any utilization thereof will be free from the proprietary rights of a third party. Xerox does not assume any responsibility or liability for any errors or inaccuracies that may be contained in the manual or have any liabilities or obligations for any damages, including but not limited to special, indirect or consequential damages, arising out of or in connection with the use of this manual or products or programs developed from its use. No part of this manual, either in whole or part, may be reproduced or transmitted mechanically or electronically without the written permission of Xerox Corporation.

Copyright © 1984 by Xerox Corporation.
All Rights Reserved.



Table of contents

1	Introduction	
1.1	Syntax notation	1-2
2	Basic data types and expressions	
2.1	A slice of Mesa code	2-1
2.1.1	Basic lexical structure	2-3
2.2	Simple declarations	2-4
2.3	The fundamental operations: assignment, equality and inequality	2-4
2.4	Basic types	2-5
2.4.1	The numeric types INTEGER and CARDINAL	2-5
2.4.2	Type BOOLEAN	2-6
2.4.3	Type CHARACTER	2-7
2.4.4	Escape conventions for literals	2-7
2.4.5	The numeric types LONG INTEGER and LONG CARDINAL	2-7
2.4.6	Type REAL	2-8
2.4.7	Relations among basic types	2-9
2.4.8	Long to short conversion	2-11
2.4.9	Predeclared identifiers	2-12
2.5	Expressions	2-12
2.5.1	Numeric operators	2-13
2.5.2	Relational operators	2-16
2.5.3	BOOLEAN operators	2-17
2.5.4	Assignment expressions	2-18
2.5.5	Operator precedence	2-19
2.5.6	Function-like operators	2-19
2.5.7	Function-like operators on types	2-20
2.6	Initializing variables in declarations	2-20

Table of contents

2.6.1	Compile-time constants	2-21
2.7	More general declarations	2-22
3	Common constructed data types	
3.1	The element types	3-3
3.1.1	Enumerated types	3-4
3.1.2	Subrange types	3-8
3.2	Arrays	3-11
3.2.1	Declaration of arrays	3-13
3.2.2	Array constructors	3-15
3.2.3	Keyword array constructors	3-16
3.3	Records	3-17
3.3.1	Field lists	3-18
3.3.2	Declaration of records	3-19
3.3.3	Qualified references	3-21
3.3.4	Record constructors	3-22
3.3.5	Default field values	3-24
3.3.6	Extractors	3-26
3.3.7	Machine-dependent records	3-28
3.4	The types <code>POINTER</code> and <code>LONG POINTER</code>	3-30
3.4.1	Constructing pointer types	3-32
3.4.2	Pointer operations	3-33
3.4.3	Long pointers	3-35
3.4.4	Automatic dereferencing	3-36
3.5	Type determination	3-37
3.5.1	Type conversion	3-38
3.5.2	Balancing *	3-40
3.5.3	Free conformance *	3-41
3.6	Determination of representation *	3-42
3.7	Extended defaults	3-45
3.8	The null value <code>NIL</code>	3-47
4	Ordinary statements	
4.1	Assignment statements	4-2
4.1.1	Assignment expressions	4-2
4.1.2	Restrictions on assignment	4-3
4.2	<code>IF</code> statements	4-3
4.2.1	<code>IF</code> expressions	4-4
4.3	<code>SELECT</code> statements	4-5
4.3.1	Forms and options for <code>SELECT</code>	4-6

	4.3.2	The NULL statement	4-7
	4.3.3	SELECT expressions	4-8
4.4		Blocks	4-8
	4.4.1	GOTO statements	4-9
	4.4.2	OPEN clauses	4-12
4.5		Loop statements	4-14
	4.5.1	Loop control	4-15
	4.5.2	GOTOS, LOOPS, EXITS, and loops	4-18
5		Procedures	
5.1		Procedure types	5-3
	5.1.1	Procedure values and compatibility *	5-5
5.2		Procedure calls	5-6
	5.2.1	Arguments and parameters	5-8
	5.2.2	Termination and results	5-8
5.3		Procedure bodies	5-9
	5.3.1	RETURN statements	5-10
	5.3.2	Operations which deal with intact parameter records	5-11
	5.3.3	Defaults in argument and result records	5-12
5.4		A package of procedures	5-13
	5.4.1	The example	5-15
	5.4.2	Invoking procedures in other modules	5-16
5.5		Nested procedures	5-17
	5.5.1	Scopes defined by procedures	5-18
5.6		Inline procedures *	5-18
6		Other data types and storage management	
6.1		Strings	6-2
	6.1.1	String literals and string expressions	6-3
	6.1.2	Declaring strings	6-4
	6.1.3	Long strings	6-5
6.2		Array descriptors	6-5
	6.2.1	Array descriptor types	6-6
	6.2.2	Long descriptors	6-8
6.3		Base and relative pointers	6-9
	6.3.1	Syntax for base and relative pointers	6-9
	6.3.2	A relative pointer example	6-10
	6.3.3	Relative pointer types	6-11
	6.3.4	Relative array descriptors	6-12
6.4		Variant records	6-13

Table of contents

6.4.1	Declaring variant records	6-15
6.4.2	Bound variant types	6-18
6.4.3	Accessing entire variant parts, and variant constructors	6-18
6.4.4	Accessing components of variants	6-20
6.4.5	Defaults and variant records	6-25
6.5	Sequences	6-25
6.5.1	Defining sequence types	6-26
6.5.2	MACHINE DEPENDENT sequences	6-27
6.5.3	Allocating sequences	6-28
6.5.4	Operations on sequences	6-29
6.5.5	StringBodies and TEXT	6-30
6.6	Dynamic storage allocation	6-30
6.6.1	Zones	6-30
6.6.2	Allocating storage	6-31
6.6.3	Releasing storage	6-32
6.6.4	Implementing uncounted zones	6-33
7	Modules, programs, and configurations	
7.1	Interfaces	7-1
7.2	The fundamentals of Mesa modules	7-4
7.2.1	Including modules: the DIRECTORY clause	7-5
7.2.2	Accessing items from an included module	7-7
7.2.3	Scopes for identifiers in a module	7-9
7.2.4	Implications of recompiling included modules	7-10
7.3	DEFINITIONS modules	7-12
7.3.1	READONLY variables	7-13
7.3.2	Default fields in interfaces	7-14
7.3.3	Inline procedures in interfaces	7-15
7.3.4	Usage hints for INLINE procedures in interfaces *.	7-16
7.4	PROGRAM modules: IMPORTS and EXPORTS	7-17
7.4.1	IMPORTS , interface types and interface records	7-17
7.4.2	Importing program modules	7-18
7.4.3	Exporting interfaces and program modules.	7-20
7.4.4	IMPORTS in DEFINITIONS modules and implicitly imported interfaces	7-21
7.5	Access control: PUBLIC and PRIVATE	7-22
7.5.1	Access attributes in declarations	7-23
7.5.2	Access attributes in TYPE definitions	7-24
7.5.3	Default global access	7-25
7.5.4	Accessing PRIVATE names of other modules *.	7-25
7.6	Exported (opaque) types	7-25
7.6.1	Interface modules	7-26

7.6.2	Client modules	7-27
7.6.3	Implementation modules	7-27
7.7	Dot notation and interface items	7-28
7.8	The Mesa configuration language, an introductory example	7-29
7.8.1	Lexicon: a module implementing <i>LexiconDefs</i>	7-30
7.8.2	<i>LexiconClient</i> : a client module	7-32
7.8.3	Binding, loading, and running a configuration: an overview	7-33
7.8.4	A configuration description for running <i>LexiconClient</i>	7-33
7.9	C/Mesa: syntax and semantics	7-34
7.9.1	IMPORTS, EXPORTS, and DIRECTORY in C/Mesa	7-35
7.9.2	Explicit naming, IMPORTS, and EXPORTS *	7-36
7.9.3	Default names for interfaces and instances *	7-38
7.9.4	Multiple exported interfaces from a single component *	7-38
7.9.5	Multiple components implementing a single interface *	7-40
7.9.6	Nested (local) configurations	7-41
7.9.7	Package creation: EXPORTS ALL	7-41
7.10	Loading and running modules and configurations	7-42
7.10.1	Making copies of modules	7-42
7.10.2	How the loader binds interfaces	7-43
7.10.3	STARTing, STOPping, and RESTARTing module instances	7-43
7.10.4	Loading and starting configurations, control modules	7-44
8	Signaling and signal data types	
8.1	Declaring and generating SIGNALS and ERRORS	8-2
8.1.1	ERROR in expressions	8-3
8.2	Control of generated signals	8-4
8.2.1	Preparing to catch signals: catch phrases	8-4
8.2.2	The scope of variables in catch phrases	8-5
8.2.3	Catching signals	8-6
8.2.4	RETRY and CONTINUE in catch phrases	8-10
8.2.5	Resuming from a catch phrase: RESUME	8-11
8.3	Signals within signals	8-12
9	Processes and concurrency	
9.1	Concurrent execution, FORK and JOIN	9-1
9.1.1	A process example	9-1
9.1.2	Process language constructs	9-2
9.2	Monitors	9-4
9.2.1	An overview of monitors	9-4
9.2.2	Monitor locks	9-5

Table of contents

9.2.3	Declaring monitor modules, ENTRY and INTERNAL procedures	9-6
9.2.4	Interfaces to monitors	9-7
9.2.5	Interactions of processes and monitors	9-8
9.3	Condition variables	9-8
9.3.1	Wait, notify, and broadcast	9-8
9.3.2	Condition variable timeouts	9-11
9.4	More about monitors	9-12
9.4.1	The LOCKS clause	9-12
9.4.2	Monitored records	9-13
9.4.3	Monitors and module instances.	9-13
9.4.4	Multi-module monitors	9-14
9.4.5	Object monitors	9-16
9.4.6	Explicit declaration of monitor locks	9-17
9.4.7	Inline ENTRY procedures.	9-17
9.5	Signals	9-17
9.5.1	Signals and processes	9-17
9.5.2	Signals and monitors	9-17
9.6	Initialization	9-19

Appendices

A	Pronouncing Mesa	A-1
B	Programming conventions	
B.1	Names	B-1
B.1.1	Capitalization	B-1
B.1.2	Qualification	B-2
B.1.3	Module naming	B-3
B.2	Types	B-4
B.3	Exceptions: SIGNALS and ERRORS	B-5
B.3.1	General	B-5
B.3.2	In DEFINITIONS modules	B-5
B.3.3	In PROGRAM modules	B-6
B.4	Module histories	B-6
B.5	Documentation of definitions modules	B-7
B.6	Module organization	B-7
B.7	Layout	B-7
B.8	Spaces	B-8
B.9	Miscellaneous	B-9

C	Mesa machine dependencies	
C.1	Numeric limits	C-1
C.2	ASCII character set and ordering of character values	C-1
D	Binder extensions	
D.1	Code packing	D-1
	D.1.1 Syntax	D-1
	D.1.2 Restrictions	D-2
D.2	External links	D-2
	D.2.1 Syntax	D-3
	D.2.2 Restrictions	D-3
E	Mesa reserved words	E-1
F	Collected grammar	F-1
Index	I-1



Introduction

This manual concentrates on the Mesa programming language. Mesa is really a programming system of which the language is but one part. Other components of the system are documented separately, as are the details of preparing, compiling, debugging and running Mesa programs.

Each chapter of this manual discusses some aspect of the language, using examples as well as descriptions of semantics and syntax. The chapters emphasize different language features and provide different levels of detail. The complete treatment of some features requires more than one chapter. Generally, earlier chapters introduce topics, and later ones supply additional detail. Titles of chapters, sections and subsections indicate the language issues with which they deal.

In each major section, information is presented at three levels:

- (1) Ordinary usage (motivation, forms and semantics), frequently with examples.
- (2) Syntax equations (when appropriate).
- (3) Fine points (if applicable): restrictions, special cases, references to later material, precise semantics, etc.

Level (1) is intended to offer a basic understanding of Mesa. Reading only first level material should be adequate to begin programming in the language. Levels (2) and (3) supply more detail and provide information about the full power of Mesa.

As a rule, these levels of discourse occur separately and in the indicated order. A section with a heading followed by an asterisk (*) deals with specialized material that can be skimmed or skipped entirely on first reading. Occasionally, fine points or syntactic details are presented within first-level material. The reader will be able to distinguish between levels by their appearance. Fine points are written in a small font, like this. Syntax equations and syntactic categories appear in the following font: **FontForSyntax**.

Any italicized word or phrase is important. If a Mesa technical term is being introduced, it will be in italics; if a term is used before being defined, it will be italicized to warn the reader that it should not be taken lightly and that it has a particular meaning in Mesa. Occurrences of a technical term, once defined, are not distinguished. Lastly, names

appearing in programs are italicized in both the program text itself and the explanations of that text.

Programming examples are indented relative to the surrounding text.

1.1 Syntax notation

Mesa's grammar is described by syntax equations written using a variation of Backus-Naur Form (or *BNF*). For those unfamiliar with BNF, an explanation follows. Reading and understanding that explanation is imperative for full use of this manual; in a first reading, details of the syntax equations can safely be skipped. Those familiar with BNF should scan this section to discover the particular variation being used.

An individual syntax equation defines a portion of the Mesa grammar. It specifies a rule for forming some class of phrases in the language. A *phrase class* has a name, e.g., **Program**, and is defined by one or more syntax equations. Phrase names are always printed in the syntax font when their use is meant to be technically accurate. For example, an **OctalDigit**, which can be any of 0, 1, 2, . . . , 7, is defined by the equation:

$$\text{OctalDigit} \quad ::= \quad 0|1|2|3|4|5|6|7$$

Each equation consists of a phrase name on the left, followed by the operator ::= (pronounced "is defined to be"), in turn followed by a *formation rule* for that phrase class. A formation rule consists of one or more *alternatives*, separated by the syntactic operator vertical bar, | (pronounced "or"). The ordering of alternatives is not critical. In the definition of **OctalDigit**, "3" is an alternative.

Each alternative is a sequence of *symbols*, where a symbol is either a phrase name (in the syntax font) or a syntactic literal. In a syntax equation, a literal symbol stands for itself. The *reserved words* of Mesa, such as **BEGIN**, appear as literals; they are always written using upper-case characters in the font shown. The digits 0, 1, 2, etc. and special characters, such as =, + and ←, also are used to form literal symbols. Some composite symbols are formed from more than one special character, e.g., =>. Spaces in syntax equations are used only to separate the items in the rules and have no special significance.

The phrase name **empty** is often used as one of the alternatives in a formation rule. It means that the rule permits an "empty" phrase as one of its alternatives (i.e., an actual phrase is optional; it may or may not occur in the result of applying the formation rule).

Comments embedded in syntax rules are preceded by a double dash, --, and appear to the right, e.g.,

$$\text{Digit} \quad ::= \quad \text{OctalDigit} | 8 | 9 \quad \text{-- a decimal digit is an OctalDigit an 8 or a 9}$$

Often, only part of the total definition of a phrase class is given. To indicate that there are other ways of forming phrases of that class, an ellipsis (...) is used as an alternative within the rule. The definition of **Statement** is distributed throughout much of the manual in this way. When a certain statement form, such as the **AssignmentStmnt**, is being discussed, the following partial rule appears:

Statement ::= **AssignmentStmt** | ... -- *this is just an example*

One can read this as, "A **Statement** is defined to be an **AssignmentStmt**, among other things."

Within a single alternative, the order of symbols is important. The alternative acts as a "template" for forming an actual phrase; literal names and literal characters are *copied*, while *substitutions* are made for the phrase names. Consider the following example:

ReturnStmt ::= **RETURN** | **RETURN Constructor**

("A **ReturnStmt** is defined to be **RETURN** or **RETURN** followed by a **Constructor**.") The second alternative means that **RETURN** and some actual phrase defined by **Constructor** occur in exactly that order.

Syntax equations can indicate recursive substitution; for example:

IdList ::= **identifier** | **identifier , IdList**

In a Mesa program, an **identifier** is basically a name. This equation defines an **IdList** to be a list of one or more names, with commas separating them if there is more than a single name in the list.

Fine points:

This result is explained as follows. The formation rule for **IdList** consists of two alternative rules:

Rule 1: (First alternative) "An **IdList** is defined to be an **identifier**," i.e., any one *name* can replace an **IdList**.

Rule 2: (Second alternative) "An **IdList** is defined to be an **identifier** followed by a comma followed by another **IdList**," i.e., *name, IdList* can replace an **IdList**.

To derive a single name, use Rule 1 as shown below. (Note: The substitutions are emphasized by writing them in *italics*.)

IdList ::= *name* (by Rule 1)

To derive two names separated by a comma:

IdList ::= *name, IdList* (by Rule 2)
name, name (by Rule 1)

To derive three names separated by commas:

IdList ::= *name, IdList* (by Rule 2)
name, name, IdList (by Rule 2)
name, name, name (by Rule 1)

To derive *n* names separated by commas, use Rule 2 *n* - 1 times and then use Rule 1.

The following syntax equation also relies on recursion:

StmtSeries ::= **empty** | **Statement** | **Statement ; StmtSeries**

The equation is read as, "A **StmtSeries** is defined to be empty, or a single statement, or a series of statements separated by semicolons; the last statement may be followed by a semicolon."

Fine point:

A trailing semicolon is possible because:

- 1) A **StmtSeries** may take the form specified by the third alternative, "**Statement ; StmtSeries**."
- 2) After some number of further substitutions using the third alternative, the recursive reference to **StmtSeries** may take the "empty" form, i.e., "... **Statement ; empty**."
- 3) **empty** is replaced by nothing at all, i.e., "... **Statement ;**".

Commas and semicolons are used as major separators for a variety of constructs in Mesa. To distinguish between such constructs, a convention is adopted that the suffix "List" on a phrase name implies a sequence separated by commas, while "Series" implies a sequence separated by semicolons. This convention is reflected in the phrase names **IdList** and **StmtSeries** above.

Mesa is a living language that has undergone many changes since its initial implementations. Extensions and refinements continue to be made. Consequently, the BNF found in this manual may not be an exact copy of the current Mesa grammar, but it is a very close approximation. The "official" grammar used by the parser of the Mesa 11.0 compiler has been reproduced in Appendix F.



Basic data types and expressions

This chapter discusses how to declare, initialize and assign values to variables. It also describes the basic types for numeric, character and boolean data, as well as the operators used to construct expressions having these types.

The Mesa language is *strongly typed*. The programmer is given a collection of predefined types and the ability to construct new ones, and is encouraged to choose or invent suitable types for each particular application. Every variable and constant in a Mesa program must be declared to have one of these types; every constant has a type; and every expression has a type derived from its components and context. All types can be deduced by static analysis of the program, and the language requires that each value be used in a way consistent with its type according to rules specified here and in chapter 3. The type of an object determines its representation and structure as well as the set of applicable operations. In addition, the type system can be used to partition the universe of objects and avoid confusion, even among classes of objects that are represented identically.

2.1 A slice of Mesa code

The example below is an excerpt from a Mesa program. It assigns to *gcd* the greatest common divisor (GCD) of a pair of integers, *m* and *n* (where *m*, *n* and *gcd* are integer variables in the program from which this excerpt was taken; we assume their values need not be preserved). The example uses the Euclidean Algorithm for finding the GCD of two numbers and works as follows:

If both *m* and *n* are zero, the GCD is zero (by convention).

Otherwise, repeat the following until *n* is zero: find the remainder of dividing *m* by *n*; set *m* to the value of *n*; then set *n* to the remainder. The final value of *m* is the GCD of the original *m* and *n* except that it may be negative; taking its absolute value gives the GCD.

Example 1. Slice of Mesa Code Using the Euclidean Algorithm

```

-- Given are integers m and n, which can be altered.           (1)
IF m = 0 AND n = 0 THEN gcd ← 0 -- by convention              (2)
ELSE                                                            (3)
  BEGIN                                                        (4)
    r: INTEGER;                                               (5)
    UNTIL n = 0                                               (6)
    DO                                                         (7)
      r ← m MOD n; -- r gets remainder of m/n                 (8)
      m ← n; n ← r; -- update variables                       (9)
    ENDLOOP;                                                  (10)
    gcd ← -- in case one of m or n was negative -- ABS[m];   (11)
  END;                                                         (12)

```

The example contains twelve lines of source code, including comments. The numbers in parentheses at the right side are for reference only and are not part of the source code. Comments begin with the symbol "--" and terminate at line endings. They may also be completely embedded within lines, in which case they both begin and end with "--".

Line (2) begins an IF statement that uses the values of m and n to select between two alternatives. If both values are zero, the assignment statement following THEN is executed; it assigns the value 0 to gcd (the character "←" is Mesa's assignment operator). If either is nonzero, the assignment is skipped and the compound statement following ELSE (lines (4) through (12) inclusive) is executed. (Distinguishing the two cases is actually unnecessary, but doing so illustrates more features of Mesa.)

The second alternative is a block, a series of declarations followed by a series of statements, all bracketed by "BEGIN" and "END." Line (5) declares a variable r of type INTEGER for use within that block. A semicolon separates the declaration from the statements that follow it.

The iteration in the algorithm is performed by the loop (UNTIL $n=0$ DO ...ENDLOOP), which contains three embedded assignment statements. The loop repeats until n is equal to zero. If it is zero at the outset, the embedded statements are not executed at all. Statements are separated by semicolons. A semicolon at the end of a statement series that is embedded in another statement (such as the series in the loop) is optional; it is permissible to write a semicolon after every statement in the series.

Within the loop, line (8) assigns to r the value of the expression " $m \text{ MOD } n$," which gives the remainder of dividing m by n . Line (9) updates m to contain the previous value of n and then updates n for the next iteration, if any. Control transfers from the end of the loop, line (10), back to line (6), where the new value of n is tested. If it is not zero, the loop is repeated; otherwise, execution continues with the first statement after the loop, line (11).

When control reaches the assignment statement in line (11), m either has its original value (if n was zero) or contains the value n had just before it became zero. The expression "ABS [m]" has the form used for calling a function and passing it one or more arguments; square brackets enclose the argument list. Normal parentheses, "(" and ")," are used only for nested expressions, e.g., " $a*(b+c/(d-e)*f)$." The assignment places the absolute value of m into gcd ; this is the correct result. At this point, the reader is urged to trace through

the example with initial values for m and n of 15 and 12, respectively; the result should be $gcd=3$.

2.1.1 Basic lexical structure

The names gcd , m , n and r in the example are called **identifiers**. The general form of an **identifier** is given by the following (informal) syntax:

An **identifier** is a sequence consisting of any mixture of upper-case letters, lower-case letters or digits, the first of which is a letter. Upper and lower case letters *are different* and do distinguish identifiers.

The following, valid identifiers are all distinct:

aBc Abc DiskCommandWord displayVector mach1 x32y40

Certain identifiers consisting entirely of capital letters are reserved for use by the Mesa language. Some, such as **IF**, are punctuation symbols; others name built-in types, such as **INTEGER**, or functions, such as **ABS**. All such words that have special meaning and are not to be defined by the programmer are called *reserved words*. It is legal for the programmer to use fully capitalized identifiers, but he risks a clash with a reserved word (possibly a new one in some future version of the language). This can be easily avoided by including at least one digit or lower case letter in any identifier. Appendix E lists the current set of reserved words.

Mesa uses the blank (or space) character to separate basic *lexical units* of the language (such as reserved words and identifiers). *Blanks are significant separators of lexical units*. They may not be embedded in identifiers, composite symbols (such as $>=$), or numeric literals (such as 1000). Blanks are meaningful in **STRING** constants (§ 6.1.1), and there is a **CHARACTER** constant for space (§ 2.4.3). As a separator, any sequence of contiguous blanks is equivalent to a single blank. A **TAB** character also behaves exactly as a blank when used as a separator.

A carriage-return character behaves as a blank for separating lexical units also, but it has one extra function: if the last part of a line is a comment, the carriage return acts as the terminator of that comment. Multiline comments (those containing carriage returns) may be inserted into source code by beginning each new line with "--." Long comments (either containing carriage returns or not) may be bracketed with matched pairs of double angle brackets (<< and >>). All characters within the brackets are ignored, with the exception of nested long comment brackets. Care should be taken to insure that any code being commented out does not contain any unmatched pairs of angle brackets within comments initiated by "--." Line breaks have no significance as statement separators. For example, the single loop statement in the example extends over a number of lines, and a semicolon is used to separate two statements in a series.

Semicolons are used for separating declarations, series of declarations from following statements, and statements in a series from one another. Semicolons act as a separator between declarations, rather than being required after each declaration. *They cannot be used with abandon, however; care is necessary when writing IF statements (§ 4.2.1) or SELECT statements (§ 4.3.1)*. Multiple statements can be written on a single line, separated by semicolons.

2.2 Simple declarations

The example (Euclidean Algorithm) contains the following declaration:

```
r: INTEGER;
```

This declares r to be a variable of type `INTEGER` (§ 2.4.1), one of Mesa's built-in types. More than one variable can be declared at the same time. For instance,

```
x, y, divisor: INTEGER;
```

declares identifiers x , y and $divisor$ as variables of type `INTEGER`. These examples reflect the two primary purposes of every declaration:

to designate one or more identifiers as variables, and

to specify their type.

A declaration always begins with a single identifier or a list of identifiers. Conventionally, "list" is used to denote a single item as well as multiple items separated by commas. An *identifier list* (`IdList`) is defined as follows:

```
IdList      ::= identifier |
              identifier , IdList
```

A declaration begins with an `IdList` followed by a *colon*. The colon is followed by a *type specification* (`INTEGER`, for instance, is a type specification).

2.3 The fundamental operations: assignment, equality and inequality

The example contains the following five assignment statements:

```
gcd ← 0
r ← m MOD n
m ← n
n ← r
gcd ← ABS[m]
```

An assignment statement has the following syntax:

```
AssignmentStmt ::= LeftSide ← RightSide | ...
LeftSide       ::= identifier | ...    -- plus forms for array indexing, etc.
RightSide      ::= Expression
```

The `RightSide` may be any expression (§ 2.5) provided that its type *conforms* to that of the `LeftSide`. "Conforms" is defined in subsection 2.4.6 and is discussed further in section 3.5; for now, it can be taken to mean: "is the same as." The `LeftSide` may be a simple variable or a component of an aggregate variable (such as an element of an array). In any event, a `LeftSide` denotes a variable, something capable of receiving values. A `LeftSide` cannot, for example, be a constant, while a `RightSide` can.

The assignment operation (\leftarrow), the equality operation ($=$) and the inequality operation (\neq) are called the *fundamental operations*. They can be applied to values of most types (including, for instance, entire arrays). The rules governing which pairs of operands may be used in a fundamental operation are detailed in section 3.5.

2.4 Basic types

The types of variables in a Mesa program fall into two broad classifications, *user-defined* types and *built-in* types. Chapter 3 describes how a programmer can define new data types using *type constructors*; this section discusses the basic, built-in types. These include several numeric types (INTEGER, LONG INTEGER, NATURAL, CARDINAL, LONG CARDINAL and REAL), a type for logical values (BOOLEAN), and a type for individual character values (CHARACTER). The built-in type STRING (for sequences of characters) is described in chapter 6.

2.4.1 The numeric types INTEGER, CARDINAL, and NATURAL

Mesa provides two single word numeric types, one with values ranging over the signed integers; the other, over the unsigned integers. Neither type completely mirrors the corresponding mathematical abstraction (the integers \mathbb{Z} or the natural numbers \mathbb{N} , respectively) because a finite representation is used for values of each type. The range of the type INTEGER is (approximately) symmetric about zero, and values of type INTEGER are represented as *signed* numbers. The range of the type CARDINAL is some finite interval of the natural numbers that includes zero, and values of type CARDINAL are represented as *unsigned* numbers. "Signed" and "unsigned" are not types; rather, they describe the *machine representation* of a numeric value. There is an additional type, NATURAL, whose range of legal values is the intersection of that for INTEGER and CARDINAL.

The programmer must choose an appropriate type for each numeric variable. CARDINALS offer a somewhat greater positive range than INTEGERS, and this is significant in a few applications, e.g., those that manipulate addresses that might be the same size as the word size. More importantly, declaring a variable to have type CARDINAL asserts that its value is always nonnegative; the compiler can use such assertions to perform more checking and to generate better code. Programmers are encouraged to declare as much information about each variable as possible; the ranges of numeric variables can be further constrained by using subrange types (§ 3.1.2).

The types INTEGER and CARDINAL are distinct and not interchangeable. They are, however, closely related. Mesa allows most combinations of these types to occur within assignments and arithmetic expressions (but not relational expressions). Care is necessary to avoid ambiguity and failures of representation when values with different representations are mixed. This is discussed further in subsection 2.4.6 and sub-subsection 2.5.1.1.

The major advantage of NATURAL is that it can be compared to either INTEGER or CARDINAL (§ 2.5.2). It also requires fewer bits to represent, so a NATURAL and a BOOLEAN can be packed into a single word inside a record.

2.4.1.1 Numeric literals

A *numeric literal* is an instance of the phrase class **number**, defined as follows:

A **number** is a sequence of digits. The digits may optionally be followed by the letter B, H or D, which in turn may optionally be followed by another sequence of digits denoting a scale factor. No spaces are allowed within numeric literals.

If D is specified explicitly, or if neither B, D, nor H appears, the **number** is treated as decimal. The letter B means the **number** is octal (radix 8). A scale factor indicates the number of zeros to be appended to the first sequence of digits; *the scale factor itself is always a decimal number*. The syntax for numeric literals in base 16 notation (hexadecimal) is one or more hexadecimal digits (0-9, A-F) followed by the character H. The first character must be a valid decimal digit (0-9) or the literal will scan as a valid Mesa identifier. Lowercase letters may also be used for a through f and for the suffix. The literals below all denote the same value:

6400 6400D 64D2 14400B 144B2 1900H 19H2

A numeric literal always denotes a non-negative number (i.e., -5 is considered to be an expression in which the unary negation operator is applied to the literal 5 to produce an **INTEGER** value). To be valid in a context requiring a **CARDINAL**, the value of the literal must be a valid **CARDINAL** number. Similarly, if an **INTEGER** is required by context, the value must be a valid (positive) **INTEGER**. (§ 2.4.5 for more details). Note that literal values of type **REAL** have a syntax discussed in subsection 2.4.6.

2.4.2 Type **BOOLEAN**

A **BOOLEAN** value can be either **TRUE** or **FALSE**; these are the only literals of type **BOOLEAN**; i.e.,

BooleanLiteral ::= **FALSE** | **TRUE**

BOOLEAN expressions are used in conditional statements (following **IF**) and in certain loop constructs. For instance, the following skeletal form describes the flow of control in Example 1:

```

IF  $m = 0$  AND  $n = 0$  THEN ...
ELSE
...
UNTIL  $n = 0$ 
DO
...
ENDLOOP;
```

The expression " $n = 0$ " is a **BOOLEAN** expression; its value is **TRUE** if the value of n is zero and **FALSE** otherwise. The expression " $m = 0$ **AND** $n = 0$ " is also a **BOOLEAN** expression; its value is **TRUE** just if both relations are. The relational and logical operators discussed in subsections 2.5.2 and 2.5.3 all yield **BOOLEAN** values.

Variables of type **BOOLEAN** can be assigned values and appear as operands (although not of arithmetic operators) just as any other Mesa variables. For example, the above program outline could validly be replaced by the following:

```

mIsZero, nIsZero: BOOLEAN;
...
mIsZero ← ( $m = 0$ ); nIsZero ← ( $n = 0$ ); -- compute whether  $m$  and  $n$  are zero
```

```

IF mIsZero AND nIsZero THEN ...
ELSE
...
UNTIL nIsZero = TRUE           -- equivalent to just nIsZero by itself
DO
...
nIsZero ← n = 0;           -- recompute whether n is zero just
                               before testing
ENDLOOP;

```

2.4.3 Type CHARACTER

A value of type **CHARACTER** represents a *single* character of text. **CHARACTER** values are ordered (according to the order specified in appendix C) and can be compared using the normal arithmetic relations. **CHARACTER** values are distinct from numbers, and they cannot be assigned to variables with numeric types. Limited arithmetic is, however, allowed on characters (§ 2.5.1.3). See also the functions **ORD** and **VAL** (§ 2.5.6).

A **characterLiteral** is written as an apostrophe (') immediately followed by a single character (which can be a blank, carriage-return, semicolon, apostrophe, or any other character) or as an octal number followed by **C**. For example:

```

lowerCaseA ← 'a;
mark ← ' ;           -- mark is set to be a blank. Here a blank is significant
endMarker ← ';      -- endMarker is set to be a semicolon
asciiCR ← 15C;      -- an Ascii Carriage Return character,

```

2.4.4 Escape conventions for literals

Mesa provides an escape convention to allow non-printing characters in character and string literals (cf. the escape convention for the language C). The escape character is ****, and the following codes are recognized:

<u>Code</u>	<u>Interpretation</u>	
\n, \N, \r, \R	Ascii.CR	
\t, \T	Ascii.TAB	
\b, \B	Ascii.BS	
\f, \F	Ascii.FF	
\l, \L	Ascii.LF	-- note that \n is LF in C .
\ddd	dddC	-- where <i>d</i> is an octal digit, $ddd \leq 377B$
\\	\	
\'	'	
\"	"	

2.4.5 The numeric types LONG INTEGER and LONG CARDINAL

For some applications, the ranges of the numeric types introduced in subsection 2.4.1 are too limited. Mesa provides both a predefined type **LONG INTEGER**, with signed representation, and a predefined type **LONG CARDINAL**, with unsigned representation, for such applications. These types offer greater ranges, but their values occupy more storage and are generally more time-consuming to manipulate than those of **INTEGER** and **CARDINAL**.

In an implementation, values of types `INTEGER` and `CARDINAL` are expected to be represented by single machine words, while values of types `LONG INTEGER` and `LONG CARDINAL` are expected to occupy two words. For this reason, `INTEGER` and `CARDINAL` will be referred to as *short numeric types*; `LONG INTEGER` and `LONG CARDINAL`, as *long numeric types*. On a machine using two's complement arithmetic and a word length of N bits, the following table indicates the range spanned by each numeric type (".." replaces the mathematician's comma in this interval notation):

<code>INTEGER</code>	$[-2^{N-1} .. 2^{N-1}]$
<code>NATURAL</code>	$[0 .. 2^{N-1}]$
<code>CARDINAL</code>	$[0 .. 2^N]$
<code>LONG INTEGER</code>	$[-2^{2N-1} .. 2^{2N-1}]$
<code>LONG CARDINAL</code>	$[0 .. 2^{2N}]$

The actual ranges for these types are given in appendix C, the machine dependencies appendix.

Long numeric constants are denoted by numeric literals defined by the phrase class `number` (§ 2.4.1.1). The allowable type of any decimal or octal literal is determined by its value, as summarized by the following table (using the conventions introduced in the preceding paragraph):

<i>Range</i>	<i>Allowable Types</i>
$[0 .. 2^{N-1}]$	<code>INTEGER</code> , <code>CARDINAL</code> , <code>NATURAL</code> , <code>LONG INTEGER</code> , <code>LONG CARDINAL</code>
$[2^{N-1} .. 2^N]$	<code>CARDINAL</code> , <code>LONG INTEGER</code> , <code>LONG CARDINAL</code>
$[2^N .. 2^{2N-1}]$	<code>LONG INTEGER</code> , <code>LONG CARDINAL</code>
$[2^{2N-1} .. 2^{2N}]$	<code>LONG CARDINAL</code>

As in the case of short numeric types, the types `LONG INTEGER` and `LONG CARDINAL` are distinct but closely related. Mesa allows most combinations of these types and the types `INTEGER` and `CARDINAL` to occur within assignments, arithmetic expressions and relational expressions, but care is necessary when this is done (§ 2.4.6 and § 2.5.1.1).

2.4.6 The numeric type `REAL`

The values of Mesa's type `REAL` are approximations of mathematical real numbers. These approximations are sometimes called *floating-point* numbers.

Mesa has adopted the proposed IEEE standard for floating-point arithmetic (see e.g., Coonen. "An implementation guide to a proposed standard for floating-point arithmetic," *Computer*, January 1980, pp. 68-79). In support of this, the language provides floating-point literals and the compiler performs a limited number of operations upon floating-point constants. A floating-point constant is known as a `realLiteral`.

Syntax

Primary	:: =	... realLiteral	-- (§ 2.5)
realLiteral	:: =	unscaledReal unscaledReal scaleFactor wholeNumber scaleFactor	
unscaledReal	:: =	wholeNumber fraction fraction	
fraction	:: =	. wholeNumber	
scaleFactor	:: =	E optSign wholeNumber e optSign wholeNumber	
optSign	:: =	empty + -	
wholeNumber	:: =	digit wholeNumber digit	

An **unscaledReal** has its usual interpretation as a decimal number. The **scaleFactor**, if present, indicates the power of 10 by which the **unscaledReal** or **wholeNumber** is to be multiplied to obtain the value of the literal.

Mesa represents **REAL** numbers by 32 bit approximations as defined in the IEEE standard. The rounding mode used to convert literals is "round-to-nearest." A literal that overflows the internal representation is an error; one that underflows is replaced by its so-called "denormalized" approximation. In Mesa, the value of the **unscaledReal** in a literal must be a valid **LONG INTEGER** when the decimal point is deleted.

No spaces are allowed within a **realLiteral**. Note that such a literal can begin, but not end, with a decimal point. Thus, the interpretation of [0...1) is unambiguous but perhaps surprising; use [0.. .1) or [0.0..0.1) instead. (See **SubRangeTC** in subsection 2.5.2.)

Operations

The compiler performs the following operations involving floating-point constants:

Unary negation (with $-0 = 0$)
ABS
 Fixed-to-Float (in "round-to-nearest" mode).

Other operations are deferred until runtime, even if all their operands are constant, so that the programmer can control the treatment of rounding and exceptions (see the proposed standard).

2.4.7 Relations among basic types

If two types are completely interchangeable, they are said to be *equivalent*. A value having a given type is acceptable in any context requiring a value of any other type equivalent to it; there is no operational difference between two equivalent types. None of the basic types discussed in section 2.4 is equivalent to another basic type.

One type is said to *conform* to another if any value of the first type can be assigned to a variable of the second type. A type trivially conforms to itself or to any type equivalent to

itself. In more interesting cases, an automatic application of a conversion function may be required prior to the assignment. Conformance and its implications are discussed further in section 3.5.

There are nontrivial conformance relations involving the types `INTEGER`, `LONG INTEGER`, `NATURAL`, `CARDINAL`, `LONG CARDINAL` and `REAL`. These relations allow certain combinations of the numeric types to be mixed, not only in assignments but also in arithmetic and relational operations (§ 2.5.2). They also permit these types to share denotations of constants (§ 2.4.5). The conformance relations can be summarized as follows:

`INTEGER`, `NATURAL`, `LONG INTEGER`, `CARDINAL` and `LONG CARDINAL` conform to `INTEGER` (note § 2.4.8).

`INTEGER`, `NATURAL`, `LONG INTEGER`, `CARDINAL` and `LONG CARDINAL` conform to `CARDINAL` (note § 2.4.8).

`INTEGER`, `NATURAL`, `LONG INTEGER`, `CARDINAL` and `LONG CARDINAL` conform to `NATURAL` (note § 2.4.8).

`INTEGER`, `NATURAL`, `LONG INTEGER`, `CARDINAL` and `LONG CARDINAL` conform to `LONG INTEGER`.

`INTEGER`, `NATURAL`, `LONG INTEGER`, `CARDINAL` and `LONG CARDINAL` conform to `LONG CARDINAL`.

`INTEGER`, `NATURAL`, `LONG INTEGER`, `CARDINAL`, `LONG CARDINAL` and `REAL` conform to `REAL`.

Pairs of numeric types not on this list do not conform; e.g., it is not possible to assign a `REAL` to a `CARDINAL`.

Particular care is required when numeric types with different representations are intermixed. Mathematically, $\mathbb{Z} \supset \mathbb{N}$; however, it is not necessarily true that `INTEGER` \supset `CARDINAL` or that `LONG INTEGER` \supset `LONG CARDINAL`. For instance, with the assumptions above, the intersection of `INTEGER` and `CARDINAL` is $[0..2^N-1]$. Within this interval, the signed and unsigned representations agree, and the interpretation of a short numeric value is unambiguous. *If a `CARDINAL` value lies in this range, it can validly be assigned to an `INTEGER` variable, and vice-versa; outside this range, the value represented by a given word depends upon whether it is viewed as a `CARDINAL` or as an `INTEGER`.* Similar considerations apply to `LONG CARDINAL` and `LONG INTEGER`.

Example:

With the assumptions above and word length $N=16$, the unsigned value 177777B and the signed value -1 are encoded by the same bit pattern.

Assignment of an unsigned value to an `INTEGER` variable, or of a signed value to a `CARDINAL` variable, implicitly invokes a conversion function, which is just an assertion that the value to be assigned is an element of `CARDINAL` \cap `INTEGER` (i.e., a `NATURAL`). If bounds checking is requested of the compiler, code will be inserted before each cross assignment to insure that the value is within range; otherwise, *it is the responsibility of the programmer to ensure that the conversion is valid*. In many cases this is not too difficult, but programmers are urged to avoid mixing signed and unsigned representations when this is possible. It almost always is.

Mesa does guarantee that `LONG T` \supseteq `T` for any type `T` and that `LONG INTEGER` \supset `CARDINAL`; thus it is always valid to assign a short numeric value to a `LONG INTEGER` variable or a short

unsigned value to a **LONG CARDINAL** variable. The properties of conversion to type **REAL** are not specified by the language.

Fine points:

A built in procedure *FLOAT* is automatically applied to convert a value from type **LONG INTEGER** to **REAL**. Short numeric values are converted first to **LONG INTEGER** and then to **REAL**.

Conversion from a short numeric value to a **LONG INTEGER** (and thus to a **REAL**) is substantially more efficient when the value has an unsigned representation.

Neither **BOOLEAN** nor **CHARACTER** conforms to any other basic type.

Examples:

```

i: INTEGER;      n: CARDINAL;    ii: LONG INTEGER;    x: REAL;

(valid)          i ← 0;
                  ii ← 0;
                  x ← n;
                  x ← ii;

(invalid)        i ← x;
                  n ← TRUE;

```

2.4.8 Long to short conversion

Mesa provides conversion from types **LONG INTEGER** and **LONG CARDINAL** to **INTEGER**, **CARDINAL**, and any subranges thereof. If you request bounds checking, the compiler will insert code to check that the values are in the proper range, otherwise the values are simply truncated to fit the destination.

In order to avoid surprises caused by loss of precision, the compiler issues a warning whenever there is an implicit shortening from a two word quantity to a smaller one. To avoid a warning, you should use an explicit range assertion (§ 3.1.2.2). For reasons of backward compatibility, no explicit range assertion is required when storing a single word quantity into a smaller subrange variable. For example:

```

i: INTEGER; c: CARDINAL; li: LONG INTEGER; S: TYPE = [0..10);
s: S;
i ← li;                -- gets a warning
i ← INTEGER [li];     -- no warning
c ← li;                -- gets a warning
c ← CARDINAL [li];   -- no warning
s ← li;                -- gets a warning
s ← S [li];           -- no warning
s ← i;                -- no warning!
s ← S [i];            -- no warning

```

In all cases, warning or not, the compiler will generate code to check the suitability of the right hand side value if bounds checking (the /b switch) is requested.

2.4.9 Predeclared identifiers

The following predeclared identifiers may be used to make programs less verbose:

```

BOOL : TYPE = BOOLEAN;
CHAR : TYPE = CHARACTER;
NAT : TYPE = NATURAL;
INT : TYPE = LONG INTEGER;

```

2.5 Expressions

Expressions are constructs describing rules of computation for evaluating variables and for generating new values by the application of operators. The overall syntactic rule for an expression is given by

```

Expression ::= Disjunction | AssignmentExpr | IfExpr | SelectExpr | ...

```

The **Disjunction** form, which is discussed in this section, includes all the numeric operations, relational operations, and **BOOLEAN** (logical) operations. An **AssignmentExpr** allows one to write multiple assignments in a single statement; it is discussed in subsection 2.5.4. The **IfExpr** and **SelectExpr** forms are discussed in chapter 4.

The basic unit from which expressions are built is called a **Primary**. This syntactic class includes references to variables, literals, function calls (chapter 5), and any arbitrary expressions embedded in parentheses:

```

Primary ::= Variable | Literal | ( Expression ) | FunctionCall | ... | realLiteral

Variable ::= LeftSide

Literal ::= number | BooleanLiteral | characterLiteral

FunctionCall ::= BuiltinCall | Call -- defined in chapter 5

```

Recall that every expression has a well-defined type in Mesa. The general rules for determining the type of an expression from the types of its constituent parts are given in section 3.5. This section outlines the types of the basic expression forms (as functions of the types of their operands). For example, the type of a **Primary** is the type of the **Variable** or **Literal** involved, or reduces to the type of the **Expression** within parentheses, or is the type of the value returned by the **BuiltinCall** (some of which are defined below) or the **Call** of a user-defined procedure (§ 5.1).

Some operators are numeric and some are **BOOLEAN**. The next sections discuss the numeric operations, the relational operations, and the operations applicable only to **BOOLEAN** values. Considered together, the operators form a single hierarchy with respect to their precedence, which is described with each operator class and summarized in subsection 2.5.5.

2.5.1 Numeric operators

The operations on numeric values are addition, subtraction, multiplication, division, modulus, and arithmetic negation. The syntax for this group of operations is

```

Factor           ::= Primary | -Primary | +Primary

Product          ::= Factor | Product MultiplyingOperator Factor

MultiplyingOperator ::= * | / | MOD

Sum              ::= Product | Sum AddingOperator Product

AddingOperator  ::= + | -

```

These operators have their usual mathematical meanings. The division operation on integers, $/$, always truncates toward zero; thus $-(i/j) = -i/j = i/-j$. The MOD operator yields the remainder of dividing one number by another (MOD is not applicable to REAL operands). MOD is defined by the relation $(i/j)*j + (i \text{ MOD } j) = i$, and the sign of the result of MOD is always the sign of the dividend. (This is the reason that line 11 of Example 1 takes the absolute value of the computed *gcd*; if $m = -12$ and $n = 8$ initially, the *gcd* would be -4 if its absolute value were not taken.)

The built-in function MIN computes the minimum value in a list of expressions; similarly, the MAX function, the maximum value. The built-in function ABS computes the absolute value of its argument. The syntax for calls on the built-in functions is

```

BuiltinCall      ::= MIN [ ExpressionList ] |
                  MAX [ ExpressionList ] |
                  ABS [ Expression ] |
                  ...
                  -- other built-in functions later

ExpressionList   ::= Expression | ExpressionList , Expression

```

For the arithmetic operators and built-in functions, the order in which the operands are evaluated is undefined, but the syntax implies a precedence ordering that controls the association of operators with their operands. In that ordering, unary negation precedes the multiplying operators, which in turn precede the adding operators. *Sequences of operators of the same precedence associate from left to right* (with the exception of the embedded assignment operator, § 2.5.4). Thus, an expression such as $a + b* - c$ does *not* specify the order of evaluation of a , b and c but does require that the operations be performed in the following order: negate c ; then multiply the result by b ; finally, add that result to the value of a .

Examples:

```
i, j, k: INTEGER; m, n: CARDINAL;
```

```

Factors:      n
              15
              (i+j+k)
              -15

```

```

MIN[i, j, k, -15]
3.1416

Products:  m*n
           i/-15
           n MOD 8
           m/n*10      -- same as (m/n)*10 because of left-associativity
           -k*(i+1)/2 MOD 3 -- same as (((-k)*(i+1))/2) MOD 3

Sums:      i+1
           -i+j
           j-i
           n-n MOD 8    -- same as n-(n MOD 8) because of precedence
           m-m/n*n     -- same as m MOD n

```

2.5.1.1 Domains of the numeric operators *

In principle, each arithmetic operator designates the corresponding mathematical function. Unfortunately, the hardware underlying any implementation of Mesa does not provide this function but only a set of related partial functions. For each operator, the compiler must choose as appropriately as possible from this set. The choice is made by considering the types of the operands.

Example:

With the usual assumptions, 177777B and -1 are represented by the same bit pattern. The value of 177777B > 0 is TRUE, but that of -1 > 0 is FALSE.

Mesa provides the operators +, -, *, /, MIN, MAX and ABS for all the numeric types. The operation MOD is defined for all numeric types except REAL; the operation of unary negation, for all but CARDINAL and LONG CARDINAL. For each of these operators, the type of the result is the same as the type of the operands. Additionally, the result of the operation is considered to have signed representation if all the operands have signed representation, and to have unsigned representation if all the operands have unsigned representation. Thus, adding two INTEGER values yields an INTEGER result, and dividing one CARDINAL by another yields a CARDINAL result.

Fine points:

Division and modulus operations on short numeric values are substantially more efficient if their operands are unsigned.

Addition, subtraction, and comparison of long numeric values are fast; multiplication and division are done by software and are relatively slow.

Although the mathematical integers (\mathbb{Z}) and real numbers are closed under all these operations (except division by zero), the subranges defining the types INTEGER, LONG INTEGER, NATURAL, CARDINAL and LONG CARDINAL generally are not. When the result of an operation falls outside the range of its assumed type, a representational failure called *overflow* or *underflow* occurs. In the current version of Mesa, *it is the programmer's responsibility to guard against overflow and underflow conditions.*

The implications of Mesa's conventions for subtraction are worth emphasizing. If both operands have valid signed representations, the result has a signed representation. If both

have only unsigned representations, the result has an unsigned representation and is considered to overflow if the first operand is less than the second.

Example:

i : INTEGER; m, n : CARDINAL;

$i \leftarrow m - n$;

-- should be used only if it is known that $m \geq n$

$i \leftarrow \text{IF } m \geq n \text{ THEN } m - n \text{ ELSE } -(n - m)$; -- a safer form (§ 3.6)

The arithmetic operations are defined for operands that all have the same type, but it is possible to mix numeric types (and thus representations) within an expression. In this case, operands are converted as necessary to the "smallest" type to which all the operands conform, the operation for that type is applied, and the result also has that type. The rule for expressions involving type REAL is easy to state:

If any operand has type REAL, the REAL operation is used.

The rules governing combination of numeric operands with differing representations involve some additional concepts and are stated in section 3.6. Again, the programmer should try to avoid such combinations when possible. (Recall that literals in INTEGER \cap CARDINAL have whatever representation is required by context.)

2.5.1.2 The operator LONG *

The built-in function LONG converts any value with a short numeric type to a long numeric type and provides explicit lengthening of pointer types to long pointer types (§ 3.4.3) or explicit lengthening of an array descriptor type (§ 6.2). A value with a signed representation is converted to LONG INTEGER; one with an unsigned representation, to either LONG INTEGER or LONG CARDINAL as required. The syntax is as follows:

BuiltinCall ::= ... | LONG [Expression]

This operation is necessary when the standard conversion rules do not give the desired result. It can also be used to emphasize the conversion.

Example:

LONG [$m * n$]

-- "short" multiplication, overflow lost

LONG [m]*LONG [n]

-- "long" multiplication

Fine points:

Lengthening a single-precision expression is substantially more efficient if that expression has an unsigned representation.

The Mesa implementation provides standard procedures (not part of the language) for performing certain multiplication and division operations in which the operands and results do not all have the same length. These procedures provide less expensive equivalents of, e.g., LONG [m]*LONG [n].

2.5.1.3 CHARACTER operators *

Limited CHARACTER arithmetic is possible and is sometimes useful for manipulating the encodings of CHARACTER values. The following arithmetic operations are defined for operands of type CHARACTER:

A CHARACTER value plus or minus a short numeric value yields a CHARACTER value.

Subtracting two CHARACTER values yields an INTEGER value.

No other arithmetic operations on characters are allowed. Since the results of character arithmetic depend upon details of the character encoding, such arithmetic should be used with discretion.

Examples:

```
c: CHARACTER; digit: INTEGER;
digit ← c - '0;
c ← c + ('A' - 'a')           -- converts lower case to upper
```

2.5.2 Relational operators

The relational operators include = and # (not equal), <, <= (less than or equal), >= (greater than or equal), >, and their negatives (e.g., NOT<, ~<, ~>=, etc.). These operators always yield BOOLEAN results, depending on the truth or non-truth of the relation expressed. The operators = and # apply to most types; the others, to any *ordered* type (i.e., to any type whose values are considered to be ordered). Ordered types include INTEGER, LONG INTEGER, NATURAL, CARDINAL, LONG CARDINAL, REAL, BOOLEAN, CHARACTER, enumerated types (§ 3.1), and subranges of ordered types (§ 3.1).

The relational operators also include the composite operator IN, which takes a numeric value as its left operand and an *interval* as its right operand. Its value is TRUE if the left value lies in the interval and FALSE otherwise. The syntax for relational operators is

Relation	:: =	Sum Sum RelationTail
RelationTail	:: =	RelationalOperator Sum Not RelationalOperator Sum IN SubRange Not IN SubRange
RelationalOperator	:: =	< <= = # > >=
Not	:: =	~ NOT
SubRange	:: =	SubRangeTC ... -- explained in chapter 3
SubRangeTC	:: =	Interval ... -- explained in chapter 3
Interval	:: =	[Expression .. Expression) (Expression .. Expression)

(Expression .. Expression]]
 [Expression .. Expression]

The extra syntax for `SubRange` and `SubRangeTC` is placed here to be consistent with later uses of the class `Interval` in chapter 3. The syntax for intervals follows mathematical notation; a square bracket indicates the inclusion of the respective end point in the interval, while a parenthesis indicates its exclusion. For example, the following intervals all designate the range from -1 to 5 inclusive:

`[-1 .. 5]` `[-1 .. 6)` `(-2 .. 6)` `(-2 .. 5]`

In the above examples, -1 is the *lower bound* of each interval; the *upper bound* is 5 . The bounds of an interval are its end points, regardless of whether the interval is written as a closed, half-open or open one. The bounds are not required to be constants. An interval with an upper bound less than its lower is said to be *empty*; no values lie in such an interval. For example, the following are all empty intervals:

`[-1 .. -2]` `[-1 .. -1)` `(-2 .. -1)` `(-2 .. -2]`

Intervals may use real numbers as endpoints. Recall from subsection 2.4.6 that a `realLiteral` can begin, but not end, with a decimal point. Consequently, `[0...1]` (three dots) is unambiguous but is better expressed as `[0.0 .. 0.1]`.

Examples:

Relations:	<code>n = 15</code>	
	<code>m # n</code>	-- or <code>m ~ = n</code>
	<code>i <= j</code>	
	<code>(i < j) = (j < k)</code>	-- = with two <code>BOOLEAN</code> operands
	<code>n IN [1 .. 5)</code>	-- <code>n >= 1</code> and <code>n < 5</code>
	<code>i NOT IN [-1 .. 5]</code>	-- only legal if <code>i</code> is signed (because -1 is)

Fine point:

The relational operators, like the arithmetic operators, denote families of hardware operations when they have numeric operands. Again, there is one operation for each numeric type. If there is a unique "smallest" type to which all the operands conform, they are converted to that type as necessary and then the comparison is performed. There is no unambiguous choice of such a type for numeric operands with different representations; an attempt to compare two such values is an error. The precise rules appear in section 3.5.

2.5.3 BOOLEAN operators

The operators `NOT` (logical negation), `AND` and `OR` apply only to `BOOLEAN` values. The syntax is

Negation	:: =	Relation Not Relation
Conjunction	:: =	Negation Conjunction AND Negation
Disjunction	:: =	Conjunction Disjunction OR Conjunction

NOT *negates* the logical value of a **BOOLEAN** expression. *p AND q* has the value **TRUE** if and only if both *p* and *q* are **TRUE**. *p OR q* is **TRUE** if at least one of *p* or *q* is **TRUE**.

When evaluating a **BOOLEAN** expression, evaluation of primaries is guaranteed to take place from left to right. In the operation **AND** or **OR**, the second operand is evaluated only if the first operand's value does not determine the value of the expression.

Fine points:

"*x AND y*" is equivalent to the `IfExpr` "IF *x* THEN *y* ELSE FALSE"; i.e., when *x* is **FALSE**, *y* is not evaluated.

"*x OR y*" is equivalent to the `IfExpr` "IF *x* THEN TRUE ELSE *y*"; i.e., when *x* is **TRUE**, *y* is not evaluated.

It is therefore safe to have expressions of the form "*x AND y*," where *y* is defined only when *x* is **TRUE**, e.g., "*x#0 AND c/x > 2*," or "*p=NIL OR p.f=0*."

Examples:

Negations:	<code>NOT i = 15</code>	-- same as <code>NOT(i = 15)</code>
	<code>~q</code>	-- <i>q must be of type BOOLEAN</i>
	<code>~(p AND q)</code>	
Conjunctions:	<code>i <= j AND j < k</code>	
	<code>p AND ~q</code>	
	<code>i = 5 AND j NOT IN [- 1..1]</code>	
Disjunctions:	<code>m > n OR m = 15</code>	
	<code>~p OR ~q</code>	

2.5.4 Assignment expressions

The assignment operation can be embedded in other expression forms. When it is, the result of the operation has the type of the `LeftSide` and the value received by the `LeftSide` in the assignment. The "`←`" operator has the lowest precedence of any operator. Its syntax is the same as that of the `AssignmentStmt`:

`AssignmentExpr ::= LeftSide ← RightSide`

If this form is used to perform multiple assignments, it is important to note that "`←`" is *right-associative*. Thus, the assignment expression `a←b←b+1` first assigns the value of `b+1` to `b` and then assigns `b`'s new value to `a`.

Examples:

Assignment Expressions:	
<code>m←15</code>	
<code>m←n←15</code>	
<code>m←n←n+1</code>	-- same as <code>m←(n←(n+1))</code>
<code>i←(j←(j+1) MOD n)*2</code>	-- all these parentheses are necessary

Rules governing assignments of numeric values when the types are not identical are summarized in subsection 2.4.7.

Fine point:

Because the order of evaluation of the primaries is not defined, expressions such as " $(i \leftarrow j) + (j \leftarrow k)$ " have unpredictable values and should not be used.

2.5.5 Operator precedence

The following table summarizes the precedences of the unary and binary operators introduced in this section. The order is from highest precedence (tightest binding of operands) to lowest; operators on the same line have the same precedence.

- , +	-- unary negative and positive
*, /, MOD	
+ , -	-- addition and subtraction
= , # , < , <= , > , >= , IN	
~ , NOT	
AND	
OR	
←	

Parentheses can be used to explicitly control the association of operands with operators.

2.5.6 Function-like operators

There are a number of unary operators whose application looks like a function call, such as the `LONG` operator (§ 2.5.1.2).

BuiltinCall	::=	... PrefixOp [Expression] VAL [Expression]
PrefixOp	::=	ABS -- (§ 2.5.1) BASE -- (§ 6.2) LENGTH -- (§ 6.2) LONG -- (§ 2.5.1.2) ORD PRED SUCC

The operators `PRED` and `SUCC` operate upon values of any ordered type except `REAL`. For numeric and character types, `SUCC [x]` and `PRED [x]` are equivalent to $x+1$ and $x-1$ respectively. For enumerated types (§ 3.1.1), the values are successor and predecessor of x in the enumeration; a bounds fault occurs if there is no such element and you requested bounds checking, otherwise you get undefined results and possibly overflow.

The operator `ORD` (ordinal) provides a `LOOPHOLE`-free mechanism (§ 3.5.1.2) for converting a character or enumerated value into a numeric value. For example, given the standard ASCII representation for characters, `ORD ['A'] = 101B`.

The `VAL` operator is the inverse of `ORD`. It must be used in a situation where the compiler can determine the type of the result from context. For example:

```
c: CHARACTER ← VAL [101B];           -- sets c to 'A'
```

All of the **PrefixOps** can be invoked using “dot notation” as well as using brackets. Thus `x.SUCC` may be used instead of `SUCC[x]`.

2.5.7 Function-like operators on types

There are several function-like operators with arguments of type **TYPE**.

```
BuiltinCall      ::= ... | TypeOp [ TypeExpression ] |
                   SIZE [ TypeExpression , Expression ]
```

```
TypeOp           ::= SIZE | FIRST | LAST | NIL
```

The operators **FIRST** and **LAST** are applicable to all element types (§ 3.1), including **INTEGER**, **NATURAL**, **CARDINAL**, and **CHARACTER**, as well as **LONG INTEGER** and **LONG CARDINAL**. When applied to the numeric types other than **REAL**, they may supply information about the range of values supported by a particular implementation. When applied to an enumerated type (§ 3.3.1) they yield the least and greatest elements, respectively, of the enumeration.

The operator **SIZE** is used to find the number of machine words occupied by an object of any type. The result is a **CARDINAL** value. The two parameter form of the operator is used to determine the number of machine words occupied by a **PACKED ARRAY** of elements of the specified type (§ 3.2 and 3.3).

The operator **NIL** returns a nil value of a **POINTER** type (§ 3.4).

If the **TypeExpression** can be parsed as an **Expression**, these operators may be written in “dot notation.” Thus **INTEGER.LAST** is equivalent to **LAST [INTEGER]**, but **LONG INTEGER.LAST** or even **(LONG INTEGER).LAST** is not allowed. For another example of a situation where this cannot be done, see the section on variant records (§ 6.4).

2.6 Initializing variables in declarations

A variable may be given an initial value in a declaration. For example, the Boolean variable *delimited* could be set initially **FALSE** by using the declaration:

```
delimited: BOOLEAN ← FALSE;
```

Variables (of the same type) can be initialized collectively:

```
n, n0: INTEGER ← -7;
```

This declares two separate integer variables *n* and *n0* and initializes each to -7.

Any expression that could be used as the **RightSide** of an assignment can be used to initialize a variable:

```
i: INTEGER ← ABS[n];           -- this will set i to 7
iSquared: INTEGER ← i*i;      -- iSquared is initialized to 49
j: INTEGER ← iSquared - i + 1;  -- j is initialized to 49 - 7 + 1 = 43
```

All initializations shown so far have taken “assignment” (or “←”) form. There is another form, the “fixed” (or “=”) initialization. For example,

```
octalRadix: INTEGER = 8;
```

This means that `octalRadix` is to have a fixed value. *It is never valid as the LeftSide of an assignment.* We call `octalRadix` a *constant* because its value can never change after it is initialized (recall that the number 8 is called a *literal*). Normally, the term “constant” will include the term “literal”; if the distinction is important, then “literal” will be used.

Initial values for fixed initialization can be arbitrary expressions. Paraphrasing the earlier example:

```
i0: INTEGER = ABS[ - octalRadix]; i0Squared: INTEGER = i0*i0;
j0: INTEGER = i0Squared - i0 + 1;
```

The initializing expression can use values that are not known at compile time. In this example, if `octalRadix` did not have fixed initialization, the values of `i0`, `i0Squared`, and `j0` would be computed and assigned at run-time. Variables are initialized in the order of appearance in a declaration, and later declarations can use variables initialized earlier, as shown by the example.

2.6.1 Compile-time constants

Wherever possible, the Mesa compiler evaluates expressions containing only constants. If a variable is initialized using the fixed form and the expression can be evaluated at compile time, then that variable has a known value. Since it can never appear as the `LeftSide` of an assignment operator, it too becomes a compile-time constant (the variables `i0`, `i0Squared`, and `j0` in the previous section are all compile-time constants).

Example:

```
beta: INTEGER = 3;
alpha: INTEGER = beta - 1;
```

In this case, `alpha` is a compile-time constant (with the value 2), since the expression `beta - 1` involves only compile-time constants. Compile-time constants need not occupy memory at run-time; the compiler can replace references to compile-time constants, such as `alpha` and `beta`, by their known values.

Fine points:

Knowledge of compile-time constant values can also be exploited when analyzing expressions, processing other declarations, or generating object code.

One side effect of this propagation of constants is that the representation of a numeric constant is known at compile-time. For instance, `alpha` above is declared to be an `INTEGER`, but because its value is 2, it may also be used as a `CARDINAL`. However, declaring the type of `alpha` determines what kind of arithmetic (signed or unsigned) will be used to compute its value, whether at compile-time or run-time (§ 2.5.1).

In certain contexts, an expression is required to yield a compile-time constant value. A (sub)expression denotes such a constant if all the operands are compile-time constants and the operation is *not* one of those listed below (current restrictions):

Any arithmetic or relational operation with operands of type `REAL`.

Application of any function (chapter 5) other than a built-in function or simple **INLINE** procedures (but no guarantees are made for **INLINE** procedures).

The @ operation (§ 3.4), except in some cases where constants have been **LOOPHOLEd** (§ 3.5.1.2) to pointer types.

2.7 More general declarations

Preceding sections have introduced all the syntactic components of a declaration. The general form is defined as follows:

```
Declaration      ::= IdList : TypeSpecification Initialization ; Declaration |
                  IdList : TypeSpecification Initialization
```

For the moment, **TypeSpecification** is defined as one of the built-in types; chapter 3 describes other forms of **TypeSpecification**.

```
TypeSpecification ::= PredefinedType | ...
```

```
PredefinedType  ::= INTEGER | CARDINAL | NATURAL
                  BOOLEAN | CHARACTER |
                  LONG INTEGER | LONG CARDINAL | REAL |
                  STRING |                -- see chapter 6
                  WORD |                  -- see fine point below
                  UNSPECIFIED |          -- see fine point below
                  MONITORLOCK |         -- see chapter 9
                  CONDITION              -- see chapter 9
```

An **Initialization** is formally defined as follows:

```
Initialization  ::= empty |
                  ← Expression |
                  = Expression |
                  ...                --other forms are given later
```

Fine points:

The predefined type **WORD** is provided to describe values on which bit-by-bit logical operations are to be performed. Currently, it is a synonym for **CARDINAL**.

The predefined type **UNSPECIFIED** is a device for bypassing most type checking. An **UNSPECIFIED** value is a single machine word, and it matches the type of any object that occupies at most a single machine word, including **INTEGER**, **NATURAL**, **CARDINAL**, **CHARACTER**, **BOOLEAN**, **UNSPECIFIED**, **STRING**, and any user-defined type (chapter 3) that fits in a single machine word.

For numeric operations, its representation is similarly fluid. If a **CARDINAL** and an **UNSPECIFIED** value are the operands of some arithmetic operation, then the **UNSPECIFIED** value is considered to be unsigned. If an **UNSPECIFIED** is combined with a signed value, it is treated as if it were signed too. If an **UNSPECIFIED** is combined with an **UNSPECIFIED**, they are both treated as signed.

Less type checking is sacrificed by using **LOOPHOLE** (§ 3.5.1.2) than by declaring variables with type **UNSPECIFIED**.



Common constructed data types

Mesa encourages the programmer to augment the collection of predefined types by *constructing* new types. Types can be defined to describe objects that are structured collections of related values (e.g., a vector of Booleans, a table, or a complex number consisting of real and imaginary components). Mesa's type system has other, perhaps less obvious applications. These include expressing some of the programmer's knowledge about a class of variables (e.g., that all take on values restricted to some known interval), separating variables with different semantics into different classes so that they cannot be confused (e.g., to avoid "comparing apples and oranges"), hiding implementation details of abstractions (e.g., to prevent the user of a table-lookup package from depending upon the internal organization of the table), and facilitating the introduction of synonyms to provide better description and improved readability.

Programmer-created types have the same status as Mesa's built-in types. They can be used to declare variables and to construct additional new types. In addition, values of most constructed types can be operands of the fundamental operations (\leftarrow , =, #).

A new type identifier is declared using the following syntax:

TypeDeclaration ::= idList : TYPE = TypeSpecification ;

Each identifier in the **idList** is thereby declared to name the type denoted by the **TypeSpecification**. If this declaration form is compared to a normal declaration, i.e.,

Declaration ::= IdList : TypeSpecification Initialization ;

it can be seen that "**TYPE**" fills the role of a **TypeSpecification**, and "**= TypeSpecification**" plays the role of **Initialization**. In fact, the newly declared identifier has type "**TYPE**" and a value (which must be constant, hence the "=") that is a **TypeSpecification**.

There are several predeclared identifiers that may be used to make programs less verbose.

```
BOOL: TYPE      = BOOLEAN;  
CHAR: TYPE     = CHARACTER;  
INT: TYPE      = LONG INTEGER;
```



```
NAT: TYPE      = NATURAL;
```

Any predefined Mesa type (§ 2.7) is a valid **TypeSpecification**; thus the following are valid type declarations:

```
SignedNumber: TYPE = INTEGER;
UnsignedNumber: TYPE = CARDINAL;
TruthValue: TYPE = BOOLEAN;
Char: TYPE = CHARACTER;
```

These *type identifiers* are now valid type specifications and can be used to declare variables:

```
i, j: SignedNumber;
n: UnsignedNumber;
b: TruthValue;
c: Char;
```

After this series of declarations, *i* and *j* have type *SignedNumber*, which is equivalent to **INTEGER**; *n* has type *UnsignedNumber*, which is equivalent to **CARDINAL**; etc. This is a trivial way of defining new types. A more interesting way uses a *type constructor* as the **TypeSpecification** and generates a truly new type, not just an additional name for an existing one. A **TypeSpecification** can be defined as

```
TypeSpecification ::= PredefinedType |
                    TypeIdentifier |
                    TypeConstructor
```

(**TYPE** itself is *not* a **TypeSpecification**; it can be used only to declare types.)

There is an important point worth emphasizing here. A **TypeSpecification** that is a **PredefinedType** or a **TypeIdentifier** denotes an existing type and yields the same type every time it is used. A declaration such as the one of *SignedNumber* introduces a synonym for the name of an existing type. Synonyms can be more descriptive and thus improve readability, but they do *not* partition the set of values. The types *SignedNumber* and **INTEGER** are fully equivalent, and values with these types can be used interchangeably. On the other hand, a **TypeConstructor** constructs a new type. The rules for equivalence and conformance of constructed types depend upon the forms of their constructors and are discussed as the constructors are introduced. In some cases, each appearance of a constructor generates a unique type, i.e., writing the same sequence of symbols twice generates two distinct, incompatible types. For this reason, programmers usually should name such a type, using a **TypeDeclaration**, and thereafter use the type's identifier. Of course, introducing an identifier for a constructed type can make a program easier to read and modify in any case.

The **PredefinedTypes** are described in chapter 2 (except for **STRING** in chapter 6 and process related types in chapter 9).

The simplest form of a **TypeIdentifier** is given by

```
TypeIdentifier ::= identifier |      -- which is a declared type
                ...                  -- other forms given in chapters 6 and 7
```

The rest of this chapter discusses the attributes and uses of some common constructed types: *enumerations*, *subranges*, *arrays*, *records*, and *pointers*. The syntax for `TypeConstructor` is

```

TypeConstructor ::= EnumerationTC | -- for enumerations
                  SubrangeTC     | -- for subranges
                  ArrayTC        | -- for arrays
                  RecordTC       | -- for records
                  PointerTC      | -- for pointers
                  LongTC         | -- for long pointers, etc
                  ProcedureTC    | -- see chapter 5
                  ArrayDescriptorTC | -- see chapter 6
                  RelativeTC     | -- see chapter 6
                  SignalTC       | -- see chapter 8
                  ProcessTC      | -- see chapter 9

```

(The suffix “TC” is to be understood as an abbreviation for “TypeConstructor.”)

Enumerations define a set of values by giving a list of *identifiers*. These identifiers can be viewed as members of an ordered set.

Subranges define types with values drawn from those of a larger, encompassing type but restricted to lie in a specified interval. The subrange takes on the characteristics of the enclosing type; for example, a subrange of `INTEGER` can be used to declare variables that behave as `INTEGER`s but are constrained to take values within some interval.

Arrays are sequences of components that are homogeneous with respect to type and are accessed by computed indices (“subscripting”). Records are sequences of components that have potentially different types and are accessed using fixed component names (“selection”). Records and arrays are Mesa’s *aggregate* data types.

Pointers are scalar values used to access data objects indirectly. A pointer value is represented by an address. Pointers can be used to build structures such as linked lists and tree structures. Long pointers are pointers capable of spanning a larger address space than ordinary pointers.

This chapter concludes with a discussion of *type determination*, the process by which Mesa decides whether an expression has an acceptable type for a given operation. This is closely related to questions of the equivalence and conformance of types.

3.1 The element types

This section describes a class of types called *element types*. Their common properties are the following:

- (1) They are ordered types; values of an element type can be operands of all the relational operators (§ 2.5.2).
- (2) They are *scalar* types; a value of an element type does not have any visible or directly accessible internal structure insofar as the language is concerned.

- (3) They can be used to declare subrange types (§ 3.1.2).
- (4) They are the only types allowed as index types of arrays (§ 3.2).

The element types are `INTEGER`, `NATURAL`, `CARDINAL`, `CHARACTER`, `BOOLEAN`, the types generated by `EnumerationTC`, and the types generated by `SubrangeTC`. Because of (3) above, this definition is recursive; subranges of subranges are allowed. The definition of the class `ElementType` is

```

ElementType      ::= INTEGER | NATURAL | CARDINAL | CHARACTER | BOOLEAN |
                  EnumerationTC |
                  SubrangeTC

```

Fine point:

Note that `LONG INTEGER` and `LONG CARDINAL`, although ordered scalar types, are not element types, i.e., one cannot use long numeric values as array indices. As a notational convenience, it is possible to declare subranges of these types provided that the resulting subrange is in fact a valid subrange of `INTEGER` or `CARDINAL`.

3.1.1 Enumerated types

Consider the following declarations and a typical assignment:

```

channelState: INTEGER;
disconnected: INTEGER = 0;
busy: INTEGER = 1;
available: INTEGER = 2;
...
channelState ← busy;

```

Suppose `channelState` is a variable that is intended to range over a set of three “states” named `disconnected`, `busy`, and `available`, which are represented by values 0, 1, and 2. These values have no real significance; 5, 6, and 7 would serve equally well. Enumerated types are well suited to such an application (where the underlying values are unimportant). The above declarations could be replaced by a single declaration of a variable with an enumerated range:

```

channelState: {disconnected, busy, available};
...
channelState ← busy;

```

The effect is the same as before; `channelState` is a variable with values ranging over the same “states,” and similar assignment statements can be used.

The enumeration has some advantages over the original declarations:

It is more convenient; the programmer does not have to provide values for `disconnected`, `busy`, and `available`.

It allows more type checking. In the `INTEGER` case, one could assign any short numeric value to `channelState`.

It helps documentation; an enumeration shows all of its possible values.

An enumerated type is constructed by specifying a list of identifiers between braces, "{...}". These identifiers are not variables, but constants of that enumeration called *identifier constants*. They represent nothing more than their own names.

The type constructor **EnumerationTC** is defined as follows:

```
EnumerationTC ::= { IdList }
```

The **IdList** supplies all the identifier constants for the enumeration, and duplication of identifiers is illegal. *Separately specified enumerations are distinct*. Every appearance of an **EnumerationTC** generates a new type that is not equivalent to, and does not conform to, any other enumeration. Thus the declarations

```
foreground: {red, orange, yellow, green, blue, violet};
background: {red, orange, yellow, green, blue, violet};
```

specify two *different* enumerations. It is illegal to assign *background* to *foreground*, despite the fact that the same identifier list appears in each declaration. Occasionally, the inability to declare any further variables with the same type can be used to advantage by the programmer. Otherwise, the best way to avoid such problems is first to declare a type and then to declare variables using the identifier of that type; for example:

```
Color: TYPE = {red, orange, yellow, green, blue, violet};
foreground: Color;
background: Color;
```

This allows the assignment of *background* to *foreground* as well as the declaration of further variables with the same type (perhaps initialized differently).

The identifier constants in two different enumerated types have no association whatsoever and do not need to be distinct from one another. To identify unambiguously the enumeration from which a constant is taken, one can, and sometimes must, *qualify* the identifier constant by giving the name of the intended enumerated type. For example, given the additional declaration

```
Fruit: TYPE = {orange, lemon};
```

Color[orange] denotes a color and *Fruit[orange]* denotes a fruit. More generally, the syntax used for this form of qualification is

```
Primary ::= ... | TypIdentifier [ identifier ] |
          TypIdentifier.identifier
```

(This adds a new case to the syntactic definition of **Primary**, which already allows an identifier constant.) The "dot" form of qualification is discussed in subsection 3.3.3.

Often qualification is not necessary; for instance, the following is permitted:

```
hue: Color;
hue ← orange;                -- the type of hue implies Color[orange]
```

In the following situations, an identifier constant need not be qualified, because the intended enumerated type is established by the context:

- as the **RightSide** of an assignment
- as an initializing **Expression**
- as a component in an array or record constructor (§ 3.2.2 and 3.3.4)
- as an argument of a procedure (chapter 5)
- as an array index (§ 3.2)
- as the *right* operand of a **Relation**, including that part of a **Relation** used to label an arm in a discrimination (§ 4.3)
- as the bounds in a **SubrangeTC** (§ 3.1.2)

The values of an enumeration are ordered. The ordering is given by the order of appearance in the **IdList** used to construct the enumerated type. The leftmost identifier has the smallest value, and values increase from left to right. The following relations all have the value **TRUE**:

```
Color[red] < Color[orange]
Color[red] < violet
hue IN [red..yellow]           -- assuming hue = orange
```

There are two additional built-in functions that are applicable to enumerations: **FIRST [TypeSpecification]** yields the smallest value of the specified enumeration; e.g., **FIRST [Color]=red**. Similarly, **LAST [TypeSpecification]** produces the greatest value in an enumeration; e.g., **LAST [Color]=violet**. It is also possible to iterate over all values of an enumeration (§ 4.5).

The predefined type **BOOLEAN** is really an enumerated type, and its definition is

```
BOOLEAN: TYPE = {FALSE, TRUE};
```

Thus, **FALSE < TRUE**, **FIRST [BOOLEAN]=FALSE**, and **LAST [BOOLEAN]=TRUE**. The **BOOLEAN** constants **TRUE** and **FALSE** may always be used without qualification since Mesa contains the predefined symbols

```
TRUE : BOOLEAN = TRUE ;
FALSE : BOOLEAN = FALSE ;
```

The operators **ORD** and **VAL** provide a **LOOPHOLE**-free mechanism for converting between numbers and enumerated type values.

```
Color: TYPE = {red, orange, yellow, green, blue, indigo, violet};
c: Color;
y: Color ← yellow;
i: CARDINAL;
```

```

c ← VAL [3];           -- assigns green
i ← ORD [y];          -- assigns 2
i ← y.ORD;            -- also assigns 2
i ← Color.green.ORD;  -- assigns 3
i ← green.ORD;        -- illegal

```

Notice that the use of the adjective `Color` is required before the identifier `green` to establish the context of the intended enumerated type to which the `ORD` operator is to be applied.

The `VAL` operator must be used in a context where the desired type is known, such as assignment, parameter passing or any of the other situations described above for which qualification is not necessary.

Fine point:

“Dot notation” is a form of qualification that was used in Mesa originally to refer unambiguously to a named component of some record (§ 3.3.3). This notational formalism has been extended to a number of other situations requiring qualification, including the denotation of an identifier constant of an enumerated type (also § 6.4.4.1 and § 7.6.5). Thus, `Color.red` is equivalent to `Color[red]`.

3.1.1.1 Machine dependent enumerations

Sometimes a programmer can enumerate the values of some type but requires control of the encoding of each value or of the number of bits used to represent the type (usually for future expansion). Machine-dependent enumerations are provided for such applications.

Syntax

```

EnumerationTC      ::= MachineDependent { ElementList }

MachineDependent   ::= empty | MACHINE DEPENDENT

ElementList        ::= Element | ElementList , Element

Element            ::= identifier |
                    identifier (Expression) |
                    ( Expression )

```

Examples

```

Status: TYPE = MACHINE DEPENDENT {off(0), ready(1), busy(2), finished(4),
broken(7)}

```

```

Tint: TYPE = MACHINE DEPENDENT {red, blue, green, (255)}      -- reserve 8 bits

```

Each `Expression` in an `EnumerationTC` must denote a compile-time constant, the value of which is a `CARDINAL`.

In an enumerated type with the `MACHINE DEPENDENT` attribute, the values used to represent the enumeration constants are assigned according to the following rules. If a parenthesized expression follows the element identifier, the value of that expression is used; otherwise, the representation of an element is one greater than the representation of

the preceding element. If you specify only a representation, the corresponding element (normally a place holder) is anonymous. If the representation of the initial element is not given, the value zero is used.

You cannot explicitly specify the representation of any element unless the attribute **MACHINE DEPENDENT** appears in the type constructor. Two element identifiers cannot be represented by the same value (either given explicitly or determined implicitly as described above). The ordering of elements determined by position in the **ElementList** must agree with the ordering determined by the (unsigned) arithmetic ordering of the representations.

Sparse Enumerations

A machine-dependent enumerated type is *sparse* if there are gaps within the set of values used to represent the constants of that type or if the smallest such value is not zero. Mesa currently takes the following position on gaps: they are filled by valid but anonymous elements of the enumerated type. These elements can be generated only by the operators **FIRST**, **LAST**, **SUCC** and **PRED** (or by the iteration forms that implicitly use these operators). For example, **SUCC** [*busy*] is an anonymous element with the representation 3.

If you use a sparse enumerated type as the index type of an array, the array itself will have components for all elements of the enumeration, including the anonymous ones. The latter are awkward to access (except through **ALL**) and may cause problems in constructors, comparison operations, etc., as well as wasted space. (For example, **ARRAY** *Tint* **OF** **INTEGER** would occupy 256 words.)

3.1.2 Subrange types

In many cases, the values of a variable are inherently range-limited. For instance, a value for *day* (of the month) lies in the range [1..31]. In other cases, the range is limited by design. For instance, a value for *year* might be limited to the range [1900..1999]. Mesa permits the user to declare such variables in the following way:

```
day: CARDINAL [1 .. 31];
year: CARDINAL [1900 .. 1999];
```

Since these intervals cover a subrange of **CARDINAL**, the variables *day* and *year* are called *subrange variables*. It is useful to think of *day* and *year* as having type **CARDINAL** with the additional constraint that values are restricted to the specified intervals.

Subrange types have a number of advantages and uses. Subrange declarations unambiguously document the range of values intended for a variable and thus aid software maintenance. The compiler is able to optimize storage allocation when dealing with range-restricted variables (for example, in arranging the fields of a record, § 3.3) and can take advantage of subrange declarations to generate more efficient object code.

The general form of a **SubrangeTC** is

```
SubrangeTC      ::= TypelIdentifier Interval |
                  Interval
```

The **Typewriter** must evaluate to an **ElementType**. Thus, one can declare types that are subranges of **INTEGER**, **NATURAL**, **CARDINAL**, **CHARACTER**, **BOOLEAN**, enumerated types, and other subrange types. For example,

```

SymmetricRange: TYPE = INTEGER [-1..1];
PositiveInteger: TYPE = CARDINAL [1..LAST [INTEGER]];
UpperCaseLetter: TYPE = CHARACTER ['A..'Z'];
DegenerateType: TYPE = BOOLEAN [TRUE..TRUE];
CoolColor: TYPE = Color(yellow..LAST [Color]); -- excludes red, orange, yellow
AthroughM: TYPE = UpperCaseLetter['A..'M]; -- subrange of a subrange

```

The *base type* for a subrange is that type of which it is a subrange and which is not itself a subrange; e.g., the base type for both *UpperCaseLetter* and *AthroughM* is **CHARACTER**.

The **Expressions** that define the end points of an interval must have types that conform to the type denoted by the **Typewriter** (or yield short numeric values if the identifier is omitted). Also, *for the purpose of defining a subrange type, the end points must be compile-time constants*.

Fine point:

It is permissible for the interval defining a subrange type to be empty, e.g., [0..0). It is not legal to use a variable of such a type, but an empty subrange is sometimes useful for specifying the bounds of an array in a record declaration (§ 3.2).

A subrange type conforms to its base type, and a base type conforms to any of its subrange types. By extension, any two subrange types with the same base types are mutually conforming (even if they do not overlap in any way). A more revealing point of view is that the value of a subrange variable has the base type as its type, and an assignment of a value to a subrange variable makes an associated assertion that the value is in the appropriate interval. A violation of such an assertion is called a *range error*. *It is the programmer's responsibility to guard against range errors*. However, the compiler has an option that applies bounds checking to insert run-time tests to detect range errors. As implied by this viewpoint, appropriate literals of the base type serve as literals of the subrange type, and any operations defined on the base type automatically extend to the subrange type (but usually without closure).

Examples:

```

n: CARDINAL [0..10]; m: INTEGER [-5..5];

m ← 0; n ← 0; -- inherited literals
n ← n + 1; -- not valid if n = 10
n ← m; -- only valid if m IN [0..5]

```

The preceding discussion implies that subrange restrictions can be ignored in answering many type-related questions; in this sense, subrange types are “weak.” Two subrange types are equivalent if their base types are equivalent and if the corresponding bounds are equal. For these types, equivalence is much stronger than conformance. Equivalence becomes important when subrange types are used in the construction of other types.

FIRST and LAST are applicable to all subrange types and yield the corresponding bound. For example, FIRST [CoolColor] = green and LAST [AthroughM] = 'M'. It is also possible to iterate over all values in a subrange (§ 4.5).

3.1.2.1 Subranges of numeric types *

The description above applies to subranges of both enumerated and numeric types. Numeric subranges introduce one further complication, which is the question of representation. Omission of the initial `Typewriter` in a `SubrangeTC` is permissible if and only if each bound in the `Interval` specifies a short numeric value. In that case, `INTEGER` or `CARDINAL` is the base type, and the choice depends upon the representations of the bounds.

A numeric subrange type has a signed representation if both bounds are elements of `INTEGER` and at least one is not an element of `INTEGER` \cap `CARDINAL`. Similarly, it has an unsigned representation if both bounds are elements of `CARDINAL` and at least one is not an element of `INTEGER` \cap `CARDINAL`. If both bounds are elements of `INTEGER` \cap `CARDINAL`, values of that subrange type are considered to have *both* representations. Any other combination of bounds is illegal.

Examples:

```
s1: [-10..10];           -- signed representation
s2: [100..33000];       -- unsigned representation (if 33000 > LAST [INTEGER])
s3: [0..10];           -- both representations
```

Fine point:

There is currently a shortcoming in the symbol table representation that requires that the lower bound of a numeric subrange be a valid `INTEGER` value. Thus `CARDINAL [40000..40005]` is not a legal subrange type.

With respect to the choice of signed or unsigned versions of arithmetic and relational operators, a quantity with both representations is treated flexibly. When combined with an unsigned value, the quantity is considered to be unsigned; the unsigned operation and result are chosen. When combined with a signed value, the quantity is considered to be signed; the operation and result are signed. The rules governing combinations of values with both representations depend upon the context in which the result is used; the default is to choose signed representation and `INTEGER` operations. The precise rules are discussed in section 3.6.

Examples:

```
i: INTEGER; n: CARDINAL;    -- plus the declarations above

(signed)    s1 + 1
            s1 + s3
            s3 - i
```

```
(unsigned)    s2 + 1
              s2 + s3
              s3 * n
```

Fine point:

The representation assumed for a literal also depends upon context. In fact, any short numeric constant *c* is treated as if its type were [*c...c*].

3.1.2.2 Range assertions *

Assignment to a subrange variable implies an assertion about the range of the expression being assigned. The programmer may make such an assertion explicitly, for any expression, by using a *range assertion*. If *S* is an identifier of a subrange type and *e* is an expression with a type *T* conforming to *S*, the Primary *S* [*e*] has the same value as *e* and is additionally an assertion that *e* IN [FIRST [*S*∩*T*] .. LAST [*S*∩*T*]] is TRUE. In addition to user defined types, the basic types **INTEGER**, **NATURAL** and **CARDINAL** may be used in range assertions.

A program that violates one of its range assertions is in error. In addition to providing documentation and (optional) run time checking, a subrange assertion affects the attributes attached to an expression. For example, an assertion of an **INTEGER** range (or a signed subrange) forces the result to be treated as a value with signed representation. This is useful for controlling the choice of an operation when the intended one cannot correctly be inferred from the operands (§ 3.6).

Examples:

```
i: INTEGER; n: CARDINAL; S: TYPE = [0..10];

CARDINAL [i]          -- i is asserted to be nonnegative
S[n]                  -- asserts n IN [0..10]
```

3.2 Arrays

Arrays are indexable collections of homogeneous components. The components of a given array will have the same type, and each component corresponds to one index value in a range of indices associated with that array. The index range of an array is itself a type called an *index type*. The index type and *component type* together determine the type of the array.

```
earningsPerQuarter: ARRAY [1..4] OF INTEGER;
```

declares a variable with a constructed array type having an index type of [1..4] and a component type of **INTEGER**. Thus, *earningsPerQuarter* is an array of four integer elements: *earningsPerQuarter*[1], *earningsPerQuarter*[2], ... , *earningsPerQuarter*[4]. *earningsPerQuarter* by itself refers to the entire array variable. (Aggregate variables and components of aggregates are generally called "variables." If a distinction is needed, the term *component* is used and always means an item contained within an aggregate.)

An index type must be an **INTEGER**, **NATURAL**, **CARDINAL**, **CHARACTER**, **BOOLEAN**, **EnumerationTC**, or **SubrangeTC** (the element types). Ordinarily, one only uses subranges of **INTEGER** or

CARDINAL as an index type. A one-to-one correspondence exists between the components of an array and the values of the index type. This allows array elements to be accessed via "indexed references." An *indexed reference* selects and accesses the component corresponding to a particular index value. In its simplest form, it consists of the name of an array followed by a bracketed **Expression** with a type conforming to the array's index type.

An index type can be specified using a type identifier:

```
Quarter: TYPE = [1..4];
profit, loss, earnings: ARRAY Quarter OF INTEGER;
thisQuarter: Quarter;
...
earnings[thisQuarter] ← profit[thisQuarter] - loss[thisQuarter];
```

The arrays *profit*, *loss*, and *earnings* have *Quarter* as their index types, and *thisQuarter* is a subrange variable with type *Quarter*.

Index types may also be enumerations or subranges thereof. For example,

```
CallType: TYPE = {longDistance, tieLine, toll, local, inPlant};
nearbyCalls: ARRAY CallType[toll..inPlant] OF CARDINAL;
...
nearbyCalls[local] ← nearbyCalls[local] + 1;
```

Components may be of any desired type. In particular, the component type may itself be an array type. This allows an approximation of multidimensional arrays, which are otherwise absent in Mesa. For example, a two-dimensional data structure can be declared and used as follows:

```
Matrix3by4: TYPE = ARRAY [1..3] OF ARRAY [1..4] OF INTEGER;
mxy: Matrix3by4;
...
mxy[3][4] ← 0;                -- clear last component.
```

In the assignment statement, *mxy* is an expression of array type (with index type [1..3] and component type **ARRAY [1..4] OF INTEGER**). *mxy*[3] is an indexed reference to the third component of *mxy*. This in turn yields an expression of array type (with index type [1..4] and component type **INTEGER**). Thus, *mxy*[3][4] is an indexed reference to the fourth component of that subarray. Because of left-associativity, *mxy*[3][4] is the same as (*mxy*[3])[4].

An *array constructor* (§ 3.2.2) consists of an optional type identifier followed by a list of values (syntactically, **Expressions**) enclosed in brackets. The list specifies values for components of an array in index order. The declaration below uses an array constructor to initialize an array that can be used as a translation table; i.e., *octalChar*[*n*] holds the character denoting octal digit *n*:

```
octalChar: ARRAY [0..7] OF CHARACTER = ['0', '1', '2', '3', '4', '5', '6', '7'];
```

Note that the number of values in the list (eight) matches the number of indices in the index type. *This is required for array constructors.* A special form using the replicator **ALL**

is available for abbreviating array constructors in which all components have the same value. For example, the following two declarations are equivalent:

```
dashes: ARRAY [0..7] OF CHARACTER ← ['-', '-', '-', '-', '-', '-', '-', '-'];
dashes: ARRAY [0..7] OF CHARACTER ← ALL ['-];
```

Array variables may also be initialized using other array values. Consider the following example:

```
freshVector: ARRAY [0..3] OF CARDINAL = ALL [0];
currentVector: ARRAY [0..3] OF CARDINAL ← freshVector;
```

In this case, *currentVector* is initialized with *freshVector*'s value, i.e., all three of *currentVector*'s elements are initially set to zero. Because the declaration of *freshVector* uses fixed initialization, assignment either to the entire array or to one of its elements is illegal.

When the operands of any of the fundamental operations (\leftarrow , $=$, $\#$) are arrays, the operation is applied on a component-by-component basis. The initialization of *currentVector* above uses assignment in this way. Similarly, the expression "*currentVector* = *freshVector*" yields the result **TRUE** if and only if all three components of each array are equal (as they are in the above example).

3.2.1 Declaration of arrays

Arrays are declared using the *array type constructor*, **ArrayTC**:

```
ArrayTC           ::= PackingOption ARRAY IndexType OF ComponentType

PackingOption    ::= empty |           -- elements word aligned
                    PACKED           -- elements potentially packed within words

IndexType       ::= ElementType |
                    TypeIdentifier

ComponentType   ::= TypeSpecification
```

Two array types are equivalent if both their index types and their component types are equivalent and if they are both packed or both unpacked (see below). An array type conforms to another if the two types are equivalent. Conforming arrays need not be declared together (unlike **RECORD** or enumerated declarations where separately declared types are unique even if they look the same (§3.2.2)). For example:

```
IndexType : TYPE = [0..10)
ArrayType1 : TYPE = ARRAY IndexType OF INTEGER ;
ArrayType2 : TYPE = ARRAY [0..10) OF INTEGER ;
ArrayType3 : TYPE = ARRAY [0..10) OF INTEGER ;
Array1 : ArrayType1 ; Array2 : ArrayType2 ; Array3 : ArrayType3 ;
```

Arrays *Array1*, *Array2*, and *Array3* all conform to one another. Thus, it is possible to assign or compare array variables with separately constructed types if those types are structurally identical.

Fine point:

In addition, one array type freely conforms to another if the component type of the first freely conforms to that of the second, the index types are equivalent, and they are both packed or both unpacked (§ 3.5). Packed array types with non-equivalent component types do not freely conform, as their elements may well occupy different numbers of bits within a word.

Declarations of initialized array variables take the form

IdList : ArrayTC Initialization

The initializing expression must have an array type that conforms to the one being declared.

The previous section describes indexed references to array components. A formal definition follows:

IndexedReference ::= **Variable** [**Expression**] |
(**Expression**) [**Expression**]

LeftSide ::= ... | **IndexedReference**

The **Variable** or parenthesized **Expression** must be of some array type, and the bracketed **Expression** must conform to the index type for that array type. An **IndexedReference** is itself part of the definition of a **Leftside** (and therefore of a **Variable**, § 2.5).

Fine points:

If you specify the **PACKED** attribute for an array type, the granularity of packing is 1, 2, 4, 8 or 16n bits and is determined by the component type of the array. Unless an array is packed, each component is "aligned," i.e., begins on a word boundary. Thus a packed array of **CHARACTER** wastes no space.

Since packed array elements are not necessarily word aligned, one cannot use the @ operator (§ 3.4) to generate the address of an element.

The value of the construct **SIZE** [T, n] is the size, in words, of the storage required by a packed array of n items of type T.

The *length* of an array is the number of its elements. For variables with an array type, the length is fixed and known at compile-time. (Dynamic arrays are possible in Mesa through the use of array descriptors, discussed in subsection 6.2.1 or sequences discussed in section 6.5.)

The **IndexType** of an array may legally be an empty interval. In this case, no storage is allocated for the array. This is useful when the array appears as the last component of a **RECORD** (§ 3.3) and the user will be obtaining storage for each record plus some number of array elements from a free storage manager. Note that [0..0] is not equivalent to [1..1], since the intervals specify different initial indices for the array. The use of sequences (§6.5) instead of dummy arrays is strongly encouraged.

Three function-like operators are relevant to arrays (and more relevant to array descriptors): **LENGTH**, **BASE**, and **DESCRIPTOR**. These are discussed in section 6.2, but a brief summary is provided below. For this summary, *arg* denotes an expression with some array type.

LENGTH (<i>arg</i>)	-- yields the number of array elements.
BASE (<i>arg</i>)	-- yields a pointer value for locating the first array element.
DESCRIPTOR (<i>arg</i>)	-- yields <i>arg</i> 's array descriptor value (consisting of base and length).

3.2.2 Array constructors

In the preceding examples, array constructors are used only for initialization. Actually, constructors for arrays may be used in any **RightSide** context. An array constructor is defined as follows:

Primary	:: =	Constructor ...
Constructor	:: =	OptionalTypeId [ComponentList] ALL [Component]
OptionalTypeId	:: =	TypeIdentifier empty
ComponentList	:: =	PositionalComponentList ... -- other forms for record constructors
PositionalComponentList	:: =	Component PositionalComponentList , Component
Component	:: =	empty -- elided component Expression -- explicit component NULL -- voided component

The **empty** components in a constructor are said to be *elided*, and **NULL** components are said to be *voided*. The values of both elided and voided components are undefined when the component type does not have a default value (§ 3.3.5). In the first form of array constructor, using **OptionalTypeId** [**ComponentList**], the number of **Expressions** plus elided or voided components must match the length implied by the array type. The type of each **Expression** must conform to the array's component type. The expressions (and elided or voided components) are taken in order to form a sequence that is the constructed array value.

Consider the following example:

```
Triple: TYPE = ARRAY [1..3] OF CARDINAL;
triplet: Triple ← Triple[11, 12, 13];
```

The declaration assigns 11 to *triplet*[1], 12 to *triplet*[2] and 13 to *triplet*[3].

When the array type is implied by context, the **TypeIdentifier** may be omitted. Thus, the declaration above could be written as

```
triplet: Triple ← [11, 12, 13];
```

Taken out of context, the constructor [11, 12, 13] is ambiguous; it could be assigned to *any* array of three numeric elements; for example:

```
trio: ARRAY {Patty, Laverne, Maxine} OF LONG INTEGER ← [11, 12, 13];
```

The second form of constructor, using ALL, is *only* valid when the array type is implied by context. The type of the Expression must conform to the array's component type. The value of the constructor is an array in which the specified value is replicated a number of times equal to the length of the array. The expression is evaluated just once. In the case of an array of arrays, the structure must be mirrored by nesting in the constructor, as shown by the following example:

```
allOnes: Matrix3by4 ← ALL [ALL [1]];
```

Fine points:

The value of an elided or voided component of an array constructor is not defined, but it will have *some* value. In particular, if the statement

```
triplet ← [1, . 3];
```

is executed after the previous assignment to *triplet*, the value of *triplet*[2] is undefined. If the element type of the array has a default (§3.7), then the elided element will have that value.

Any array constructor in which all components are compile-time constants is a compile-time constant. Also, selection from an array that is a compile-time constant using a constant index yields a compile-time constant.

3.2.3 Keyword array constructors

If the index type of the array is an enumeration or a subrange thereof, one can use a *keyword array constructor* where the individual element positions are named. The acceptable keywords are the constants appearing in the enumeration. In the case of a subrange, the endpoints must be defined by expressions involving only those constants, the operators FIRST, LAST, SUCC and PRED, and identifiers equated by declaration to such expressions. It is possible to specify the same optional default value (§3.7) for each component of an array by assigning a default value (including NULL) to the component type of the array at the time that component type is declared. If this is done, keyword items can be omitted, the corresponding elements receive the default value. For example,

```
Element : TYPE = {red, green, blue};
BoolFalse : TYPE = BOOLEAN ← FALSE;
Set : TYPE = PACKED ARRAY Element OF BoolFalse;
s, t : Set;
s ← [red : TRUE];                -- equivalent to [TRUE, FALSE, FALSE]
t ← [red : TRUE, blue : TRUE]    -- equivalent to [TRUE, FALSE, TRUE]
```

The rules for when components of an array or record constructor may be omitted, elided, or voided are explained more fully in subsection 3.3.5.

3.3 Records

A *record* is an aggregate that allows a group of related data items of different types to be packaged together. In the definition of a record type, the type of each individual component must be supplied, as in the following example:

```
MilitaryTime: TYPE = RECORD [hrs: [0..24), mins: [0..60)];
oldTime, newTime: MilitaryTime;
```

Here, *MilitaryTime* is a newly defined type, and *oldTime* and *newTime* are record variables of that type. *MilitaryTime* is a two-component record type, where the first record component is named *hrs* and the second *mins*. Every *MilitaryTime* record contains both components, but different record objects have their own values for these components.

A constructor of a record type contains a *field list* after the word **RECORD**. Each element in the list specifies one (or more) components of the record. For *MilitaryTime*, the field list is [*hrs*: [0..24), *mins*: [0..60)]. The component names, *hrs* and *mins*, are called *field names*. They are used to refer to components in any *MilitaryTime* record. For instance, the first component of *oldTime* may be selected using the *qualified reference*, "*oldTime.hrs*."

One can construct an entire record value using a *record constructor*. For instance, the constructors below yield *MilitaryTime* values with *hrs* components that have the value 13 and *mins* components that have the value of the expression "*y + 1*":

```
MilitaryTime[13, y + 1]
MilitaryTime[hrs: 13, mins: y + 1]
MilitaryTime[mins: y + 1, hrs: 13]
```

The second constructor is an example of a *keyword constructor*, since it specifies the name of the component (e.g., as "*hrs*:") with which a value is to be associated. The third example shows that record components need not be specified in order of field declaration provided that keywords notation is used in the record constructor.

A default value (§3.3.5) can be specified for any field in the definition of a record type. The default is used in constructing records of that type when no value is specified in the constructor. Defaults are useful for suppressing detail and ensuring initialization of fields. In the following example, the two constructors have the same value:

```
Datum: TYPE = RECORD
[
  value: INTEGER,
  nReads: CARDINAL ← 0,
  nWrites: CARDINAL ← 1
];

Datum[x]
Datum[value: x, nReads: 0, nWrites: 1]
```

The basic operations on (non-variant) record values include the fundamental operations (=, #, ←), qualification, the unambiguous reference to a named component of some record, and extraction, the expansion of record objects and assignment of their components

to individual variables in a single statement. Variant records are records of the same type that do not necessarily contain the same components. They are discussed in chapter 6.

3.3.1 Field lists

There are two kinds of field lists, depending on whether the fields are “named” or “unnamed.” (Field lists used to construct multi-component record types are almost always named.)

Syntax equations:

```

FieldList           ::= [ UnnamedFieldList ] | [ NamedFieldList ]

UnnamedFieldList   ::= TypeSpecification |
                    TypeSpecification , UnnamedFieldList

NamedFieldList     ::= IdList : FieldDescription DefaultOption |
                    NamedFieldList , IdList : FieldDescription
                    DefaultOption

FieldDescription    ::= TypeSpecification

DefaultOption      ::= empty | ← DefaultSpecification -- see subsection 3.3.5

```

Examples:

```

[i: INTEGER, b: BOOLEAN, c: CHARACTER] -- a named field list
[INTEGER, BOOLEAN, CHARACTER]         -- a similar, but unnamed field list
[f1: CHARACTER, f2, f3: INTEGER]       -- components listed and declared together
[f1: CHARACTER, f2: INTEGER, f3: INTEGER] -- equivalent to the previous

```

Note that if one field is named, *all* must be named. Also, field names must be unique within a given field list. (The same identifiers may be used as field names in other field lists, however, or as names of declared variables.)

Field descriptions in a named field list contain a type specification, indicating the type of the field. Any type may be specified, including an array type or (some other) record type.

Fine points:

A field’s type specification must not imply an infinite nesting of records. For instance, the following is illegal:

```

A: TYPE = RECORD {b: B};
B: TYPE = RECORD {a: A};

```

Field lists occur in constructors of types other than records, such as PROCEDURES (chapter 5), SIGNALS (chapter 8), and in variant record specifications (chapter 6).

Unnamed field lists are normally used when component names would be ignored if they were present. This is common for functions that return single-component results. Unnamed field lists are sometimes used in specifying the input parameters for procedure variables that are to be set to one of several actual

procedures. (However, an unnamed field list does not allow Calls using such a procedure variable to refer to the parameters by name.)

3.3.2 Declaration of records

The type constructor `RecordTC` is defined as follows:

```
RecordTC      ::= RECORD FieldList |
                ...           -- plus variant records (chapter 6)
```

where `FieldList` is defined in the previous section. Separately declared record types are unique, *even if they look the same*. Every appearance of a record constructor creates a new type that is not equivalent to, and does not conform to, any other record type. In the example:

```
RecType1: TYPE = RECORD [a,b: INTEGER];
rec1: RecType1;

RecType2: TYPE = RECORD [a,b: INTEGER];
rec2: RecType2;

rec3: RECORD [a,b: INTEGER];
rec4: RECORD [a,b: INTEGER];
```

the record variables `rec1`, `rec2`, `rec3`, and `rec4` all have different, non-conforming types. None of these can be assigned to any of the others (despite the similarity of their components). It is, of course, legal to assign to a component any value with a conforming type. For example:

```
rec1.a ← rec2.b ← rec3.a ← 5;
rec4.a ← rec1.a; rec4.b ← rec1.b;
```

Any single-component record type conforms to the type of its single component, but not vice versa. The automatic conversion in this case requires no computation.

Example:

```
Bundle: TYPE = RECORD [value: INTEGER];
recVar: Bundle;
intVar: INTEGER;
...
intVar ← recVar;           -- means intVar ← recVar.value
intVar ← recVar + 1;       -- operand conversion
recVar ← Bundle[intVar];   -- a constructor
recVar.value ← intVar;
```

This conversion simplifies dealing with functions that return single-component records (chapter 5). It also provides a way of partitioning a set of variables that can be checked by the type system. In the example above, a direct assignment of `intVar` to `recVar` is invalid. Furthermore, no other single-component record type, such as

```
Prime: TYPE = RECORD [value: INTEGER];
```

can be confused with *Bundle*; assignment of a *Bundle* value to a *Prime*, or a *Prime* to a *Bundle*, is illegal. Either a *Bundle* or a *Prime* can, however, appear as a numeric operand. Defining *Bundle* and *Prime* as synonyms for `INTEGER` would not provide this additional checking.

Because of the uniqueness of constructed record types, record variables are typically declared in two steps: first the record type, then the record variables. The general form is:

```

identifier : TYPE = RecordTC;           -- define record type.
IdList : identifier Initialization ;      -- same identifier as just defined

```

Record variables can also be declared directly:

```

IdList : RecordTC Initialization ;

```

This form is not very useful because the (anonymous) record type is not available for purposes such as declaring other records of the same type or writing constructors.

The **Initialization** shown in these general forms applies to the entire record variable, not to individual components. Any **Initialization** must have the proper record type. Initialization of record variables is shown in the next example.

```

noon:           MilitaryTime = [hrs:12, mins:0];
midnight:      MilitaryTime = [hrs:0, mins:0];
time:           MilitaryTime ← midnight;      -- start time at midnight.

```

Fine points:

The Mesa compiler packs record components into machine words. The components may be arranged in an order that differs from the left-to-right order of the fields in the type constructor. All records of the same type have the same component arrangement.

Normally, the user is unconcerned with the actual arrangement of record components. When component arrangement is important, the user may specify "MACHINE DEPENDENT" records (§ 3.3.6).

Except in **MACHINE DEPENDENT** records, components are packed for storage efficiency. Some fields may be aligned (to the beginning of a word boundary) and some may not. Components occupying a full word or more are always aligned: arrays, `INTEGERs` and pointers, for example. Subrecords may or may not be aligned, depending on their size. Packed arrays that don't fit completely within a word are aligned, even if there would have been space in the preceding word for some of the elements. As an example:

```

FourBits: TYPE = PACKED ARRAY [0..4] OF BOOLEAN;
Record: TYPE = RECORD(a, b, c, d: FourBits);      -- fits in a single word

```

The function-like operator **SIZE** (§ 2.5.6) is often used to find the number of machine words occupied by a record of some type.

3.3.3 Qualified references

Qualification is used to refer unambiguously to a named component of some record. The general form (which extends the definition of a `LeftSide`) is

```
QualifiedReference ::= Variable . identifier |
                  ( Expression ) . identifier
```

```
LeftSide          ::= ... | QualifiedReference
```

The field name is said to be "qualified by" the record value (the `Variable` or `Expression`) to the left of the dot. The operator associates from left-to-right in the case of multiple qualification. For example:

```
Latitude:        TYPE = RECORD [degs:[0..360], mins, secs:[0..60)];
Longitude:       TYPE = RECORD [degs:[-90..90], mins, secs:[0..60)];
Position:        TYPE = RECORD [latitude: Latitude, longitude: Longitude];
somePosition: Position;
```

Some of the possible qualified references to components of `somePosition` are listed below:

<u>Qualified Reference</u>	<u>Refers To</u>
<code>somePosition.latitude</code>	1st sub-record
<code>somePosition.longitude</code>	2nd sub-record
<code>somePosition.latitude.degs</code>	1st component of 1st sub-record
<code>somePosition.longitude.secs</code>	3rd component of 2nd sub-record

The association order for qualification means that names must occur in the proper sequence; e.g., `somePosition.mins.longitude` is incorrect. Also, a qualified reference must be complete, i.e., names may not be skipped (as in `somePosition.secs`, which would be ambiguous in any event).

Qualified references and indexed references have the same precedence (the highest possible) and may be intermixed. For example:

```
recordOfArrays: RECORD [a,b: ARRAY [0..100] OF CARDINAL];
arrayOfRecords:[1..5] OF RECORD [i1,i2,i3: CARDINAL];
...
arrayOfRecords[5].i3 ← recordOfArrays.a[0];    -- ("last" gets "first")
```

Fine point:

Qualification briefly opens up a given "name scope." For instance, in the record qualification, `rec.x`, the qualified name, `x`, must name a field of `rec` and selects that field. Scope is treated more fully in chapter 7.

3.3.4 Record constructors

A record constructor assembles a record value from a set of component values. In the following example, a constructor is used as a **RightSide** of an assignment.

```
MonthName: TYPE = {Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec};
```

```
Date: TYPE = RECORD
```

```
  [
    day: [1 .. 31],
    month: MonthName,
    year: [1900 .. 2000)
  ];
```

```
birthDay: Date;
now: [1900..2000) ← 1984;
birthDay ← Date[25, Apr, now - 33];
```

This constructor yields a record value with type *Date*. The record assigned to *birthDay* contains the following component values:

<u>Component</u>	<u>Value</u>	
<i>day</i>	25	
<i>month</i>	<i>Apr</i>	
<i>year</i>	<i>now</i> - 33	(which is 1951)

A **Constructor** is a **Primary** and may not be used as the **LeftSide** of an assignment.

Record constructors are of two kinds: *keyword constructors* and *positional constructors*. Within both kinds, the component value for a particular field may either be supplied or be omitted. If it is omitted, the value of the field is determined by the **DefaultOption** appearing in the declaration of the field (§ 3.3.5).

Syntax equations:

```
Primary ::= ... | Constructor

Constructor ::= OptionalTypeId [ ComponentList ]

OptionalTypeId ::= TypeIdentifier | empty

ComponentList ::= KeywordComponentList |
PositionalComponentList

KeywordComponentList ::= KeywordComponent |
KeywordComponentList , KeywordComponent

PositionalComponentList ::= Component |
PositionalComponentList , Component

KeywordComponent ::= identifier : Component
```

Component	:: = empty	-- elided component
	Expression	-- explicit component
	NULL	-- voided component

The initial `TypeIdentifier`, if present, must name the type of the record being constructed.

In *keyword constructors*, the correspondence between constructor components and record components is strictly "by name." *Keyword names may not be repeated in a constructor*, but the order is irrelevant. For example, the following keyword constructors are equivalent:

```
Date[day: 25, month: Apr, year: now - 33]
Date[month: Apr, day: 25, year: now - 33]
```

All of these keyword constructors specify values for all the components. In the following example, the first keyword constructor *elides* the *month* component (the place for the component value is specified, but no value is given); the second *voids* the *month* component (by specifying `NULL` instead of a value):

```
Date[day: 25, month: , year: now - 33]           -- month is elided
Date[day: 25, month: NULL, year: now - 33]      -- month is voided
```

The distinction between an elided and a voided field arises in the treatment of defaults (§ 3.3.5). Since the declaration of *Date* specifies no default value for *month*, both of these examples construct records with a second component that has an undefined value.

In a *positional constructor*, the correspondence between constructor components and record components is strictly "by position." The first constructor component corresponds to the first record component, the second value to the second component, etc. Positional constructors may be used for both records and arrays (§ 3.2.2). It does not matter whether or not fields are named in the definition of the record type. The following three constructors are equivalent:

```
Date[day: 25, month: , year: now - 33]          -- value of month is undefined (elided)
Date[25, , now - 33]                            -- value of 2nd component is undefined (elided)
Date[25, NULL, now - 33]                       -- value of 2nd component is undefined (voided)
```

Positional constructors may elide or void components as shown above, and trailing fields (only) can be omitted by omitting the final commas. The positional constructor `[]` is considered to omit, not elide, its first component.

Keyword and positional notations may not be mixed in a single constructor. The order of evaluation of components is not specified for either kind of constructor.

The initial `TypeIdentifier` in a constructor may be omitted when the constructor is used as:

- the `RightSide` of an assignment (unless the `LeftSide` is an extractor, § 3.3.6)
- an expression in an `Initialization`
- a component of an enclosing record or array constructor
- an argument of a procedure

the right operand of a **Relation**.

In other cases, an initial **TypeIdentifier** must appear. It is never incorrect to supply the identifier, and sometimes doing so improves readability.

Fine point:

Any record constructor in which all components are compile-time constants is a compile-time constant. Also, a field selected from a record that is a compile-time constant is itself a compile-time constant.

3.3.5 Default field values

The definition of most types, including those in this chapter, allows a default specification for the type being defined. The default specification for a type is most useful when that type is used as a component of a record or as the component type of an array (§3.7).

The definition of a record type is unique in that it may specify a default value for each field of the record overriding, if necessary, any default specification the particular component type may have already had. All these default specifications are optional; if present, they can be used in constructing records and arrays when no values for the corresponding fields are specified in constructors. However, there are different ways of not specifying a value for a field in a constructor. An elided field, as discussed in the preceding section, supplies the field, but supplies no value. An omitted field is simply not present at all. In a keyword constructor, a field is omitted by omitting the keyword entirely; in a positional constructor, trailing fields (only) can be omitted by omitting the final commas. The positional constructor “[]” is considered to omit, not elide, its first component. Fields in an array constructor, however, can only be omitted if the index type of the array is an enumeration or a subrange thereof (§3.2.3). A field is voided by specifying the word **NULL** or **TRASH** as the value of the field (§3.7). A discussion of the semantics of omit, elide, and void follows the default specification given below.

In the following example, all constructors have the same value.

Interval: **TYPE = RECORD**

```
[
  range: INTEGER,
  origin: INTEGER ← 0,
  direction: {up, down} ← up
];
```

<i>Interval</i> [range: 10, origin: 0, direction: up]	-- all fields specified
<i>Interval</i> [range: 10, origin: , direction:]	-- origin, direction elided
<i>Interval</i> [range: 10]	-- origin, direction omitted
<i>Interval</i> [10]	-- origin, direction omitted (positional form)

The syntax for specifying defaults in a **NamedFieldList** or for any type of declaration (§3.7) follows:

```

DefaultOption      ::= empty | ← DefaultSpecification

DefaultSpecification ::= empty |
                       Expression |
                       TRASH |
                       NULL |
                       Expression | TRASH |
                       Expression | NULL

```

Note: In the final two lines, the vertical bar denotes itself and is embedded within an alternative.

In a **DefaultSpecification**, the **Expression** must have a type that conforms to the type of the corresponding field.

Suppose that R is a record type with a field v of type T or that v is simply a variable that has been defined to have type T . The above syntax allows five forms for the **DefaultOption** in the declaration of v . No matter which form is used, a constructor of an R may explicitly specify a value for the field v . The various options control whether the existence of the field must be made evident in the constructor, whether an explicit value must be supplied and, if not, what action is taken. The options are interpreted as follows:

- (1) $v: T$
In a constructor, the value of v can be left undefined, but that must be indicated explicitly, by eliding or voiding the field. v cannot be omitted. This rule also applies to an unnamed field.
- (2) $v: T \leftarrow$
Every constructor must supply an explicit value (not **NULL**) for v .
- (3) $v: T \leftarrow e$
If a constructor elides or omits v , the value of the expression e is used; voiding the field is not permitted.
- (4) $v: T \leftarrow \text{TRASH}$
As in (1) above, except that the constructor may omit v entirely. If the field is omitted, elided or voided, its value is undefined.
- (5) $v: T \leftarrow e | \text{TRASH}$
As in (3) above, except that a constructor may explicitly void v . If the field is omitted or elided, the value of e is used; if it is voided, its value is undefined.

If the first or second form is used, the field cannot be omitted from a constructor; these forms are useful when such omission is likely to indicate a programming error. Omission is permitted by the other forms, which differ in the default action for an omitted or elided field. These forms are appropriate when a field has some common and meaningful default value (the third and fifth cases) or, alternatively, is relevant only in unusual circumstances (the fourth case). The last three forms are particularly suitable for extending the definition of a record type; constructors in existing programs need not be modified.

In the fourth and fifth cases, **NULL** can be used in place of **TRASH** with the same results.

Fine points:

The second form of field declaration guarantees that the field *v* has a well-defined value. Constructors cannot void this field, and declaration of a record of this type (or allocation by **NEW**, § 6.6.2) must include a constructor that sets this field.

If the **Expression** form of a default specification is used, that expression is evaluated at the time of construction but in the context of the declaration of the record, i.e., the expression is treated as a parameterless procedure invoked by evaluation of the constructor (see chapter 5).

The default value of a field cannot be specified in terms of other fields in the same record. Default values for fields of record types defined in **DEFINITIONS** modules (§ 7.3.2) must be compile-time constants.

Examples:

```
R: TYPE = RECORD
  [
    v1: CARDINAL,
    v2: CARDINAL ←,
    v3: CARDINAL ← 3,
    v4: CARDINAL ← TRASH,
    v5: CARDINAL ← 5 | NULL
  ];
-- the following are valid
R[v1: 1, v2: 2]           -- means R[v1: 1, v2: 2, v3: 3, v4: TRASH, v5: 5]
R[v1: , v2: 2, v5: ]     -- means R[v1: , v2: 2, v3: 3, v4: TRASH, v5: 5]
R[v1: 1, v2: 2, v5: NULL] -- means R[v1: 1, v2: 2, v3: 3, v4: TRASH, v5: NULL]

-- the following are not valid
R[]                       -- neither v1 or v2 may be omitted
R[v1: 1, v2: NULL, v3: NULL] -- neither v2 nor v3 may be voided
```

3.3.6 Extractors

Extractors are used to “explode” record objects and assign their components to individual variables in a single statement. For example, the extractor below assigns the components of *birthDay* (defined in subsection 3.3.4) to the variables *dd*, *mm*, and *yy*, in that order:

```
dd: [1..31]; mm: MonthName; yy: [1900..2000];
[dd, mm, yy] ← birthDay;
```

This has the same effect as the following three separate assignments, except that *birthDay* is evaluated only once:

```
dd ← birthDay.day; mm ← birthDay.month; yy ← birthDay.year;
```

An extractor resembles a constructor in form, but there are some important differences:

The “components” of an extractor specify **LeftSides**, not **Expressions**.

Extractors always begin with a left bracket, never with a **TypeIdentifier**.

The type of the record value assigned to an extractor must be known to the compiler. This means that the following (rather useless) statement is invalid because the constructor's type cannot be determined:

```
[dd, mm, yy ] ← [25, Apr, 1943];           -- invalid
```

The statement should specify the type of the constructed value:

```
[dd, mm, yy ] ← Date[25, Apr, 1943];       -- valid
```

Extractors, like constructors, may use keywords. This allows an extractor to be written without regard to the record's component order. For instance, the following statements are equivalent to the first one in this section:

```
[day: dd, month: mm, year: yy ] ← birthDay;
[month: mm, day: dd, year: yy ] ← birthDay;
```

Extractors may elide or omit *any* item, in which case the corresponding record component is not assigned. The extractors shown below are equivalent:

```
[day: dd, month: , year: yy ] ← birthDay;   -- month elided
[day: dd, year: yy ] ← birthDay;           -- month omitted
[dd, , yy ] ← birthDay;                   -- 2nd component elided
```

A positional extractor may omit trailing components without supplying trailing commas. The *year* component of *birthDay* is omitted below.

```
[dd, mm] ← birthDay;
```

Syntax equations:

```
AssignmentStmt  ::= ... | Extractor ← RightSide

Extractor       ::= [ KeywordExtractList ]
                  [ PositionalExtractList ]

KeywordExtractList ::= KeywordExtract |
                        KeywordExtract , KeywordExtractList

KeywordExtract  ::= identifier : ExtractItem

PositionalExtractList ::= ExtractItem |
                        ExtractItem , PositionalExtractList

ExtractItem     ::= empty |           -- component is ignored
                  LeftSide          -- component is assigned to LeftSide
```

The identifiers in a **KeywordExtractList** must be field names for the record type. Note that an extraction list can be empty, in which case the effect is to discard a record value.

Fine point:

An **ExtractItem** may itself be an **Extractor**. This allows extraction from records embedded within records. In particular, this form is useful in situations where a single-component record is not

automatically converted to its single component, e.g., extraction from a record that is the only component of a return record.

3.3.7 Machine dependent records

Machine-dependent records are provided for situations in which the exact position of each field is important. (A first time reader is referred to section 6.4 for a complete discussion of variant records.) You can explicitly specify word- and bit-positions in the declaration of the record type.

Syntax

```

VariantFieldList      ::= CommonPart FieldId : Access VariantPart |
                          VariantPart |
                          NamedFieldList |
                          UnnamedFieldList |
                          empty

CommonPart            ::= NamedFieldList , | empty

NamedFieldList       ::= NamedField | NamedFieldList , NamedField

NamedField           ::= FieldIdList :
                          Access TypeSpecification DefaultOption

FieldIdList          ::= FieldId | FieldIdList , FieldId

FieldId              ::= identifier | identifier ( FieldPosition )

Tag                  ::= FieldId | ...

FieldPosition        ::= Expression : Expression .. Expression |
                          Expression

```

Examples

```

InterruptWord: TYPE = MACHINE DEPENDENT RECORD [
  channel (0: 8..10): [0..nChannels),          -- nChannels <= 8
  device (0: 0..7): DeviceNumber,
  stopCode (0: 11..15): MACHINE DEPENDENT {finishedOK(0), errorStop(1), powerOff(3)},
  command (1: 0..31): ChannelCommand];

```

```

Node: TYPE = MACHINE DEPENDENT RECORD [
  type (0: 0..15): TypeIndex,
  rator (1: 0..13): OpName,
  rands (1: 14..47): SELECT valence (1: 14..15): * FROM
    nonary = > [],
    unary = > [left (2): POINTER TO Node],
    binary = > [left (2), right (3): POINTER TO Node]
  ENDCASE]

```

An identifier with an explicitly specified **FieldPosition** can occur only in the declaration of a field of a record defined to have the **MACHINE DEPENDENT** attribute. If the position of any

field of a record is specified, the positions of all must be. Each **Expression** in a **FieldPosition** must denote a compile-time constant, the value of which is an unsigned integer.

The first expression appearing in a **FieldPosition** specifies the (zero-origin) record-relative index of the word containing the start of the field; the second and third specify the indices (zero-origin) of the first and last bits of the field with respect to that word. The second and third expressions may specify a bit offset greater than the word size if the word offset is adjusted accordingly. Similarly, the difference between the second and third expressions may exceed the word size. If the bit positions are not specified, a specification of $0..n*WordSize-1$ is assumed, where n is the minimum number of words required by the type of the field. Note that in the preceding example, in the field position for *left*, a single word quantity could be given equivalently as (1: 16..31) or as (2: 0..15).

Each field must be at least wide enough to store any value of the corresponding type. Values are stored right-justified within the fields. The current implementation of Mesa imposes the following additional restrictions on the sizes and alignment of fields:

A field smaller than a word (16 bits) cannot cross a word boundary.

Any field occupying a word or more must begin at bit zero of a word and have a size that is a multiple of the word size.

A variant part (§ 6.4) may begin at any bit position (as determined by its tag field); the tag must be the first thing in the variant part.

If the sizes of all variants of a record type are less than a word, those sizes must be equal; otherwise, the size of each variant of the type must be a multiple of the word length.

In the definition of a machine-dependent record type, explicitly specified field positions must not overlap. For a variant record type, this requirement applies to the variant part (including the tag) considered in conjunction with the fields of the common part; the tag and fields particular to each variant must lie entirely within the variant part.

The order of fields in a record type declaration need not agree with the order of those fields in the representation of the record; however, no gaps are permitted. For variant records, the fields of at least one variant (including the tag field) must fill the position specified for the variant part.

Fine point:

An earlier, purely positional, form of machine dependent record remains for compatibility. The new form is encouraged because it provides better documentation and is usually easier to use. An example of the purely positional form:

InterruptWord: **TYPE = MACHINE DEPENDENT RECORD**

```
[
  device: DeviceNumber,
  channel: [0..7],
  stopCode: {finishedOk, errorStop, powerOff},
```

```
command: ChannelCommand
;
```

In this case, the user takes full responsibility for component arrangement. Components are positioned exactly as given, from left to right in machine words. In general, "fill" components are needed to ensure that no field crosses a word boundary (unless it starts on one). Components (such as *ChannelCommand*) may themselves be aggregates occupying more than one word

It is also the user's responsibility to "fill out" the record to a full word if the record crosses a word boundary. (*InterruptWord* might be correct for a 16-bit machine, but not for a machine having a larger word length).

3.4 The types `POINTER` and `LONG POINTER`

`POINTERS` and `LONG POINTERS` provide efficient indirect access to objects that reside in virtual memory. All Mesa processors provide a single large, uniformly addressed virtual memory organized as an array of words. Virtual addresses occupy two words and are represented by values of type `LONG POINTER`.

Within a distinguished region of the virtual memory, called the main data space, data may be referenced using short pointers; these addresses occupy a single word and are represented by values of type `POINTER`. The main data space (MDS) is a contiguous region of 64K words of virtual memory in which system data structures and all local and global frames of executing processes reside. The purpose of the main data space is to allow commonly used data structures to be referenced by single word pointers. The use of the main data space for general allocation of other data structures is strongly discouraged.

Dynamically allocated data structures typically reside outside the main data space and are accessed via long pointers. (Short) pointers may be used by programs that indirectly address local and global variables, although for generality, such accesses are often done by long pointers as well. (Short) pointers are also used within the operating system. We begin our discussion of pointer types with (short) pointers for two reasons: long pointer types are constructed from pointer types using the type constructor `LONG`, and all the standard operations that can be applied to pointers can be applied to long pointers as well.

A pointer may refer to only one specific type of item. For instance, the following pointer provides access only to objects of type `INTEGER`:

```
intPtr: POINTER TO INTEGER;
```

Another pointer might be specified to point only to `BOOLEAN` objects:

```
boolPtr: POINTER TO BOOLEAN;
```

These are different types of pointers since they have different *reference types*, `INTEGER` and `BOOLEAN`. Furthermore, since `INTEGER` and `BOOLEAN` are incompatible types, these pointer types are also incompatible; i.e., assignment of *boolPtr* to *intPtr*, or vice versa, is disallowed.

A pointer value is represented by the address of some data object, called the pointer's *referent*. The postfix operator \uparrow may be applied to a pointer value of any type to yield that value's referent. The process of "following" a pointer to its referent is called *dereferencing*.

A dereferenced pointer designates a variable. When the pointer is declared as above, the variable can be used as a **LeftSide** or as a **Primary**. Thus $intPtr \uparrow$ and $boolPtr \uparrow$ are variables of type **INTEGER** and **BOOLEAN** respectively. The statement

```
boolPtr  $\uparrow$   $\leftarrow$  (intPtr  $\uparrow$  = 0);
```

is executed by following $intPtr$ to obtain a **INTEGER** value, testing that value, and assigning the result to the **BOOLEAN** variable referenced by $boolPtr$.

Sometimes a pointer is created simply to identify an object or to allow indirect access to a value that is not to be modified. Mesa provides *readonly* pointers for such applications. A value with a readonly pointer type cannot be used to update its referent. For example, the declaration

```
ROintPtr: POINTER TO READONLY INTEGER;
```

declares a readonly pointer. $ROintPtr \uparrow$ is a **Primary** with type **INTEGER** but not a valid **LeftSide**.

Any type specification is permitted as the reference type of a pointer type. The pointers declared below reference a named record type.

```
Person: TYPE = RECORD
  [
    age: [0..200],
    sex: {male, female},
    party: {Democratic, Republican}
  ];
candidate1, candidate2: Person;
winner, loser: POINTER TO Person;
```

Pointers to record objects may be used to qualify field names. If record $candidate1$ is the referent of $winner$, then qualifications such as

```
winner.age           winner.sex           winner.party
```

select the corresponding components of $candidate1$. However, if $candidate2$ were the referent, these same qualifications would select components of $candidate2$. When applied to a pointer, the operation of selection implies dereferencing. For example, $winner.age$ specifies dereferencing $winner$ to obtain a record variable of type *Person* and then performing normal field selection on that record. Thus $winner.age$ is an abbreviation of $winner \uparrow .age$.

It is common to define a record type containing components that are pointers referencing objects with the same record type. For example, the type declared as follows:

```
FamilyMember: TYPE = RECORD
  [
    someone: Person,
    mother, father: POINTER TO FamilyMember
  ];
```

might be used to create a tree of related persons in which the relations are expressed directly by pointer linkages.

The fundamental operations ($=$, $\#$, \leftarrow) applied to pointer values deal with the pointers themselves, *not* with their referents. In the examples:

```
winner ← loser;
winner ↑ ← loser ↑;
```

the first sets *winner* to point to the same *Person* as *loser*; the second assigns the referent of *loser* to the referent of *winner*, and thus has a quite different effect.

The full set of relational operators can be applied to pointers declared to be *ordered*; for example:

```
orderedPtr: ORDERED POINTER TO Person;
```

The ordering is determined by the memory addresses that represent the pointers, not by the properties of the referents. Pointers not declared to be ordered can be only be compared using the operators $=$ and $\#$.

There is one pointer literal, `NIL`. It conforms to *any* unordered pointer type and denotes a pointer value that has no valid referent. For example:

```
IF intPtr = NIL THEN boolPtr ← NIL;
```

A pointer with value `NIL` should not be dereferenced; the result is undefined.

Pointer values are most commonly obtained from allocators that provide and manage storage for a class of objects. The unary prefix operator $@$ also generates pointers. When applied to a variable with type *T*, it yields a pointer to that variable with type `POINTER TO T`; for example:

```
winner ← @candidate1;
```

Pointer generation should be done with caution; it is possible for the resulting pointer to outlive the referenced object. A non-`NIL` pointer value with no valid referent is said to be a *dangling reference*. The language does not prevent dereferencing such a pointer, but doing so produces an undefined result. *It is the user's responsibility to avoid dereferencing a dangling (or uninitialized) reference.*

3.4.1 Constructing pointer types

The type constructor for pointers is defined as follows:

```
PointerTC      ::= Ordered Base POINTER TO ReadOnly TypeSpecification |
                  Ordered Base POINTER Interval TO ReadOnly TypeSpecification

Ordered        ::= empty | ORDERED

Base           ::= empty | BASE
```

ReadOnly ::= empty | READONLY

The **TypeSpecification** in a **PointerTC** specifies the reference type of the pointer type. Two pointer types are equivalent if their reference types are equivalent and if their attributes **ReadOnly** and **Ordered** are specified identically. Thus equivalent pointer types can be constructed in separate places, but they must have the same structure. One pointer type conforms to another if the two reference types are equivalent, if either the **ReadOnly** attributes are identical or the second is **READONLY** and the first is not, and if either the **Ordered** attributes are identical or the first is **ORDERED** and the second is not. The **Base** attribute is ignored in determining conformance (base pointers are discussed in section 6.3).

In the following examples, the first type in each pair conforms to the second, but the second does not conform to the first:

POINTER TO <i>FamilyMember</i> ORDERED POINTER TO <i>Person</i> ORDERED POINTER TO <i>Date</i>	POINTER TO READONLY <i>FamilyMember</i> POINTER TO <i>Person</i> POINTER TO READONLY <i>Date</i>
---	---

Fine points:

If one pointer type conforms to another, it conforms freely (§ 3.5.3). Conformance of pointer types is extended by the following rule: one pointer type conforms freely to another if the second is **READONLY**, the reference type of the first conforms freely to the reference type of the second, and the **Ordered** attributes satisfy the restriction above.

The second form of **PointerTC** constructs a subrange of a pointer type. Subranges of pointers have the usual properties of subranges; e.g., a pointer subrange type and its base type mutually conform (but not freely). The values of a subrange pointer are restricted to the given interval (and can potentially be stored in smaller fields). Subrange pointer types are not recommended for general use. They are intended primarily for constructing relative pointer types (§ 6.3) which, unlike the subrange types, do not allow dereferencing without relocation.

The attribute **BASE** specifies that values with that pointer type are to be used as base values for relocating relative pointers (§ 6.3). Such values may also be used as ordinary pointers.

3.4.2 Pointer operations

The general form of an indirect reference is:

IndirectReference ::= Variable ↑ |
 (Expression) ↑

LeftSide ::= ... | **IndirectReference**

The postfix operator ↑ performs explicit dereferencing of the pointer expression it follows. Its precedence is the same as indexing and qualification (the highest possible), and these operations can be intermixed. For example:

```
group: ARRAY [0..10] OF POINTER TO FamilyMember;
...
group[i] ↑ .mother ↑ .someone      -- (((group[i] ↑).mother) ↑).someone
```


If p is an arbitrary pointer expression, then $p \uparrow$ can be read as “ p 's referent” or “referent of p .” Application of the \uparrow operator produces a variable that may be used as a **Primary**. Unless p is a readonly pointer, $p \uparrow$ (or any of its components) may also be used as a **LeftSide**. The definition of conformance implies that an ordinary pointer can be assigned to a readonly pointer, but not vice versa. Thus, the referent of a readonly pointer is not necessarily immutable; i.e., its value might change during the lifetime of the **READONLY** pointer. The Mesa language only prevents updates of the object through those pointers to it that are declared to be readonly.

The syntax used for address generation is

Primary **:: =** ... | **@ LeftSide**

The prefix operator **@** produces the address of its operand. If x is a variable of type T whose access path does not involve a long pointer, the value of $@x$ is a pointer to x , and its type is **POINTER TO T** . Otherwise $@x$ is a long pointer to x , and its type is **LONG POINTER TO T** . $@x$ can be read as “address of x .” The operand for **@** must be a valid **LeftSide** (it cannot be a constant or an arbitrary expression, for instance). The operator's precedence is lower than that of \uparrow ; e.g., $@x \uparrow$ is equivalent to $@(x \uparrow)$ (or simply x).

Fine points:

If a variable is declared with fixed form (“=”) initialization, its address may be taken with the **@** operator if and only if the pointer to which it is assigned has the **READONLY** attribute.

There are variables that cannot be the referents of pointers and thus cannot be the operands of **@**. In addition, a pointer value is represented by a word address. Therefore, a referent must lie on a word boundary; an object having this property is called *aligned*. Variables are aligned except in the following cases:

Elements of packed arrays are not aligned.

Any component of a record that occupies less than a single word is not aligned (but arrays, even if packed, are always aligned unless they are small enough to fit entirely within a partial word in the record).

Care must be taken so that a pointer to a declared variable does not exist longer than the variable to which it points. Consider the following example (which assumes familiarity with procedures, local variables and global variables):

```

pointer1, pointer2: POINTER TO INTEGER; -- two global variables
...
RiskyProc: PROCEDURE [ i: INTEGER ] = -- i is a local variable
  BEGIN
    local: INTEGER;                    -- and so is local
    ...
    pointer1 ← @i;                      -- risky: i will disappear upon RETURN
    pointer2 ← @local;                  -- likewise
    ...
    RETURN                              -- the "risky" pointers are valid up to this point, but
    END;                                -- NOT after this statement is executed.

```

After the **RETURN** statement is executed, local storage is released for other purposes; thus the pointers will reference unpredictable data when that storage is reused. *One should use pointers with referents existing at least as long as the pointers.*

Another common mistake is to return the address of a local variable as follows:

```
IncorrectProc: PROCEDURE RETURNS[ POINTER TO INTEGER ] =
  BEGIN
    local: INTEGER;
    ...
    RETURN [@local];           -- WRONG, local is deallocated upon RETURN.
    ...
  END;
```

Pointers that are declared to be **ORDERED** may be used as operands of all the relational operators (§ 2.5.2). For this purpose, they behave as *unsigned* numeric values. The definition of conformance implies that an ordered pointer can be assigned to an unordered pointer variable, but not vice versa. **NIL** is *not* a valid ordered pointer constant, and the relation of its value to other pointer values is undefined. Also, the **@** operator always produces an unordered pointer value.

The following fine points cover pointer capabilities that should be used with caution (and avoided when possible). Some of these capabilities circumvent normal type-checking, and may result in unpredictable results if used.

The type **POINTER TO UNSPECIFIED** (or simply **POINTER**) has the following two properties, which are almost unrelated: It can be dereferenced to yield a value of any type whose size is a single word, and it conforms to any other pointer type, and conversely.

Limited arithmetic can be performed on pointers, but programmers are encouraged to use **BASE** and **RELATIVE** pointers (chapter 6) if the purpose of the arithmetic is simple relocation. A short numeric value added to, or subtracted from, a pointer produces another pointer with the same type. Also, the difference of two pointer values with equivalent types is a **CARDINAL**.

3.4.3 Long pointers

Long pointer types are constructed as follows:

```
LongTC ::= LONG TypeSpecification
```

Long pointers may be created by lengthening (short) pointers as described below. In particular, **NIL** is automatically lengthened to provide a null long pointer when required by context. The standard operations on pointers (dereferencing, assignment, testing equality, comparing ordered pointers) all extend to long pointers.

Both automatic and explicit lengthening (using the operator **LONG**) are provided for pointer types, and the type **POINTER TO T** conforms to (but is not equivalent to) the type **LONG POINTER TO T**. Lengthening an expression with the first of these types produces a value with the second; i.e., the reference type and the **Base**, **Ordered** and **ReadOnly** attributes are unchanged.

The operator @ applied to a variable of type T produces a pointer of type LONG POINTER TO T if the access path to that variable itself involves a long pointer and of type POINTER TO T otherwise.

Fine points:

Two conforming pointer types conform freely only if both are long pointers or both are not.

NIL is lengthened in a standard way and has a universal representation. All other pointers are lengthened in a hardware dependent way.

If either operand in a pointer addition or subtraction is long, all operands are lengthened and the result is long.

Examples:

```
R: TYPE = RECORD [ f: T, ... ];
p, q: POINTER TO R;
pp, qq: LONG POINTER TO R;
pT: POINTER TO T;
ppT: LONG POINTER TO T;
```

-- the following are valid.

```
pp ← qq; pp ← NIL; pp ← p;
pp = qq, pp = NIL, pp = q;
pT ← @p.f; ppT ← @pp.f;
ppT ← @p.f;
```

-- long comparisons

-- pointer lengthened

-- the following are not valid.

```
pp = ppT;
p ← pp; pT ← @pp.f;
```

-- incompatible types

-- no automatic shortening

3.4.4 Automatic dereferencing

Automatic dereferencing converts a *pointer* RightSide of type POINTER TO T into one of type T if that RightSide is followed by dot qualification (§ 3.3.3), a bracketed array index, or a bracketed argument list (the last two are syntactically identical). For example, in the following two statements, the LeftSides are equivalent:

```
winner.party ← Democratic;
winner ↑ .party ← Republican;
```

Automatic multilevel dereferencing is possible. Given the following declarations, the three final assignment statements have the same effect:

```
actualArray: ARRAY [0..20] OF INTEGER;
arrayPtr: POINTER TO ARRAY [0..20] OF INTEGER ← @actualArray;
arrayFinger: POINTER TO POINTER TO ARRAY [0..20] OF INTEGER ← @arrayPtr;
actualArray[1] ← 3;
arrayPtr[1] ← 3;
arrayFinger[1] ← 3;
```

-- arrayPtr ↑ [1] ← 3

-- arrayFinger ↑ ↑ [1] ← 3

Fine points:

The pointer attribute **BASE** inhibits automatic dereferencing in the context of subscript or argument brackets (§ 6.3).

A pointer expression following **OPEN** or **WITH** (§ 4.4.2 and 6.4.4) will be dereferenced an arbitrary number of times (not just once) to obtain an expression designating a record.

3.5 Type determination

Every expression in a Mesa program has a type that can be deduced by static analysis of the program text. Such analysis is called *type determination*. The language imposes constraints on the type of each expression according to the context in which it is used. A program that does not violate any of these constraints is *type-correct*; every valid Mesa program must be type-correct.

In principle, every variable and every expression has an *inherent type* derived from its structure. The inherent type of a variable is established by declaration; the form of a literal implies its type, and each operator produces a result with a type that is a function of the types of the operands. Inherent types of some expression forms are listed below:

<u>Expression</u>	<u>Inherent Type of Expression</u>
34	[34..34] which has base types INTEGER and CARDINAL (§ 3.1.2.1)
NIL	POINTER TO UNSPECIFIED
$x < y$	BOOLEAN
x	declared type of x
$array[i]$	type specified for the components of $array$
$@x$	POINTER TO type of x
$(x \leftarrow e)$	type of x

The type rules in Mesa take two general forms, which are the following:

The exact type required by the context is known, and a given type must conform to it. The required type is called the *target type*.

The exact type required is not implied by context, but a relation that must be satisfied by a set of types is known. The process of satisfying that relation is called *balancing*.

Situations in which the target type is known are simpler and more common; they will be discussed first.

All assignment-like contexts establish a target type for the expression to be assigned. These contexts include not only assignment itself (where the target type is the type of the **LeftSide**) but also initialization, record construction (where the target type for each component expression is the declared type of the corresponding field), array construction, parameter list construction, and the like.

Example:

```
LType: TYPE = RECORD [c: CType];
lVar: LType;
```

```

...
IVar ← anyExp;           -- target type of anyExp is LType
IVar ← LType[c: someExp]; -- target type of someExp is CType ...
IVar.c ← someExp;        -- ... which is more obvious here

```

The following rule applies to assignments:

There is never any automatic dereferencing or type conversion of any kind for the LeftSide of an assignment, and the inherent type of the LeftSide is the target type of the right side. (Of course, a LeftSide may contain subexpressions, such as array subscripts, that are themselves right sides and subject to conversion.)

Certain other contexts imply a target type. For example, the target type for an array subscript is the index type of the array. Also, the target type of the expression following IF, WHILE, etc., is BOOLEAN.

If the inherent type of an expression is equivalent to the target type, the use of that expression is type-correct. If it is not equivalent, it may still be possible to obtain conformance by applying various *type conversions*, which are sometimes called *coercions*. In Mesa, there is at most one sequence of conversions that can be applied automatically to convert a value from one type to another. When implicit conversion from the inherent type to the target type is impossible, the program is in error; e.g., assigning a BOOLEAN value to an INTEGER variable is never valid.

Fine points:

When the target type is well defined, certain expression forms may be abbreviated. Identifier constants need not be qualified, and explicit identification of the type of a constructor is optional. The abbreviated constructs have no inherent type when viewed out of context, and they cannot be used in situations requiring implicit conversion. For example,

```

Color: TYPE = {red, orange, yellow, green, blue, violet};
i: CARDINAL ← Color.green.ORD      -- qualification of green is required

```

An Extractor never has an inherent type; the extraction is controlled by the inherent type of the RightSide, which therefore cannot be abbreviated or converted. For example,

```

r: RECORD [ inner: RECORD [ f1, f2: INTEGER]];
[ i, j ] ← r.inner;      -- the field selection cannot be omitted

```

3.5.1 Type conversion

There are four automatic type conversions that can be applied to establish type conformance. All have been discussed in preceding sections. They are the following:

- (1) A value with a subrange type may be converted to a value with its base type, and vice versa (§ 3.1.2).
- (2) A value with a single-component record type may be converted to a value with the type of that component (§ 3.3.2).

- (3) A value with a short numeric, pointer or array descriptor type may be lengthened to a value with the corresponding long type (§ 2.4.5, 2.5.1.2, 3.4.3, 3.2.1 (and 6.2)).
- (4) A value with any numeric type may be converted to type REAL (§ 2.4.6).

The first of these is a somewhat special case; as mentioned in subsection 3.1.2, it is more accurate to view this as a pair of conversions that are applied unconditionally when evaluating, and assigning to, a subrange variable.

Examples:

```

r: RECORD [ f: INTEGER];
i: INTEGER;
ii: LONG INTEGER;
...
i ← r;           -- i ← r.f
ii ← r;          -- ii ← LONG[r.f]

```

Fine points:

A number of the conversions used to achieve conformance require computation and cannot be applied recursively to the constituents of constructed types. For example, **INTEGER** conforms to **LONG INTEGER**, but **ARRAY IndexType OF INTEGER** does not conform to **ARRAY IndexType OF LONG INTEGER**. subsection 3.5.3 discusses the concept of "free" conformance and the rules governing such cases.

There is one other automatic conversion, dereferencing, that is applied only in certain syntactic contexts (§ 3.4.4). It is never applied automatically to achieve type conformance in an assignment.

3.5.1.1 **ISTYPE** predicate and **NARROW** operator

Let T be a type and x be a variable. The expression **ISTYPE** [x, T] has the value **TRUE** if the type of x is T and **FALSE** otherwise. The **ISTYPE** operator is useful for describing the meaning of the **NARROW** operator as explained below. **ISTYPE** cannot be used with numeric variables to test for interval membership; use **IN** instead.

NARROW is used to restrict the type of variable.

NARROW [x, T] allows a value x to be viewed as a value of type T , and succeeds if and only if **ISTYPE** [x, T] is **TRUE**. It may be thought of as approximately the following Mesa code:

```
IF ISTYPE [x, T] THEN LOOPHOLE [x, T] ELSE ERROR Runtime.NarrowFault;
```

(See chapter 8 for **ERROR**'s, sub-subsection 3.3.1.2 for **LOOPHOLE** syntax, and the *Pilot Programmer's Manual* for details of the *Runtime* interface.)

NARROW, when used with the **ISTYPE** predicate (e.g.

```
IF ISTYPE [u, Type1] THEN {v1: Type1 ← NARROW [u]; ... };
--where u is a variant record)
```

can be used to discriminate a variant record in a situation where only one tag value is reasonable (§ 6.4.4.1). Such use of **NARROW** causes the compiler to generate a run-time test

of the variant tag. Observe that if the type *Type1* can be determined from context, it need not be supplied as an explicit parameter to the **NARROW**.

The **NARROW** operation can also be applied at compile-time to view a value of some opaque type *T* (§ 7.6) to be of the concrete type *T'*. This is particularly useful when the opaque value is embedded within some larger composite type.

An expression of the form **NARROW** [*x*, *T*] cannot appear on the left of an assignment operator. One can usually narrow a suitable pointer and dereference it as a left-hand side. For example:

```
x: S;
NARROW [x, T] ← valueOfTypeT;           -- illegal
NARROW [@x, POINTER TO T] ↑ ← valueOfTypeT; -- okay
```

3.5.1.2 Unchecked type conversion: the **LOOPHOLE** operator

Sometimes it is necessary to subvert Mesa's type checking, particularly in programs that manipulate low-level representations of objects. A **Primary** with the form

```
LOOPHOLE [Expression, TypeSpecification]
```

has the same value as the **Expression** (viewed as a sequence of bits) and the type denoted by **TypeSpecification**. This "conversion" never requires any computation. The only restriction is that values with the inherent type of **Expression** must be represented in the same number of machine words as values of the type **TypeSpecification**. If **Expression** is a valid **LeftSide**, then **LOOPHOLE**[**Expression**, **TypeSpecification**] is also. When the target type is well-defined, the **TypeSpecification** may be omitted. For example:

```
b: BOOLEAN; n: CARDINAL;
n ← LOOPHOLE [b, CARDINAL];           -- to discover the representation
n ← LOOPHOLE [b];                     -- also acceptable
LOOPHOLE [n, BOOLEAN] ← b;           -- as a LeftSide
```

Since **LOOPHOLE** bypasses most checking, its use should be limited as much as possible.

3.5.2 Balancing *

Many of Mesa's operators are generic; i.e., the operation performed depends upon the types of the operands. Examples are the fundamental operators **=** and **#**, which accept two operands with arbitrary (but compatible) types and produce a **BOOLEAN** result. In this case, neither operand has a defined target type. Instead, it is necessary to find some type to which the inherent type of each operand conforms: any automatic type conversions are applied to the operands as necessary to produce values of that type, and the operation is then performed. The common type is the "least upper bound," i.e., the one requiring the fewest conversions.

Examples:

```
R: TYPE = RECORD [ f: INTEGER];
RR: TYPE = RECORD [ ff: LONG INTEGER];
i: INTEGER;
```

```

ii: LONG INTEGER;
r1, r2: R;
rr: RR;
...
i = ii           -- LONG[ i ] = ii
r1 = r2         -- compared as records
r1 = i          -- r1.f = i
r1 = rr        -- LONG[r1.f] = rr.f

```

Balancing is also applied to IF expressions (§ 4.2.1), SELECT expressions (§ 4.3.3), and the arithmetic and relational operators.

Fine points:

Many generic operators do not propagate the target type of the expression in which they appear; instead, the operands are balanced and combined to produce a result that is converted further if necessary. For example,

```

ii ← i + r;           -- ii ← LONG[ i + r.f ]
ii ← LONG[ i ] + r;  -- ii ← LONG[ i ] + LONG[ r.f ]

```

The current version of Mesa does not fully implement balancing when lengthening (or conversion to REAL) is required. The restrictions are:

Operands of MIN and MAX and the alternatives of conditional expressions are lengthened to match the expression's target type, if any, and otherwise to match the type of the first operand.

The endpoints of an interval in the right operand of IN are lengthened to match the type of the left operand, but the left operand is never lengthened.

The expressions selecting the arms of a selection (§ 4.3) are lengthened to match the type of the selecting expression, but that expression is never lengthened.

3.5.3 Free conformance *

A number of the conversions used to achieve conformance require computation and cannot be applied recursively to establish the conformance of types constructed from pairwise conforming types. For example, INTEGER conforms to REAL, but the conversion from INTEGER to REAL transforms the representation. Thus a POINTER TO INTEGER and POINTER TO REAL cannot validly have the same referent, and these types do not conform.

The relation of *free conformance* is less restrictive than strict type equivalence but is defined so that it can be computed recursively. Loosely speaking, one type freely conforms to another if a value of the first can always be used as a value of the second without any computation or run-time check of validity. The relations of equivalence, free conformance and conformance are not independent. Equivalence always implies free conformance; if two types are equivalent, each freely conforms to the other. Also, free conformance implies conformance; if one type freely conforms to another, the first also conforms to the second.

Of the automatic conversions discussed in subsection 3.5.1, only a restricted form of the first (subrange conversion) can be applied to establish free conformance. The restriction (which arises from the representation of subrange values in Mesa) is the following:

The subrange type $T[i..j]$ conforms freely to T if $i = \text{FIRST}[T]$ and to $T[i..k]$ if $j \leq k$.

If automatic conversion (1) of subsection 3.5.1 must be applied in any other circumstance or if application of conversion (2), (3) or (4) of that section is required to establish the conformance of two types, they do not conform freely.

Of the constructed types discussed in this chapter, array and pointer types also have rules for free conformance less restrictive than equivalence. To summarize:

One array type conforms freely to another if the index types are equivalent and the component type of the first freely conforms to the component type of the second (§ 3.2.1).

One pointer type freely conforms to another whenever the first pointer type conforms to the second as defined in subsection 3.4.1.

Free conformance is also important for procedure types (§ 5.1) and variant records (§ 6.4).

In the following pairs of types, the first conforms to the second (but does not freely conform):

[0..100)	[0..10)
[5..10)	[0..10)
INTEGER	REAL
POINTER TO <i>Person</i>	LONG POINTER TO <i>Person</i>

In the following pairs, the first type freely conforms to the second (but is not equivalent):

POINTER TO [0..10)	POINTER TO READONLY [0..100)
POINTER TO READONLY [0..10)	POINTER TO READONLY [0..100)
ARRAY [0..10) OF [0..10)	ARRAY [0..10) OF CARDINAL

Fine point:

Note that **POINTER TO [0..10)** does not conform to **POINTER TO [0..100)** so that the following is illegal:

```
p: POINTER TO [0..10); q: POINTER TO [0..100);
...
q ← p; q ↑ ← 99;           -- now p ↑ = 99
```

3.6 Determination of representation *

This section discusses the rules used by Mesa for choosing between signed and unsigned versions of the numeric operations. These rules assume that there are conversion functions (taking the form of range assertions for short numeric types, § 3.1.2.2) that convert values from **CARDINAL** to **INTEGER** (from **LONG CARDINAL** to **LONG INTEGER**) and vice versa. In both directions, the "conversion" amounts to an assertion that the value is an element of $\text{INTEGER} \cap \text{CARDINAL}$ ($\text{LONG INTEGER} \cap \text{LONG CARDINAL}$). *Such assertions must be verified by the programmer.*

For any arithmetic expression, the *inherent representations* of the operands and the *target representation* of the result are used to choose between the signed and unsigned versions of the arithmetic and relational operators.

The target type determines the target representation. The preceding section describes the derivation of target types; in addition, a range assertion establishes the asserted type as the target type of its operand. If all valid values of the target type are nonnegative, the target representation is unsigned; otherwise, it is signed. The arithmetic operators propagate target representations unchanged to their operands, but the target representation of an operand of a relational operator is undefined. The target representation is also undefined in all other cases in which the target type is undefined. Thus each (sub)expression has at most one target representation.

The inherent representation of a **Primary** is determined by its type (if a variable, function call, etc.), by its value (if a compile-time constant), or explicitly (if a range assertion). Possible inherent representations are signed and unsigned; in addition, a compile-time constant in `INTEGER \cap CARDINAL` or a **Primary** with an inherent type that is a subrange of `INTEGER \cap CARDINAL` is considered to have *both* inherent representations. Inherent representations of operands are propagated to results as described below.

The operation denoted by a generic operator is chosen by considering first the inherent representations of its operands, next the target representation, and finally a preferred default. If the operation cannot be disambiguated in any of these ways, the expression is considered to be in error. The exact rules follow:

If the operands have exactly one common inherent representation, the operation defined for that representation is selected (and the target representation is ignored).

If the operands have no common inherent representation but the target representation is well-defined, the operation yielding that representation is chosen, and each operand is "converted" to that representation (in the weak sense discussed above).

If the operands have both inherent representations in common, and if target representation is well-defined, it selects the operation. If the operands have both inherent representations in common but the target representation is ill defined, the signed operation is chosen.

If the operands have no representation in common and the target representation is ill-defined, the expression is in error.

In all cases, the inherent representation of the result is determined by the selected operation.

The unary operators require special mention. Unary minus converts its argument to a signed representation if necessary and produces a signed result.

Example:

If m and n have unsigned representation, both the following are legal and assign the same bit pattern to i , but the first overflows if $m < n$.

```
i ← m − n; i ← IF m ≥ n THEN m − n ELSE −((n − m));
```

ABS is a null operation on an operand with an unsigned representation; it always yields a value with unsigned representation. The target representation for the operand of **LONG** (or of an implied lengthening operation) is unsigned.

Examples:

```
i, j: INTEGER; m, n: CARDINAL; s, t: [0..77777B]; b: BOOLEAN
```

-- the statements on each of the following lines are equivalent.

```
i ← m + n; i ← INTEGER [m + n]           -- unsigned addition
i ← j + n; i ← n + j; i ← j + INTEGER [n] -- signed addition
i ← s + t; i ← INTEGER [s] + INTEGER [t]    -- signed (overflow possible)
n ← s + t; n ← CARDINAL [s] + CARDINAL [t]  -- unsigned (overflow impossible)
s ← s − t; s ← CARDINAL [s] − CARDINAL [t]  -- unsigned (overflow possible)
b ← s − t > 0; b ← INTEGER [s] − INTEGER [t] > 0 -- signed (overflow impossible)
```

```
i ← −m; i ← −INTEGER [m]
```

```
i ← m + n*(j + n); i ← INTEGER [m] + (INTEGER [n]*(j + INTEGER [n]))
n ← m + n*(j + n); n ← m + (n*CARDINAL [j] + n)
i ← m + n*(s + n); i ← INTEGER [m] + (n*CARDINAL [s] + n)
```

```
b ← s IN [t − 1 .. t + 1]; b ← INTEGER [s] IN [INTEGER [t − 1] .. INTEGER [t + 1]]
FOR s IN [t − 1 .. t + 1] ...; FOR s IN [CARDINAL [t − 1] .. CARDINAL [t + 1]] ...
```

The following statements are incorrect because of representational ambiguities.

```
b ← i > n; b ← i + n IN [s .. j]
```

```
SELECT i FROM m => ...; t => ...; ENDCASE
```

Fine point:

When an **INTEGER** is lengthened, its inherent type is **LONG INTEGER**. When a **CARDINAL** or **NATURAL** is lengthened, its inherent type is **LONG INTEGER** and **LONG CARDINAL**.

```
LongCardinal: LONG CARDINAL;
```

```
LongInteger: LONG INTEGER;
```

```
Integer: INTEGER;
```

```
Cardinal: CARDINAL;
```

The following statements are valid:

```
IF LongInteger < Integer THEN ...;
IF LongCardinal < Cardinal THEN ...;
IF LongInteger < Cardinal THEN ...;
```

The following statement is invalid:

```
IF LongCardinal < Integer THEN ...;
```

3.7 Extended defaults

As previously explained, you can associate a default initial value with a type (not just with a field of a record). If a type is constructed from other types using one of Mesa's type operators (e.g., `RECORD`), the default value for that type is determined by the default values of the component types and by rules associated with each operator and by any default specification for the record type itself. When you declare a named type, you have the option of explicitly specifying a default for that type.

With this extension, you will find that uses of defaults in Mesa generally fall into two classes. Default values for fields of records and arrays make the corresponding constructors more concise and more convenient to use. On the other hand, the usual reason for associating a default initial value with a type is to ensure that storage allocated for that type is well-formed, i.e., that any variable of such a type always has a meaningful value. There is some interaction between these uses; the default value of a record type is partly determined by any default values specified for its fields, and a record field may inherit its default value from the type of that field. The details appear below.

The rules for inheritance of defaults are designed to provide the following property (currently not quite preserved by sequence (§ 6.5) or variant record types (§ 6.4.5)): if a type *T* has been given a non-NULL default value, any type derived from *T* will have a defined and non-NULL default value for any embedded component of type *T*. Because of the potential cascading effect implied by this, you should carefully consider the relative costs and benefits of specifying a default, especially one that does not include NULL as an alternative.

Defaults are ignored in determining equivalence and conformance of types. Thus, it is possible to have two compatible types with different default initializations.

Specification of Default Initialization

None of the built-in types (`INTEGER`, `CARDINAL`, `NATURAL`, `BOOLEAN`, `CHARACTER`, `STRING` and `REAL`) has a default initial value. All of the transfer types `PROCEDURE`, `PROGRAM`, `SIGNAL`, `ERROR`, and `PROCESS` have a default initial value of `NIL`.

The following rules determine the default initial value of a type designated by an expression involving a type operator:

The default initial value for a type constructed using `RECORD` (or `ARRAY`) is defined field-by-field (or element-by-element). For each field (element), it is the default value for that field if there is one; otherwise, it is the default initial value for the type of that field (element) or is undefined if there is no such default.

Types constructed using other operators have no implied default initialization.

The default initial value of a type designated by a declared type identifier *T* depends upon the form of the declaration of *T*, as follows:

T: `TYPE = TypeExpr`;

T receives all the attributes of *TypeExpr* including any default.

T : TYPE = *TypeExpr* ← e ;

T receives all the attributes of *TypeExpr* except that its default initial value is e .

Examples

```

Flag: TYPE = BOOLEAN ← FALSE;
Rec1: TYPE = RECORD [f: Flag];           -- default value is [f: FALSE]
Rec2: TYPE = RECORD [f: Flag] ← [ ];    -- ditto (the field defaults)
Rec3: TYPE = RECORD [f: Flag] ← [TRUE]; -- explicit default
Rec4: TYPE = Rec3;                       -- default value is [f: TRUE]
Rec5: TYPE = Rec3 ← [f: FALSE];         -- default value is [f: FALSE]

```

Any **DefaultSpecification** is acceptable in a type declaration (§ 3.3.5). A declaration giving a type T a **NULL** default cannot, however, equate T to a type with a default that does not include **NULL**. A default appearing in a type definition within a **DEFINITIONS** module must be either **NULL** or an expression with a compile-time constant value.

Default values associated with types are used

to initialize local variables of procedures and programs, in the absence of explicit initialization,

to initialize variables that are dynamically allocated using **NEW** (§ 6.6.2), in the absence of explicit initialization (see below),

to construct records (except argument and result records), in the absence of an explicit value for a field in the constructor and of a default value for that field in the record declaration,

to construct arrays, in the absence of an explicit value for an element (see below).

Defaulted Array Elements

Elements in an array constructor may be voided or elided. Omission of elements is permitted in a keyword constructor (see below) but not in a positional constructor. The empty constructor (**[]**) is a keyword constructor with all items omitted. An elided or omitted element receives the default value for the type of the components of the array (if any); the value of a voided element is undefined.

ALL abbreviates a positional constructor of the appropriate length; thus **ALL []** elides all elements (defaulting if possible) and **ALL [NULL]** voids all positions.

Fine point:

The examples below illustrate the generality and complexity that default values may have. In the example below, the assignment statement assigns 2 to *group*[1], 4 to *group*[2], 3 to *group*[3], 4 to *group*[4], and 3 to *group*[5].

```

GradeRange: TYPE = CARDINAL[0..4] ← 3;
GradeType: TYPE = ARRAY[0..4] OF GradeRange
group : GradeType;

```

```
group ← [2,4, ,4, ];           --two fields elided
```

A final example:

```
Subrange: TYPE = CARDINAL [0..4] ← 2 | TRASH;
Array: TYPE = ARRAY [0..4] OF Subrange ← [1, 1, ,TRASH, 1 ] | TRASH;
Record: TYPE = RECORD
    [
        r1: CARDINAL ← TRASH,
        r2: Subrange ← 4, --overrides previous default specification
        r3: Subrange,
        r4: Array,
        r5: CARDINAL ← 77 | TRASH
    ] ← [ , ,3, [0,1,2,3,4], ];
Color: TYPE = {red, orange, blue, white} ← ;
NestedRec: TYPE = RECORD
    [
        n1: Color,
        n2: Record ← [r1: 10, r2: , r3: TRASH, r4: [3,3,3,3,3], r5: 2]
    ];
Ptr: TYPE = LONG POINTER TO Color ← NIL;

v1: Array;
v2: Record;
v3: NestedRec ← [red];
v4: Color ← red;
```

--the following are valid

```
v1 ← [2,4, ,4, ];           -- means Array[2,4,2,4,2]
v2 ← [r3: TRASH, r5: 4];     -- means Record[r1: TRASH, r2: 4, r3: TRASH,
                             --                    r4: [1,1,2,TRASH,1], r5: 4 ]
v2 ← [r3: 1, r4: [1,2,3,4, ], r5: 6]; -- means Record[r1: TRASH, r2: 4, r3: 1,
                             --                    r4: [1,2,3,4,2], r5: 6 ]
v2 ← [];                    -- means Record[r1: TRASH, r2: 4, r3: 2,
                             --                    r4: [1,1,2,TRASH,1], r5: 77 ]
v2 ← [r1: TRASH, r4: TRASH, r5: 6]; -- means Record[r1: TRASH, r2: 4, r3: 2,
                             --                    r4: TRASH, r5: 6 ]
v3 ← [n1: red, n2: [r4: [4, ,4, ,4]]]; -- means NestedRec[n1: red,
                             -- n2: [r1: TRASH, r2: 4, r3: 2, r4: [4,2,4,2,4], r5: 77]
```

3.8 The null value NIL

In Mesa, null values are available for all address-containing types. An address-containing type is one constructed using **POINTER**, **DESCRIPTOR**, **PROCEDURE**, **PROGRAM**, **SIGNAL**, **ERROR**, **PROCESS**, **UNCOUNTED ZONE** or a **LONG** or subrange form of one of the preceding. The built-in type **STRING** is address-containing.

Fine point:

A relative pointer or relative descriptor type is not considered to be address-containing in Mesa.

Null values are denoted as follows:

If T designates any address-containing type, `NIL[T]` denotes the corresponding null value.

Whenever T is implied by context, `NIL` abbreviates `NIL[T]`.

If T is not implied by context, `NIL` means `NIL [POINTER TO UNSPECIFIED]` and thus matches any `POINTER` or `LONG POINTER` type.

A fault will occur if you attempt either to dereference a null value or to transfer control through a null value.



Ordinary statements

Statements are the units of action in Mesa; they control the flow of execution and the updating of variables. This chapter treats *ordinary* statements: those statements having wide applicability (such as assignment statements); later chapters cover the remaining statements. The following syntax lists the phrase names of all the statement forms covered in this chapter:

Statement ::= AssignmentStmt | IfStmt | SelectStmt | NullStmt |
Block | GotoStmt | LoopStmt | ExitStmt | ...

Some statements have expression counterparts, with the same general purposes but slightly different constraints. For instance, assignment can be performed by an expression as well as a statement. The expression forms covered in this chapter are

Expression ::= ... | AssignmentExpr | IfExpr | SelectExpr

In Mesa, certain statement forms such as the IF statement contain other statements. These statements in turn may contain still other statements, and so forth. Consequently, the term "statement" should be understood to encompass the large and small alike.

The dynamic successor of a statement embedded within another depends upon the embedding form. For simplicity, the discussion assumes that most statements occur in the middle of a hypothetical series of statements. Execution paths within a statement are described for each form of control statement, and the successor is described in terms of a postulated "*Next-Statement*." *Next-Statement* represents nothing more than completion of a given statement; another statement may or may not appear at that point in an actual program.

Although execution of a statement can be aborted prior to its normal completion, the discussion of statement sequencing also assumes normal termination of each statement unless otherwise stated.

In the examples, *Stmt-0*, *Stmt-1*, *Stmt-2*, etc. denote arbitrary statements, the details of which are irrelevant.

4.1 Assignment statements

Syntax:

```
AssignmentStmt ::= LeftSide ← RightSide |
                Extractor ← RightSide
```

The **RightSide** must be an expression with a type conforming to the type of the left-hand side. The left-hand side must be a valid recipient of data such as a declared variable or a component. For assignment *statements*, a left-hand side may also be an extractor (§ 3.3.6):

Examples:

```
i ← 3; a ← b + c;
birthDay.month ← Apr; birthTable[Tom].year ← 1955;
[mm, dd, yy] ← birthDay;           -- an extractor as the LeftSide
```

4.1.1 Assignment expressions

Assignment operations may be carried out by expressions, as well as by assignment statements. The syntax for an assignment expression is:

```
AssignmentExpr ::= LeftSide ← RightSide |
                Extractor ← RightSide
```

Assignment expressions can be used for performing multiple assignments in a single statement, and for saving the value of an intermediate expression without having to write a separate statement:

```
x2 ← x1 ← x0 ← v;           -- set x0, x1, and x2 to the value in v
array[j ← j + 1] ← x [i];   -- j is changed while changing the array
component
```

Evaluation of the first statement proceeds as if it were written:

```
x2 ← (x1 ← (x0 ← v))
```

Note that $x2 \leftarrow (\dots)$ is an assignment *statement*. The assignment expression, $x0 \leftarrow v$, yields the value assigned to $x0$, this becomes the **RightSide** value for the other assignment expression, and so on.

The difference between an assignment expression and an assignment statement is that the expression yields a value (in addition to performing assignment).

An assignment statement with an **Extractor** as its left side may be used as an expression. This allows multiple extractions, among other things. The value of such an assignment is the value of its right side.

An **AssignmentExpr** is an **Expression**. Its type is the type of the **LeftSide**, and its value is the value actually assigned (possibly after type conversion) of the **RightSide**. The assignment operator has the *lowest possible precedence*. As a stylistic rule, an assignment expression embedded in another expression is enclosed in parentheses.

Fine point:

In an expression such as the following:

$$a[k \leftarrow k + 1] + b[k];$$

the order of evaluation is undefined, and the embedded assignment may be executed either before or after evaluation of $b[k]$. Such use of embedded assignments should be avoided.

4.1.2 Restrictions on assignment

The assignment operations defined upon certain types have been restricted so that variables of those types can be initialized (either explicitly or by default) when they are created but cannot subsequently be updated. A variable is considered to be created at its point of declaration or, for dynamically allocated objects, by the corresponding **NEW** operation.

The following types have restricted assignment operations:

MONITORLOCK (§ 9.2)

CONDITION (§ 9.3)

any type constructed using **SEQUENCE** (§ 6.5)

any type constructed using **ARRAY** in which the component type has a restricted assignment operation.

any type constructed using **RECORD** in which one of the field types has a restricted assignment operation.

Note that the restrictions upon assignment for a type do not impose restrictions upon assignment to component types. Thus selective updating of fields of a variable may be possible even when the entire variable cannot be updated; e.g., the *timeout* field of a **CONDITION** variable can be updated by ordinary assignment. Also, you may apply the operator **@** to obtain the address of the entire variable in such a situation.

4.2 IF statements

An **IF** statement is a control statement that functions as a two-way switch:

IfStmnt ::= **IF** Predicate **ThenClause** **ElseClause**

Predicate ::= **Expression**

ThenClause ::= **THEN** **Statement**

ElseClause ::= **empty** | **ELSE** **Statement**

A simple **IF** statement is shown below.

```
IF  $v = 0$  THEN WriteString["Done."][ELSE  $v \leftarrow v - 1$ ;  
Next-Statement
```

The **BOOLEAN** expression ($v = 0$) is called the **Predicate** of the **IF** statement. The **Predicate** is evaluated first, and if **TRUE**, the **Statement** in the **ThenClause** is executed (in this case a call on the procedure *WriteString*). Upon its completion, execution continues at *Next-Statement*. If the **Predicate** value is **FALSE**, the **Statement** in the **ElseClause**, " $v \leftarrow v - 1$ ", is executed; if there is no **ElseClause** control goes directly to *Next-Statement*.

Other examples:

```
IF ( $flag = on$ ) AND  $i$  IN [ $m..n$ ] THEN  $i \leftarrow i + iDelta$  ELSE  $i \leftarrow m$ ;  
  
IF  $winner \sim = NIL$  THEN  
  BEGIN  
    totalAge  $\leftarrow$  totalAge +  $winner.age$ ;           -- this Statement is a block (§ 4.4)  
    IF  $winner.party = Democratic$  THEN  $demoScore \leftarrow demoScore + 1$   
    ELSE  $gopScore \leftarrow gopScore + 1$ ;  
  END;                                           -- end of the ThenClause
```

Note that a semicolon cannot follow a **ThenClause** when an **ElseClause** is present.

If the **Statement** in a **ThenClause** is a second **IF** statement, then the outer **IF** may have an **ElseClause** only if the inner one does; i.e., an **ElseClause** "belongs" to the innermost possible **IF**. For example:

```
IF  $a \geq 0$  THEN  
  IF  $a > 0$  THEN  $b \leftarrow 1$            --  $a > 0$  means set  $b$  to 1  
  ELSE  $b \leftarrow 0$ ;                 --  $a = 0$  means set  $b$  to 0  
  ...                                   -- no action if  $a < 0$ 
```

It is recommended that "IF...THEN IF" combinations be avoided entirely unless the second **IF** has an **ElseClause**. Often, a single **IF** statement is sufficient. For example, let $p1$ and $p2$ be arbitrary predicates. Then the following two statements have identical effect:

```
IF  $p1$  AND  $p2$  THEN  $Stmt$ ;           -- recommended form (§ 2.5.3)  
  
IF  $p1$  THEN IF  $p2$  THEN  $Stmt$ ;       -- longer form
```

Fine point:

If the **Predicate** is a compile-time constant, the compiler does not produce object code for the text that would never be executed. This also holds for **IF** expressions.

4.2.1 IF expressions

The **IF** statement has a counterpart that is an expression. Its syntax is similar to that of an **IfStmt**:

```
IfExpr ::= IF Predicate THEN Expression ELSE Expression
```

There are two differences between an **IfExpr** and an **IfStmt**:

The clauses of an IF expression contain expressions, not statements;

An IF expression must have an ELSE-clause.

Examples:

```
slope ← IF y = 0 THEN max ELSE x/y;           -- avoid division by zero.
```

```
b ← IF a >= 0 THEN (IF a > 0 THEN 1 ELSE 0) ELSE - 1;
```

Evaluation of an IF expression begins with evaluation of the **Predicate** (in the first example, $y = 0$). If it is **TRUE**, the expression in the **ThenClause** (i.e., max) is evaluated, and its value becomes the value of the IF expression. If the predicate is **FALSE**, the **ElseClause** expression (i.e., x/y) is evaluated, and its value becomes the value of the IF expression. The second example sets the value of b to -1 , 0 , or $+1$, depending on whether a is negative, zero, or positive, respectively.

The **ThenClause** and **ElseClause** expressions must conform to some common type (possibly after type conversion, as outlined in subsection 3.5.1). The type to which they conform is the IF expression's inherent type.

An IF operator has the same precedence as an assignment operator, i.e., the lowest possible precedence. IF expressions should be enclosed in parentheses when embedded in other expressions.

4.3 SELECT statements

The **SELECT** statement chooses for execution at most one statement from an ordered list of statements. The choice is based upon the relation between a given expression and expressions associated with each selectable statement. Thus, this statement form permits multiway branching, not just the two way branching of an IF statement.

A **SELECT** statement is shown below. The separator " $= >$ " should be read as "chooses." The entire statement may be read as follows: "Select, using x 's value, from the comparisons preceding the substatements. First, (x 's value) 'equal to zero' chooses *Stmt-1*. Second, 'in subrange m through n ' chooses *Stmt-2*. Third, 'less than m ' chooses *Stmt-3*. Otherwise, choose nothing."

```
SELECT x FROM
  = 0           => Stmt-1;
  IN [m..n]    => Stmt-2;
  < m          => Stmt-3;
ENDCASE
```

The next four sections cover various forms of **SELECT**, their precise syntax, and the expression counterpart of the **SELECT** statement. The term "**SELECT**," used by itself, includes both statement and expression forms.

4.3.1 Forms and options for SELECT

Syntax equations:

```

SelectStmt      ::= SELECT LeftItem FROM           -- (the head)
                  StmtChoiceSeries                -- (the arms)
                  ENDCASE FinalStmtChoice |       -- (the foot)
                  . . .

LeftItem        ::= Expression

StmtChoiceSeries ::= empty |
                  TestList => Statement |
                  TestList => Statement ; StmtChoice Series

FinalStmtChoice ::= empty |
                  => Statement

TestList        ::= Test | TestList , Test

Test            ::= Expression |                  -- no operator implies an equality test
                  RelationTail

```

Example:

```

i:[0..5];
. . .
SELECT i FROM
  0      => i ← i + 1;          -- i = 0
  < 3    => BEGIN j ← i; i ← i - 1 END; -- i = 1 or i = 2
  = 5    => i ← 0;            -- i = 5
  ENDCASE => i ← 2;          -- i = 3 or i = 4 (none of the above)
Next-Statement

```

In the execution of a **SELECT** statement, the **LeftItem** is evaluated first; a sequence of comparisons then follows. Each arm of the **SELECT** statement begins with one or more **Tests**. The **Expression** in each **Test** is evaluated and compared with the value of the **LeftItem**. The evaluation occurs in order, from left to right, and continues until a comparison succeeds or the **TestList** for that particular arm is exhausted. If a test succeeds, control passes immediately to the statement following the **TestList** in that arm (no further **Tests** are evaluated, even in that same list). If all **Tests** in a given arm fail, the next arm in the series is tried. After a test succeeds and its associated statement is executed, control passes to *Next-Statement*. Thus, at most, one statement can be chosen in a given execution of a **SELECT** statement.

When combined with the **LeftItem** (perhaps with an implied "="), each **Test** must be a valid **Relation**. The type of the **Expression** in a **Test** must conform to the type of the **LeftItem**. If a **Test** uses "IN Subrange," the base type of the subrange must conform to the type of the **LeftItem**.

A single **SELECT** arm may specify more than one test:

```

SELECT i*j + k FROM
  1, IN [7..10]    => Stmt-1;      -- values: 1, 7, 8, 9, 10
  2, 5, > 10       => Stmt-2;      -- values: 2, 5, 11, 12, ...
ENDCASE;

```

A final choice may be appended to a **SELECT** to handle all remaining cases; it follows **ENDCASE**. For example:

```

PriorityState: TYPE = RECORD [i0, i1, i2, i3: BOOLEAN];
oldState, newState: PriorityState;
...
SELECT TRUE FROM                                -- picks the first TRUE state
  oldState.i0   => Stmt-0;
  oldState.i1, newState.i0 => Stmt-1;
  oldState.i2, newState.i1 => Stmt-2;
  oldState.i3, newState.i2 => Stmt-3;
ENDCASE      => Stmt-99;

```

If this **SELECT** statement does not choose one of the first four statements, the final statement (*Stmt-99*) is executed.

Fine points:

If all **SELECT** arms (or those in some contiguous subseries) specify constant values in each **Test**, the compiler can produce code using a "jump table" for efficient selection (It only does this if the set of interesting values is sufficiently "dense").

The other alternatives for **SelectStmt** apply to variant records and are discussed in chapter 6.

The expressions in the **SELECT** arms are lengthened as necessary to match the type of the **LeftItem**, but the **LeftItem** is never lengthened.

4.3.2 The NULL statement

The **NULL** statement, which serves only as a placeholder, is often useful as the statement in an arm of a **SELECT** statement:

```

NullStmt      ::= NULL

```

For example:

```

SELECT currentChar FROM
  IN ['0..'9] => Stmt-1;      -- handle digits.
  IN ['A..'Z] => Stmt-2;      -- handle capital letters.
  IN ['a..'z] => Stmt-3;      -- handle small letters.
  SP         => NULL;        -- ignore blanks.
ENDCASE     => Stmt-99;      -- handle all other chars.

```

4.3.3 SELECT expressions

The **SELECT** statement has an expression counterpart. There are three differences between the expression and statement forms of **SELECT**:

- (1) The choices in each arm must be expressions, not statements.
- (2) The arms are terminated by commas, not semicolons.
- (3) **ENDCASE** must be followed by “=>” and a final (expression) choice.

Its syntax is defined by

```

SelectExpr      ::= SELECT LeftItem FROM      -- (the head)
                  ExprChoiceList           -- (the arms)
                  ENDCASE => Expression    -- (the foot)
                  | ...

ExprChoiceList  ::= empty |
                  TestList => Expression |
                  TestList => Expression, ExprChoiceList

```

LeftItem and **TestList** are defined in subsection 4.3.1.

For example:

```

pt: INTEGER;                -- point on a line.
lo, hi: INTEGER ← 0;        -- bounds for a line segment, initially a null segment
...
PointPosition: TYPE = {leftMargin, rightMargin, inside, outside, degenerate};
position: PointPosition;
...
position ← SELECT pt FROM
  IN [lo..hi]                => inside,
  NOT IN [lo..hi]           => outside,
  < hi                       => leftMargin,          -- = lo but ≠ hi
  > lo                       => rightMargin,         -- = hi but ≠ lo
  ENDCASE                   => degenerate;          -- = lo and = hi

```

A **SELECT** expression is executed just as a **SELECT** statement, except that the selected arm yields a value, which becomes the value of the **SELECT** expression as a whole. The inherent type of a **SELECT** expression is the one to which all the expressions in the arms conform (§ 3.5.3).

A **SELECT** operator has the same precedence as an assignment operator, i.e., the lowest possible precedence. **SELECT** expressions should be enclosed in parentheses when embedded in other expressions.

4.4 Blocks

A block is a way of packaging a series of statements so that they can be used where only a single statement is permitted syntactically. In its simplest form a block is a pair of “brackets,” **BEGIN** and **END**, with a series of statements (of any form) between them. The general syntax is:

```

Block          ::= BEGIN
                OpenClause      -- optional; subsection 4.4.2
                EnableClause    -- optional; subsection 8.2.1
                DeclarationSeries -- optional
                StatementSeries
                ExitsClause      -- optional; subsection 4.4.1
                END

StatementSeries ::= empty |
                Statement |
                Statement ; StatementSeries

DeclarationSeries ::= empty | DeclarationSeries Declaration

```

The bracket pair `{ }` can be used any place the bracket pair `BEGIN END` can be used (but not conversely).

Fine point:

A semicolon terminates every declaration and therefore is not mentioned as a separator here.

In the following `IF` statement, a block takes the place of the single `Statement` normally allowed in a `ThenClause`:

```

IF lo > hi THEN
  BEGIN                               -- Exchange lo and hi.
    temp: INTEGER ← lo;
    lo ← hi;
    hi ← temp;
  END

```

A semicolon must separate each statement in the `StatementSeries` but is optional after the last statement.

The optional `DeclarationSeries` in a block introduces new identifiers, such as `temp` above, with scope smaller than an entire procedure (or module) body. Scope is discussed further in subsections 4.4.2 and 5.5.1 and in chapter 7.

Fine point:

During the execution of a Mesa program, frames are allocated at the procedure and module level only (§ 5.2). Any storage required by variables declared in an internal `Block` (one used as a `Statement`) is allocated in the frame of the smallest enclosing procedure or module. When such internal blocks are disjoint, the areas of the frame used for their variables overlay one another.

Ordinarily, when a block is executed, every statement in its `StatementSeries` is executed, and *Next-Statement* is the successor of the entire block. It is possible, however, to jump out of a block, as described in the next section on `GOTOS`.

4.4.1 GOTO statements

A more general form of a block allows a series of labeled statements to be written immediately preceding its `END`. One can jump to any one of these statements from within

the block only, using a **GOTO** statement. There are two consequences of this way of constraining the **GOTO**:

A **GOTO** may only jump forward in the program, never backward.

A **GOTO** may only jump out of a block, never into one.

The syntax for the **ExitsClause** of a block and for the **GOTO** statement is the following:

```

ExitsClause      ::= empty |
                    EXITS ExitSeries |
                    EXITS ExitSeries ;    -- optional final semicolon

ExitSeries      ::= empty |
                    LabelList = > Statement |
                    LabelList = > Statement; ExitSeries

LabelList       ::= Label | LabelList , Label

Label           ::= identifier

GotoStmnt       ::= GOTO Label | GO TO Label

```

A simple example:

```

IF input.status # open THEN
  BEGIN
  ...
  IF input.fileHandle = defaultInput THEN GOTO useDefault;
  ...                                     -- processing for non-default file
  IF input.fileNumber = ttyNumber THEN GOTO fileIsDefault;
  IF input.length = 0 THEN GOTO newFile;
  ...                                     -- compute number of pages in the file
  EXITS
    useDefault, fileIsDefault = >      -- multiple labels are allowed
    BEGIN input ← ttyInput; pages ← maxPages END;
    newFile = > pages ← 0;
  END;    -- end of the ThenClause and the IF statement
  Next-Statement

```

The **Labels** in this example are *useDefault*, *fileIsDefault*, and *newFile* (it is helpful to view the labels as the names of conditions or reasons for which the block is being left). If any one of the **GOTOS** is executed, control transfers immediately to the statement labeled with the identifier used in the **GOTO**. The normal successor of any one of the labeled statements is *Next-Statement*, which is also the normal successor of the last statement in the main body of the block (i.e., the one just before **EXITS**).

Since one block can appear within the body of another, a **GOTO** can jump directly out of one (or more) blocks to the **ExitsClause** of an enclosing block. For example,

```

BEGIN-- outer block
...
  BEGIN-- inner block
  ...

```

```

    IF  $i = iMax$  THEN GO TO endOfArray;      -- jump to end of outer compound
    ...
    END;                                     -- end of inner
    ...
     $i \leftarrow i + 1$ ;
    EXITS
        endOfArray =>  $i \leftarrow 0$ ;
    END;                                     -- end of outer
    Next-Statement

```

If the **GOTO** statement is executed, control jumps to the exit labeled *endOfArray*. The chosen statement ($i \leftarrow 0$) is executed and control then goes to *Next-Statement*. The identifiers used as **Labels** are only known inside the block in which they appear, and it is possible to use the same identifier as a label in a number of blocks. If this is done in nested blocks, a **GOTO** naming that identifier will always go to the statement with that label in the smallest enclosing block. Generally, using the same label in nested blocks is a bad idea.

Since Mesa allows declarations in any block, it is possible to declare a procedure (§ 5.5) within the scope of the **Labels** of a block. Jumping out of a procedure into a surrounding block is disallowed. Such a result may be obtained, however, by use of the **SIGNAL** machinery (see chapter 8). For example, the following is *illegal*:

```

BEGIN
...
p: PROCEDURE =
    BEGIN
        GOTO panicExit; -- illegal -- ...
    END;
... p[]; ...
EXITS
    panicExit => ...
END

```

The desired result is achieved with the following program (see chapter 8 for a description of signals and catch phrases):

```

BEGIN
Panic: SIGNAL = CODE;
...
p: PROCEDURE
    BEGIN
        ... SIGNAL Panic; ...
    END;
... p[!Panic => GOTO panicExit]; ...
EXITS
    panicExit => ...;
END

```

A statement in an **ExitsClause** may contain a **GOTO**, but the label in the **GOTO** can only refer to labels in surrounding blocks, not to labels in the same **ExitsClause** as the **GOTO**. For example, the following is legal:

```

BEGIN                                -- outer
...
  BEGIN                              -- inner
  ...
  EXITS
  endOfFileReached = > BEGIN ... GOTO outOfData END;
  END;                                -- end of inner
...
  EXITS
  outOfData = > ...;
END                                  -- end of outer

```

4.4.2 OPEN clauses

An **OPEN** clause allows more convenient reference to the fields of a record. In the simplest form, it allows *fieldname* as an abbreviation for *recordname.fieldname*. If the name of the record is complicated (e.g., *candidateList[tableOfObjects[i]]*), this can make programs much more readable. The programmer should be cautioned, however, that this is merely a syntactic shorthand; the code generated is actually *recordname.fieldname* and *recordname* is recomputed each time. Thus, in the example above, if *i* or *tableOfObjects* is changing within the scope of the **OPEN**, each reference to a field can potentially access a different element of *candidateList*. Similarly, an **OPEN** clause may simplify access to an interface which it opens (§ 7.2.2.2). The syntax for **OPEN** follows:

```

OpenClause      ::= empty | OPEN OpenList;  -- note the terminal semicolon

OpenList       ::= OpenItem | OpenList , OpenItem

OpenItem       ::= AlternateName : Expression |
                  Expression

AlternateName   ::= identifier

```

The *scope* of an **OPEN** clause (the portion of the program over which the synonym can be used) is the body of a block or loop, including the optional **EXITS** clause (§ 4.5). The following diagram summarizes the scope of the various parts of a **Block**. The scope of each phrase extends over others with greater indentation.

```

BEGIN
  OpenClause
    EnableClause
      DeclarationSeries
        StatementSeries
    ExitsClause
END

```

An **OpenItem** using an **AlternateName** allows a simple identifier to replace an expression as the designator of some record object. For example, the two blocks below are equivalent:

```

PersonChain: TYPE = RECORD [ p: POINTER TO Person, next: POINTER TO PersonChain ]
candidateList: POINTER TO PersonChain;          -- Person is defined in section 3.4

```

```

BEGIN OPEN c: candidateList.p;
IF c.party = Republican AND c.age < 30 THEN youngRepublicans ←
youngRepublicans + 1;

IF c.sex = Female THEN women ← women + 1;
...
END

BEGIN
IF candidateList.p.party = Republican AND candidateList.p.age < 30 THEN
  youngRepublicans ← youngRepublicans + 1;
IF candidateList.p.sex = Female THEN women ← women + 1;
...
END

```

The **OPEN** statement does not provide a general renaming capability; it merely allows more convenient access to the fields of a record. Each **Expression** in an **OpenList** must either have a record type or be a pointer to a record. When the **AlternateName** form is used, the alternate identifier always designates the opened record, even if the **Expression** is a pointer to that record.

The form of **OpenClause** without an **AlternateName** allows access to the fields of a record object as though they were simple variables. For example, using this feature in the above example allows omission of the "c."s:

```

BEGIN OPEN candidateList.p;
IF party = Republican AND age < 30 THEN youngRepublicans ←
youngRepublicans + 1;

IF sex = Female THEN women ← women + 1;
...
END

```

*Note: if the **AlternateName** form is used, qualification of record fields using the alternate name is mandatory.*

Besides record objects, it is possible to open a *module* (chapter 7) or an interface to simplify access to the identifiers available from the module, or items from an interface. However, the use of the **AlternateName** form of **OpenClause** is most strongly encouraged in these cases because the resulting code is more clear.

If an **OpenClause** contains multiple **OpenItems**, the opened expressions might refer to records having some selector names the same. In the example below, *x* is a selector name for two records, *recVar* and *recVar.subRecord*. An unqualified occurrence of *x* is taken to be the *x* component of the *rightmost* opened record (*recVar.subRecord*). To refer to an earlier opened record, explicit qualification is necessary (the **AlternateName** form should be used).

```

i, j: INTEGER;
RecordType: TYPE = RECORD
[
  a, b, x: INTEGER,
  subRecord: RECORD [x, y: INTEGER]
];
recVar: RecordType;

```

```

BEGIN OPEN r1: recVar, recVar.subRecord;
i ← r1.a + r1.b * r1.x; j ← x - y;
END;

```

The above block is equivalent to:

```

BEGIN
i ← recVar.a + recVar.b * recVar.x; j ← recVar.subRecord.x - recVar.subRecord.y;
END;

```

Fine points:

The range of text affected by an **OpenItem** includes any further items in the **OpenList**. The **OpenClause** itself may use implied qualification or alternate names (from earlier **OpenItems**). Thus, in the above example, one could have said

```
OPEN r1: recVar, r1.subRecord;
```

rather than

```
OPEN r1: recVar, recVar.subRecord;
```

Opened expressions are evaluated *at each use*, whether used implicitly or explicitly under an alternate name. This is essential for dealing with relocating allocation schemes. To avoid confusion, however, it is recommended that ordinary pointers be updated before entering the statement sequence headed by an **OpenClause**. In that way, names in the statement sequence will remain consistent, i.e., will apply to the same objects throughout. Consider an extension of the above example:

```

a: ARRAY [0..10] OF RecordType;
BEGIN OPEN r: a[i];                                -- assume i IN [0..9]
  r.x ← 2;
  i ← i + 1;
  j ← r.x;                                          -- reevaluation of OPENed expression with new
                                                    -- value of i, this r.x is not the one just stored into
END;

```

```

p: POINTER TO RecordType;
p ← . . .                                         -- p gets a new value
BEGIN OPEN r: p ↑;
  r.y ← 2;
  p ← . . .                                       -- p gets a new value
  j ← r.y;                                        -- field of newly pointed to record
END;

```

4.5 Loop statements

In Mesa, a loop is a statement containing a series of statements that are to be executed repeatedly. *All the ways of controlling how many times a loop should be repeated include the ability to repeat it zero times: i.e., to bypass it entirely.* Example 1 in section 2.1 contains the following loop statement:

```

UNTIL  $n = 0$ 
DO
   $r \leftarrow m \text{ MOD } n;$            --  $r$  gets remainder of  $m/n$ 
   $m \leftarrow n; n \leftarrow r;$ 
ENDLOOP

```

“UNTIL $n=0$ ” is the *loop control* for this loop. A variety of loop controls are available in Mesa: they include control by a Boolean expression, as above, and control by iteration over a subrange, as in the following example:

```

FOR  $i$  IN  $[0..N)$  DO  $a[i] \leftarrow a[i] + b[i]$  ENDLOOP

```

This will execute the assignment N times, with i taking the values 0, 1, ..., $N-1$ on successive iterations. If $N = 0$, the assignment is not executed at all.

The formal syntax of loop statements is

```

LoopStmt      ::= LoopControl      -- optional; may be empty
                DO
                OpenClause          -- optional (§ 4.4.2)
                EnableClause        -- optional (§ 8.2.1)
                DeclarationSeries
                StatementSeries
                LoopExitsClause     -- optional; may be empty
                ENDLOOP

```

The portion between **DO** and **ENDLOOP** is the *body* of a loop. Subsequent sections discuss the forms of **LoopControl**, the **LoopExitsClause** and **GOTOS** in loops.

The scopes of identifiers introduced in the various components of a loop are summarized by the following diagram (cf. **Block**, § 4.4.2):

```

LoopControl
  DO
    OpenClause
      EnableClause
        Declaration Series
          StatementSeries
            LoopExitsClause
  ENDLOOP

```

As in the case of a block, any exit labels are visible within the **EnableClause**, and any catch phrase (§ 8.2.1) in the **EnableClause** is not enabled within the **LoopExitsClause**.

4.5.1 Loop control

The syntax for **LoopControl** is

```

LoopControl    ::= IterativeControl ConditionTest  -- either may be empty
ConditionTest  ::= empty | WHILE Expression | UNTIL Expression

```

If both the **IterativeControl** and the **ConditionTest** are missing from a loop, it will repeat indefinitely (unless terminated by an embedded **GOTO** or **EXIT**, § 4.5.2).

If a **LoopControl** includes a **ConditionTest**, the boolean expression in the test is (re)evaluated before each execution of the loop body, including the first. If the **ConditionTest** *succeeds*, the body of the loop is executed; if it *fails*, the loop is finished (*terminates conditionally*) and control continues at *Next-Statement* (or at a **FINISHED** clause, § 4.5.2). A **WHILE** test succeeds if the value of the expression is **TRUE**. In the following example, *i* has the values 1, 2, 3, ..., 9 in successive executions of the body of the loop, and the value 10 when *Next-Statement* is reached (assuming that there are no other assignments to *i*):

```

i ← 1;                               -- this statement is not part of the loop
WHILE i < 10
  DO ... i ← i + 1;    ... ENDLOOP;
Next-Statement

```

An **UNTIL** test succeeds if the value of the expression is **FALSE**: i.e., it is the opposite of **WHILE**. The following loop is equivalent to the one above:

```

i ← 1;                               -- this statement is not part of the loop
UNTIL i ≥ 10
  DO ... i ← i + 1;    ... ENDLOOP;
Next-Statement

```

An **IterativeControl** provides a way of executing a loop (no more than) a computed number of times. It may be followed by a **ConditionTest**. It optionally updates a specified **ControlVariable** prior to each iteration so that, e.g., statements in the body have access to (a simple function of) the number of iterations. A loop that finishes by satisfying the implicit test associated with an **Iteration** or a **Repetition** is said to *terminate normally*.

```

IterativeControl  ::= empty | Repetition | Iteration | Assignment
Repetition       ::= THROUGH LoopRange
Iteration        ::= FOR ControlVariable Direction IN LoopRange |
                   FOR ControlVariable: TypeExpression Direction IN LoopRange
LoopRange        ::= SubrangeTC | TypeIdentifier |
                   BOOLEAN | CHARACTER
Direction        ::= empty | DECREASING
Assignment       ::= FOR ControlVariable ← InitialExpr , NextExpr |
                   FOR ControlVariable: TypeExpression ← InitialExpr , NextExpr
ControlVariable  ::= identifier
InitialExpr      ::= Expression
NextExpr         ::= Expression

```

In the **Repetition** form of **IterativeControl**, a **LoopRange** specifies how many times the loop body is to be executed. For example,

```

THROUGH [1..100] DO ... ENDLOOP

```

executes the body 100 times. A loop range can have any element type (§ 3.1) or any subrange of **LONG INTEGER** or **LONG CARDINAL**. The bounds of a subrange can be arbitrary expressions and do not have to be compile-time constants (as they do in a **SubrangeTC**).

A **Repetition** and a **ConditionTest** may be combined in a single loop control. For example,

```
THROUGH [low..high] WHILE lineIsConnected DO . . . ENDLOOP
```

Normal termination occurs after $high - low + 1$ iterations; conditional termination can occur sooner if *lineIsConnected* is **FALSE** prior to some iteration. Note that if $low > high$, the interval [*low*..*high*] is empty and the loop body is not executed.

Iteration and **Assignment**, the two forms of **IterativeControl** that include a **ControlVariable**, begin with the keyword **FOR**. If the first option of either **Iteration** or **Assignment** is used, then the *control variable* must be a variable declared separately in the program. Its type becomes the target type for the various expressions in the remainder of the **IterativeControl**. The forms of **Iteration** and **Assignment** with “:TypeExpression” declare a new control variable. That variable cannot be explicitly updated (except by the **FOR** clause itself). Its scope is the entire **LoopStmt** introduced by the **Iteration** or **Assignment** including any **LoopExitsClause**. Note, however, that the value of a control variable used in an **Iteration** is undefined in the **FinishedExit** (§ 4.5.2).

An **Iteration** steps through a subrange much as a **Repetition**, which is described above. In addition, it may specify a **Direction**: whether to begin at the lower bound of the range and step up (**empty**) or at the upper bound and step down (**DECREASING**). In any case, the size of the step is always one; for (a subrange of) an enumerated type, this really means stepping from an element to its successor (if the direction is increasing) or to its predecessor (if the direction is **DECREASING**). The control variable is assigned the current control value each time around the loop.

When a loop terminates *normally*, the final value of the control variable is *not* defined. The only way to ensure that the control variable's final value is well defined is to terminate the loop conditionally or *forcibly* (e.g., using **EXIT** or **GOTO**, § 4.5.2).

The following examples shift the components of an array *vec* (indexed from [0..**LENGTH**[*vec*])) left or right one position, leaving one element unchanged:

```
FOR i IN [1..LENGTH [vec])
DO
  vec[i - 1] ← vec [i];           -- "Left-shift" vec's elements.
ENDLOOP;

FOR i DECREASING IN [1..LENGTH [vec])
DO
  vec [i] ← vec [i - 1];         -- "Right-shift" vec's elements
ENDLOOP;
```

In the second case, *i* is initially set to the value **LENGTH** [*vec*] - 1 and decremented by one for each subsequent iteration. During the last execution of the loop, *i* has the value 1.

Bounds expressions in a **LoopRange** are evaluated *exactly once*, before the first execution of the loop body. Subsequent alteration of variables used in those expressions does not affect the number of iterations. When an **Iteration** is combined with a **ConditionTest** in a single loop control, the control variable is updated and tested before the **ConditionTest** is evaluated.

In an **Assignment**, the value of the **InitialExpr** is assigned to the control variable prior to the first iteration. Before each subsequent iteration, the **NextExpr** is (re)evaluated and assigned to that variable. There is no implicit test associated with an **Assignment** as there is for an **Iteration**; thus, the user must either use a **GOTO** (§ 4.5.2) to terminate the loop or include a **ConditionTest** in the **LoopControl** with the **Assignment**. As with an **Iteration**, the control variable is updated for each iteration before any **ConditionTest** is evaluated. This form is useful for scanning a list structure, as in the following example:

```

NodeLink: TYPE = POINTER TO Node;
node, headOfList: NodeLink;
Node: TYPE = RECORD
  [
    listValue: SomeType,
    next: NodeLink      -- either NIL (end of list) or pointer to next element
  ];
...
FOR node ← headOfList, node.next UNTIL node = NIL
  DO ... ENDLOOP;

```

Fine point:

The control variable can be altered within a loop, but this is not recommended. An iterative loop control updates the variable according to its current value. If the statement sequence assigns a new value to the control variable, the expected series of values may be disrupted (by omission or duplication). For control variables declared in the **Iteration** or **Assignment**, altering the control variable is not allowed.

4.5.2 GOTOS, LOOPS, EXITS, and loops

A loop may be *forcibly* terminated by a **GOTO** (or an **EXIT**, see below). The **LoopExitsClause** serves the same purpose as the **ExitsClause** in a **Block**; there are just three differences:

- (1) The **LoopExitsClause** is bracketed by **REPEAT** and **ENDLOOP** instead of **EXITS** and **END**;
- (2) The **LoopExitsClause** may contain a final statement labeled with the keyword **FINISHED**; this statement is executed if the loop terminates normally or conditionally, but *not* if it is forcibly terminated.
- (3) There is a special case of the more general **GOTO**, called **EXIT**, which simply terminates a loop forcibly without giving control to any statement in the **LoopExitsClause**.

There is another kind of **GOTO** statement, **LOOP**, which does not terminate the loop but skips the remainder of the loop body in the current iteration.

Syntax equations:

```

LoopExitsClause  ::= empty | REPEAT LoopExits
LoopExits        ::= ExitSeries |
                  ExitSeries ; |
                  FinishedExit |
                  ExitSeries ; FinishedExit

```

```

FinishedExit      ::= FINISHED => Statement |
                   FINISHED => Statement;

LoopCloseStmt    ::= LOOP

ExitStmt         ::= EXIT

```

The **LOOP** statement is used when there is nothing more to do in the iteration, and the programmer wishes to go on to the next repetition, if any. For example,

```

stuff: ARRAY [0..100] OF PotentiallyInterestingData;
Interesting: PROCEDURE [PotentiallyInterestingData] RETURNS [BOOLEAN];
i: CARDINAL;

FOR i IN [0..100] DO
  -- some processing for each value of i
  ...
  IF ~Interesting[stuff[i]] THEN LOOP;
  -- process stuff[i]
  ...
ENDLOOP;

```

The example used in the previous section to illustrate **ConditionTests** can be rewritten using a **GOTO** and a **LoopExitsClause** as follows:

```

i ← 1;
DO
  IF i ≥ 10 THEN GOTO quit;           -- first statement in the body
  ... i ← i + 1; ...
REPEAT
  quit = > NULL;                     -- do nothing but exit the loop
ENDLOOP;
Next-Statement

```

Frequently, forcible loop termination requires no special processing in the **LoopExitsClause**. The **EXIT** statement simplifies this case by not requiring a labeled statement in that clause; in fact, no **LoopExitsClause** need be present. The above example can be further rewritten to use **EXIT**, as follows:

```

i ← 1;
DO
  IF i ≥ 10 THEN EXIT;                -- first statement in the body
  ... i ← i + 1; ...
ENDLOOP;
Next-Statement

```

An **EXIT** is less general than a **GOTO**. For instance, if one has a loop nested within another and wants to exit from both, **EXIT** cannot be used because it terminates only the inner loop. A **GOTO** can jump to the **ExitsClause** of any enclosing loop or block. The **ExitsClause** of either a block or a loop is considered to be outside of the block or loop. Thus, an **EXIT** can appear in any **ExitsClause** (provided there is an outer loop), and it causes forcible termination of the smallest surrounding loop.

The following example shows a typical loop that is terminated only by execution of an `EXIT` statement.

```

BufIndexType: TYPE = [1..max];
buf: ARRAY BufIndexType OF INTEGER;
i, x: BufIndexType;
...
FOR i ← x, (IF i = max THEN 1 ELSE i + 1)    -- Starting at point x,
DO
...                                             -- do something and then
IF buf[i] = 0 THEN EXIT;                       -- quit on a "clear" entry, or
buf[i] ← 0;                                    -- clear until one is found.
ENDLOOP;

```

The `NextExpr`, "IF *i* = *max* THEN 1 ELSE *i* + 1," makes *buf* behave as a ring buffer.

Sometimes one must detect normal (as opposed to forcible) termination of a loop, perhaps to take some "finishing" action. A final labeled statement with the label `FINISHED` (which may *not* appear as the identifier in a `GOTO`) provides this facility. For example,

```

FOR i IN [0..nEntries] DO
  IF a[i] = x THEN GO TO found;
  REPEAT
    found = > old ← TRUE;
    FINISHED = >
      BEGIN
        a[i ← nEntries] ← x;
        nEntries ← nEntries + 1;
        old ← FALSE
      END;
  ENDLOOP;

```

The `FINISHED` exit is taken if and only if the loop terminates normally or conditionally (i.e., when the loop range is exhausted in the case above). Upon entry to a `FINISHED` exit, the value of the **ControlVariable** is undefined. Note that if an `EXIT` statement is executed, the `FINISHED` statement is *not* executed.



Procedures

Procedures provide one of the most important abstraction mechanisms in Mesa. The definition of a procedure assigns a name to a function or action. The computation performed by a procedure is specified by a series of statements and can be expressed in terms of *parameters* of the procedure. In addition, a procedure can produce one or more values, called its *results*. To invoke or *call* a procedure, the programmer simply names it and supplies *arguments* corresponding to the parameters. He need not be concerned with the internal workings of the procedure and can use its meaningful name to denote the function or action.

The GCD computation in section 2.1 is of limited use as it stands because it depends upon (and changes) variables m , n and gcd declared somewhere in its environment. It can usefully be packaged as a procedure with parameters m and n , as in:

```
Gcd: PROCEDURE [m, n: INTEGER] RETURNS [CARDINAL] =
    BEGIN
    r: INTEGER;
    UNTIL n = 0
    DO
    r ← m MOD n; m ← n; n ← r;
    ENDOLOOP;
    RETURN [ABS [m]]
    END;
```

The parameters of a procedure constitute the fields of a record, called the *parameter record* of the procedure. When calling a procedure, the arguments are evaluated and assembled into an *argument record* using a constructor (§ 3.3.4). "Applying" a procedure value to that argument record invokes the procedure. Consider the procedure call $Gcd [x + 1, y]$. This evaluates $x + 1$ and y , constructs an argument record from these values, and then calls procedure Gcd , passing it the argument record.

Within the procedure, the argument record is assigned to the parameter record, and fields of the parameter record are accessed as simple variables (i.e., that record is **OPENED**). Thus, the effect of the call above is to assign the value of $x + 1$ to m and the value of y to n before the statements within Gcd are executed.

A procedure may return values to the point of its call. These *results* constitute a *result record*. There can be any number of results, and their types may differ. Within a procedure, a **RETURN** statement assembles the results into a record and then returns control to the caller. The procedure *Gcd* returns a result record with one component, of type **CARDINAL**. Thus, the form *Gcd*[*x* + 1, *y*] is an expression with a record type. Because of the automatic conversion from a single-component record to the component (§ 3.5.1), it can also be used in any context accepting a value of type **CARDINAL**.

The following assignment has an effect similar to that of the entire example in section 2.1:

```
gcd ← Gcd[m, n]
```

Note that arguments are always passed *by value* in Mesa. The arguments *m* and *n* (for which declarations must exist at the point of call in this case) are completely distinct from the parameters *m* and *n*, and execution of *Gcd* does not change the values of the former.

A procedure declaration, as illustrated above, defines an *actual procedure*. It introduces an identifier, supplies some *procedure type* for that identifier, and defines the computation to be performed by specifying a block called the *procedure body*. Such a declaration uses fixed form initialization and closely parallels the declaration of an ordinary variable with = initialization, such as,

```
octalRadix: CARDINAL = 8;
```

Other declaration forms may also be used, including procedure variables with values that can be updated to designate different actual procedures. *In Mesa, procedures are full-fledged data objects.*

A procedure type is defined by specifying its parameter and result records. For example, the type of *Gcd* is

```
PROCEDURE [m, n: INTEGER] RETURNS [CARDINAL]
```

Procedure types constructed with different parameter and result records are different. Thus, the type system helps to ensure that, even when procedure variables are used, a proper argument record is constructed for each procedure call (i.e., that the number and types of the arguments are correct), and that the result record is used correctly in the text surrounding the procedure call.

Since a procedure body is a block, it may contain *declarations*. These declare *local variables* for that procedure. Local variables are created when the procedure is called, may be directly accessed only from within it, and are destroyed when the procedure returns. Within a procedure body, the named fields of the parameter and result records are also considered local variables; they have the same lifetime and can be referenced without qualification. The local variables of *Gcd* are *m*, *n* and *r*.

Because this local storage is allocated and released dynamically, any Mesa procedure can be invoked recursively and used in a reentrant fashion. Thus, the following alternative declaration of *Gcd*, which directly mirrors a recursive definition of the greatest common divisor, is valid:

```

Gcd: PROCEDURE [m, n: INTEGER] RETURNS [CARDINAL] =
    BEGIN
    RETURN [IF n = 0 THEN ABS [m] ELSE Gcd[n, m MOD n]]
    END;

```

Fine points:

Although both versions of `Gcd` compute the same function, the recursive one is potentially extravagant in its use of time and space, especially since an iterative version is so easy to derive. This demonstrates an advantage of procedural abstraction: the second definition of `Gcd` could be replaced by the first without effect on any caller of `Gcd`.

The compiler detects cases of simple tail recursion such as that above and converts the procedure to an iterative one. Examples in section 5.4 demonstrate more appropriate uses of recursion.

A procedure body may also access variables declared outside the actual procedure. Such variables are *nonlocal* to the procedure; they exist longer than any single invocation of the procedure and must be defined in the enclosing program text.

Mesa also has extensive facilities supporting the separate compilation of packages of procedures and variables; these packages are called *modules* (chapter 7). These facilities allow one module to name and use the procedures in another, but the type-correct usage of argument and result records is still checked at compile-time.

If a procedure is called from many places, the “packaging” of code provided by the procedure body makes a program more compact. Procedure calls and returns, however, introduce some runtime overhead. If a procedure is called from exactly one place, that overhead is unnecessary. If it is called from time-critical code or if the body of the procedure is very simple, the overhead can be unacceptable. Mesa provides *inline* procedures for such applications. The call of an inline procedure is replaced by a modified copy of its body. This mechanism eliminates most of the overhead, but retains many of the advantages of procedures, such as introducing structure, improving readability, and isolating detail.

The foregoing discussion is only an introduction to procedures. The rest of this chapter provides further detail.

5.1 Procedure types

Procedure types are constructed by the syntactic form `ProcedureTC`, which is defined as follows:

```

ProcedureTC      ::= PROCEDURE ParameterList ReturnsClause
ParameterList   ::= empty | FieldList
ReturnsClause   ::= empty | RETURNS ResultList
ResultList      ::= FieldList
FieldList       ::= -- (§ 3.3.1)

```

The **ParameterList** and **ResultList** are **FieldLists** and define record types. If either is missing, the corresponding record type is "empty." A procedure type is fully determined by its parameter and result record types.

Default specifications are permitted for fields of a **ParameterList** or a **ResultList**, even if the **FieldList** is an unnamed one.

Fine points:

The form `[]` is permitted in the declaration of a **ParameterList** or **ReturnsClause** (it is equivalent to omitting the list or clause).

Note that constructors of procedure types require specification of the field lists; it is not possible to use a separately defined record type to specify an entire parameter or result record.

These records, unlike regular records, are not packed; every component is aligned (begins on a word boundary) to allow efficient passing of arguments and results.

PROC is acceptable as a short form of **PROCEDURE**.

A few typical procedure types are shown below:

PROCEDURE	-- takes no arguments; returns no results
PROCEDURE [<i>x</i> : INTEGER, <i>flag</i> : BOOLEAN]	-- takes two arguments
PROCEDURE RETURNS [<i>i</i> : INTEGER]	-- returns a single value
PROCEDURE RETURNS [<i>i</i> : INTEGER, <i>b</i> : BOOLEAN]	-- returns two results
PROCEDURE [<i>x</i> : INTEGER] RETURNS [<i>y</i> : INTEGER]	-- takes and returns one value

These are all distinct types; none conforms to any of the others.

Values with procedure types are allowed in Mesa; one may have procedure variables, arrays of procedures, records with components that have procedure types, and procedures with procedure parameters or results. The fundamental operations =, # and ← may be applied to procedure values with conforming types.

Constructors of procedure types appear most commonly in the declarations of actual procedures, but they may occur wherever a **TypeSpecification** is valid. Thus, a **ProcedureTC** can appear in such constructs as:

A variable declaration:

```
ErrorHandler: PROCEDURE [which: ErrorCode] ← DefaultHandler;
```

A type declaration:

```
ListProc: TYPE = PROCEDURE [in: List] RETURNS [out: List];
First, Rest, Last: ListProc;
```

A field list (notice the parameters *lessThan* and *swap*):

```
Sort: PROCEDURE [
  first, last: CARDINAL,
  lessThan: PROCEDURE [CARDINAL, CARDINAL] RETURNS [BOOLEAN],
```

```
swap: PROCEDURE [CARDINAL, CARDINAL]
];
```

An array declaration:

```
tOps: ARRAY OpNames OF PROCEDURE [T, T] RETURNS [T];
```

5.1.1 Procedure values and compatibility *

Equivalence and conformance of procedure types are defined in terms of relations between fields of their **ParameterLists** and **ResultLists**. If the number of parameters or results differs, one procedure type neither conforms to, nor is equivalent to, another. Otherwise, corresponding fields, matched according to position, are considered. Two procedure types are *equivalent* if, for each pair of fields, the names are identical (or both are unnamed), the types are equivalent, and both **DefaultOptions** are empty. One field is *compatible* with another if the names are identical or if either is unnamed, and if the types are equivalent. A procedure type *conforms* to another if all corresponding fields are compatible. Default specifications do not affect conformance.

All the assignments in the following example are valid because the types of the procedures conform:

```
Handle: TYPE = POINTER TO Person;
SignedNumber: TYPE = INTEGER;
Int: TYPE = INTEGER;
ProcA: PROCEDURE [h: Handle, v: SignedNumber];
ProcB: PROCEDURE [h: Handle, v: Int];
ProcC: PROCEDURE [POINTER TO Person, INTEGER];
ProcA ← ProcB; ProcA ← ProcC;
ProcC ← IF flag THEN ProcA ELSE ProcB;
```

Fine points:

In the current version of Mesa, the name of the component of a single-element parameter or result record is ignored when comparing two procedure types for conformance.

If one procedure type conforms to another, it also conforms freely (§ 3.5.3). Free conformance of procedure types is defined by the following less restrictive rule: One field is compatible with another if the names are identical or either is unnamed, and if the type of the first freely conforms to the type of the second. One procedure type freely conforms to another if, for the **ParameterList**, each field of the *second* is compatible with the corresponding field of the *first* and, for the **ResultList**, each field of the *first* is compatible with the corresponding field of the *second*.

In the following example, recall that *Handle* conforms freely to *ReadOnlyHandle* but not vice versa:

```
ReadOnlyHandle: TYPE = POINTER TO READONLY Person;
ProcX: PROCEDURE [ in: ReadOnlyHandle ] RETURNS [ out: Handle ];
ProcY: PROCEDURE [ in: ReadOnlyHandle ] RETURNS [ out: ReadOnlyHandle ];
ProcZ: PROCEDURE [ in: Handle ] RETURNS [ out: ReadOnlyHandle ];
```



```

-- valid assignments
ProcY ← ProcX; ProcZ ← ProcY; procZ ← ProcX;

-- invalid assignments
ProcX ← ProcY; ProcX ← ProcZ; ProcY ← ProcZ;

```

In determining the conformance of two procedure types, default specifications are ignored. Thus, it is possible to assign a procedure value to a procedure variable with differently specified defaults. In a procedure call, the type of the variable appearing in the call, not the declaration of the actual procedure, determines the treatment of defaults. Thus, the initializing declarations in the following example are valid. Note that the declaration of *Proc2* declares a procedure constant that is indistinguishable from *Proc1* except for the default value of its argument.

```

Which: TYPE = {proc1, proc2, proc3};

Proc1: PROCEDURE [ p: Which ← proc1 ] =
  BEGIN
  ...
  END;

Proc2: PROCEDURE [ p: Which ← proc2 ] = Proc1;
Proc3: PROCEDURE [ p: Which ] ← Proc1;

-- some calls
Proc1 [ ];           -- equivalent to Proc1 [proc1 ]
Proc2 [ ];           -- equivalent to Proc1 [proc2 ]
Proc3 [proc3 ];     -- note that Proc3 [ ] is not legal

```

5.2 Procedure calls

The syntax for calling a procedure is

```

CallStmt      ::= Variable | Call
Call          ::= Variable [ ComponentList ] | ...

```

where the **Variable** has some procedure type. Other forms of **Call**, discussed in chapter 8, specify “catch phrases” for dealing with signals (or errors) that are generated because of the call.

In a procedure call, the arguments are packaged into a record. Therefore, a procedure call may use all the syntax for record constructors in passing arguments. Components (arguments) may be specified using either keyword or positional notation. Arguments not explicitly specified may be supplied by default. The following calls of *Gcd* are equivalent:

```
Gcd[x+1, y]  Gcd[m: x+1, n: y]  Gcd[n: y, m: x+1]
```

Fine point:

The orders of evaluating the items in constructors (including argument lists) and the operands of infix operators (except **AND** and **OR**) are subject to change. In particular, programs that assume a left-to-right order of procedure calls in these contexts (e.g., `Divide[Pop[], Pop[]]`) are unlikely to work correctly.

If the **ReturnsClause** in a **ProcedureTC** is not empty, then its **ResultList** specifies the number and types of the results returned by a procedure of that type. It may be a named or an unnamed **FieldList** (§ 5.3.1 discusses the **RETURN** statement).

Procedures that return results must be called from within **Expressions** that use the results in some way. Such *function references* are valid **Expressions**. Procedures that do not return results are used in call statements. A procedure that does not return results is called by simply writing a **CallStmt** as a statement by itself. For example,

```

group: ARRAY [1..N] OF POINTER TO Person;
Younger: PROCEDURE [first, second: CARDINAL] RETURNS [BOOLEAN] =
    BEGIN RETURN [group[first].age < group[second].age] END;
Exchange: PROCEDURE [first, second: CARDINAL] =
    BEGIN
    t: POINTER TO Person;
    t ← group[first]; group[first] ← group[second]; group[second] ← t
    END;
...
Sort [first: 1, last: N, lessThan: Younger, swap: Exchange];

```

A call statement is ordinarily used to obtain *side effects*. Most often, these take the form of changes to variables that are not local to the invoked procedure, but they may also involve input or output. A function may also have side effects as well as return results. On occasion, only the side effects are important, and the user wishes to ignore the returned results. An easy way to do this is to assign the result record to an empty extractor:

```
[ ] ← F[x];           -- call F and discard its result record.
```

A call that supplies *no arguments* is written with an empty constructor, “[]”.

Fine point:

When a procedure call with no arguments is itself a statement, the empty brackets may be omitted. When such a call is used as an expression, the empty brackets are mandatory; otherwise, the value of the expression is the value of the procedure variable, not the value of its results. For example, consider the two procedure variables in the following:

```

Proc1: PROCEDURE RETURNS (INTEGER);
Proc2: PROCEDURE RETURNS (INTEGER);
...
-- here the program assigns values to the procedure variables
IF Proc1 = Proc2 THEN ...
-- compare the procedure variables
...
IF Proc1 [=Proc2] THEN ...
-- compare their results (integers)

```

At the time a call occurs, a specific *activation* is executing, the *caller's* activation. The effect of a call is to suspend execution of the caller, to create a new activation of the called procedure (including new storage for all parameters and local variables), and to begin execution of that activation. An important consequence of this structuring of procedure control is that all Mesa procedures are inherently capable of being *recursive* and *reentrant*.

5.2.1 Arguments and parameters

Arguments are values supplied at call-time; parameters are variables that are local to a given activation. The association of arguments with their parameters amounts to assignment, much as if the following were written:

```
InRec: RECORD [arg1: Type1, arg2: Type2, ... ];
...
InRec ← [arg1: val1, arg2: val2, ... ];      -- in the caller

...
param1: Type1;
param2: Type2;
[ param1, param2, ... ] ← InRec;           -- in the called procedure
```

This is not just an idle analogy. The semantics of assignment accurately describe how arguments are associated with parameters. The following are direct consequences of this:

An argument of a procedure need only *conform* to its parameter, just as for assignment.

All arguments are passed by value in Mesa: i.e., the value of an argument, not its address, is assigned to the parameter. Of course, this value itself can be an address (for example, if *Type1* were *POINTER TO TypeX*).

5.2.2 Termination and results

A procedure terminates by executing a **RETURN** statement, which constructs a (perhaps empty) result record. The return operation then terminates execution of the current procedure activation and restarts the caller from the point at which it was suspended by the call. As part of the return, storage for the parameters and local variables of the returning procedure is released.

Since the value of a procedure is its result record, the components of that record can be assigned to variables using an extractor. Alternatively, any single component (if named) can be referenced by a field selector. The procedure *ReturnExample* returns three integer results and may be used as indicated:

```
ReturnExample: PROCEDURE [option: [1..4]] RETURNS [a, b, c: INTEGER] =
  BEGIN . . . -- body defined in § 5.3.1 -- . . . END;

x, y, z: INTEGER;
case:[1..4];
...
x ← ReturnExample[case].a;           -- get a component only
[b: y, c: z] ← ReturnExample[case];  -- extractor a subset of the results
x ← (ReturnExample[case].c + 1) MOD 10; -- use c component
```

If a procedure returns an empty result record, the call does not have a value and can only be used as a statement.

If a procedure returns a single-component result record, extraction and selection are valid. In addition, the component may be (and usually is) accessed directly because of the automatic coercion from a single-component record to its single component. In the

following example, the first two calls of *Gcd* are valid and equivalent; the third illustrates typical use within an expression:

```
gcd ← Gcd[m, n];           -- (coercion)
[ gcd ] ← Gcd[m, n];      -- (explicit extraction)
relPrime ← Gcd[m, n] = 1; -- (coercion)
```

Fine points:

In the declaration of *ReturnExample*, `[a, b, c: INTEGER]` defines a unique type for the result record. Because of the conformance rule for record types (§ 3.3.2), it is impossible to declare a variable with that type. If a procedure is to return a record value with a particular type *T*, it must return a single-component record where that component is a record of type *T*.

For similar reasons, the result record of *G* below is not directly acceptable as the argument record of *F*.

```
F: PROCEDURE [ x, y: INTEGER];
G: PROCEDURE [ i: INTEGER] RETURNS [x, y: INTEGER];
```

With these declarations, the call `F[G[i]]` is not legal, but see subsection 5.3.2 for syntax `(APPLY(F, G[i]))` that allows this operation. It would also be legal with the following declarations:

```
T: TYPE = RECORD [ x, y: INTEGER];
F: PROCEDURE [ in: T ];
G: PROCEDURE [ i: INTEGER] RETURNS [ out: T ];
```

Note, however, that *F* takes and *G* returns only a single value now, one of type *T*.

5.3 Procedure bodies

An actual procedure declaration looks like the declaration of a procedure variable followed by a special kind of = initialization, a **ProcedureBody**. The **TypeSpecification** appearing in the declaration determines the type of the body as well as that of the procedure identifier. It may be any **TypeSpecification** equivalent to a **ProcedureTC**. **ProcedureBody** is a special form of initialization defined as follows:

```
Initialization      ::= ... | = ProcedureBody | ← ProcedureBody
ProcedureBody      ::= InlineOption Block -- see section 4.4 for Block
InlineOption       ::= empty | INLINE
```

If the attribute **INLINE** appears, the procedure body is an *inline* one; any call of the procedure is replaced by a modified copy of the body (§ 5.6).

Only a procedure initialized with = to a **ProcedureBody** is called an *actual procedure*; its meaning cannot change because it cannot be assigned to. If, however, it is initialized to a **ProcedureBody** using ← initialization, its value can be changed by assignment, and it is considered a procedure variable. Initialization using ← is not permitted for an inline procedure.

In addition to other statement forms, a procedure body can contain **RETURN** statements (described in the next section). There is an implicit **RETURN** at the end of each procedure body if one does not appear explicitly.

A **ProcedureBody** defines a *scope* for declarations. Identifiers declared within it are local to the procedure and are unknown outside it. There must be no duplicates among the names in a procedure's **ParameterList**, **ResultList**, and local variables. Names in the **ParameterList** can be used to write a keyword constructor (§ 3.3.4) in a call of a procedure. Similarly, names in the **ResultList** can be used in keyword extractors (§ 3.3.6) and as qualifiers (§ 3.3.3) to access the results returned by a procedure. Within a procedure, any named fields of parameter and result records act just as local variables; the former are initialized with the values of the actual parameters. A **ParameterList** for an actual procedure should be a *named* field list so that the procedure body can reference the parameters.

Fine point:

Although the parameters and results act as local variables within the block that is the procedure body, the scopes are slightly different. The scope of the named parameters and results includes any **OpenClause**, **EnableClause** or **ExitsClause** of that block; the scope of the local variables does not (§ 4.4.2).

5.3.1 RETURN statements

There are several basic forms of **RETURN** statements, two of which are discussed in this section: **RETURN** and **RETURN** followed by a constructor. When either form is executed, control returns to the point from which the procedure was called. In addition, the **RETURN** can supply results in the form of a constructor conforming to the type of the procedure's **ResultList**:

ReturnStmt ::= **RETURN** | **RETURN** [**ComponentList**] | . . .

There may be any number of **RETURN** statements in a procedure body. The form of a **RETURN** statement depends upon the **ReturnsClause** in the definition of the procedure type. There are three cases to be considered:

no ReturnsClause (empty result record)
 an *unnamed* field list as the **ResultList**
 a *named* field list as the **ResultList**

If there is *no ReturnsClause*, the **ReturnStmt** must be just "**RETURN**." An explicit **RETURN** statement can be omitted at the end of the procedure in this case.

If an *unnamed* field list is used for the **ResultList**, each **ReturnStmt** must include a positional constructor. That constructor must match the field list exactly, with one component for every field: *omission, elision and voiding are not allowed* (unless the **ResultList** has defaults). In this case, there is no implied return at the end of the procedure.

If the **ResultList** is a *named* field list, either form of **ReturnStmt** is acceptable. If no explicit constructor appears, the current values of the named result variables define the value of the result record. An explicit constructor may use either positional or keyword notation; again, omission, elision and voiding are disallowed (unless the **ResultList** has defaults). A **RETURN** statement is optional at the end of the procedure; if omitted, an implicit **RETURN** of the result variables is provided. Examples follow:

```

ReturnExample1: PROCEDURE [ option: [1..4 ]] RETURNS [ a, b, c: INTEGER ] =
  BEGIN
  a ← b ← c ← 0;
  SELECT option FROM
    1 => RETURN [ a: 1, b: 2, c: 3]; -- keyword parameter list
    2 => RETURN [ 1, 2, 3];         -- positional version of option 1
    3 => RETURN;                   -- a = b = c = 0
  ENDCASE => b ← 4;
  c ← 9;
  END;                               -- implicit return: a = 0, b = 4, c = 9

ReturnExample2: PROCEDURE [g: INTEGER] RETURNS [ INTEGER ← 3, INTEGER ← 4 ] =
  BEGIN
  SELECT g FROM
    0 => RETURN [ , 2];           -- elide first result; returns [3, 2]
    1 => RETURN [ 8, ];          -- elide second result; returns [8, 4]
    2 => RETURN [ , ];           -- elide both results; returns [3, 4]
    3 => RETURN [ 5];            -- omit second result; returns [5, 4]
    5 => RETURN [ ];             -- omit both results; returns [3, 4]
  ENDCASE ;
  END;                               -- implicit return: [3, 4]

```

5.3.2 Operations which deal with intact parameter records

Procedures in Mesa logically take and return “records.” These so called parameter and result records differ from other records in the language in several ways: they are unpacked, and, as indicated earlier, type equivalence is somewhat more relaxed, they must be assignable field by field in order to be assignable—fields must have assignable types and identical names (if there is more than one field and both records have named fields). The reason for requiring names to match is to disallow exporting *Proc*: PROCEDURE [*from*, *to*: POINTER] to an interface where the declaration is *Proc*: PROCEDURE [*to*, *from*: POINTER] (§ 7.4).

Consider the following declarations:

```

Proc1: PROCEDURE [x: CARDINAL] RETURNS [a: T1, b: T2] = ...
Proc2: PROCEDURE RETURNS [a: T1, b: T2] = ...
Proc3: PROCEDURE [a: T1, b: T2] = ...

```

Note that *Proc1* and *Proc2* have compatible return records.

The following statement is legal inside *Proc1*:

```
RETURN Proc2[]; -- call Proc2 and then pass its return record along to my caller.
```

You can even say something like:

```
RETURN (IF boolexp THEN Proc2[] ELSE Proc1[x + 2]); -- note the tail recursion in
-- the ELSE case.
```

The return record of a procedure may also be passed along to another procedure as its argument record. The syntax is as follows:

```

APPLY [Proc3, Proc2[]];           -- the return record of Proc2 is compatible with
                                -- Proc3's parameter record.

```

Just as the expression after **RETURN** must have a type compatible with the return record of the procedure, the second parameter of the **APPLY** must have a type compatible with the argument record of the procedure that is the first parameter to the **APPLY**. In the case of the **RETURN**, parentheses may be required for certain complicated expressions.

The **APPLY** and **RETURN** constructs also work for the argument and **RESUME** records of **SIGNALS** and **ERRORS** (§ 8.2.5).

5.3.3 Defaults in argument and result records

You may specify default values for the fields of argument and result records. Such default values must be constructed from constants or variables that are declared outside of the procedure type definition. In particular, you cannot use either a value of another field of the same record or, in the case of a result record, a value from the associated argument record to define such a default.

You may omit a field in the constructor of an argument or result record only if the definition of that record specifies an explicit default value for the field. Default initial values associated with the *types* of such fields are *not* inherited. (This protects you from assigning a value to a return variable and then forgetting to mention it in a **RETURN** statement, causing the default for its type to be returned.) On the other hand, protection against ill-formed storage *is* inherited; you may not void or elide a field unless the type of that field allows a **NULL** initialization.

Defaults that you specify in the declaration of a result record serve two purposes. Since the fields of such a record can be used as local variables within the procedure body, a default specification affects the initialization of those variables; in addition, it allows abbreviation in the constructors of the corresponding return records. The precise rules are:

Upon entry to a procedure, each field of the result record is initialized with the default value specified for that field, if any; otherwise, it is initialized with the default initial value for the type of that field, if there is one; otherwise, its initial value is undefined.

If a **RETURN** is followed by an explicit constructor, the default specifications appearing in the declaration of the result record control the values of any omitted or elided fields, *even if other assignments have been made to the result variables within the procedure body*. If the **RETURN** either stands by itself, without such a constructor, or is implicit, the return record is constructed using the current values of the result variables.

Examples

```

T: TYPE = INTEGER ← 1;
Proc1: PROC [i: INTEGER ← 0, j: T];
Proc2: PROC RETURNS [m: T, n: INTEGER ← 2] = {
  -- m initialized to 1 (from T), n to 2
  Proc1[j: 3];           -- Proc1[i:0, j:3];
}

```

```

Proc1[i: 3];           -- illegal (j does not default to 1)
...
m ← 4; n ← 5;
... RETURN;          -- returns [4, 5]
... RETURN [m, n];   -- also returns [4, 5]
... RETURN [m];      -- returns [4, 2]
... RETURN [NULL, n]; -- illegal (declaration of T disallows voiding of m)
... RETURN [, n];    -- illegal (m does not default to 1 or 4)
... RETURN [6, 7];   -- returns [6, 7]
... };               -- implicitly returns [4, 5]

```

5.4 A package of procedures

This section contains an example of a simple module, *BinaryTree*, which is designed to create and manage a data base structured as a binary tree. It is typical of the ways in which related procedures are packaged together. The example illustrates many of the issues discussed in the previous sections and introduces the use of modules and interfaces in Mesa.

The binary tree implemented by the example is a data structure containing nodes linked by pointers. Any node points to at most two others (its *sons*), and a node is pointed to by exactly one other node (its *parent*). A special *root* node exists and is referenced by a pointer not in the tree. Every node also contains a *value*, which for simplicity in the example is just an `INTEGER`. When the program starts, the tree is empty, and any call to *SeekValue* will return a count of zero.

The nodes in this particular binary tree are records with four components:

```

value           -- an integer value (with unspecified interpretation),
count          -- the number of duplications of the value in the data base,
left           -- pointer to a "left" son node (or NIL), and
right          -- pointer to a "right" son node (or NIL).

```

There are rules of association between the values and the nodes:

The first supplied value is entered into the root node.

A given value exists in only one node; duplications are counted.

If node E points to "left" son L, then all values in the subtree rooted at L are *less than* the value in E. If node E points to "right" son G, then all values in the subtree rooted at G are *greater than* the value in E.

When the module is started, the tree is initialized to be empty. Thereafter, the module itself executes no code, but its procedures can be called to alter the tree that it manages. For instance, other modules call *PutNewValue* to insert new values into the tree.

PutNewValue calls another of *BinaryTree*'s procedures, *FindValue*, which traverses the tree seeking a node that already has a given value. *FindValue* may find such a node, or it may fail by reaching a higher-valued node with a `NIL` left son or a lower-valued node with a `NIL` right son. If *FindValue* finds a node with the given value, *PutNewValue* increments that node's *count*. Otherwise, *PutNewValue* sets up a new node and attaches it to the node returned by *FindValue*.

This strategy is chosen for simplicity, but it can be a poor way to construct a binary tree. For instance, if the values are entered in strictly decreasing order, the tree becomes a linear list of left nodes. To find the lowest-valued node, *every* node must be examined.

The reader should read the explanation following the example in conjunction with the example itself.

Example 2. A Package of Procedures

```

1:  DIRECTORY
2:      Storage: USING [publicZone],
3:      OrderedTable: USING [UserProc ];
4:
5:  BinaryTree: PROGRAM IMPORTS Storage EXPORTS OrderedTable =
6:  BEGIN
7:
8:  -- type definitions and compile-time constants
9:  Node: TYPE = LONG POINTER TO BinaryNode;
10: BinaryNode: TYPE = RECORD [
11:     value: INTEGER, count: CARDINAL, left, right: Node ];
12:
13: -- a global variable
14: root: Node;
15:
16: -- public (exported) procedures:
17: SeekValue: PUBLIC PROCEDURE [val: INTEGER] RETURNS [count: CARDINAL] =
18: BEGIN
19:     node: Node;
20:     found: BOOLEAN;
21:     [found, node] ← FindValue[val];           -- see if it is in the tree
22:     RETURN [IF found THEN node.count ELSE 0 ]
23: END;
24:
25: PutNewValue: PUBLIC PROCEDURE [val: INTEGER] =
26: BEGIN
27:     node, nextNode: Node;
28:     alreadyInTree: BOOLEAN;
29:     -- Use FindValue to find where to put val:
30:     [alreadyInTree, node] ← FindValue [val ];
31:     IF alreadyInTree THEN node.count ← node.count + 1
32:     ELSE BEGIN -- name "external" UNCOUNTED ZONE publicZone by qualification
33:         -- allocate and initialize new node
34:         nextNode ← Storage.publicZone.NEW [BinaryNode ← [val, 1, NIL, NIL]];
35:         IF root = NIL THEN root ← nextNode
36:         ELSE IF val < node.value THEN node.left ← nextNode ELSE node.right ← nextNode;
37:     END;
38: END;

```

```

39:
40:  EnumerateValues: PUBLIC PROCEDURE [userProc: OrderedTable.UserProc] =
41:    BEGIN
42:      -- a local procedure (§ 5.6)
43:      Walk: PROCEDURE [node: Node] RETURNS [keepGoing: BOOLEAN] =
44:        BEGIN -- walk through the tree in order by increasing value using recursion
45:          RETURN [ node = NIL -- don't examine empty (sub) trees
46:            OR (
47:              Wal[node.left] -- enumerate the lesser-valued nodes first
48:              AND userProc [node.value, node.count] -- enumerate this node
49:              AND Walk [node.right] -- then enumerate the greater-valued nodes
50:            ) ]
51:        END; -- of Walk
52:      [ ] ← Walk [root]; -- just start enumerating at the root
53:    END;
54: -- a procedure that is private to this module
55:  FindValue: PROCEDURE [val: INTEGER] RETURNS [inTree: BOOLEAN ← FALSE, node: Node] =
56:    BEGIN
57:      nextNode: Node ← root; -- always start at the root
58:      IF root = NIL THEN RETURN [FALSE, NIL];
59:      UNTIL nextNode = NIL
60:        DO
61:          node ← nextNode;
62:          nextNode ← SELECT val FROM
63:            < node.value => node.left,
64:            > node.value => node.right,
65:          ENDCASE => NIL;
66:        ENDOLOOP;
67:      RETURN [ val = node.value, node ]
68:    END;
69:
70: -- mainline statements
71:  root ← NIL; -- make tree initially empty
72:  END.

```

5.4.1 The example

Each line of the source code in Example 2 is numbered for convenient reference; other than that, the code could be compiled as it stands.

The body of a **PROGRAM** module resembles a procedure body: **BEGIN**, followed by declarations, then some statements, and finally **END**. The declarations and statements are both optional, but it would be unusual to omit the declarations.

In this example, the module *BinaryTree* declares five actual procedures: *SeekValue* (lines 17-23), *PutNewValue* (lines 25-38), *EnumerateValues* (lines 40-53), *Walk* (lines 43-51) and *FindValue* (lines 55-68). It also declares two types (*Node* and *BinaryNode*), and a single global variable (*root*). The scope of these declarations is the entire body of the module (lines 6-72). For example, *PutNewValue*, *EnumerateValues* and *FindValue* all reference the global variable *root*.

When a module is created and started (chapter 7), the global variables are created and any statements in its body are executed. *BinaryTree* has just one such statement (line 71), which creates the initial empty tree by assigning *NIL* to *root*. Storage for activations of modules is not released when control reaches the end of the main body. Global variables

such as *root* continue to exist and may be used to retain data shared by the actual procedures in the module.

The procedure *EnumerateValues* has two major distinguishing features: it takes a procedure value as a parameter, and it contains the declaration of a nested procedure (*Walk*). For each node in the tree, *EnumerateValues* calls the procedure value *userProc* that it received as an argument, passing it the value in that node and its replication count. If *userProc* returns **TRUE**, the enumeration of the values continues; if it returns **FALSE**, *EnumerateValues* terminates and returns to its caller. The values are generated in order from least to greatest.

The nested procedure *Walk* is recursive and traverses the tree by first traversing the left subtree, then visiting the root, and finally traversing the right subtree. This postorder traversal delivers the values in increasing order (the reader should convince himself that it does). The expression in lines 47-50 depends upon the definitions of **AND** and **OR** (§ 2.5.3) to terminate the traversal as soon as *userProc* returns **FALSE**. The first procedure call occurs only if *node* is not **NULL**; the second only if the first is called and returns **TRUE**; the third only if the first and second are called and both return **TRUE**. Section 5.6 treats local procedures in more detail.

5.4.2 Invoking procedures in other modules

The **DIRECTORY** section at the beginning of a module lists the interfaces used in that module. The identifier *Storage*, for example, must be the name of a (**DEFINITIONS**) module. Such modules allow the independent development of interface definitions and the sharing of such definitions. *Storage* and *OrderedTable* are said to be *included* by *BinaryTree*. The optional **USING** clause provides compiler-checked documentation of exactly which identifiers are used in a module but defined in the associated interface.

The **IMPORTS** list (§ 7.4.1) on line 5 allows *BinaryTree* to access a variable (*publicZone*) defined in the interface *Storage*, which has the following (skeletal) form:

```
Storage: DEFINITIONS =
  BEGIN
  ...
  publicZone: UNCOUNTED ZONE;
  ...
  END.
```

The example uses explicit qualification (dot notation) to name *publicZone* (line 34).

The **EXPORTS** list (§ 7.4.3) names the single interface *OrderedTable*, which is defined as follows:

```
OrderedTable: DEFINITIONS =
  BEGIN
  -- types
  UserProc: TYPE = PROCEDURE [val: INTEGER, count: CARDINAL] RETURNS [continue:
  BOOLEAN];
  -- the interface
  SeekValue: PROCEDURE [val: INTEGER] RETURNS [count: CARDINAL];
  PutNewValue: PROCEDURE [val: INTEGER];
```

```

EnumerateValues: PROCEDURE [userProc: UserProc];
END.

```

Other modules access the **PUBLIC** procedures in *BinaryTree* (*SeekValue*, *PutNewValue* and *EnumerateValues*) by importing this interface (just as *BinaryTree* imports *Storage*); they have no other access to *BinaryTree*. For example, *FindValue* is private to *BinaryTree*, so it is only called from within the module (lines 21 and 30). The definition of the type *UserProc* is included in the interface so that it is publicly available for defining procedures to be passed to *EnumerateValues*. Note that *BinaryTree* also obtains the definition of this type from the interface (line 40).

5.5 Nested procedures

Actual procedures may be declared within procedure bodies. A *nested procedure* is one declared within (and local to) some enclosing procedure. Nesting of procedure declarations restricts the scope of the names of the inner procedures. In addition, the enclosing procedure establishes an environment for the inner; this is especially useful when the inner procedure is passed as a parameter.

The value of a nested procedure (and any activation of that value) is “tied” to the local variables of the enclosing procedure and, indirectly, to the local variables of the procedure or module in which the enclosing one is declared. An activation of the nested procedure references those variables available at its point of declaration. A different activation of the enclosing procedure declares a nested procedure with a *different* value, one with its nonlocal variables tied to that other instance of the enclosing procedure.

The following example uses the interface *OrderedTable* defined in subsection 5.4.2 and illustrates a typical application of a nested procedure.

```

AverageValue: PROCEDURE RETURNS [INTEGER] =
  BEGIN
    sum, n: INTEGER;
    AddValue: OrderedTable.UserProc =      -- a nested procedure
      BEGIN
        n ← n + count; sum ← sum + count*val;
        RETURN [continue: TRUE]
      END;
    sum ← n ← 0;
    OrderedTable.EnumerateValues [AddValue]
    RETURN [IF n = 0 THEN 0 ELSE (IF sum < 0 THEN sum - (n/2) ELSE sum + (n/2))/n ]
  END;

```

The procedure *AverageValue* computes the average value of the *value* fields in the binary tree. It declares and initializes a pair of local variables (*n* and *sum*) that are updated by the nested procedure *AddValue*, but must have a greater lifetime than any individual activation of *AddValue*. Note that a similar effect could be achieved here by making *n* and *sum* global variables; the suggested solution restricts their scope (and thus, the opportunity for accidental misuse).

Execution of *AverageValue* involves a second nested procedure, the procedure *Walk* within *EnumerateValues*. The latter’s parameter *userProc* serves a purpose similar to that of *sum* or *n* in *AverageValue*. Since there is nothing to prevent a recursive call of

EnumerateValues from some actual procedure corresponding to *userProc*, making *userProc* a global variable in the module *BinaryTree* could be disastrous.

Fine point:

Because a nested procedure is tied to an activation of the enclosing procedure (even when it references nonlocal variables), the value of a nested procedure should not be assigned to a variable with a lifetime greater than that of the enclosing procedure instance.

In a sense, all procedures are “local” procedures. They are either local to some enclosing procedure or local to some module (recall that static variables are local to the module declaring them). This nesting can continue to an arbitrary number of levels. (The level is important only to the extent that it influences name scopes, a topic covered in the next section.)

5.5.1 Scopes defined by procedures

Each procedure body defines a new scope for names declared in that procedure. Such names represent variables that are local to the body. The scope for a local variable is such that:

- (1) the local variable is unknown *outside* of that procedure body, and
- (2) a non-local variable is unknown *inside* the procedure *if its name matches some local variable's name*.

Within a procedure body, a block (§ 4.4) can be used to further restrict the scope of a local variable. In the following example, scopes for the procedures are indicated by comments:

```

SomeModule: PROGRAM =
  BEGIN
    var: INTEGER;           -- the var of INTEGER type is used here
    ...
    OuterProc: PROCEDURE
      BEGIN
        var: BOOLEAN       -- the var of BOOLEAN type is used here
        ...
        LocalProc: PROCEDURE =
          BEGIN
            var: CHARACTER; -- the var of CHARACTER type is used here
            ...
            END;
            ...
            END;
            ...
            END.
            ...
            END.

```

5.6 Inline procedures *

An actual procedure is said to be *inline* if the attribute **INLINE** appears before the body in the declaration of that procedure. Any call of the procedure is replaced by an inline *expansion*, which is a modified copy of the procedure's body. The code of the procedure and any storage required for local variables are merged with the code and storage of the calling procedure

or module. Thus, **INLINE** procedures can be used to eliminate the overhead of a procedure call and return (usually at the cost of a longer object program).

The rules for expanding an **INLINE** ensure that the presence or absence of the **INLINE** attribute has no effect upon the meaning of a program. Execution of the expansion must always produce a result with the same logical behavior as the result of applying the following operations:

- (1) For each argument, create a uniquely named variable local to the caller, and initialize that variable with the value of the argument.
- (2) If there is a result record with named fields, enclose the body of the **INLINE** procedure in a block containing a declaration of each such field.
- (3) In the resulting block, replace each reference to a field of the parameter list by the identifier introduced in the first step for the corresponding argument.

Any global variables of the procedure body refer to the corresponding variables accessible at its point of declaration, not the point of call.

Fine points:

A catch phrase can be attached to the call of an inline procedure (§ 8.2.1). The arguments are evaluated outside the scope of the catch phrase.

The Mesa compiler attempts to discover many of the common cases in which "call by name" is equivalent to the "call by value" substitution described above. When it discovers such a case, the argument is substituted directly for the corresponding parameter.

The attribute **INLINE** is never mandatory. Deleting **INLINE** is always valid, but adding it is not. No inline procedure can be recursive, either directly or indirectly through a chain of inline procedure calls. Consider a procedure *Proc* declared as follows:

```
Proc: PROCEDURE [v: INTEGER] RETURNS [INTEGER] = INLINE
  BEGIN
    RETURN [v*v + 3*v + 1]
  END;
```

Because of its **INLINE** attribute, *Proc* cannot be used in any of the following situations:

When *Proc* itself is the operand of one of the fundamental operations of assignment (*procVar* ← *Proc*, *GeneratorProc*[*Proc*], etc.) or comparison (*Proc* = *AnotherProc*).

When *Proc* itself is used as an alternative in a conditional expression, e.g.,
(**IF** *predicate* **THEN** *Proc* **ELSE** *AnotherProc*)[*x*].

When *Proc* is the operand of **FORK** (§ 9.1).

When *Proc* is to be exported to an interface (§ 7.4.3).

Fine points:

Since arguments are evaluated before procedures are called, usage such as *Proc[Proc[x]]* does not make *Proc* recursive.

Additional restrictions apply when an inline procedure is declared in a **DEFINITIONS** module (§ 7.3.3).



Other data types and storage management

This chapter introduces a number of new data types: strings, array descriptors, sequences, and zone types. Relative pointers are also discussed, and the definition of record types is extended to include variant records.

In Mesa, the type `STRING` is really "POINTER TO *StringBody*"; a *StringBody* contains a packed array of characters, a *maxlength* field giving the length of that array, and a *length* field indicating how many of the characters are currently significant.

A *sequence* is an indexable collection of items, all of which have the same type. In this respect, a sequence resembles an array; however, the maximum length of the sequence is specified at run-time when the object containing that sequence is allocated, and this maximum length cannot subsequently be changed. It is the programmer's responsibility to keep track of the number of items in the sequence which have been assigned meaningful values.

An *array descriptor* describes the location and length of an array. For ordinary arrays, these are fixed at compile-time. Values of array descriptor *type*, however, have location and length items that can vary. These array descriptors may represent arrays that are dynamic, but they may also represent ordinary arrays. For efficiency, users often pass array descriptors to procedures instead of passing the entire arrays themselves.

Relative pointers require the addition of a *base* pointer to obtain an absolute pointer. This allows data structures with internal references that are independent of memory location.

Variant records contain a set of common fields and a variant portion with a specified set of different possible interpretations.

Dynamic variables in Mesa are allocated in *zones*. Zones are not necessarily associated with fixed areas of storage; rather they are objects characterized by procedures for allocation and deallocation. There is a standard system zone, but programs that allocate substantial numbers of similar dynamic variables can often improve performance by segregating each kind into its own zone. `NEW` is used to allocate a dynamic variable from a zone, and `FREE` to release it.

6.1 Strings

In Mesa, a **STRING** represents a finite, possibly empty, sequence of characters. Associated with a string are the following:

<i>text</i>	a PACKED ARRAY of characters.
<i>maxlength</i>	the maximum number of characters that <i>text</i> can hold.
<i>length</i>	the number of significant characters in <i>text</i> ; may vary from zero up to <i>maxlength</i> .

STRING is a predefined type in Mesa. Each program contains the following relevant pre-declarations:

```
STRING: TYPE = POINTER TO StringBody;
StringBody: TYPE = MACHINE DEPENDENT RECORD [
    length: CARDINAL,
    maxlength: -- read only-- CARDINAL,
    text: PACKED ARRAY [0..0) OF CHARACTER];
```

Suppose *s* is a **STRING** variable. Then *s.length* and *s.maxlength* refer to the first two components of the *StringBody* currently pointed to by *s*. The type *StringBody* is "built into" the Mesa language so that the *i*th character of the *text* array, *s.text[i]*, may be abbreviated *s[i]*. The index type of *text* in the declaration is used only to specify a starting index of 0. It is better to think of a particular **STRING** as having an index type [0..*s.maxlength*).

The value of *s.maxlength* is assigned when a *StringBody* is created and is a constant: *it may not appear as a LeftSide in the user's program*. However, *s.length* can be used as a **LeftSide**. In fact, the user is responsible for setting and changing the length when appropriate (i.e., *s.length* is meant to reflect the "meaningful" length of the character sequence). Suppose, for instance, that *s* initially points to a *StringBody*, having no significant characters, i.e., *s.length* = 0. Then the user might append characters as follows:

```
s[s.length] ← anotherChar; s.length ← s.length + 1;
```

(Actually, characters are seldom appended in this manner. *The recommended practice is to use string-handling procedures provided by the Mesa system*. These are documented in the *Mesa Programmer's Manual* and the *Pilot Programmer's Manual*.)

Since strings in Mesa are actually pointers to string bodies, several strings may refer to the *same* body. Therefore, a change to that structure would manifest itself in all such strings. Keep the following in mind:

*When an item has type **STRING**, think "pointer to string-body".*

Fine point:

While the programmer cannot assign to the *maxlength* field with an assignment statement, it can be set (along with the length) in a constructor, e.g.,

```

AllocateString: PROCEDURE [maxLength: CARDINAL] RETURNS [s: STRING] =
  BEGIN
    s ← AllocateWords [SIZE[StringBody] + (maxLength + 1)/2]
    s ↑ ← StringBody [length: 0, maxlength: maxLength, text: ];
  END;

```

This is the way to initialize a *StringBody* when the space for it comes from some general storage allocator. Note that the *text* field cannot be set with the constructor since the **ARRAY** is of length zero in the declaration. See also subsection 6.5.5.

6.1.1 String literals and string expressions

A string literal is a sequence of characters enclosed in quotation marks, "...". A quotation mark within a string constant is represented by a pair of quotation marks (""). Here are some examples of string literals:

```

"The first example contains
some embedded
carriage-returns."

"A single quote mark (') isn't a double quotation mark( "")...."
"!
""
-- an empty string

```

A string literal is an **Expression** of type **STRING**. Its value is a constant *pointer* to a constant **StringBody** in which:

```

length      = number of characters given, and
maxlength   = length

```

The fundamental operations are defined for string **Expressions**. *The fundamental operations deal with string expressions as pointer values*; e.g., \leftarrow assigns one string pointer to another string pointer, $=$ compares two strings for the same pointer value, and $\#$ compares two strings for different pointer values.

Fine point:

The *StringBody* of a string literal is normally allocated in the global frame of the module in which the literal appears (it is copied there from the code segment when the module is **STARTed**). Unfortunately, such strings can consume substantial amounts of space in the (permanent and unmovable) global frame area. The programmer can indicate that the *StringBody* should be allocated in the global frame by following the literal by a **G** (e.g., "abc"G).

If a string literal is followed by **L** (e.g., "abc"L), the *StringBody* is allocated in the local frame of the innermost procedure enclosing the literal; the *StringBody* is initialized by copying it from the code segment whenever an instance of that procedure is invoked. As a corollary, the space is freed and the *StringBody* disappears when the procedure returns. This allows smaller global frames, but it is important to insure that pointers to local string literals are not assigned to **STRING** variables with lifetimes longer than that of the procedure.

If there is a "global" copy of a string literal, saying "...L is a no-op. Since it is already in the global frame, it is not also copied into the local frame.

6.1.2 Declaring strings

String variables are declared like ordinary variables, but there is one additional form of initialization (for strings only):

```
Initialization      ::= ... ← [ Expression ] | = [ Expression ]
```

The **Expression** must be a compile-time constant expression of type **CARDINAL**. At run-time, Mesa creates a *StringBody* with *maxlength* equal to this **Expression**'s value, *length* equal to zero, and *text* uninitialized. The declared string variable is then set to point to this *StringBody*. If an **IdList** is declared with this form of initialization, all of the listed variables initially point to the same *StringBody*.

Some examples:

```
currentLine, nowLine: STRING ← [256];
stringBuffer: STRING ← [stringMax + someExtra];
```

This would allocate two *StringBodys* in the local frame of the program or procedure containing the declarations. The strings *currentLine* and *nowLine* would point to one *StringBody*, with *maxlength* 256. The string *stringBuffer* points to the other *StringBody*. (Note that *stringMax* and *someExtra* must be compile-time constants.) Since the initialization is done with "←", it is legal to assign new pointer values to these string variables. The space for the string bodies is recovered when the procedure returns; care should be exercised to avoid assigning such a string pointer to a variable with a longer lifetime.

The following are examples of fixed form string initialization:

```
whatWasThat: STRING = "Eh?";
goofed: STRING = whatWasThat;
```

In this case, Mesa would allocate and fill in a *StringBody* for the string literal "Eh?". *whatWasThat* and *goofed* would be compile-time constants having the same string value: i.e., they would both point to the same *StringBody*. In fact, any other references to the same string literal will point to the same *StringBody*. For example:

```
huh: STRING = "Eh?";
answer: STRING ← "Eh?";
```

The strings *huh* and *answer* would point to the same *StringBody* as *whatWasThat*.

Fine point:

Since string literals can be assigned to string variables (pointer assignment), it is possible to modify the text of a literal. Doing this can lead to significant confusion. For example:

```
s1: STRING ← "abcdefg";
...
s1[2] ← 'x';
...
WriteString["abcdefg"] -- will write "abxdefg"
```

String variables can be declared with ← initialization or without any initialization:

```

stdErrorMsg: STRING ← "It seems that we have made a mistake."
firstReply, reply: STRING ← "Yes";
oldBuffer, newBuffer: STRING;
...
IF quickDialog THEN stdErrorMsg ← whatWasThat;
...
IF reply[0] = "?" THEN
  IF firstReply[0] = "?" THEN HelpaLot[]
  ELSE HelpaLittle[];
...
oldBuffer ← newBuffer ← stringBuffer1;
...
IF stringBuffer1 # stringBuffer2 THEN newBuffer ← stringBuffer2;

```

Fine point:

The Mesa system contains procedures you should use when allocating blocks of data. These procedures are helpful for applications involving an arbitrary number of strings or strings of arbitrary length. The procedures are documented in the *Mesa Programmer's Manual* and *Pilot Programmer's Manual*.

6.1.3 Long strings

A **STRING** is just a pointer, so **LONG STRING** is also a predefined type:

```
LONG STRING: TYPE = LONG POINTER TO StringBody;
```

It is perhaps curious to note that declaring a **LONG** string says nothing about its actual or potential *maxlength*.

As indicated earlier, Mesa has a facility for allocating string bodies in the local frame of a procedure by saying *localString*: **STRING** ← [exp]. You can also say *longLocalString*: **LONG STRING** ← [exp] as well (exp must be constant in either case). Note that you need not declare a local string to be **LONG** if you are only passing it to other procedures; the compiler will lengthen the pointer for you.

Most system routines deal with **LONG STRINGS** since most *StringBodys* are allocated dynamically from **UNCOUNTED ZONES** (§ 6.6.1).

6.2 Array descriptors

A full description of an array contains several items of information. Consider a typical array declaration:

```
schedule: ARRAY [0..999] OF Date;
```

The following things are known about *schedule*:

```

base = @schedule[0],
index type = [0..999] (a subrange of INTEGER or CARDINAL),
minIndex = 0,

```

```

length = 1000,
component type = Date

```

All of these items except *base* are compile time constants, and the value of *base* is the address of a fixed place in the frame, chosen by the compiler. Mesa provides a mechanism for dynamic arrays, where the *base* and *length* can vary at run-time. The implementation does not allow for a variable *minIndex*. Dynamic arrays are implemented by *array descriptors*. Array descriptors are present in Mesa mainly for backward compatibility. When possible, use sequence-containing types for allocation of arrays with dynamically computed size; use array descriptor types only for parameter passing when necessary.

6.2.1 Array descriptor types

An *array descriptor* type is constructed much like an array type:

```

DescriptorTC      ::= = DESCRIPTOR FOR ReadOnlyOption ArrayTC |
                   DESCRIPTOR FOR ReadOnlyOption PackingOption ARRAY OF
                   TypeSpecification

ReadOnlyOption    ::= = empty | READONLY

PackingOption     ::= = empty | PACKED

```

For example,

```

events: DESCRIPTOR FOR ARRAY [0..999] OF Date;

```

If **READONLY** is specified, the contents of the array cannot be changed via the descriptor.

The value of *events* is an array descriptor (a record-like object containing items similar to those described previously for *schedule* except that the *base* is not fixed). The next declaration specifies an array descriptor in which the *base* and the *length* are variable:

```

history: DESCRIPTOR FOR ARRAY OF Date;

```

Indexing can be used to access components of *events* and *history* as if they were actual arrays instead of array descriptors. Since no index type is specified for *history*, it has an *indefinite* index type *starting at zero with no specified upper bound*. This is equivalent to declaring the index type as [0 . . 0).

Two array descriptor types are equivalent if they specify equivalent types for their array components and if they have equivalent index-sets (or if both index-sets are unspecified). Note that **DESCRIPTOR FOR ARRAY [0..2] OF T** and **DESCRIPTOR FOR ARRAY [1..3] OF T** are different types, even though the lengths and element types are the same. Expressions of equivalent descriptor types may be compared for equality (= or #).

The rules for assignment are somewhat more relaxed. If *a1* has type **DESCRIPTOR FOR ARRAY OF T**, and *a2* has type **DESCRIPTOR FOR ARRAY [0..10] OF T**, then the assignment *a1* ← *a2* is legal, but the assignment *a2* ← *a1* is not.

In any case, for assignments and comparisons, both operands must be array descriptors, and it is the descriptors themselves, *not* the arrays that they describe which are the values

operated on. It would be an error to attempt to assign *events* to *schedule* because the first is a descriptor and the second is an actual array.

There are three function-like operators relevant to array descriptors: **DESCRIPTOR**, **BASE**, and **LENGTH**. **DESCRIPTOR** returns an array descriptor result and has three distinct forms which are treated syntactically as built in functions:

```

BuiltInCall      ::= = DESCRIPTOR [ Expression ] |
                  DESCRIPTOR [ Expression , Expression ] |
                  DESCRIPTOR [ Expression , Expression , TypeSpecification ] |
                  BASE [ Expression ] | LENGTH [ Expression ] | . . .

LeftSide         ::= = Expression . BASE |
                  Expression . LENGTH | . . .

```

The first form takes an argument of some array type, e.g.,

```
events ← DESCRIPTOR[schedule];
```

The result is an array descriptor for *schedule*. The second form needs two arguments:

```

base: POINTER TO UNSPECIFIED      -- address of the minIndex component
length: CARDINAL                  -- number of components

```

This form is usually assigned to an array descriptor variable which was declared without an explicit index type.

In those rare situations where the compiler cannot deduce the component type of the descriptor from context, a form of the **DESCRIPTOR** construct is provided which takes three arguments. The third one is a **TypeSpecification**, the component type.

The following example provides a fresh array of 64 *Dates*:

```

Allocate: PROCEDURE [blkSize: CARDINAL] RETURNS [POINTER TO UNSPECIFIED];
...
history ← DESCRIPTOR[Allocate [64*SIZE[Date]], 64];

```

The expressions **BASE**[] and **LENGTH**[] take one argument (of array descriptor or array type). **BASE** yields the base of the described array, and **LENGTH** yields its length. For example:

```

events ← DESCRIPTOR[schedule];           -- describe the entire array
events ← DESCRIPTOR[BASE[schedule], 5];  -- describe the first 5 elements

```

One can assign to the individual fields of an array descriptor by using the "dot notation," forms of **BASE** and **LENGTH**. For example:

```
events.LENGTH ← 4;                       -- describe only 4 elements
```

There is no special form for constructing **DESCRIPTORS** for packed arrays. The **PACKED** attribute is deduced from context. In the two or three argument form of **DESCRIPTOR** for

packed arrays, the second argument (the `LENGTH`) is the number of elements. At present, the `DESCRIPTOR` operator cannot be applied to packed arrays which occupy less than a word.

It is usually more efficient to pass array descriptors as arguments, rather than arrays. Since arguments are passed by value, an array argument causes a copy of the entire array to be made twice (once to put it into an argument record, and once to copy it into a local variable in the called procedure). The next example shows a case in which array descriptors must be used, since passing by value would not work:

```

Table: TYPE = DESCRIPTOR FOR ARRAY OF INTEGER;
SortInPlace: PROCEDURE[localTable: Table];      -- sorts in situ
thisArray: ARRAY [0..this] OF INTEGER;
thatArray: ARRAY [0..that] OF INTEGER;
anyTable: Table ← DESCRIPTOR[thisArray];
...
SortInPlace[anyTable];                          -- sorts thisArray
...
SortInPlace[DESCRIPTOR[thatArray]];             -- sorts thatArray

```

A *StringBody* (§ 6.1) contains an array, *text*, of characters. One must be careful when constructing a `DESCRIPTOR` for this array. Recall that the bounds of *text* are `[0..0]`. This declaration is used since the actual length of *text* varies from `STRING` to `STRING`. For this reason, the "one argument" form should not be used to construct a `DESCRIPTOR` for *text*.

```

textarray: DESCRIPTOR FOR PACKED ARRAY OF CHARACTER;
s: STRING;

textarray ← DESCRIPTOR[s.text];                 -- LENGTH[textarray] is incorrect
textarray ← DESCRIPTOR[BASE[s.text], s.length]; -- correct

```

6.2.2 Long descriptors

The `BASE` portion of an array descriptor is essentially a pointer. Just as the language allows the type `LONG POINTER`, it also allows the type `LONG DESCRIPTOR`. The syntax is straightforward:

```

TypeConstructor ::= ... | LongTC

LongTC           ::= LONG TypeSpecification

TypeSpecification ::= ... | DescriptorTC

```

All the standard operations on array descriptors (indexing, assignments, testing equality, `LENGTH`, etc.) extend to long array descriptors. The type of `BASE[desc]` is long if the type of *desc* is long. The `LENGTH` of an array descriptor is a `CARDINAL`, whether the descriptor (i.e. its `BASE`) is `LONG` or short.

Long array descriptors are created by applying `DESCRIPTOR []` to an array that is only accessible through a long pointer, or by applying `DESCRIPTOR [,]` or `DESCRIPTOR [, ,]` to operands the first of which is long. When a short array descriptor is assigned to a long one, the pointer portion is automatically lengthened. Alternatively, an array descriptor can be explicitly lengthened by the operator `LONG []`. Consider the following examples:

```

d: DESCRIPTOR FOR ARRAY OF T;
dd: LONG DESCRIPTOR FOR ARRAY OF T;
i, n: CARDINAL;
pp: LONG POINTER TO ARRAY [0..10) OF T;

dd ← DESCRIPTOR [pp ↑ ];           -- descriptor for the entire array
dd ← DESCRIPTOR [pp, 5];         -- descriptor for half of the array
dd ← d;                          -- automatic lengthening
pp ← BASE [dd];                  -- BASE of long is long
n ← LENGTH [dd];                -- LENGTH is always a CARDINAL

```

6.3 Base and relative pointers

Mesa provides relative pointers, i.e., pointers that are *relocated* by adding some base value before they are dereferenced. Relocation has the further effect of mapping a value with some pointer type into a value with a possibly different pointer type. Relative pointers are expected to be useful in such applications as the following:

Conserving Storage. Relative pointers can adequately identify objects stored within a zone of storage if the base of that zone is known from context. If the zone is of known and relatively small maximum size, fewer bits are needed to encode the relative pointers. Since a relative pointer and the corresponding base value can have different lengths, relative pointers can be shorter than absolute pointers to the same objects. Overall storage savings are possible when all the base values can be contained in a small number of variables shared among many different object references.

Providing Movable Storage Zones. If all interobject references within a storage zone are encoded as zone-relative pointers, the zone itself can be organized to contain only location-independent values. Moving the zone, possibly via external storage, requires only that a set of base pointers be updated.

Designating Record Extensions. Sometimes it is convenient to extend a record by appending information (especially variable-length information) to it. Pointers stored in, and relative to the base of, the extended record provide location independent and type-safe access to the extensions.

6.3.1 Syntax for base and relative pointers

The syntax for base and relative pointer type constructors is as follows:

```

PointerTC ::= ORDERED BaseOption POINTER OptionalInterval
               PointerTail
BaseOption ::= empty | BASE
TypeConstructor ::= ... | RelativeTC
RelativeTC ::= TypeIdIdentifier RELATIVE TypeSpecification

```

In a **PointerTC**, a nonempty **OptionalInterval** declares a subrange of a pointer type, the values of which are restricted to the indicated interval (and can potentially be stored in

smaller fields). Normally, such a subrange type should be used only in constructing a relative pointer type as described below, since its values cannot span all of memory.

A **BaseOption** of **BASE** indicates that pointer values of that type can be used to relocate relative pointers. Such values behave as ordinary pointers in all other respects with one exception: subscript brackets never force implicit dereferencing. Subscript brackets are used together with relative pointers to relocate relative pointers (see below). The attribute **BASE** is ignored in determining the assignability of pointer types.

A **RelativeTC** constructs a *relative pointer* or *relative array descriptor* type. The **TypeIdentifier** must evaluate to some (possibly long) pointer type which is the type of the base, and the **TypeSpecification** must evaluate to a (possibly long) pointer or array descriptor type.

The referent of a relative pointer is specified by using subscript-like notation in which the type of the "array" is the base type and that of the "index" is the relative pointer type. Thus if *base* is a base pointer and *offset* is a relative pointer (to *T*), the form

base[*offset*]

denotes an expression of type *T*, and the value of that expression is $(\text{LOOPHOLE}[base] + offset) \uparrow$.

6.3.2 A relative pointer example

Consider the *BinaryTree* example from section 5.4. In this program, an ordered table is stored as a binary tree. The tree is stored in the following Mesa data structure:

```
Node: TYPE = POINTER TO BinaryNode;
BinaryNode: TYPE = RECORD [value: INTEGER, count: CARDINAL, left, right: Node];
```

Suppose that the *BinaryNodes* are allocated from a contiguous region of memory. If the programmer now wishes to put the current state of the ordered table on secondary storage, it is not sufficient to simply write out the region of memory containing the *BinaryNodes*. This is because the data would make sense only if read back into exactly the same place in memory, a restriction that is difficult to live with. The difficulty stems from the absolute pointers used in the nodes. The problem can be solved by changing the definition of *Node*. If the *BinaryNodes* are allocated from a region of type *TreeZone*, let

```
TZBase: TYPE = LONG BASE POINTER TO TreeZone;
Node: TYPE = TZBase RELATIVE POINTER TO BinaryNode;
```

The procedure *FindValue* would be written as follows:

```
nullNode: Node = <some value never allocated (§ 6.3.3)>;
tb: TZBase ← ...;
root: Node ← nullNode;                                -- list is initially empty
...
FindValue: PROCEDURE [val: INTEGER] RETURNS [inTree: BOOLEAN ← FALSE, node: Node] =
BEGIN
  nextNode: Node ← root;
  IF root = nullNode THEN RETURN [FALSE, nullNode];
```

```

UNTIL nextNode = nullNode DO
  node ← nextNode;
  nextNode ← SELECT val FROM
    < tb[node].value = > tb[node].left,
    > tb[node].value = > tb[node].right,
  ENDCASE = > nullNode;
ENDLOOP;
RETURN [val = tb[node].value, node];
END;

```

The other procedures of *BinaryTree* can easily be rewritten to use the new definition of *Node*. The compiler would aid in the translation, since any unrelocated dereferencing of a *Node* would be a compile-time error.

This new implementation of *BinaryTree* has the feature that the *TreeZone* could be moved around in memory, or written and read on secondary storage, and only the base pointer *tb* need be updated to reflect the new position of the *TreeZone*.

6.3.3 Relative pointer types

An important topic to consider is the interaction of the relative pointer constructs with the type machinery of Mesa.

A *RelativeTC* constructs a relative pointer type whenever both the *TypeIdentifier* and the *TypeSpecification* evaluate to pointer types. Let a *RelativeTC* be

TypeIdentifier RELATIVE **TypeSpecification**,

where

TypeIdentifier is of type

[LONG] BASE POINTER [*SubRange_b*] TO [READONLY] *T_b*,

TypeSpecification is of type

[LONG] [ORDERED] [BASE] POINTER [*SubRange_r*] TO [READONLY] *T_r*,

and the brackets indicate optional attributes. Relative pointer values must be relocated before they are dereferenced. If *base* and *offset* are base and relative pointers respectively, *offset* ↑, *offset.field*, etc. are compile-time errors.

If the *TypeSpecification* says **READONLY**, a relocated pointer cannot be a **LeftSide**.

The base type must be designated by an identifier (rather than a *TypeSpecification*) to avoid syntactic ambiguities. Note that the form

LONG **TypeIdentifier** RELATIVE **TypeSpecification** -- *wrong*

does not have the effect of lengthening the base type and furthermore is always in error, since **LONG** cannot be applied to a relative type. The type designated by the *TypeSpecification* can be lengthened (to give a relative long pointer) using the form

TypeIdentifier RELATIVE LONG TypeSpecification .

Short relative pointers are never made long automatically. With respect to other operations (assignment, testing equality, comparison if ordered, etc.), relative pointers behave like ordinary pointers. In particular, the amount of storage required to store such a pointer is determined by the **TypeSpecification**.

Fine points:

In some applications, there is no obvious type for the base pointer, i.e., it might not be possible or desirable to describe a storage zone using a Mesa type declaration. In such cases, a declaration such as

```
BaseType: TYPE = BASE POINTER TO RECORD [UNSPECIFIED]
```

generates a unique type that will not be confused with other base types.

The declaration of a relative pointer does not associate a particular base value with that pointer, only a basing type. Thus some care is necessary if multiple base values are in use. Note that the final type of the relocated pointer is largely independent of the type of the base pointer. Sometimes this observation can be used to help distinguish different classes of base values without producing relocated pointers with incompatible types. Consider the following declarations:

```
baseA: BaseA;
baseB: BaseB;
OffsetA: TYPE = BaseA RELATIVE POINTER TO T;
OffsetB: TYPE = BaseB RELATIVE POINTER TO T;
offsetA: OffsetA;
offsetB: OffsetB .
```

If *BaseA* and *BaseB* are distinct types (see the preceding point), so are *OffsetA* and *OffsetB*. Expressions such as *baseA [offsetB]* and *offsetA ← offsetB* are then errors, but *baseA [offsetA]* and *baseB [offsetB]* have the same type (*T*).

The base type must have the attribute **BASE**. Conversely, the attribute **BASE** always takes precedence in the interpretation of brackets following a pointer expression. Consider the following declarations:

```
p: POINTER TO ARRAY IndexType OF . . .;
q: BASE POINTER TO ARRAY IndexType OF . . . .
```

The expression *p[e]* will cause implicit dereferencing of *p* and is equivalent to *p ↑ [e]*. On the other hand, *q[e]* is taken to specify relocation of a pointer, even if the type of *e* is *IndexType* and not an appropriate relative pointer type. In such cases, the array must (and always can) be accessed by adding sufficient qualification, e.g., *q ↑ [e]*; nevertheless, users should exercise caution in using pointers to arrays as base pointers.

Mesa currently supplies no special mechanisms for constructing relative pointers. It is expected that such values will be created by user-supplied allocators that pass their results through a **LOOPHOLE** or from pointer arithmetic involving **LOOPHOLES**. **FIRST** and **LAST** may also be used to create relative pointers in certain cases:

```
Base: TYPE = LONG BASE POINTER TO RECORD[UNSPECIFIED];
IntPtr: TYPE = POINTER [0..200) TO INTEGER;
Node: TYPE = Base RELATIVE IntPtr;
nullNode: Node = FIRST[Node];                -- Never allocate this value
```

Note that relative pointers don't have a built-in NIL like other pointers. It is up to the programmer to create his own null value like that above. Another popular null value is `LAST[Node]`. One can only use `FIRST` or `LAST` if the pointer has a subrange specification. A subrange of `[0..CARDINAL.LAST]` can be used if the entire word is to be used for the relative pointer.

For more information about base and relative pointers, read about Pilot memory management and the system supplied interface *Zone* in the *Pilot Programmers Manual*.

6.3.4 Relative array descriptors

A `RelativeTC` constructs a relative array descriptor type whenever the `TypeIdentifier` evaluates to a pointer type and the `TypeSpecification` evaluates to an array descriptor type. Let a `RelativeTC` be

`TypeIdentifier` **RELATIVE** `TypeSpecification`,

where `TypeIdentifier` is of type

`[LONG]` **BASE POINTER** `[SubRangeb]` **TO** `[READONLY]` `Tb`,

and `TypeSpecification` is of type

`[LONG]` **DESCRIPTOR FOR** `[READONLY]` **ARRAY** `Ti` **OF** `Tc`,

and the brackets indicate optional attributes. Relative array descriptor values must be relocated before they are indexed. The relocation yields an expression with type

ARRAY `Ti` **OF** `Tc`.

Relative array descriptor types are entirely analogous to relative pointer types; indeed, values of such types can be viewed as array descriptors in which the base components are relative pointers. If the `TypeSpecification` says `READONLY`, the relocated array (or its elements) cannot be a `LeftSide`.

In the constructor of a relative array descriptor type, the `TypeSpecification` must evaluate to a (possibly long) array descriptor type.

In the notation introduced above, a reference to an element of the described array has the form

`base[offset][i]`

where `i` is the index of the element.

Currently, relative array descriptor values must be constructed using `LOOPHOLES`.

6.4 Variant records

Section 3.3 discussed "ordinary" record types, where every record object of a single type has the same number and types of components. Such records are not always adequate for programming applications. For example, in the symbol table for a compiler, all the records could have certain components in common: some standard linkage, a string representing the symbol, and a category field indicating whether the symbol stands for an operator,

constant, variable, label, etc. Different categories of symbols would then need further components that were not the same in all the records.

Variant records are designed for such applications: a variant record consists of an optional *common part* followed by a *variant part*. The common part contains components that are common to all records of this type. The variant part contains the components of each variant of the record.

The specification of a variant record type has the appearance of an ordinary record specification: `RECORD [field list]`. If the record has any common components, these are specified first; then the variant part is specified. Subsection 6.4.1 discusses variant record declaration more completely.

The variant part really represents a set of alternative *extensions* to the common part. The record type as a whole can be viewed as follows:

Common Part	Variant Part
Field list for the common part ----	---- field list for variant 1
	---- field list for variant 2
	---- ...
	---- field list for variant n

Each individual variant is identified by one or more *adjectives*. Suppose record type *ClassRec* is declared to have a set of variants named *class1*, *class2*, and *class3*. Then variables could be declared as follows:

```

someClass: ClassRec;           -- sometimes one class, sometimes another

firstClass: class1 ClassRec;    -- strictly a class1 ClassRec
secondClass: class2 ClassRec;  -- strictly a class2 ClassRec
thirdClass: class3 ClassRec;   -- strictly a class3 ClassRec

```

Types like *class3 ClassRec* are *bound variant types*. *ClassRec* and *class3 ClassRec* are both type specifications, but the latter is bound to a particular variant. A variable which is declared as a bound variant contains a definite variant; these components can be accessed as if they were common components.

The field list for any variant may itself have a variant part (and a variant in that part may have its own variant part, etc.). It is possible to have a type like *small class3 ClassRec* (i.e., the field list for the *class3* variant has a variant part which, in turn, has a *small* variant).

The record, *someClass*, presents a problem. During the course of execution, *someClass* might contain a *class1*, *class2*, or *class3* variant record. (Mesa allocates enough storage to hold the largest variant specified for *ClassRec* type records.) The problem is determining which variant applies at a given time.

To decide which kind of variant a record object contains, some form of tag is needed. This tag is normally specified as part of the record, in which case every such record object will contain an "actual tag" denoting the variant it represents. Instead of storing a simple tag, it may be possible to "compute" the tag value whenever it is needed (possibly by inspecting some values in the common part). Such *computed tags* are much less safe than explicit

ones. For instance, you could refer incorrectly to a "class2" component of *someClass* when it held a *class1* variant record. The result would be undefined.

It is possible to construct an entire variant for the variant part (§ 6.4.3) by qualifying a constructor (for that variant) with the variant's name (an adjective, in other words). Suppose for example that *ClassRec* has common components *c1* and *c2* followed by a variant part named *vp*, and that the *class1* variant has components *x* and *y*. Then the record constructor below constructs an entire *class1* variant:

```
ClassRec[c1: val1, c2: val2, vp: class1 [x: val3, y: val4] ]
```

Components of an unbound variant can be accessed using the record's tag value (whether actual or computed). A variation of **SELECT** beginning with the keyword **WITH** is used for this purpose (§ 6.4.4). An example follows (given that *ClassRec* has a computed tag):

```
WITH someClass SELECT currentTag FROM
  class1    => Stmt-1;      -- someClass is a bound class1 variant here
  class2    => Stmt-2;      -- someClass is a bound class2 variant here
  class3    => Stmt-3;      -- someClass is a bound class3 variant here
ENDCASE;
```

6.4.1 Declaring variant records

Variant records, like ordinary records, are usually declared in two steps:

```
identifier : TYPE = RecordTC ;           -- define record type
...
IdList : TypeIdentifier Initialization ;   -- declare the records
```

Initialization for variant records (§ 6.4.3) is similar to that for ordinary records. The (now complete) definition of **RecordTC** follows. It extends the partial definition given in subsection 3.3.7 and includes machine-dependent record types:

```
RecordTC           ::= MachineDependent RECORD [ VariantFieldList ]

MachineDependent   ::= empty | MACHINE DEPENDENT

VariantFieldList   ::= CommonPart FieldId : Access VariantPart |
                    VariantPart |
                    NamedFieldList |
                    UnnamedFieldList |

CommonPart         ::= empty | NamedFieldList ,

VariantPart        ::= SELECT Tag FROM VariantList ENDCASE

Access             ::= empty | PUBLIC | PRIVATE    -- see section 7.5.

Tag               ::= FieldId : Access TagType |
                    COMPUTED TagType |
                    OVERLAID TagType
```

```

TagType          ::= TypeSpecification | *

VariantList      ::= empty | Variant |
                  Variant VariantList

Variant          ::= FieldIdList => [ VariantFieldList ], |
                  FieldIdList => NULL ,

```

The **TypeSpecification** in **TagType** must be equivalent to some enumeration or enumerated subrange type. If the **CommonPart** is not empty, it must be a **NamedFieldList**. If there is no **CommonPart**, the **VariantPart** itself need not be named.

The following example shows many of the possible variations resulting from the above syntax definitions. It is unnecessarily complex for the application, but does show a number of features. It would be worthwhile to parse the declaration yourself using the definitions given above. The example might be used to describe the various "accounts" in a bank; there would be a table of such entries, one per account.

```

Service: TYPE = {savings, checking, depositBox};
Account: TYPE = RECORD
[
  number: CARDINAL,
  specifics: SELECT type: Service FROM
    savings    => [term: [30..365], intRate: PerCent, balance: Money],
    checking   =>
      [
        balance: Money,
        monthlyFee: SELECT COMPUTED {free, notfree} FROM
          notfree => [monthlyFee: Money],
          free    => [],
        ENDCASE
      ],
    depositBox => [fee: Money, dueDate: Date, paid: BOOLEAN],
  ENDCASE
  -- no variant can be attached to the ENDCASE
];

```

Each arm of a **VariantPart** specifies a single variant. An arm may be empty (as in the case of a *free checking Account*) if that variant needs no components of its own. *Although the syntax above states that all the arms must end with a comma, the one before the ENDCASE is, in fact, optional.*

Fine point:

In the declaration of a variant record, the form **NULL** may be used instead of **[]**.

The adjectives are identifier constants from some enumeration. Their type can be given explicitly, or implicitly as an enumeration whose members are the adjectives used in the variant part. In any case, the enumerated type is the "tag's" type for a variant part. There are three possible forms for the tag, and they represent:

an actual tag with an explicit enumerated type (e.g., *type* in *Account*),
 an actual tag implicitly defined (e.g., *easyTag* in *NoCommon* below), or
 a computed tag (e.g., the *monthlyFee* for a *checking Account*).

If an actual tag is used, it is allocated in the common part of the record and may be accessed and used like any other common component, *but it may not appear as a LeftSide, since that would compromise the type-safeness of such variant records.* The only way an actual tag with an explicit enumerated type can be set is with a record constructor which constructs the entire variant part (§ 6.4.3). Not all possible values from the tag's enumeration type have to be used as adjectives preceding the "*= >*" in the **Variant** declaration, some values may be omitted. However, if one uses an "omitted" value as a qualifier for a variant record constructor, a compilation error results (since, of course, there is no variant for that value).

An asterisk, "****", is used to indicate that the type of an actual tag is being defined implicitly by the set of adjectives naming the variants in that tag's variant part. For example, consider the record declaration below:

```
NoCommon: TYPE = RECORD
  [ -- no common part
    variantPart: SELECT easyTag: * FROM
    i   => [compl: INTEGER],
    j, k => [x, compl: STRING],
    ENDCASE];
```

The implicit type of *easyTag* is *{i, j, k}*. Note: you can't declare variables of the same type as *easyTag*. Thus, declaring a tag with an explicit enumerated type is often the preferred method.

Computed tags are always unnamed. In fact, they are not really tags at all: when one needs to know which variant a record with a computed tag contains, some *computation* must be done. Exactly how the variant "tag" is computed is strictly up to the program using it. For instance, to determine whether a *checking Account* was *free* or not, the program might look at some property of the *Account number* (such as whether it was odd or even).

An **OVERLAID** tag is a special case of a computed tag. The differences occur in the ways in which fields of the record are accessed (§ 6.4.4).

Fine point:

Special care must be exercised when declaring a **MACHINE DEPENDENT** variant record. Recall that **MACHINE DEPENDENT** records can contain no "holes" between fields. For variant records, this leads to the following rules: If the minimum amount of storage required for each variant is a word or less, each variant must be "padded" to occupy the same number of bits as the longest. Otherwise, each variant must occupy an integral number of words.

6.4.2 Bound variant types

The declaration of a variant record specifies a type, as usual. This is the type of the whole record. The variant record type itself defines some other types: one for each variant in the record. Consider the following example:

```

StreamType: TYPE = {disk, display, keyboard};
StreamHandle: TYPE = POINTER TO Stream;
Stream: TYPE = RECORD [
  Get: PROCEDURE [StreamHandle] RETURNS [Item],
  Put: PROCEDURE [StreamHandle, Item],
  body: SELECT type: StreamType FROM
    disk = > [
      file: FilePointer,
      position: Position,
      SetPosition: PROCEDURE [POINTER TO disk Stream, Position],
      buffer: SELECT size: * FROM
        short = > [b: ShortArray],
        long = > [b: LongArray],
      ENDCASE],

    display = > [
      first: DisplayControlBlock,
      last: DisplayControlBlock,
      height: ScreenPosition,
      nLines: [0..100] ],
    keyboard = > [],
  ENDCASE];

```

The record type has three main variants: *disk*, *display*, and *keyboard*. Furthermore, the *disk* variant has two variants of its own: *short* and *long*. The total number of type variations is therefore six, and they are used in the following declarations:

```

r: Stream;
rDisk: disk Stream;
rDisplay: display Stream;
rKeyb: keyboard Stream;
rShort: short disk Stream;
rLong: long disk Stream;

```

The last five types are called *bound variant types*. The rightmost name must be the type identifier for a variant record. The other names are **Adjectives** modifying the type identified to their right. Thus, *disk* modifies the type *Stream* and identifies a new type. Further, *short* modifies the type *disk Stream* and identifies still another type. Names must occur in order and may not be skipped. For instance, *short Stream* would be wrong since *short* does not identify a *Stream* variant.

A *ragged variant record* is a variant record type whose arms do not all occupy the same number of words. For example

```

Variant: TYPE = RECORD [SELECT tag: * FROM
    short = > [c: CARDINAL],
    long = > [lc: LONG CARDINAL],
ENDCASE];

```

It is illegal to use the "=" or "#" operation to compare two variables whose types are unbound ragged variant records. The reason is that such a comparison would produce an incorrect result if garbage bits at the ends of the records were different, but all other bits were identical. Indeed, if a small variant is allocated to exact size, the trailing bits aren't really there, and the comparison might attempt to read from an illegal address.

However, it is legal to use "=" or "#" to compare two bound variants, or to compare a bound variant with an unbound one.

Furthermore, it is illegal to use "=" or "#" on any RECORD or ARRAY type which directly or indirectly contains an unbound ragged variant record.

The formal definition of **TypIdentifier** can now be completed (it was only partially defined in chapter 3):

```

TypIdentifier    ::= ... | Adjective TypIdentifier |
                    TypIdentifier . Adjective
Adjective      ::= identifier

```

where **Adjective** is an adjective of the variant part in the type specified by **TypIdentifier**. Note that the recursive use of **TypIdentifier** in the first line allows a sequence of adjectives. Sub-subsection 6.4.4.1 discusses the use of dot notation with variant records.

6.4.3 Accessing entire variant parts, and variant constructors

Mesa does not allow an entire variant part to occur on the right of an assignment. The only way to assign to an entire variant part is via a constructor, not by copying the variant part of an already initialized record.

This section considers accesses to entire variant records (e.g., for initialization) and the variant part of the record as a whole. The next section covers accesses to individual components in a variant part.

The actual tag, *type*, in the *body* variant part may be accessed by qualification:

```

IF r.type = keyboard THEN Stmt-1;

```

It is also possible to construct values of a variant record type. The syntax of a constructor for a variant part is the same as that of a normal constructor except that the identifier preceding the "[" must be present and must be one of the adjectives used in defining the variant. For example, some of the following declarations use constructors to initialize the variables:

```

myDisplay: display Stream ← [myGet, myPut, display[d1,d,h,8] ];
yourDisplay: display Stream ← myDisplay;

```

```
currentStream: Stream ← myDisplay;
s: Stream ← [SysGet, SysPut, disk[fp, 0, SysSetPos, long[al ] ] ];
```

The *keyboard* variant of *Stream* is a **NULL** variant; so there are no components for that variant in a *keyboard* constructor:

```
rKeyb ← Stream[Get: Kget, Put: Kput, body: keyboard[ ] ];
```

A side effect of assigning a bound variant value to a variable is that the actual tag of the record is also changed. *This is the only way to change the variant contained in a variable (except in the case of a **COMPUTED** tag).* This restriction ensures type-safeness. For example, the following assignment changes the *type* tag for *r*:

```
r.body ← keyboard[ ];
```

If one is assigning a completely bound variant value, *bv*, say (which could be a constructor, of course) in an **AssignmentExpr** (§ 2.5.4.), then the type of the **AssignmentExpr** is the type of *bv*, not the type of the **LeftSide**, which might not be a bound variant.

6.4.4 Accessing components of variants

When a record is a bound variant, the components of its variant part may be accessed as if they were common components. For example, the following assignments are legal:

```
rDisplay.last ← rDisplay.first;
rDisk.position ← rShort.position;
```

Components of an unbound variant may not be accessed in this way. If a record is not a bound variant (e.g., *r* in the previous section), the program needs a way to decide which variant it is before accessing variant components. More importantly, however, this must be type-safe. For this reason, the process of discriminating among possible variants and then accessing within a variant part is combined in one syntactic form, called a *discrimination*, which is a generalization of **SELECT**.

A *discrimination* closely mirrors the form of **SELECT** used to *declare* a variant part. However, the arms in a discriminating **SELECT** contain statements or **Expressions**, and, within a given arm, the discriminated record value *is viewed as a bound variant*. Therefore, within that arm, its variant components may be accessed. The syntax equations for variant record declaration follow:

```
SelectStmt      ::= ... | Select Variant
SelectVariant   ::= WITH OpenItem SELECT TagItem FROM
                  ChoiceSeries
                  ENDCASE FinalStmtChoice |
                  ...
ChoiceSeries     ::= empty |
                  AdjectiveList = > Statement |
                  AdjectiveList = > Statement ; ChoiceSeries
TagItem         ::= empty |                               -- the actual tag is used
                  Expression                             -- compute the tag value
```

```

FinalStmtChoice ::= empty | => Statement
SelectExprVariant ::= WITH OpenItem SELECT TagItem FROM
                    ChoiceList
                    ENDCASE => Expression |
                    ...
ChoiceList ::= empty |
            AdjectiveList => Expression |
            AdjectiveList => Expression , ChoiceList
OpenItem ::= Expression |
           AlternateName : Expression

```

The value discriminated is the one given in the **WITH** clause, which behaves just like an **OPEN** clause (§ 4.4.2) to simplify naming the record value in the arms of the **SELECT**. The following example discriminates on *r*:

```

WITH strm: r SELECT FROM
  display =>
    BEGIN
      strm.first←strm.last;
      strm.height←73;
      strm.nLines←4;
    END;
  disk => WITH strm SELECT FROM
    short => b[0]←10;
    long => b[0]←100;
  ENDCASE;
ENDCASE => strm.body ← disk[GetFp["Alpha"], 0, SysSetPos, short[ ]];

```

In this example, the tag implicitly used for the outer **WITH** clause is *type*.

First, suppose *r* contains a variant record of *display Stream* type. Then the first arm is chosen by this **SELECT**. Within it, *strm* (but not *r*) is considered a record of *display Stream* type; so all components of the *display* variant may be accessed in the statement chosen by that arm (as they are in the example).

Suppose *r* contains a variant record of *disk Stream* type. Then the actual tag has the value *disk*, and the second arm is chosen. In this example, only one of the *disk* components is accessed, its variant part. The inner **SELECT** uses variant record *strm*. Within the outer arm, Mesa knows that *strm* (but not *r*) is a record of *disk Stream* type by definition. Consequently, the tag implicitly used for this **SELECT** is the tag specified for the *disk Stream* type (namely, *size*). This is why *strm: r* was used in the **WITH** clause instead of just *r*. Use of *strm: r* ensures that within the *disk* arm, the identifier *strm* will have type *disk Stream* and be available for use by the inner **WITH** clause.

If the tag value of *size* is *short*, then the chosen arm accesses component *b* in the *short disk Stream* variant record; if it is *long*, then the chosen arm accesses component *b* in the *long disk Stream* variant record.

However, the **ENDCASE** for the inner **SELECT** could have accessed components that are common to a *disk Stream* (*file*, *position*, *SetPosition*, variant part *buffer*, and actual tag

size, plus all the original common components: *Get*, *Put*, variant part *body*, and actual tag *type*).

Suppose, lastly, that *r* does not contain a variant record of *display Stream* or *disk Stream* type: the outer **ENDCASE** statement will be chosen. This statement accesses the common component *body* (the *entire* variant part is considered a common component), and gives the record a specific variant type (*short disk Stream*) by wholesale assignment. An **ENDCASE** may only access common components; *it may not access components of variants in the given type*.

If the labels on an arm of a discrimination identify more than one variant structure, the record is not considered to be discriminated within that arm and only the common fields are accessible (cf. **ENDCASE**).

Since the outer variant part of *Stream* was declared using an actual tag, the tag's value is obtained from the record itself, and no **Expression** follows the keyword **SELECT**. (Both **SELECTS** above have this form.)

The **Expression** in the **WITH** clause's **OpenItem** must represent either a variant record or a pointer to a variant record (e.g., *r* in the above). The alternate name is essentially a synonym for that **Expression** (e.g., *strm* in the above). If it is a pointer, however, the alternate name designates a record value, *not* a pointer value in each arm of the **SELECT**. In the following example, the *display* arm is correct, and the *disk* arm is in error:

```
rp: StreamHandle;
proc: PROCEDURE [StreamHandle];
WITH sRec:  rp SELECT FROM
    display = > proc[@sRec];           -- correct
    disk    = > proc[sRec];           -- wrong
ENDCASE;
```

An open item with no alternative name opens a name scope so that components can be accessed with implicit qualification (as in the inner **SELECT** of the first example), *but then no further levels of WITH...SELECT using the same record can be done within such a WITH...SELECT*. In addition, there is no way to name the bound variant: only the alternative name can be used for that purpose. The type of the open item's **Expression** indicates the nature of the variant part, including whether the tag is an actual or computed tag, its enumerated type, and the names of each variant (i.e., the adjectives) in the variant part.

If a *computed tag* had been used, the program would have to supply an **Expression** following **SELECT** to determine the variant. This **Expression**'s value would have to be an adjective in the applicable variant part. For example, assume that *tbl[i]* in the following has type *checking Account* (§ 6.4.1); then this is a legal (if not very sophisticated) discrimination for it:

```
WITH this: tbl[i] SELECT (IF (this.number MOD 2) = 0 THEN free ELSE notfree) FROM
    free      = > NULL;
    notfree   = > AddToBill [this.monthlyFee];
ENDCASE;
```

The record value in a **WITH** clause must not represent a completely bound variant (which is really not a variant at all). For example, a valid discrimination for a *disk Stream* record, *aDiskStream*, follows:

```
WITH aDiskStream SELECT FROM
  short => b[0]←10;
  long  => b[0]←100;
ENDCASE;
```

It would be illegal to rewrite this as follows:

```
WITH alt: aDiskStream SELECT FROM           -- wrong!
  disk => WITH alt SELECT FROM
  short => b[0]←10;
  long  => b[0]←100;
ENDCASE;

ENDCASE;
```

An **OVERLAID** record is a special case of a computed variant record in that there is no explicit tag field in the record. The fields of the individual variants may be accessed using a "computed" **WITH** construct in the same manner as a **COMPUTED** record. In addition, any field name of a variant that is unambiguous (i.e., it appears in only one variant) can be referenced without discrimination. In essence, the programmer is telling the compiler "When I use a fieldname, you can trust me that the record has the proper variant." Consider the following example:

```
TrustMe: TYPE = RECORD[
  SELECT OVERLAID * FROM
    one => [c: CHARACTER, i: CARDINAL, next: POINTER TO TrustMe],
    two => [b: BOOLEAN, next: POINTER TO TrustMe],
    three => [s: STRING],
  ENDCASE];
t: TrustMe;
t.c           -- legal
t.b           -- legal
t.next       -- illegal, both variants one and two contain such a field
```

Fine point:

In the declaration of *TrustMe* above, the two *next* fields were of the same type, but occupied different positions within the record. Even if they did occupy the same position, one could still not refer to *t.next*. The ambiguity is one of variant, not of value.

6.4.4.1 Dot notation for discriminated variant types

If *V* is a type expression designating some variant record type with variant *a*, then *V.a* is a type expression designating the discriminated type. Thus forms such as

```
Stream.disk           Stream.disk.long           Stream.disk.long[80]
```

are equivalent to

*disk Stream**long disk Stream**long disk Stream*[80]

The dot form is preferred because it is similar to that of other expressions.

Mesa also allows *Stream*[*disk*] and *Stream*[*disk*][*long*], but use of this form is not encouraged.

Fine point:

Certain type attributes may also be obtained using dot notation. These include **SIZE**, **FIRST**, **LAST**, and **NIL**. There is a restriction, however. In order for *T.SIZE* to parse, *T* must be syntactically valid as an expression, not just a type expression, e.g., you can't say *disk Stream.SIZE*, or even (*disk Stream*).**SIZE**, but you can say *Stream.disk.SIZE* (or **SIZE** (*disk Stream*)).

6.4.4.2 Discriminating SELECT

The Mesa construct for discriminating a variant record, **WITH exp SELECT FROM ...**, has the semantics of opening **exp** (possibly dereferenced, if **exp** is a pointer) by name within the arms of the **SELECT**. This leads to inefficient code if **exp** is complicated, and to potential chaos if **exp** is changed within one of the arms. Mesa has a form of discrimination where the variant record (or pointer thereto) is copied before the selected statement is executed. Thus the form

```
WITH v SELECT FROM
  v1: T1 => S1;
  v2: T2 => S2;
  ...
  vn: Tn => Sn;
  ENDCASE => Se;
```

is equivalent to

```
u: T = v;
IF u # NIL AND ISTYPE [u, T1] THEN {v1: T1 ← NARROW [u]; S1}
ELSE IF u # NIL AND ISTYPE [u, T2] THEN {v2: T2 ← NARROW [u]; S2}
...
ELSE IF u # NIL AND ISTYPE [u, Tn] THEN {vn: Tn ← NARROW [u]; Sn}
ELSE Se;
```

where *T* is the type of *v*. The tests against **NIL** are omitted if *T* does not have a **NIL** value.

Note that this form always copies the discriminated value. Thus

```
r: LONG POINTER TO Stream;          -- where Stream has variants including disk
...
WITH r SELECT FROM
  x: LONG POINTER TO Stream.disk => {... x ...};
  -- x is a copy of r of type LONG POINTER TO Stream.disk
  ...
  ENDCASE;
```

WITH r ↑ SELECT FROM

```

x: Stream.disk => {... x ...};    -- x is a copy of r ↑ of type Stream.disk
...
ENDCASE;

```

Contrast these with the other form of variant record discrimination, which does not copy the discriminated value and reevaluates the discriminating expression each time that it is used within the selected statement:

```

WITH x: r SELECT FROM
  disk => {... x ...};           -- x is a synonym for r ↑, with type Stream.disk
...
ENDCASE;

```

6.4.5 Defaults and variant records

You may specify a default for the entire variant part in the declaration of a variant record type by writing a default specification after the `ENDCASE`. In the absence of such a specification, the default value of that part, including the tag, is undefined.

The default initial value of a discriminated variant record type has a tag value corresponding to the discriminating adjective, and defaults for the other fields of the variant part are those implied by the fields selected by that tag. In particular, the declaration or allocation of a variable with discriminated record type sets the tag correctly. For example,

```

VRec: TYPE = RECORD [
  common: INTEGER ← 0,
  variant: SELECT tag: * FROM
    red => [r1: BOOLEAN ← FALSE],
    green => [g1: INTEGER ← 0]
  ENDCASE ← red[TRUE] | NULL];    -- default for whole variant part

v: VRec;  -- initial value is [common: 0, variant: red[r1: TRUE]]
v1: VRec ← [common: 10];         -- initial value is [common: 10, variant: red[r1: TRUE]]
v2: VRec ← [variant: NULL];      -- tag and variant part are undefined, common = 0
v3: VRec ← NULL;                -- illegal (declaration of common does not allow NULL)
rv: red VRec;                   -- initial value is [common: 0, variant: red[r1: FALSE]]
gv: green VRec;                 -- initial value is [common: 0, variant: green[g1: 0]]

```

6.5 Sequences

A *sequence* in Mesa is an indexable collection of objects, all of which have the same type. In this respect, a sequence resembles an array; however, you cannot specify the length of the sequence when its type is declared, only when an instance of that type is created. Mesa provides sequence-containing types for applications in which the size of an array is not known at compile time. Note however, that only a subset of a more general design for sequences has been implemented. The contexts in which a sequence type may appear are somewhat restricted, as are the available operations on them. The subset provides enough functionality to accommodate most uses of sequences, but you may encounter some annoying restrictions.

A sequence type is considered to be a union of some number of array types, just as the variant part of a variant record type is a union of an enumerated collection of record types. This has the following consequences:

A sequence type can be used only to declare a field of a record. At most one such field may appear within a record, and it must occur last.

A sequence-containing record has a tag field that specifies the length of the sequence, and thus the set of valid indices for its elements.

To access the elements of a sequence, you use ordinary indexing operations; no discrimination is required.

Fine point:

In this sense, all sequences are overlaid, but simple bounds checking is sufficient to validate each access.

The use of sequence-containing variables is more restricted than is currently enforced for variant records. The length of a sequence is fixed when the object containing that sequence is created, and it cannot subsequently be changed. In addition, Mesa imposes the following restrictions on the uses of sequences:

You cannot embed a sequence-containing record within another data structure. You must allocate such records dynamically and reference them through pointers. The **NEW** operation is convenient for this purpose (§ 6.6.2).

You cannot derive a new type from a sequence-containing type by fixing the length; i.e., there is no analog of a discriminated variant record type.

There are no constructors for sequence-valued components of records, nor are such components initialized automatically.

The following sections describe sequences in more detail.

6.5.1 Defining sequence types

You may use sequence types only to declare fields of records. A record may have at most one such field, and that field must be declared as the final component of the record:

Syntax

```

VariantPart      ::= ... |
                  PackingOption SEQUENCE SeqTag OF TypeSpecification

SeqTag           ::= identifier : Access BoundsType |
                  COMPUTED BoundsType

BoundsType       ::= IndexType

TypeSpecification ::= ... |
                  TypIdentifier [ Expression ]

```

The **TypeSpecification** in **VariantPart** establishes the type of the sequence elements. The **BoundsType** appearing in the **SeqTag** determines the type of the indices used to select from those elements. It is also the type of a tag value that is associated with each particular sequence object to encode the length of that object. For any such object, all valid indices are smaller than the value of the tag. If T is the **BoundsType**, the sequence type is effectively a union of array types with the index types

$$T[\text{FIRST}[T] \dots \text{FIRST}[T]], T[\text{FIRST}[T] \dots \text{SUCC}[\text{FIRST}[T]]], \dots T[\text{FIRST}[T] \dots \text{LAST}[T]]$$

and a sequence with tag value v has index type $T[\text{FIRST}[T]..v]$. Note that the shortest sequence is empty, and the longest has $\text{LAST}[T] - \text{FIRST}[T] - 1$ elements. For example, if a sequence has a **BoundsType** of **CARDINAL**, it could have at most $\text{LAST}[\text{CARDINAL}] - 1$ elements.

If you use the first form of **SeqTag**, the value of the tag is stored with the sequence and is available for subscript checking. In the form using **COMPUTED**, no such value is stored, and no bounds checking is possible.

Examples:

```
StackRep: TYPE = RECORD [
    top: INTEGER ← - 1,
    item: SEQUENCE size: [0..LAST [INTEGER]] OF T]

Number: TYPE = RECORD [
    sign: {plus, minus},
    magnitude: SELECT kind: * FROM
        short = > [val: [0..1000)],
        long = > [val: LONG CARDINAL],
        extended = > [val: SEQUENCE length: CARDINAL OF CARDINAL]
    ENDCASE]

WordSeq: TYPE = RECORD [SEQUENCE COMPUTED CARDINAL OF WORD]
```

Fine point:

The final example is a recommended method for imposing an indexable structure on raw storage.

If S is a type containing a sequence field, and n is an expression with a type conforming to **CARDINAL**, both S and $S[n]$ are **TypeSpecifications**. They denote different types, however, and the valid uses of those types are different, as described below.

6.5.2 MACHINE DEPENDENT sequences

You may declare a field with a sequence type within a **MACHINE DEPENDENT** record. Such a field must come last, both in the declaration and in the layout of the record, and the total length of a record with a zero-element sequence part must be a multiple of the word length. If you explicitly specify bit positions, the size of the sequence field also must describe a zero-length sequence; that is, it must account for just the space occupied by any tag field.

Examples:

```

Node: TYPE = MACHINE DEPENDENT RECORD[
  info (0: 0..7): CHARACTER,
  sons (0: 8..15): SEQUENCE nSons (0: 8..15): [0..256] OF POINTER TO Node]

```

```

CharSeq: TYPE = MACHINE DEPENDENT RECORD [
  length (0): CARDINAL,
  char (1): PACKED SEQUENCE COMPUTED CARDINAL OF CHARACTER]

```

6.5.3 Allocating sequences

If S is a record type with a sequence as its final component, $S[n]$ is a type specification describing a record with a sequence part containing exactly n elements. The expression n must have a type conforming to **CARDINAL**. Its value need *not* be a compile-time constant; however, you can use specifications of this form only to allocate sequence-containing objects (as arguments of **NEW**) or to inquire about the size of such objects (as arguments of **SIZE**). In particular, you cannot use $S[n]$ to define or construct a new type or to declare a variable.

The value of the expression **SIZE**[$S[n]$] has type **CARDINAL** and is the number of words required to store an object of type S having n components in its sequence part.

The value of the expression z .**NEW** [$S[n]$] has type **POINTER TO S** (or **LONG POINTER TO S**, depending upon the type of the zone z , § 6.6.2). The effect of its evaluation is to allocate **SIZE**[$S[n]$] words of storage from the zone z and to initialize that storage as follows:

Any fields in the common part of the record receive their default values.

The sequence tag field receives the value **SUCC** ^{n} [**FIRST** [T]], where T is the type of the tag.

The elements of the sequence part have undefined values.

To supply initial values for the fields in the common part, you may use a constructor for type S in the call to **NEW**. There are currently no constructors for sequence parts, however, and you must void the corresponding field. In any case, you must explicitly program any required initialization of the elements of the sequence part. In Mesa, this is true even if the element type has non-NULL default value.

Examples:

```

ps: POINTER TO StackRep ← z.NEW[StackRep[100]];  -- s.top = -1
pn: POINTER TO Node ← z.NEW[Node[degree[c]] ← [info: c, sons: NULL]]
pxn: POINTER TO extended Number ← z.NEW[extended Number[2*k]]

```

Note that n specifies the maximum number of elements in the sequence part and must conform to **CARDINAL** no matter what **BoundsType** T_i appears in the **SeqTag**. The value assigned to the tag field is **SUCC** ^{n} [**FIRST** [T_i]]. A bounds fault occurs if this is not a valid value of type T_i , i.e., if $n \geq \text{LAST}[T] - \text{FIRST}[T] + 1$, and you have requested bounds checking.

If **FIRST** [T_i] = 0, **SUCC** ^{n} [**FIRST** [T_i]] is just n , i.e., the interpretation of the tag is most intuitive if T_i is a zero-origin subrange. Usually you will specify a **BoundsType** (e.g., **CARDINAL**) with

a range that comfortably exceeds the maximum expected sequence length. If, however, some maximum length N is important to you, you should consider using $[0..N]$ as the **BoundsType**; then the value of the tag field in a sequence of length n ($n \leq N$) is just n and the valid indices are in the interval $[0..n)$.

6.5.4 Operations on sequences

You can use a sequence-containing type S only as the argument of the type constructor **POINTER TO**. Note that the type of $z.\text{NEW}[S[n]]$ is **POINTER TO** S . The operations defined upon a sequence-containing type are

- ordinary access to fields in the common part
- readonly access to the tag field (if not **COMPUTED**)
- indexing of the sequence field
- constructing a descriptor for the components of the sequence field (if not **COMPUTED**).

There are no other operations upon either a sequence-containing record or the sequence type embedded within the record. In particular, you cannot assign or compare sequences or sequence-containing records (except by explicitly programming operations on the components).

You may use indexing to select elements of the sequence-containing field of a record by using ordinary subscript notation, e.g., $s.\text{seq}[i]$. The type of the indexing expression i must conform to the **BoundsType** appearing in the declaration of the sequence field and must be less than the value of the tag, as described above. The result designates a variable with the type of the sequence elements. A bounds fault occurs if the index is out of range and the sequence is not **COMPUTED** and you have requested bounds checking.

By convention, the indexing operation upon sequences extends to records containing sequence-valued fields. Thus you need not supply the field name in the indexing operation. Note that both indexing and field selection provide automatic dereferencing.

Examples:

```
ps ↑ .item[ps.top]  ps.item[ps.top]  ps[ps.top]  -- all equivalent
```

You may apply the **DESCRIPTOR** operator to the sequence field of a record; the result is a descriptor for the elements of that field. The resulting value has a descriptor type with index and component types and **PACKED** attribute equal to the corresponding attributes of the sequence type. By extension, **DESCRIPTOR** may be applied to a sequence-containing record to obtain a descriptor for the sequence part. The **DESCRIPTOR** operator automatically dereferences its argument (multiple times if necessary) until it finds something of a type to which it applies.

You cannot use the single-argument form of the **DESCRIPTOR** operator if the sequence is **COMPUTED**. The multiple-argument form remains available for constructing such descriptor values explicitly (and without type checking).

In any new programming, you should consider the following style recommendation: use sequence-containing types for allocation of arrays with dynamically computed size; use array descriptor types only for parameter passing.

Examples:

```
DESCRIPTOR[pn ↑]  DESCRIPTOR[pn.sons]      -- equivalent
```

6.5.5 StringBodies and TEXT

The type *StringBody* illustrates the intended properties and uses of sequences. For backward compatibility, it is not defined as a sequence. (The declarations of the types **STRING** and *StringBody* are given in section 6.1.)

The operations upon sequence-containing types have, however, been extended to *StringBody* so that its operational behavior is similar. In these extensions, the common part of the record consists of the field *length*, *maxlength* serves as the tag, and *text* is the collection of indexable components (packed characters). Thus `z.NEW [StringBody[n]]` creates a *StringBody* with *maxlength* = *n* and returns a **STRING**; if *s* is a **STRING**, *s[i]* is an indexing operation upon the text of *s*, `DESCRIPTOR[s]` creates a **DESCRIPTOR FOR PACKED ARRAY OF CHARACTER**, but whose length is the *maxlength* of the string rather than the more useful `DESCRIPTOR[BASE[s.text], s.length]`.

Fine point:

There are two anomalies arising from the actual declaration of *StringBody*: *s.text[i]* never uses bounds checking, and `DESCRIPTOR[s.text]` produces a descriptor for an array of length 0.

Type TEXT

The type **TEXT**, which describes a structure similar to a *StringBody* as a true sequence, is predeclared in Mesa. Its components *length* and *maxLength* are declared to have a type compatible with either signed or unsigned numbers (but with only half the range of **INTEGER** or **CARDINAL**).

```
TEXT: TYPE = MACHINE DEPENDENT RECORD [
    length (0): NATURAL ← 0,
    text (1): PACKED SEQUENCE maxLength (1): NATURAL OF CHARACTER]
```

6.6 Dynamic storage allocation

In Mesa, you can use *zones* to perform dynamic allocation and deallocation of variables. You are responsible for managing the storage and guarding against dangling pointers. Associated features handle certain routine aspects of allocation and deallocation (such as computing sizes), provide proper default initialization of newly allocated variables, and reduce the number of **LOOPHOLES** required to deal with an allocator.

6.6.1 Zones

A zone need not be associated with any specific storage area: it is just an object characterized by procedures for allocation and deallocation, as described below. The

storage managed by a zone is said to be *uncounted*. In a zone, object management is the responsibility of the programmer, who must explicitly program the deallocation.

To use a zone, you must have available procedures that manage the zone and implement the required set of operations. Most users will use a standard package such as those described in the *Pilot Programmer's Manual*.

A zone object has a value and a type. You will normally obtain a zone value from a package that implements zones. Typically, such a package constructs a zone (and perhaps an initial storage pool) according to user-supplied parameters.

Mesa provides two types for zones, **UNCOUNTED_ZONE** and *MDSZone*. Transactions with objects having these types are generally in terms of **LONG_POINTER** and **POINTER** values respectively.

Fine points:

Syntactically, **UNCOUNTED_ZONE** is a type constructor. *MDSZone* is a predeclared identifier; you may think of it as a synonym for **MDS_RELATIVE_UNCOUNTED_ZONE** (which you currently cannot write directly).

You may declare variables having zone types; fixed initialization is recommended. Zone types may also be used to construct other types.

6.6.2 Allocating storage

The operator **NEW** allocates new storage of a specified type, initializes it appropriately, and returns a pointer to that storage. The **NEW** operation is considered an attribute of a zone, which must be specified explicitly.

Syntax

```

Primary      ::= ... |
                Variable . NEW [ TypeSpecification Initialization OptCatch ] |
                ( Expression ) . NEW [ TypeSpecification Initialization OptCatch ]

Initialization ::= empty | ← InitExpr | = InitExpr

OptCatch     ::= empty | ! CatchSeries
  
```

The value of the **Variable** or **Expression** identifies the zone to be used, either directly or after an arbitrary number of dereferencing operations. The **TypeSpecification** determines the type of the allocated object. If an **InitExpr** is provided, it must conform to the specified type and its value is used to initialize the new object; otherwise, the default value associated with that type (if any) is used. Any signals raised or propagated by the allocation procedure will activate a **CatchSeries** attached to **NEW**. Signals raised during the evaluation of the **InitExpr** will not be caught, but you can enclose the statement in a block with an **ENABLE** (§ 8.2.1). The initialization is done before the value of the expression is available, so the value of the pointer is not available to the catch code. For example:

```

r ← z.NEW[T ← buggyinit];      -- r is not set until after buggyinit finishes
  
```

The value of the **Primary** is a pointer to the newly allocated object. The type of that pointer depends upon the type of the zone and the form of the **Initialization**. If the argument of **NEW** is some type T , the type of the result is

LONG POINTER TO T , if the type of the zone is equivalent to **UNCOUNTED_ZONE**
POINTER TO T , if the type of the zone is equivalent to **MDSZone**.

If you specify fixed (=) initialization, the result is a read-only pointer with type **LONG POINTER TO READONLY T** or **POINTER TO READONLY T** respectively.

The **InitExpr** cannot be the special form for string body initialization ([**Expression**]). You can, however, allocate string bodies with dynamically computed sizes as described in subsection 6.5.5. If you do so, the **Initialization** must be empty.

6.6.2.1 Allocating variant records

If the type expression T is a bound variant type (§ 6.4.2), then only enough space is allocated to hold that particular variant. In addition, the variant tag is initialized to the proper value (along with any default field values). For example, given the declaration of *Account* from subsection 6.4.1:

a : **LONG POINTER TO Account** ← **z.NEW[Account.savings]**;

allocates sufficient storage to hold the *savings* variant of *Account*, and sets the tag of $a \uparrow$ to be *savings*. There is some danger in this operation; since a is an undiscriminated pointer, Mesa will allow a program to assign a new value to $a \uparrow$. If a new value is a larger variant, the assignment will write beyond the end of the space allocated by the original **NEW**.

6.6.3 Releasing storage

Uncounted zones have a **FREE** operation; when applied to an object, **FREE** releases the storage allocated for that object.

Syntax

Statement ::= ... |
Variable . **FREE** [**Expression OptCatch**] |
(Expression) . **FREE** [**Expression OptCatch**]

The zone used in a **FREE** operation is determined as described for **NEW**; it must be the zone from which the variable was originally allocated. The argument of **FREE** is the *address* of a pointer to the variable to be deallocated; **FREE** sets the pointer to **NIL** and deallocates the storage for the variable *in that order*. This can lead to dereferencing **NIL** if the zone is a field of a record pointed to by the freed pointer.

$p.z$.**FREE**[$@p$]; -- will set p to **NIL** before evaluating $p.z$

Any signals raised or propagated by the deallocation procedure will activate a **CatchSeries** on a **FREE**.

It is the responsibility of a implementor of an **UNCOUNTED ZONE** to treat freeing of **NIL** as a no-op. Thus the following two statements are equivalent, but the one with the explicit test will be considerably faster if the probability of **NIL** is high.

```
IF p # NIL THEN z.FREE[@p];
z.FREE[@p];
```

6.6.4 Implementing uncounted zones

This section describes the assumptions currently made by the compiler about the user-supplied implementations of uncounted zones. These assumptions are compatible with the style of "object-oriented" programming that has proven successful in a number of applications. You need to read this section only if you are designing a storage management package.

An **UNCOUNTED ZONE** dealing with **LONG POINTER** values is represented by a two word value, which the compiler assumes to be a long pointer compatible with the following skeletal structure:

```
UncountedZoneRep: TYPE = LONG POINTER TO MACHINE DEPENDENT RECORD [
  procs (0:0..31): LONG POINTER TO MACHINE DEPENDENT RECORD [
    alloc (0): PROC [zone: UncountedZoneRep, size: CARDINAL] RETURNS [LONG POINTER],
    dealloc (2): PROC [zone: UncountedZoneRep, object: LONG POINTER]
    -- possibly followed by other fields-- ],
  data (2:0..31): LONG POINTER -- optional, see below
  -- possibly followed by other fields-- ];
```

If *z* is an **UNCOUNTED ZONE**, the code generated for $p \leftarrow z.\text{NEW}[T]$ is equivalent to

```
 $p \leftarrow z \uparrow .procs \uparrow .alloc[z, \text{SIZE}[T]]$ 
```

and the code generated by $z.\text{FREE}[@p]$ is equivalent to

```
{temp: LONG POINTER  $\leftarrow p$ ;  $p \leftarrow \text{NIL}$ ;  $z \uparrow .procs \uparrow .dealloc[z, temp]$ }.
```

Remember that it is the responsibility of *dealloc* to treat an input value of **NIL** as a no-op.

Within this framework, you may design a representation of zone objects appropriate for your storage manager. In general, you should create an instance of a *UncountedZoneRep* for each instance of a zone. The record designated by the *procs* pointer can be shared by all zones with the same implementation. The *data* pointer normally designates a particular zone and/or the state information characterizing that zone. Note that the compiler makes no assumptions about the designated object and does not generate any code referencing the *data* field. The extra level of indirection provided by that field is not obligatory; you may replace it with state information contained directly in the *UncountedZoneRep*.

The compiler assumes a similar (but single word) representation for an *MDSZone* value; the skeletal structure is as follows:

```
MDSZoneRep: TYPE = POINTER TO MACHINE DEPENDENT RECORD [
  procs (0:0..15): POINTER TO MACHINE DEPENDENT RECORD [
    alloc (0): PROC [zone: MDSZoneRep, size: CARDINAL] RETURNS [POINTER],
```



```
dealloc (2): PROC [zone: MDSZoneRep, object: POINTER]  
-- possibly followed by other fields-- ],  
data (1:0..15): POINTER -- optional  
-- possibly followed by other fields-- ];
```



Modules, programs, and configurations

In Mesa, large programs are constructed by linking or binding together individual *modules*. A module is the basic unit of compilation and also the smallest self-contained executable program unit. Most of this chapter deals with how modules are combined to build large systems.

There are two kinds of modules. *Definitions* modules serve primarily as “blueprints” or specifications for how the parts of a system will fit together. They provide a common set of definitions that can be referenced by other modules being compiled. The second kind of module, called a *program*, contains actual data and executable code. Program modules can be loaded and interconnected to form complete systems.

The Mesa compiler translates, or *compiles* a program module's source code (which is a text file) into an *object module*. An object module is a binary file containing object code, symbol table information, and data structures to be used in *binding* this module together with others. Compiling a definitions module produces symbol table information only, which may then be used in compiling other modules of either type.

7.1 Interfaces

An *interface* is a connection among programs: it allows code in one module to access parts of other modules, specifically procedures, signals (chapter 8) and variables. Interfaces are defined by *definitions* modules (§ 7.3), which contain declarations for public items and allow the compiler to do type checking across inter-module references. The interface, considered as a record, also provides the data structure for binding programs together.

The procedures, signals, and variables that implement a given abstraction are often collected in a single interface. For example, an interface for an allocator might consist of declarations of procedures for allocating and freeing blocks of storage, and pointers to shared blocks of storage. The data types and constants required by these procedures (for parameters and return values) are usually defined in the same definitions module. Such *non-interface* type definitions are available for use when compiling other modules, but are not considered part of the interface specified by that definitions module.

A program module containing references to variables, procedures, or signals defined by some definitions module must *import* the interface corresponding to the definitions

module. Actual *binding* occurs later when the compiled module is coupled with other compiled program modules (§ 7.7).

The actual implementation of an interface is usually provided by a single program module, although it may be realized by a group of modules, each supplying a part. In any case, if a program module implements all or part of the interface specified by a definitions module, it is said to *export* that interface. The compiler checks that the procedures and variables in the implementing module(s) are type-compatible with those in the exported interface (§ 5.2).

A program's object module contains a set of *interface records*, one for each *imported* interface, and a set of *export records*, one for each exported interface (a single program module can implement more than one interface). *Binding* a group of modules together into a system involves the association of imported interface records with exported interface records for all the modules in the group. Until this binding has been done, imported items are dangling references; an attempt to use one will result in an error at run-time.

The following definitions module, *IO*, provides a minimal (and unrealistic) interface to a computer terminal:

```

IO: DEFINITIONS =
BEGIN
-- Interface definitions
  ReadChar: PROCEDURE RETURNS [CHARACTER];
  ReadLine: PROCEDURE [input: STRING];    -- reads from terminal into input

  WriteChar: PROCEDURE [output: CHARACTER];
  WriteLine: PROCEDURE [output: STRING];

  IOImpl: PROGRAM;

-- Non-interface definitions
  CR: CHARACTER = 015C;                    -- an ASCII Carriage-Return character
END. -- IO

```

The interface record for *IO* is imported by the following *CopyDriver* program module. The program reads lines from the terminal and retypes them. When the user types a line beginning with a period, it writes a parting message and stops.

```

DIRECTORY
  IO;

CopyDriver: PROGRAM IMPORTS IO =
BEGIN
input: STRING = [256];                    -- 256-character string to hold input lines typed by
                                          -- user

-- the main body of the program starts here:

DO                                          -- infinite loop; only left by EXIT
  IO.ReadLine [input];                    -- read a line into input
  IF input [0] = '.' THEN EXIT;           -- quit if first character is a period
  IO.WriteLine [input];                   -- otherwise copy it back to the user
ENDLOOP;

```

```

IO.WriteLine ["End of example."];           -- final output
IO.WriteChar [IO.CR];                       -- leave terminal on a new line
END.                                         -- CopyDriver

```

The skeleton of a module that implements the *IO* interface follows. It **EXPORTS** *IO* and **IMPORTS** nothing.

```

DIRECTORY
  IO;

IOImpl: PROGRAM EXPORTS IO =
-- this module implements the procedures of the interface specified by IO.
BEGIN
terminalState: {off, on, hung} ← off;       -- initial state of the terminal
ReadChar: PUBLIC PROCEDURE RETURNS [CHARACTER] = BEGIN ... END;
ReadLine: PUBLIC PROCEDURE [input: STRING] = BEGIN ... END;

WriteChar: PUBLIC PROCEDURE [output: CHARACTER] = BEGIN ... END;
WriteLine: PUBLIC PROCEDURE [output: STRING] = BEGIN ... END;
END. -- IOImpl

```

After compiling the above modules, the next step is *binding* them together. Binding is the process of matching import records to export records.

A separate language, C/Mesa, is used to describe binding. This language has a syntax similar to Mesa's, but is much smaller. C/Mesa "programs" are processed by a program called the Binder. The C/Mesa source code is called a *configuration*; binding a configuration results in an object file, called a *binary configuration description* (BCD) file. An object file produced by the Mesa compiler is actually a very simple BCD file containing object code and binding information for just one module.

Object files can be loaded and run. (Actually, it is the individual modules in the BCD that are loaded). The loading process supplies actual memory addresses where required, then the modules in the BCD can be started (§ 7.9). The following configuration describes a system of three modules: *CopyDriver*, *IOImpl*, and *CopyControl*.

```

Copier: CONFIGURATION
  CONTROL CopyControl =
  BEGIN CopyDriver; IOImpl; CopyControl; END.

```

This configuration specifies how the *CopyDriver*, *IOImpl*, and *CopyControl* object modules are to be bound together; just listing the names is all that is usually required in a configuration. After binding, the loader can load the complete system using the BCD for *Copier*. *CopyControl* is named as the **CONTROL** module for the BCD, so starting the loaded object file would actually result in starting *CopyControl*, which follows:

```

DIRECTORY
  IO, CopyDriver;

CopyControl: PROGRAM IMPORTS CopyDriver, IO =
BEGIN
START IO.IOImpl;           -- so its variables (e.g., terminalState) are initialized
START CopyDriver;         -- to initialize its variables and run its main body code
END.

```

This example is simple, but *Copier* and *CopyControl* would still be simple even if the system had 50 modules instead of just two. For this example, they seem like excess baggage, but for a larger system, they are invaluable because:

- (a) they describe exactly how the various modules are bound together and initialized;
- (b) C/Mesa allows Mesa's compile-time checks on types to extend to binding time;
- (c) loading and linking with this scheme can be very efficient.

We can now give the details of Mesa **DEFINITIONS** and **PROGRAM** modules. Section 7.8 discusses C/Mesa and the binding process.

7.2 The fundamentals of Mesa modules

The complete syntax for a module is the following:

```

CompilationUnit ::= Directory
                    ModuleName : ModuleHead = GlobalAccess
                    ModuleBody

Access           ::= empty | PUBLIC | PRIVATE

Directory       ::= empty | DIRECTORY IncludeList ; | DIRECTORY ;

ExportsList    ::= empty | EXPORTS |                               -- see subsection 7.4.3
                    EXPORTS IdList

FileName       ::= stringLiteral                               -- see subsection 6.1.1

GlobalAccess   ::= Access                                       -- see subsection 7.4.3

ImportsList    ::= empty | IMPORTS |                               -- see subsection 7.4.1
                    IMPORTS InterfaceList

IncludeList    ::= IncludeItem |
                    IncludeList, IncludeItem

IncludeItem    ::= identifier UsingClause |
                    identifier:FROM FileName UsingClause |
                    identifier:TYPE UsingClause |
                    identifier:TYPE identifier UsingClause

ModuleBody     ::= Block .                                       -- section 4.4
                    -- note the terminating period

ModuleHead     ::= DEFINITIONS LocksClause ImportsList ShareList |
                    ProgramTC ImportsList ExportsList ShareList

InterfaceItem ::= identifier | identifier : identifier

InterfaceList ::= InterfaceItem | InterfaceList , InterfaceItem

LocksClause   ::= empty |                               -- see subsection 9.4.1
                    LOCKS Expression |
                    LOCKS Expression USING identifier : TypeSpecification

ModuleName    ::= identifier | IdList

```

```

ProgramTC      ::= PROGRAM ParameterList ReturnsClause |
                   MONITOR ParameterList ReturnsClause LocksClause
ShareList     ::= empty | SHARES IdList           -- see subsection 7.5.4
UsingClause   ::= empty | USING [IdList] | USING []

```

Fine point:

ImportsList (Exports List): The form **IMPORTS(EXPORTS)** may be used when the corresponding list is empty. This reflects the view that **IMPORTS** and **EXPORTS** introduce formal parameter lists, even though their punctuation omits [and].

A definitions module may perform two functions: it may define an interface, and it may contain declarations of constants and types. Definitions modules are further discussed in section 7.3.

The text of a program module *X* implicitly defines a *frame type*, **FRAME** [*X*]. Values of this type are created dynamically by loading *X* and can only be accessed indirectly: i.e., a program may have variables of type **LONG POINTER TO FRAME** [*X*], but never of type **FRAME** [*X*]. A program's frame contains storage for its variables, along with some system overhead. This frame is sometimes referred to as the module's *global frame*.

We will first deal with the initial syntactic unit which is common to all modules (the **Directory** clause), then with **DEFINITIONS** modules as a whole. After these sections there is a complete example including

- a **DEFINITIONS** module,
- a **PROGRAM** module that implements it,
- a client program that uses it, and
- a configuration that binds the programs together into a system.

7.2.1 Including modules: the **DIRECTORY** clause

A given module may *include* previously compiled modules for the following reasons:

It might need to use some of the symbols defined by those modules.

It may need to import the interfaces defined by those modules.

It might refer to instances of such modules in order to **START** them, to make new instances of them, or to access their data.

A module that includes another is called a *client* of that module. Suppose module *A* is included in module *B*. This means that when compiling *B*, the compiler must have access to *A*'s object file in order to obtain the information needed by *B*. In normal programming practice, only definitions modules are included in other modules. Only in rare circumstances will a program module be included in another module. *Including a module is not simply an insertion of text from one module into another—it is important to read these sections carefully to use this capability correctly.*

The following is a simple but complete **DEFINITIONS** module:

```
Simple: DEFINITIONS =
  BEGIN
    limit: INTEGER = 86;
    Range: TYPE = [-limit..limit];
    Pair: TYPE = RECORD[first, second: Range];
    PairPtr: TYPE = POINTER TO Pair;
  END.
```

Suppose that the above source code is contained in a file named "Simple.mesa." After compilation, anyone who has a copy of the object file for *Simple* (which will be named "Simple.bcd" by the compiler), may then include it in other modules. The ".bcd" portion of the file name stands for binary configuration description (§ 7.8 .3). The ".bcd" part of the name need not be specified in the directory section (see below).

A module that includes other modules begins with a **Directory** clause, which performs two functions:

- (1) It associates a Mesa identifier with the name of a file containing an object module (usually the identifier and the name of the object module are the same).
- (2) It checks that the given identifier matches the **ModuleName** of the module in the object file.

An **IncludeItem** entry in a directory clause defines the following things about the included module:

the local name by which the including program refers to it,

the external name of the file that contains its object module, and

the **ModuleName** that it must have.

The system provides convenient defaults for the common case in which all of these names are simply derivable from one another. The more complex forms may be used to handle unusual cases. The following table illustrates the various forms of an **IncludeItem** and the resulting names:

IncludeItem form	local name	file name	ModuleName
name	name	name.bcd	name
name: TYPE	name	name.bcd	name
name: TYPE moduleName	name	name.bcd	moduleName
name: FROM "fileName"	name	fileName.bcd	name
name: FROM "fileName.ext"	name	fileName.ext	name

The **name: TYPE moduleName** form can be used to include several different object files with the same **ModuleName**. In the **FROM** form, any characters may be used in the **fileName** field, not just Mesa identifier characters. The file name described here may be overridden when the compiler is invoked—see the *Mesa User's Guide* for details.

7.2.1.1 Enumerating items from an included module: the USING clause

A module may list the symbols it expects to access from an included module in the **USING** clause of an **IncludItem**. If a **USING** clause is present, it must list all of the symbols to be included; the compiler will not allow access to any symbol not in the list. Warnings will be issued for symbols appearing in the list that are not referenced in the module. The **USING** clause thus accurately documents the symbols that are defined in each included module.

Here is an example of a **DIRECTORY** with a **USING** clause:

```
DIRECTORY
    Simples USING [Range, Pair],
    String;
```

A module with this **DIRECTORY** statement would be allowed to use the symbols *Range* and *Pair* defined in *Simple*, but would not be allowed to use any other symbols defined there. Access to symbols defined in *String* is not restricted.

The **USING** clause only allows and restricts access to symbols. Actual references to the symbols must be made in one of the ways described below. The form **USING []** is permitted and is used to emphasize that the only use made of the module is to export items to it.

7.2.2 Accessing items from an included module

An identifier *p* defined in a definitions module *Defs* can be named in an including module *User* in one of two ways:

Qualification: *p* can be named as *Defs.p* in *User*.

OPEN clauses: In the scope of a clause of the form "**OPEN** *Defs*", the simple name *p* suffices.

The remainder of this section gives more detail on these methods.

7.2.2.1 Qualification

As explained above, an identifier in an included module may be named using qualification, as in *Simple.Pair*. *Simple.Pair* means "the item named *Pair* in *Simple*." As a rule, qualification provides more readable code than using **OPEN**. In the following example, qualification is the only access method used:

```
DIRECTORY
    Simple;
```

Table: DEFINITIONS =

```
    BEGIN
    limit: INTEGER = 256;           -- this has no connection with Simple.limit
    Index: TYPE = [0..limit];
    PairTable: TYPE = ARRAY Index OF Simple.Pair;
    StringTable: TYPE = ARRAY Index OF STRING;
    END.
```


A module that includes *Table* may use the symbols *Table* itself defines; however, to use *Pair*, the module would have to include *Simple* (as in the next section's example). Declared symbols in the included module do not include record component names: they are part of a record's type specification and can be used wherever the record type is known.

A qualified name may denote a type defined in an included module e.g., the type *Simple.Pair*. Thus, the syntax for **TypeIdentifier** includes the case

```
TypeIdentifier      ::= ... | identifier . identifier
```

7.2.2.2 OPEN clauses

The following program *TableUser* includes both *Simple* and *Table*. It accesses names from *Simple* by qualification, but uses an **OPEN** clause to access items from *Table*:

```
DIRECTORY
    Simple USING [Pair],
    Table;

TableUser: PROGRAM =
BEGIN OPEN Table;           -- (Notice the OPEN-clause.)
vIndex: INTEGER ← limit;    -- this is Table.limit because of the OPEN
vString: StringTable;
vPair: PairTable;
StoreString: PUBLIC PROCEDURE[s: STRING, v: Index] =
    BEGIN
        vString[v] ← s;
    END;
StorePair: PUBLIC PROCEDURE[t: Simple.Pair] RETURNS [ok: BOOLEAN] =
    BEGIN
        ok ← vIndex ≤ limit;
        IF ok THEN vPair[vIndex] ← t;
    END;
END.
```

In the scope of the **OPEN** clause, the names *limit*, *StringTable*, *PairTable*, and *Index* are those in *Table*. The scope of these **OPEN** clauses follows the same rules as the **OPEN** clauses for records described in subsection 4.4.2., including the ability to specify an alternative name to use in the scope of the **OPEN**. In fact, a single **OPEN** clause can contain **OpenItems** that open either modules or records.

Table could have been in an **OPEN** clause anywhere that one is permitted. This feature can be used to help the readers of a program. For example, if the names from *Table* were only needed in the procedure *StoreString*, we could put an "**OPEN Table**" on its **BEGIN** rather than on the **BEGIN** for the whole module. This would localize the region of the program where a reader would have to consider whether an identifier is from an included module or not. Note that qualification is still required to reference *Simple.Pair* even though *Pair* appears in the **USING** clause.

7.2.3 Scopes for identifiers in a module

The use of identifiers appearing in modules falls into two broad categories, *defining occurrences* (e.g., to the left of the ":" in a declaration), and *name references* (such as the appearance of a name in an **Expression**). *Scope rules* determine which defining occurrence goes with a given reference. In Mesa, these rules are *lexical*, i.e., they depend only on the textual structure of the module.

A *name scope* is a contiguous region of a module (e.g., everything between a **BEGIN . . . END** pair, or between a [. . .] pair) and may contain other name scopes nested within it. The parts of a name scope which are not also in some nested name scope are called a *simple name scope*, and they obey the following rule:

Within a simple name scope, there can be at most one defining occurrence of a given identifier.

An important corollary of this rule is that a given identifier is either undefined in a simple name scope or it has exactly one meaning.

New name scopes are created by the following:

OPEN clauses
 declarations within a **Block** or **Loop**
 enumerated types and their subrange types
 loop control variables
 record types that use named field lists
 procedure types that use named parameters or results
 actual procedures
 exit regions for loops and compound statements
 the heads and arms of discriminating **SELECT** statements

The meaning of a name reference, which is not itself qualified, is determined by the following rule:

Use the *innermost* enclosing name scope that defines the referenced identifier. (If none of the scopes do so, the identifier is undefined and there is an error.)

In a qualified name reference, the qualification supplies the name scope for the qualified identifier. For example,

<i>Simple.Pair</i>	-- (§ 7.2.2.1) qualification by module name; name scope is the definitions module <i>Simple</i>
<i>rDisk.Get</i>	-- (§ 6.4.2) record qualification; name scope is the type <i>disk Stream</i> , the type of <i>rDisk</i>
<i>winner.party</i>	-- (§ 3.4) pointer qualification; name scope is the type <i>Person</i> , the reference type of <i>winner</i>

The rule of scope is simple for a qualified reference:

The qualified identifier is associated with its definition in the specified scope. (If there is no such defined name, the qualified identifier is undefined and there is an error.)

OPEN clauses may introduce multiple name scopes, which are nested, outer to inner, in order from left to right. Consider the following revision of the earlier *TableUser* module:

```
DIRECTORY
    Simple,
    Table;
TableUser: PROGRAM =
BEGIN OPEN Simple, Table;
...
StorePair: PUBLIC PROCEDURE [t: Pair] RETURNS[ok: BOOLEAN] =
...
```

Notice that we no longer need qualification for the parameter *t* of procedure *StorePair*. The process by which the referent of *Pair* is found is the following: look for such a definition in the current module declarations; then try the next outer scope, which according to the **OPEN** was *Table*; then look in the next (and outermost) scope given by the **OPEN** (*Simple*) at which point a defining occurrence is found. Notice that a reader might be uncertain where *Pair* was obtained from, since he would have to perform this process of scanning name scopes mentally. *This is why the use of OPEN (without renaming) is considered bad programming practice.*

Localizing the scope of identifiers from included modules is so important that the following naming guidelines are recommended:

- (1) Place a **USING** clause on items in the **DIRECTORY**. This collects in one place a list of all symbols referenced from each included module. The list is always accurate because the compiler checks it on each compilation. (There is a utility program that will generate the **USING** clauses automatically.)
- (2) Use explicit qualification as the normal way of naming an external item.
- (3) It is acceptable to **OPEN** an interface or record over a limited scope (e.g., a procedure or a block) within which identifiers from that interface are used heavily. Records should only be opened using the renaming form, e.g., **OPEN** *rwln: recordWithLongName*. A renaming **OPEN** is also the recommended form for opening included modules. For example,

```
BEGIN OPEN DUD: DebugUsefulDefs;
DUD.LongCopyRead[ ... ];
...
DUD.LongCopyWrite[ ... ];
END
```

7.2.4 Implications of recompiling included modules

Consider a set of modules *Adefs*, *User1*, and *User2* where *Adefs* is included in *User1* and *User2*. For simplicity, assume *User1* and *User2* include only module *Adefs*. Suppose *Adefs*

and *User1* have already been compiled, but *Adefs* is *recompiled* for some reason before *User2* is compiled. Then *User1* must also be recompiled.

In general, recompiling *Adefs* will invalidate the current version of *User1*. This is obvious when *Adefs* undergoes significant change between compilations, but it may also be true when seemingly innocuous changes are made. In fact, if *User1* uses record or enumeration types defined by *Adefs*, the current version of *User1* is invalidated when *Adefs* is recompiled even if *no* changes are made to its source code!


For example, suppose *Adefs* defines **RECORD** type *Account* which is used by *User1* as the type of *r1* and by *User2* for *r2*. Normally, one would expect these records to have the same type. If events occur as follows, however, they will *not*:

Adefs is compiled.
User1 is compiled including (old) *Adefs*.
Adefs is recompiled.
User2 is compiled including (new) *Adefs*.

The record types for *r1* and *r2* will differ because of the way Mesa guarantees uniqueness for record types. The compiler associates a "time stamp" (e.g., time of compilation) with each record type. Old *Adefs* defined *Account* at one time and new *Adefs* defined it a later time; this makes them different (non-equivalent) record types which only "look" the same.

Consider the case where *Defs1* is included in *Defs2*, and *Defs2* is included in *User1*. (For simplicity, assume that *Defs2* includes only *Defs1* and *User1* only *Defs2*.) Suppose that *Defs1* is recompiled and then *Defs2* is recompiled. Then *User1* should also be recompiled. The reason for this is the uniqueness of record types defined in *Defs2* and used by *User1*.

The (re)compilations of *Defs1*, *Defs2*, and *User1* must occur in a specific order: *Defs1*, *Defs2*, and then *User1*. Suppose, however, that *User1* included *Defs2* plus another module *Defs3*, and that *Defs3* included *Defs1*. The diagram below illustrates these dependencies. Included modules are above the modules that include them. The rule for avoiding errors due to incorrect compilation order is that: a module may not be (re)compiled until all the modules above it have been.



Thus, *User1* should be recompiled after *Defs2* and *Defs3* have *both* been (re)compiled. The order in which *Defs2* and *Defs3* are compiled is unimportant, however. *Moral: There is an important partial order defined on modules by their inclusion relations.*

7.3 DEFINITIONS modules

Generally, a **DEFINITIONS** module contains a set of related definitions. There may be definitions for compile-time constants, types, procedures, signals, and other variables. These definitions are used both by the program(s) that implement those procedures and variables and by programs that only wish to use the items. Separating definitions from implementations allows programs that call those procedures to be independent of changes in implementation. The items in a **DEFINITIONS** module fall into two classes:

noninterface items: definitions of **TYPES** and compile-time constants, and

interface items: definitions of procedures, signals, programs, and other variables.

Interface items are variables which are “defined” in a **DEFINITIONS** module, but must be exported by some implementing program before they actually exist. Interface items are further divided into two categories:

transfer items: interface items of type **PROCEDURE**, **SIGNAL**, **ERROR**, or **PROGRAM**, and

nontransfer variables: interface items of all other types.

Transfer items are further divided into two cases:

transfer constants: transfer items that are constants at run time, and

transfer variables: transfer items whose value varies—either actually or potentially—at run-time.

Thus every item declared in a **DEFINITIONS** module is either a noninterface item, a nontransfer variable, a transfer constant, or a transfer variable. Transfer items are usually run time constants, and Mesa has special provisions for optimizing this important, common case.

Declarations of interface items in definitions modules are like declarations of variables in program modules, with two additional options.

```

Declaration      ::= IdList : Access ReadOnlyOption EntryOption
                    TypeSpecification Initialization |
                    IdList : VAR TypeSpecification | . . .

ReadOnlyOption   ::= empty | READONLY

```

The **VAR** option is used to distinguish between transfer constants and transfer variables. In a *program* module, a declaration such as

```
SampleProc: PROCEDURE [i: INTEGER];
```

declares a procedure variable. In a **DEFINITIONS** module, however, its effect is to *define* the name and type of a procedure transfer constant of the interface specified by that definitions module. Section 7.6 contains an example of a **DEFINITIONS** module that defines procedure interfaces. In the same manner, **SIGNALS**, **ERRORS**, and **PROGRAMS** can be declared

in a **DEFINITIONS** module as transfer constants of the interface. A program definition as an interface item is discussed in subsection 7.4.3.

The **VAR** option is used when you want a transfer item in an interface to be a variable at run time instead of a constant. To continue the above example, a declaration such as

```
SampleProc: VAR PROCEDURE [i: INTEGER];
```

creates a variable *SampleProc* that may be changed at run time. The **VAR** option is ignored if applied to nontransfer variables, since they are variables already.

The **READONLY** option is used to prevent importers from changing a variable. Additionally, it declares the item to be a variable—just like **VAR**. **READONLY** can be used only in an interface. This option is explained further in the next section.

In addition to providing common access to data structures, interface variables make default fields (§ 7.3.2) and inline procedures (§ 7.3.3) more useful in **DEFINITIONS** modules; that is, interface variables can be used within these definitions to access nonconstant information in the exporter. Interface variables used for this purpose must be **PUBLIC** in the exporter but can be declared with the attribute **PRIVATE** in the **DEFINITIONS** module. This practice is strongly recommended to prevent unintended direct sharing of the variables by clients of the interface. Note that code within an interface cannot update a **READONLY** component of that interface.

7.3.1 READONLY variables

The **READONLY** option can be attached to a variable only in an interface. If this option is specified, importers can read the variable but not update it; otherwise, importers are able to read and update the variable freely. Note that a **READONLY** variable is not necessarily constant; it can be changed from within a program module exporting the variable. Consider the following example of nontransfer interface variables:

```
Defs: DEFINITIONS =  
  BEGIN  
    var1: T;  
    var2: READONLY T;  
    . . .  
  END.
```

```
Impl: PROGRAM EXPORTS Defs =  
  BEGIN  
    var1: PUBLIC T;  
    var2: PUBLIC T;  
    . . .  
    var1 ← e1; . . . var2 ← e2;  
    . . .  
  END.
```

```

User: PROGRAM IMPORTS Defs =
  BEGIN
  . . .
  Defs.var1 ← e3;
  . . .
  IF Defs.var2 = e4 THEN . . .
  . . .
  END.

```

In this example, the exporter *Impl* provides storage for the nontransfer interface variables *var1* and *var2*. Within *Impl*, both are ordinary writable variables, e.g., *var2* can be updated. By importing *Defs*, the client *User* gains access to the storage for *var1* and *var2* but cannot change *var2*.

When the compiled versions of the above three modules are bound and loaded, the links for *var1* and *var2* in *User* will be set to the addresses of the storage for *var1* and *var2* provided by *Impl*. Although *Defs.var1* and *Defs.var2* are referenced indirectly through the pointers contained in the links, those pointers are invisible to the importer and are always automatically dereferenced. This is the way all imported nontransfer variables are handled. If there is no exporter for one of the variables, say *var1*, then the invisible pointer will be *NIL*, and any access to the variable will generate a runtime fault.

Imported nontransfer variables are reached via long pointers (rather than short pointers, as in earlier versions of Mesa). If *User* contained the expression *@Defs.var1* then this expression has type *LONG POINTER TO T* (§ 3.4.2). Therefore, within *User*

```
p: POINTER TO T ← @Defs.var1;
```

will not compile but

```
p: LONG POINTER TO T ← @Defs.var1;
```

will.

Also, observe that

```
p: LONG POINTER TO T ← @Defs.var2;
```

is illegal because *Defs.var2* is *READONLY*, but that

```
p: LONG POINTER TO READONLY T ← @Defs.var2;
```

is allowed.

7.3.2 Default fields in interfaces

One valuable use of interface variables is in default values for procedure argument records (§ 5.1 and 5.2). Default arguments can contain references to procedures (*INLINE* or otherwise), signals, and interface variables that are components of the same interface. These references are bound to values in the same instance of the imported interface as the one supplying the procedure definition itself. References to non-constant components of

other interfaces require that those interfaces be imported by the **DEFINITIONS** module and all its users (§ 7.4.4). For example, interface *Defs1* could contain:

```
globalQ: PRIVATE Queue;           -- an interface variable
Add: PROCEDURE [i: Item, q: Queue ← globalQ];
```

This declaration allows users to *Add* items to *globalQ*, but not to access the variable directly. Thus the statement

```
Defs1.Add [myItem];
```

is equivalent to

```
Defs1.Add [myItem, Defs1.globalQ];
```

Fine point:

If a program imports two instances of *Defs1*, say *D1* and *D2*, the defaulted value for *q* will be from the same instance as the called procedure. In other words,

```
D1.Add [myItem]      is equivalent to      D1.Add [myItem, D1.globalQ],
D2.Add [myItem]      is equivalent to      D2.Add [myItem, D2.globalQ],
```

In a program module, default arguments of exported procedures (and signals, etc.) require special attention. Because of the assignment rule, the **DefaultOptions** specified in the exporter and in the interface are *not* required to agree. If the implementor and the clients are to behave identically with respect to defaults, the same **DefaultSpecification** must appear twice.

For example, if the interface has the declaration

```
Proc: PROCEDURE [x: T ← e1];
```

and the exporter has the declaration

```
Proc: PUBLIC PROCEDURE [x: T ← e2] = BEGIN . . . END;
```

then within the exporter, *Proc[]* means *Proc[e2]*; within an importer of the interface, *Proc[]* means *Proc[e1]*. The language requires only that the types of *e1* and *e2* be compatible with *T*; if they should also provide the same value, the programmer must ensure this. To avoid this pitfall, it is recommended that exporters not declare such redundant defaults; to obtain the default within the exporter, import the interface and invoke as, e.g., *Defs.Proc[]*. This, however, will result in slightly less efficient code.

7.3.3 Inline procedures in interfaces

An **INLINE** procedure can be declared within a **DEFINITIONS** module. Any caller of that procedure must import an instance of the corresponding interface.

Within a **DEFINITIONS** module, the body of an **INLINE** procedure can contain references to procedures (**INLINE** or otherwise), signals and interface variables that are components of the same interface. These references are bound to values in the same instance of the imported interface as the one supplying the inline procedure itself. References to non-constant

components of other interfaces require that those interfaces be imported by the **DEFINITIONS** module and any program that uses that **INLINE** procedure (§ 7.4.4).

Interface components with the attribute **PRIVATE** are visible to the bodies of inline procedures declared within the same **DEFINITIONS** module. An inline procedure referencing such components can be imported into a **PROGRAM** module in which those components are not visible. Interface components used only as default arguments of imported procedures or as variables whose only occurrence is in the body of an **INLINE** procedure should not be mentioned in the corresponding **USING** clauses. For example, suppose that an interface contains the following declarations:

```

DomainFault: SIGNAL;
Proc: PROCEDURE [CARDINAL] RETURNS [T];

N: CARDINAL = 100;

Table: PRIVATE ARRAY [0 . . N) OF POINTER TO T;

IProc: PROCEDURE [i: CARDINAL] RETURNS [T] = INLINE
BEGIN
  IF i ~IN [0 . . N) THEN ERROR DomainFault;
  RETURN [IF Table[i] = NIL THEN Proc[i] ELSE Table[i] ↑]
END;
```

Note that the body of an **INLINE** procedure (*IProc*) can contain references to constants (*N*), interface procedures (*Proc*), signals (*DomainFault*) and interface variables (*Table*) in the same interface. If *Defsl* is an instance of this interface, *Defsl.IProc* references *Defsl.Table*, calls *Defsl.Proc*, etc. An importer cannot reference *Table* directly because of the **PRIVATE** attribute but can reference it indirectly through *IProc*.

It is not legal for a **PROGRAM** module *Prog1* to call an inline procedure defined in some other **PROGRAM** module *Prog2*, even if *Prog1* imports a **LONG POINTER TO FRAME** [*Prog2*].

7.3.4 Usage hints for **INLINE** procedures in interfaces *

Expansion of inline procedures can cause internal data structures of the compiler to grow rapidly; indiscriminate use of the **INLINE** attribute can substantially degrade compiler performance or cause tables to overflow. The current compiler has been organized so that inline expansion is particularly efficient, and incurs little added overhead, in the following circumstances:

The **INLINE** procedure, with an arbitrarily complex body, is defined within a **PROGRAM** module and called *exactly once* in that same module. Thus introducing named procedures for clarifying and structuring a program can be cheap when such procedures are called only once.

The **INLINE** procedure, defined either in a **DEFINITIONS** module or a **PROGRAM** module and called an arbitrary number of times, is very simple, with no local variables, no *named* output parameters, and no side effects.

The debugger cannot currently set breakpoints within, or display the expanded source text of, an **INLINE** procedure (although it can display the local variables resulting from the

expansion). Debugging can be easier if the `INLINE` attribute is used only as needed and is specified after initial testing has been successfully completed.

7.4 PROGRAM modules: IMPORTS and EXPORTS

A `PROGRAM` module may contain

- definitions of constants and types,
- declarations of variables,
- actual procedures and signals (chapter 8), and
- executable statements of its own (i.e., not part of procedure bodies within it).

After compilation, a program module contains a set of interface records, one for each imported interface. An interface record consists of a collection of links, one for each interface item referenced by the program module. Links provide the binding mechanism which allows access to other modules' procedures and variables.

At run time a module has a *global frame* that provides storage for its variables. Storage for links may be allocated either within the global frame or with the modules' code segment at the programmer's discretion. Binding and loading fill in the links with procedure descriptors, signal codes, pointers to program frames and pointers to exported variables.

7.4.1 IMPORTS, interface types and interface records

The `IMPORTS` list for a program declares the interface records that the program needs and associates them with `DEFINITIONS` modules (called *interface types*). *Interface records and interface types are different!* When using an interface type, a program may access only non-interface elements (the data types and constants required by the procedures, signals, and variables specified in the `DEFINITIONS` module), but when using an interface record a program can access both interface and non-interface elements.

The names of interface records and interface types are declared in a `PROGRAM` module's `IMPORTS` list:

-- Example of an explicitly renamed interface record

```
DIRECTORY Defs1;
```

```
Prog1: PROGRAM IMPORTS IntRecName: Defs1 = ...
```

The identifier preceding a ":" in an `IMPORTS` list names an interface record and the name following the ":" names an interface type. Omitting the name of an interface record in an `IMPORTS` list and giving only the name of an interface type means that the record's name is the same as the type's. For example, writing

-- Example of an implicitly named interface record

```
DIRECTORY Defs2;
```

```
Prog2: PROGRAM IMPORTS Defs2 = ...
```

is the same as writing

```
DIRECTORY Defs2;
```

```
Prog2: PROGRAM IMPORTS Defs2: Defs2 = ...
```

Within the body of *Prog2*, *Defs2* refers to an interface record, not an interface type. In fact, it is impossible to refer to the interface type *Defs2* in this case. Since interface records permit access to all items included in the named **DEFINITIONS** module, implicit naming of interface records is adequate unless you import multiple instances. In the body of *Prog1*, the reference *Defs1.x* is valid provided *x* is a non-interface element of *Defs1* (*IntRecName.y* can refer to either an interface or non-interface element *y* of *Defs1*).

Sometimes one needs to have access to more than one instance of an interface record at run time. For example, the Mesa compiler needs to access one instance of a symbol table package for the program that it is compiling, and at least one for the symbol tables for modules included by that program. This can be done by importing a number of distinct interface records for a single interface type, as in the following:

```
DIRECTORY SymDefs;
```

```
PartOfCompiler: PROGRAM IMPORTS mainSym: SymDefs, auxSym: SymDefs =
  BEGIN ... END.
```

Within the body of *PartOfCompiler*, one would access an interface element of *SymDefs* named *LookUp* for the main symbol table as *mainSym.LookUp*, and for the auxiliary symbol table as *auxSym.LookUp*.

Fine point:

An interface may have interface aliases; i.e., multiple identifiers (preceding "DEFINITIONS"). In the **DIRECTORY** clause of another module, the interface can be referenced using any one of the (equivalent) identifiers.

7.4.2 Importing program modules

Any module can include a program module *X* by naming *X* in its directory. One can use *X* to declare program *variables* of type **LONG POINTER TO FRAME [X]**. **FRAME [X]** is not a valid type because frames cannot be embedded in other structures.

A module can import a program *X* by naming it in its **IMPORTS** list. For example,

```
DIRECTORY XProg1, XProg2;
```

```
Prog: PROGRAM IMPORTS frame1: XProg1, XProg2 =
  BEGIN ... END.
```

This has an effect similar to declaring

```

frame1: LONG POINTER TO FRAME [XProg1] = ...;
XProg2: LONG POINTER TO FRAME [XProg2] = ...;  -- illustrates effect only, this is
                                                -- illegal as a Mesa statement

```

The declaration for *XProg2* is not a valid Mesa statement because `FRAME [XProg2]` here refers to the interface record defined in the imports list, not the interface type defined in the `DIRECTORY` clause.

Such imported program constants are the analogs of interface records. You can access variables with a frame pointer, as well as the types and compile-time constants of the module. Also, you can execute the mainline code (if any) of a module instance corresponding to a frame pointer using `START` and `RESTART` (§ 7.10.3) and create additional instances of it using `NEW` (§ 7.10.1). You cannot, however, access the transfer constants of a module to which you have only a `POINTER TO FRAME`, since the symbol table of the imported program module does not contain sufficient information to allow the compiler to generate the necessary procedure call.

Accessing values in a program frame as described above treats the frame as a record with its variables as its components. The price paid for such close coupling with a program is that the importer must be recompiled whenever the program is.

You cannot get `TYPES` from a `PROGRAM` by simply naming the program as an interface type: you can, though, get `TYPES` from a `PROGRAM`'s interface record (i.e., you can't get types if the program name simply appears in your `DIRECTORY` clause, but you can get types if you `IMPORT` the program). There is a trick to allow access to a `PROGRAM`'s types without actually importing an instance at run-time. The strategy

```

DIRECTORY XProg1;

Prog: PROGRAM =
  BEGIN
  ...
  XProg1.TypeName1;
  ...
  END.

```

will fail. The compiler is unwilling to access *TypeName1* since *XProg1* is an interface type. One may accomplish the desired result by declaring a "dummy" `LONG POINTER TO FRAME [ModuleName]` and using it to access the types. For example:

```

DIRECTORY XProg1;

Prog: PROGRAM =
  BEGIN
  ...
  XProg1Types: LONG POINTER TO FRAME [XProg1] = NIL;
  --no code is generated above, the pointer isn't even dereferenced

```

```

foo: XProg1Types.TypeName1;           -- valid declaration
...
END..

```

Note that if you are in fact importing module *ModuleName* (as is done in the first example of this section), then you already have **LONG POINTER TO FRAME** [*ModuleName*] whose name is *ModuleName* (either *frame1* or *XProg2* in the first example).

7.4.3 Exporting interfaces and program modules

The **EXPORTS** list for a program declares the interface records and interface types to which the program exports items. The name of an exported interface record is the same as the name of the interface type (which is the name of the **DEFINITIONS** module). A program exporting an interface is often called an *implementor* of that interface.

A variable or constant is exported from a program module to an interface if (1) the module exports the interface, (2) the item is declared **PUBLIC** in the program and, (3) the name of the item in the program matches the name of an item in the interface. The exported item must conform to the declaration of the item in the interface; the compiler will check this. Note that an item may be exported to several interfaces simultaneously.

As described in subsection 7.4.2, an exported **PROGRAM** is analogous to an exported interface. All items declared **PUBLIC** in a program module are exported to the module itself, as well as to any matching exported interfaces.

Exported transfer constants (§ 7.3) must be defined with fixed (“=”) initialization. For example, if an interface declared

```

ConstantProc: PROCEDURE [ . . . ] RETURNS [ . . . ]; -- a interface transfer constant

```

then a program module could initialize and export *ConstantProc* by

```

ConstantProc: PUBLIC PROCEDURE [ . . . ] RETURNS [ . . . ] = BEGIN . . . END;

```

Exported transfer and nontransfer variables (§ 7.3) should be defined with assignment (“←”) initialization. For example, if an interface declared

```

VariableProc: VAR PROCEDURE [ . . . ] RETURNS [ . . . ]; -- a transfer variable

```

then a program module could initialize and export *VariableProc* by

```

VariableProc: PUBLIC PROCEDURE [ . . . ] RETURNS [ . . . ] ← AnotherProc;

```

Fine points:

READONLY exported transfer and nontransfer variables may be defined with fixed initialization if they do not evaluate to compile time constants (including transfer constants). Fixed initialization of such items with compile time constants is not currently implemented.

If a transfer constant is assigned the value of another transfer constant (remember that “=” initialization is required), the assigned value must be a procedure whose body is declared within the current module; it cannot be a procedure constant imported from some interface.

It is important to note that assignment initialization is not performed until the exporting module is **STARTED**, and *reference to an interface variable does not cause a start trap in the exporter* (§ 7.9.4).

7.4.4 IMPORTS in DEFINITIONS modules and implicitly imported interfaces

Recall from section 7.2 that a definitions module can contain an **ImportsList**. One interface (*Defs2*) must import another (*Defs1*) if *Defs2* requires access to an interface item of *Defs1*, e.g., a procedure (**INLINE** or otherwise) or interface variable in *Defs1*. All interfaces mentioned in the **ImportsList** of a **DEFINITIONS** module must be unnamed.

An *implicitly imported interface* is one from which imported values are required for binding the variables of another explicitly imported interface.

An imported interface is a *principal instance* if it is the only instance of that interface imported by a module or if it is not renamed in the imports list.

Consider the example

```

Defs2: DEFINITIONS
  IMPORTS Defs1 =           -- an unnamed instance of Defs1
  BEGIN
  . . .
  Proc2: PROCEDURE [ . . . ] = INLINE
    BEGIN
    . . . IF Defs1.var # v THEN Defs1.Proc[ . . . ]; . . .
    END;
  . . .
  END.

Prog1: PROGRAM
  IMPORTS
    Defs1,                 -- an unnamed instance of Defs1
                           -- the principle instance of Defs1
    Defs2,                 -- an unnamed instance of Defs2
                           -- the principle instance of Defs2
    AnotherDefs1: Defs1    -- an named instance of Defs1
    YetAnotherDefs1: Defs1 = -- a second named instance of Defs1
  BEGIN
  . . .
  Defs2.Proc2[ . . . ]; -- expansion references Defs1.Proc, not AnotherDefs1.Proc
  . . .
  END.

Prog2: PROGRAM
  IMPORTS Defs2 =           -- an unnamed instance of Defs2
                           -- the principle instance of Defs2
                           -- Defs1 is implicitly imported, its
                           -- principle instance is created automatically

```

```

BEGIN
...
Defs2.Proc2[ . . . ];
...
END.

```

In the example above, *Defs1* is explicitly imported by *Prog1* and implicitly imported by *Prog2*. In Mesa, the variables of *Defs1* that are used by *Defs2* are bound to the principal instance of *Defs1* in a program module. Furthermore, if a program module imports no instances of *Def1*, a principal instance will be created automatically. If module *M* has no other reason to mention *Defs1*, then *Defs1* need not appear in either the **DIRECTORY** or the **IMPORTS** list of *M*. This is the case in *Prog2*. Explicitly importing a principal instance of *Defs1* in such a situation is not an error, and you must do so if

you plan to use positional notation to specify the imports of *M* in a C/Mesa configuration description, since the positions of automatically created interface instances are not defined, or

you already import more than one instance of *Defs1*, each of which is named.

In a C/Mesa configuration, principal instances of interfaces are *not* supplied automatically; you must import them explicitly if they cross (sub)configuration boundaries.

7.5 Access control: **PUBLIC** and **PRIVATE**

Every name defined in a module possesses an **Access** attribute, either **PUBLIC** or **PRIVATE** (the module in which a name is defined is called its *home module*). These are used to determine whether a name may be referenced when its home module is included by some other module. A **PUBLIC** name can always be used by an including module; a **PRIVATE** name cannot generally be used, except by modules which specify that they **SHARE** the included module. Sharing modules are called *privileged modules*; others are called *non-privileged modules*. A variable's home module is privileged.

Generally speaking, an **Access** may be specified

- (a) anywhere a name can be declared. This includes normal declarations, named field lists (for records or parameter lists), preceding **SELECT** in a record's variant part, and the declaration for an actual tag in a variant part.
- (b) preceding the **TypeSpecification** in a type definition.

In addition, an **Access** may be specified

- (c) at the beginning of a module (the **GlobalAccess**), to provide a default **Access** for any of the module's identifiers which do not have any explicit access.

The syntax in the following section is intended to supersede earlier definitions of the same constructs only by showing where attributes may be inserted. Otherwise, the earlier versions are correct. Each syntax definition is followed by examples of its use.

7.5.1 Access attributes in declarations

The following three subsections deal with the placement of **Access** options in declarations, field lists, and in variant records.

7.5.1.1 Declarations

The form of **Declaration** specifying an **Access** for its declared names is as follows:

```
Declaration ::= IdList : Access TypeSpecification Initialization ;
```

Examples:

```
q1, q2: PUBLIC INTEGER ← 0;
Mine: PRIVATE TYPE = {yes, no, maybe};
```

Mine can only be used in (i.e., seen from) privileged modules. In non-privileged modules, *mine* effectively does not exist.

7.5.1.2 Field lists

The forms for specifying **Access** in a **NamedFieldList** (§ 3.3.1) are as follows:

```
NamedFieldList ::= IdList : Access FieldDescription |
NamedFieldList , IdList : Access FieldDescription
```

```
FieldDescription ::= TypeSpecification |
TypeSpecification ← Expression
```

Example:

```
blk: PUBLIC RECORD [
  a: INTEGER,
  b: PRIVATE INTEGER ← 1234,
  c, d: BOOLEAN,
  e: PRIVATE BOOLEAN ];
```

A non-privileged module could only access components *a*, *c*, and *d* in this case, and then only using qualified references such as *blk.a*. Within a non-privileged module, extractors and constructors cannot be employed for a record type with any **PRIVATE** components.

7.5.1.3 Variant parts and tags in variant records

The forms for specifying **Access** in a **VariantFieldList** or **Tag** (§ 6.4.1) are as follows:

```
VariantFieldList ::= CommonPart identifier : Access VariantPart |
VariantPart |
NamedFieldList |
UnnamedFieldList
```



```

CommonPart      ::= empty |
                  NamedFieldList ,

VariantPart     ::= SELECT Tag FROM
                  VariantList      -- same as in subsection 6.4.1
                  ENDCASE

Tag             ::= identifier : Access TagType |
                  COMPUTED TagType |
                  OVERLAID TagType

TagType        ::= TypeSpecification | *

```

Example:

```

VarRec: PUBLIC TYPE = RECORD [
  link: LONG POINTER TO VarRec,           -- public common component
  vp1: SELECT tg1: PRIVATE Etype FROM     -- public variant, private tag
    adj1 => [
      youGet: ThisItem,
      iGet: PRIVATE SELECT tg2: * FROM    -- a private variant part
      ...
    ]
  ENDCASE ]
  adj2 => ...
ENDCASE ];

```

Suppose a non-privileged module has a record of type *VarRec*. Then it could access variant part *vp1* but neither tag *tg1* nor variant part *iGet*. This only prevents it from referring to *tg1* by qualification; it may still use a discriminating **SELECT** (which implicitly accesses *tg1*) for records of type *VarRec*. Thus, an *adj1* arm of such a **WITH . . . SELECT** could access component *youGet*. However, it would be unable to access component *iGet* in any case.

Notice that the only way that the tag of a variant can be changed is by writing a variant constructor (§ 6.4.3).

7.5.2 Access attributes in TYPE definitions

The form for defining a **TypeIdentifier** with an explicit **Access** is as follows:

```

Declaration      ::= . . . | IdList : TYPE = Access TypeSpecification ;

```

Example:

```

OurType: PUBLIC TYPE = PRIVATE RECORD [comp1: INTEGER, comp2: BOOLEAN];

```

A non-privileged module could declare records of type *OurType*, but it could not access the record components. The module could, however, pass values of type *OurType* as parameters, receive them as results from procedures, and use them as operands of a fundamental operation (\leftarrow , $=$, $\#$).

Note: Only *names* specified within the defined type are affected by this form of attribute specification. Consequently, it is intended for use only when defining record types and is

just a factorization: the **PRIVATE** could have been written after each inner colon; also, specific fields can be made accessible by writing **PUBLIC** internally, as shown below:

```
AlmostPrivateType: PUBLIC TYPE = PRIVATE RECORD [
  comp1: PUBLIC INTEGER,           -- overrides outer PRIVATE
  comp2: BOOLEAN];
```

An **Access** applied to a type specification can be specified as **PUBLIC**, but is pointless (if *OurType* is **PUBLIC** then the type of its components would be **PUBLIC** by default; if *OurType* is **PRIVATE** then their access attribute is irrelevant).

7.5.3 Default global access

As in subsection 7.5.1, if a declaration specifies an **Access** for a name, then that unilaterally determines its **Access**. If not, the given item receives a *default Access*. The default may be specified by the programmer in the **GlobalAccess** for a module; otherwise one is assumed. For a program module, the default is **PRIVATE**; for a **DEFINITIONS** module, it is **PUBLIC**. For example,

```
M1: PROGRAM = PUBLIC           -- specified GlobalAccess
  BEGIN
  ...
  END.

M3: PROGRAM =                 -- PRIVATE (by default)
  BEGIN
  ...
  END.
```

7.5.4 Accessing **PRIVATE** names of other modules *

A module may be made privileged to use **PRIVATE** items in an included module by using a **SHARES** clause. This clause lists the included module names whose **PRIVATE** symbols it needs to access. Consider the *Friend* module that follows:

```
DIRECTORY
  Special, Standard, Private;

Friend: PROGRAM SHARES Private, Special =
  BEGIN
  ...
  END.
```

In this case, *Friend* can use **PRIVATE** symbols defined by *Private* and *Special* but not the **PRIVATE** symbols of *Standard*. There is no particular significance to the ordering of module names listed after **SHARES**. Any kind of module may use **SHARES** (but it ought to be one that is "a friend," to say the least).

7.6 Exported (opaque) types

An *exported type* (often called an *opaque type*) is a **TYPE** that is declared in an interface and subsequently bound to some *concrete type* supplied by a module exporting that interface.

This is analogous to the treatment of constant procedures (and other transfer constants) in interfaces, where the implementations (i.e., the procedure bodies) do not appear in the interface but are defined separately. The advantages are twofold:

The internal structure of the type is guaranteed to be invisible to clients of the interface.

There are no compilation dependencies between the definition of the concrete type and the interface module. The definition of that type can be changed and/or recompiled at any time (perhaps subject to a size constraint; see below) without requiring recompilation of either the interface or any client of the interface.

The uses of an exported type are the same as those of any other type, e.g., to construct other types. The type provided by the interface is constant but has no accessible internal structure. There are two other important differences between exported procedures and exported types.

The first is a restriction necessary to ensure type safety across module boundaries. Different exporters of an interface can supply different implementations of a procedure in that interface if they are exported to different interface records. This is not true for exported types; *all exporters of a exported type must supply the same concrete type*, which is called the *standard implementation* of that exported type. Because of this restriction, clients can safely interassign values with exported type T , no matter how obtained. In addition, any exporter of T may convert a value of type T to a value of the concrete type it uses to represent T , and conversely.

The second difference is that it is not necessary to import an interface to access an exported type defined within it or to distinguish among values of such a type coming from different imported instances. This is another consequence of the fact that all interfaces must use a single implementation of the exported type.

7.6.1 Interface modules

An exported type is declared in a definitions module using one of the following two forms:

```
T: TYPE;
T: TYPE [Expression];
```

The first introduces a type T , no properties of which are known to the interface or to any client of the interface. In particular, the size of T is not known; this is adequate and desirable if the interface and clients deal only with values of type **LONG POINTER TO T** .

The second form specifies the size of the values used in the representation of the type. The value of **Expression**, which must be a compile-time constant with a positive integer value, gives this size in units of machine words. Supplying the size of an exported type is a shorthand for exporting a set of fundamental operations (creation, \leftarrow , $=$, and $\#$) upon that type. The eventual concrete type must be one with the *standard implementations* of these operations, which are defined as follows:

```
creation      allocate the specified number of words, with no initialization
```

←	copy an uninterpreted block of words
=, #	compare uninterpreted block of words

A type with non-NULL default value does not match the standard creation operation. Such a type cannot be exported with known size. Likewise, a variant record with arms that are not all the same length fails the comparison criterion. You should therefore consider writing your interfaces in terms of **LONG POINTER TO *T***, where *T* is a completely opaque exported type and not subject to these restrictions.

7.6.2 Client modules

An importing module or *client* has no knowledge of the type *T* beyond those properties specified in the interface. If the size is not specified, no operations on *T* are permitted. If the size is available from the interface, **SIZE [*T*]** is legal; also declaration of variables (including record fields and array components of type *T*) and the operations ←, =, # are defined for type *T*.

7.6.3 Implementation modules

An implementor exports a type *T* to some interface *Defs* by declaring the type with the **PUBLIC** attribute, and the value that is the concrete type; e.g., in

```
T: PUBLIC TYPE = S;
```

S specifies the concrete type. If the size of *T* appears in the interface, the definition of *T* in the exporter must specify a type with that size and the standard fundamental operations (the compiler checks this).

Within an exporter, *Defs.T* and *T* conform freely and are assignment compatible. Otherwise, *Defs.T* is treated opaquely there and is *not* equivalent to *T*. You should therefore attempt to write an exporting module entirely in terms of concrete types. Consider the following example:

Interface Module (*Defs*):

```
T: TYPE;
H: TYPE = LONG POINTER TO T;
R: TYPE = RECORD [f: H, ...];
Proc1: PROC [h: H];
Proc2: PROC [r: LONG POINTER TO R];
...
```

Exporting Module:

```
T: PUBLIC TYPE = RECORD [v: ...];
P: TYPE = LONG POINTER TO T;           -- to the concrete type
Proc1: PUBLIC PROC [h: P] = {... h.v ...};
Proc2: PUBLIC PROC [r: LONG POINTER TO Defs.R] = {
  q: P = r.f;                          -- convert to pointer to concrete type
  ... q.v ...};                          -- legal (indirect) reference to r.f.v
...
```

If the type of *h* were *Defs.H* in the implementation of *Proc1*, the reference to *h.v* would be illegal. By defining a type such as *P* and using it within the exporter instead of *H*, you can

avoid most such problems. Note that *Proc1* is still exported with the proper type. This strategy of creating concrete types in one-to-one correspondence to interface types involving *T* fails for record types such as *R* (because of the uniqueness rule for record constructors). In this example, you must use *Defs.R* to define the type of *r* in the implementation of *Proc2*, but a reference to *r.f.v* is illegal. In such cases, one may use redundant assignments, such as the one to *q*, or the **NARROW** operator as below.

Exporting Module:

```
... NARROW [ r.f, LONG POINTER TO T ].v ...
```

This narrowing of a (pointer to)* the opaque type into a (pointer to)* the concrete type is legal inside the exporting module, and generates no code. It is usually only necessary when dealing with opaque record fields defined in included definitions modules.

7.7 Dot notation and interface items

Mesa has a limited facility to allow an object oriented programming style—one in which the imported procedures that may be applied to a variable are deduced from context.

Consider the following example:

```
Defs: DEFINITIONS =
BEGIN
  T: TYPE = RECORD [ ... ];
  OpaqueT: TYPE;
  P: PROCEDURE [self:T, ... ]
  Q: PROCEDURE [self: OpaqueT, ... ];
  ...
END.

DIRECTORY
  Defs USING [P,Q, ... ],
  ...;
Prog1: IMPORTS Defs, AnotherDefs: Defs =
BEGIN
  x: LONG POINTER TO LONG POINTER TO T;
  y: OpaqueT;
  ...
  x.P[args];
  ...
  y.Q[otherArgs];
  ...
END.
```

Assume that *x* and *y* have type ([**LONG**] **POINTER TO**)* *T* where *T* is declared in interface *Defs* to be either an opaque or record type and * denotes zero or more consecutive occurrences of [**LONG**] **POINTER TO**.

In *Prog1*, dot notation can be used to achieve a form of object notation. If *x* (when dereferenced) has no field named *P*, then the expression

$x.P[args]$

is interpreted as

$Defs.P[x, args]$

Mesa will use the principal instance (§ 7.4.4) of the interface of the definitions module in which x 's type is declared, call procedure P of this interface and pass x as the implicit first parameter to P . If there is a directory entry for $Defs$ with a **USING** clause then P must be contained in it.

The interpretation of $y.Q[otherArgs]$ is similar.

Notes:

The expression $x.P$ abbreviates $x.P[]$ (but one can always supply the empty brackets to force a particular interpretation, e.g.,

$(Defs>ReturnsAProc(x))[args]$ must be written as $x>ReturnsAProc[[]args]$ if $ReturnsAProc$ has no other arguments or they are all defaulted.

If $args$ uses keyword notation, the rewritten form is

$Defs.P[self: x, args]$

where "self" is the identifier of the positionally first formal parameter of P .

Principal instances of imported interfaces are not created automatically in a context where dot notation would require the examination of an interface which was not explicitly declared (c.f. § 7.4.4).

7.8 The Mesa configuration language, an introductory example

This section discusses C/Mesa, the Mesa configuration language, first by example and then more rigorously by syntactic definition and detailed semantics. It ends with a number of detailed examples which explore some of the intricacies of C/Mesa.

We first present an example consisting of three Mesa modules:

An interface (a **DEFINITIONS** module),

an implementor for it (a **PROGRAM** module),

and a client for the implementation (also a **PROGRAM** module).

The example is a program for collecting a set of strings in a sorted binary tree. Since it is written principally to show the relationships among definitions, implementors, and clients, several important functional details were omitted. For example, it allocates storage from the common system zone but doesn't have any way to deallocate it. It has a single tree rooted in the global frame of *LexiconImpl*, when it might be more useful to pass in the tree address as a parameter.

Following the program listings is a sequence of sample configurations for systems constructed from them. The line numbers in the left margin are provided for ease of reference and are not part of the source code. First the interface:

```
d1: Lexicon: DEFINITIONS =
d2: BEGIN
d3: FindString: PROCEDURE [LONG STRING] RETURNS [BOOLEAN];
d4: AddString: PROCEDURE [LONG STRING];
d5: PrintLexicon: PROCEDURE;
d6: END.
```

7.8.1 *Lexicon*: a module implementing *LexiconDefs*

The following module, *LexiconImpl*, implements the *Lexicon* interface. That is,

- (a) *LexiconImpl* declares **PUBLIC** procedures *FindString*, *AddString*, and *PrintLexicon*, which have procedure types conforming to their counterparts in the **DEFINITIONS** module;
- (b) *LexiconImpl* **EXPORTS** the interface *Lexicon*.

LexiconImpl **IMPORTS** three interfaces: *Heap*, *IO*, and *String*. The **USING** clauses of the **DIRECTORY** note which procedures and variables are referenced from each.

Details on Mesa system interfaces are contained in the *Pilot Programmer's Manual* and *Mesa Programmer's Manual*. A fictitious interface *IO* is used in the example for simplicity. An actual Mesa program would use one of several output interfaces, depending upon where it wanted to send the output.

The code for *LexiconImpl* follows.

```
i1:  DIRECTORY
i2:    Heap USING [systemZone],
i3:    IO USING [WriteLine],
i4:    Lexicon USING [],
i5:    String USING [AppendString];
i6:
i7:  LexiconImpl: PROGRAM
i8:    IMPORTS Heap, IO, String
i9:    EXPORTS Lexicon =
i10:
i11:  BEGIN
i12:
i13:    Node: TYPE = RECORD [llink, rlink: NodePtr, string: LONG STRING];
i14:    NodePtr: TYPE = LONG POINTER TO Node;
i15:    Comparative: TYPE = {lessThan, equalTo, greaterThan};
i16:
i17:    root: NodePtr ← NIL;
i18:
i19:    FindString: PUBLIC PROCEDURE [s: LONG STRING] RETURNS [BOOLEAN] =
i20:      BEGIN RETURN [SearchForString[root, s]]; END;
i21:
```

```

i22:   SearchForString: PROCEDURE [n: NodePtr, s: LONG STRING]
i23:   RETURNS [found: BOOLEAN] =
i24:   BEGIN
i25:   IF n = NIL THEN RETURN [FALSE];
i26:   SELECT LexicalCompare[s, n.string] FROM
i27:     lessThan => found ← SearchForString[n.llink, s];
i28:     equalTo => found ← TRUE;
i29:     greaterThan => found ← SearchForString[n.rlink, s];
i30:   ENDCASE;
i31:   RETURN [found];
i32:   END;
i33:
i34:   AddString: PUBLIC PROCEDURE [s: LONG STRING] =
i35:   BEGIN InsertString[root, s]; END;
i36:
i37:   InsertString: PROCEDURE [n: NodePtr, s: LONG STRING] =
i38:   BEGIN
i39:     NewNode: PROCEDURE RETURNS [n: NodePtr] =
i40:     BEGIN
i41:       n ← Heap.systemZone.NEW [Node ←
i42:         [string: Heap.systemZone.NEW [StringBody[s.length]], llink: NIL, rlink: NIL]];
i43:       String.AppendString[n.string, s];
i44:       RETURN;
i45:     END;
i46:
i47:     IF n = NIL THEN root ← NewNode[ ]           -- then just return
i48:     ELSE
i49:       SELECT LexicalCompare[s, n.string] FROM
i50:         lessThan => IF n.llink # NIL THEN InsertString[n.llink, s]
i51:           ELSE n.llink ← NewNode[ ];
i52:         equalTo => NULL;                               -- already there; just return
i53:         greaterThan => IF n.rlink # NIL THEN InsertString[n.rlink, s]
i54:           ELSE n.rlink ← NewNode[ ];
i55:       ENDCASE;
i56:     END;
i57:
i58:   LexicalCompare: PROCEDURE [s1, s2: LONG STRING] RETURNS [c: Comparative] =
i59:   BEGIN
i60:     n: CARDINAL = MIN [s1.length, s2.length];
i61:     i: CARDINAL;
i62:     FOR i IN [0..n] DO
i63:       SELECT LowerCases[s1[i]] FROM
i64:         < LowerCase[s2[i]] => RETURN [lessThan];
i65:         > LowerCase[s2[i]] => RETURN [greaterThan];
i66:       ENDCASE;
i67:     ENDOLOOP;
i68:     c ← SELECT s1.length FROM
i69:       < s2.length => lessThan,           -- s1 is shorter than s2
i70:       > s2.length => greaterThan,       -- s1 is longer than s2
i71:       ENDCASE => equalTo;             -- lengths are the same
i72:     RETURN [c];
i73:   END;
i74:
i75:   lower: PACKED ARRAY CHARACTER ['A .. 'Z] OF CHARACTER =
i76:     ['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];
i77:

```



```

i78:   LowerCase: PROCEDURE [c: CHARACTER] RETURNS [CHARACTER] =
i79:       BEGIN RETURN [IF c IN ['A..'Z] THEN lower[c ] ELSE c]; END;
i80:
i81:   PrintLexicon: PUBLIC PROCEDURE =
i82:       BEGIN PrintNode[root] END;
i83:
i84:   PrintNode: PROCEDURE [n: NodePtr] =
i85:       BEGIN
i86:           IF n = NIL THEN RETURN;
i87:           PrintNode[n.llink];
i88:           IO.WriteLine[n.string];
i89:           PrintNode[n.rlink];
i90:       END;
i91:   END.

```

7.8.2 LexiconClient: a client module

LexiconClient is a client for *LexiconImpl* and `IMPORTS Lexicon`. It is also a client for the interface *IO* (and also uses the constant *CR* defined in *IO* in section 7.1). The program provides a simple terminal interface to a user for testing *LexiconImpl*.

```

c1:   DIRECTORY
c2:       IO USING [CR, ReadChar, ReadLine, WriteChar, WriteLine],
c3:       Lexicon USING [AddString, FindString, PrintLexicon];
c4:
c5:   LexiconClient: PROGRAM IMPORTS IO, Lexicon =
c6:
c7:       BEGIN OPEN IO, Lexicon;
c8:
c9:       s: LONG STRING ← [80];
c10:      ch: CHARACTER;
c11:      DO -- loop until stopped by user typing q or Q (last case below).
c12:          WriteChar[CR]; WriteLine["Lexicon Command: "];
c13:          ch ← ReadChar[ ];
c14:          WriteChar[ch]; -- Echo the character (ReadChar doesn't).
c15:          SELECT ch FROM
c16:              'f, 'F = >
c17:              BEGIN
c18:                  WriteLine["ind: "]; -- terminal will read: "find":
c19:                  ReadLine[s ]; -- s will contain the string read from the terminal
c20:                  IF FindString[s ] THEN WriteLine[" -- found"]
c21:                  ELSE WriteLine[" -- not found"];
c22:              END;
c23:              'a, 'A = >
c24:              BEGIN
c25:                  WriteLine["dd: "]; -- terminal will read: "add":
c26:                  ReadLine[s ];
c27:                  AddString[s ];
c28:              END;
c29:              'p, 'P = >
c30:              BEGIN
c31:                  WriteLine["rint lexicon"]; -- terminal will read: "print lexicon"
c32:                  WriteChar[CR]; PrintLexicon[ ];
c33:              END;

```

```

c34:      'q, 'Q = >
c35:      BEGIN
c36:      WriteLine["uit"]; WriteChar[CR];  -- terminal will read: "quit"
c37:      STOP;
c38:      END;
c39:      ENDCASE = > WriteLine["Commands are Find,Add,Print lexicon,and Quit"];
c40:      ENDOLOOP;
c41:
c42:      END.

```

7.8.3 Binding, loading, and running a configuration: an overview

A *configuration description* is a recipe written in C/Mesa that describes how a set of Mesa modules are to be joined together to form a configuration. This joining is accomplished by *binding* the configuration and results in an object file, also called a *binary configuration description* (BCD).

The simplest (or atomic) object file is the object module for a Mesa program module. Thus, the Mesa compiler produces the simplest object files, and the Mesa *binder* produces more complex object files from simpler ones. Indeed, a configuration may combine both atomic and non-atomic object files together into a single, new object file. For these reasons, the object modules produced by the Mesa compiler have the same form of names as the output of the binder, i.e., names of the form "BasicName.bcd."

Once a object file has been created, it can be loaded and run.

Loading is a sequence of two actions. The first makes an *instance* of the configuration by allocating a global frame for each atomic module in the object file. Each frame has space for the module's *static variables* (those declared in the main body of the module) and some extra space for information used by the Mesa system. Imported procedures and variables are accessed via *links*. Space for these links is allocated either in the frame or in the code of the module.

The second part of loading completes the binding process by filling in the links for each module instance in the configuration. Some of these links will be connected to procedures and variables in the same configuration. Others will be connected to procedures and variables in the running system into which the configuration is being loaded. Unsatisfied imported links in the running system which are satisfied by exports of the configuration being loaded are filled in at the same time.

Once a configuration is loaded, each module instance in it has all its interfaces bound. However, no code has been executed in the instances, so global variables are not initialized, and no main body has executed. **STARTing** (§ 7.10.3) an instance executes any code for initializing static variables and also executes its main body code. For correct operation, this must occur before any of its procedures are used or before any of its global variables are referenced. If a module is not explicitly **STARTed** before one of its procedures is called, then a trap occurs, and it is automatically started. Subsequent procedure calls will not repeat this trap and auto-initialization. Section 7.10 details how these mechanisms generalize for configurations.

7.8.4 A configuration description for running LexiconClient

The following configuration will bind *LexiconImpl*, *LexiconClient*, and other necessary modules and can be used to start the client program running. The comments to the right of each module name indicate which interfaces are imported and exported by that particular

module; they are not part of *Config1*. This configuration is completely *self-contained*: all the needed imports are satisfied by exports from modules that are part of the configuration.

```

Config1: CONFIGURATION
      CONTROL LexiconClient =
      BEGIN
      Fsp;           --           EXPORTS Heap
      IOImpl;       --           EXPORTS IO
      StringImpl;  --           EXPORTS String
      LexiconImpl; -- IMPORTS Heap, IO, String   EXPORTS Lexicon
      LexiconClient; -- IMPORTS IO, Lexicon
      END.

```

To see that this configuration is completely self-contained, notice that *LexiconClient* imports *IO*, which is exported by *IOImpl*, and imports *Lexicon*, which is exported by the instance of *LexiconImpl*. Similarly, the import requirements of the other instances are satisfied by some exported interface in *Config1*.

7.9 C/Mesa: syntax and semantics

The following is the complete syntax for C/Mesa. It bears a strong resemblance to Mesa itself, but this grammar describes a completely separate language. A phrase class beginning with a **C** indicates a syntactic unit that is unique to C/Mesa. All the other units have the same *syntax* (but not necessarily exactly the same semantics) as they do in Mesa.

```

ConfigDescription ::= CDirectory           -- optional
                       CPacking
                       Configuration .      -- note the final period

CDirectory          ::= -- same as in Mesa, only no USING clauses

Configuration      ::= identifier : CHead =
                       CBody

CExports           ::= empty | EXPORTS ItemList | EXPORTS ALL

CExpression       ::= CPrimary | CExpression THEN CRightSide

CLeftSide         ::= Item | [ ItemList ]

CBody              ::= BEGIN CStatementSeries END

CHead              ::= CONFIGURATION CLinks Imports CExports ControlClause

ControlClause     ::= CONTROL idlist | empty

CLinks             ::= empty | LINKS : CODE | LINKS : FRAME

CPacking           ::= empty | CPackSeries ;

CPackList         ::= PACK IdList

```

CPackSeries	:: = CPackList CPackSeries ; CPackList
CPrimary	:: = CRightSide CPrimary PLUS CRightSide
CRightSide	:: = Item Item [] CLinks Item [IdList] CLinks
CStatement	:: = CLeftSide ← CExpression CRightSide Configuration
CStatementSeries	:: = CStatement CStatementSeries ; CStatementSeries : CStatement
Imports	:: = empty IMPORTS ItemList
Item	:: = identifier identifier : identifier
ItemList	:: = Item ItemList , Item

We will use the term “component” to refer to the parts of a configuration; i.e., for both atomic modules and configurations containing several modules. When necessary, the kind of component will be expressly given.

Similarly, we will use the term “interface” to stand for an interface record or a module instance (if used in discussing imports or exports), and we will distinguish as necessary. However, “interface” will *never* include or imply the term “interface type” (§ 7.4.1).

Lastly, we will need to distinguish between instances of components and their prototypes (the object files) from which such instances are made. Hence, a *program prototype* is the object file for a Mesa program module, and a *configuration prototype* is the analog for configurations. If the term *prototype* is used by itself, it includes both cases.

The **CPacking** and **CLinks** clauses in the syntax are directives to the Mesa Binder. **CPacking** identifies modules whose code should be packed together for swapping purposes. **CLinks** specifies for a module or a configuration whether links to imported interfaces should be stored in the frame or in the code. The use and implications of these optional clauses is described in Appendix D.

7.9.1 IMPORTS, EXPORTS, and DIRECTORY in C/Mesa

For completely self-contained, simple configurations like *Config1*, a configuration description is primarily just a list of component names. An instance of each named component will be part of the configuration, and if a component imports any interfaces, they will be supplied by those exported from other components of the configuration.

Configurations need not be self-contained, however, and may themselves import interfaces to be further imported by their components. In this way, subsystems can be constructed with some imported interfaces unbound. Loading such a configuration or naming it as a component in another configuration will supply the necessary interfaces. Furthermore, a configuration can make exported interfaces available for importation by other modules and configurations. For example, the interfaces *Heap*, *IO*, and *String* needed by *Config1* would normally be supplied by a pre-existing Mesa system

configuration. Therefore, it is really not necessary to include instances of *Fsp*, *IOImpl*, and *StringsImpl* in *Config1*. Instead, it can just import them:

```
c2.1:  Config2: CONFIGURATION
c2.2:      IMPORTS Heap, IO, String
c2.3:      CONTROL LexiconClient =
c2.4:      BEGIN
c2.5:      LexiconImpl;
c2.6:      LexiconClient;
c2.7:      END.
```

The imports clause in a configuration serves the same purpose as in a program module. The rule for importing is: If some component named in a configuration imports *SomeDefs*, and *SomeDefs* is not exported by a component in the configuration, then it must be imported. For example, *Heap* did not have to be imported into *Config1*, but it did have to be imported into *Config2*.

The rule for exports is simpler: If a component in a configuration exports an interface, that interface may also be exported from the entire configuration. It does not *have* to be exported, however. Thus, the programmer controls what is exported from a configuration and what is hidden from external view.

None of the sample configurations given so far have had a **DIRECTORY** section. This is because the default association of a component named *Prog* is to a file named "Prog.bcd" in which the **ModuleName** is also *Prog*. When this is the case, the programmer does not need to supply a **DIRECTORY** part; one would be needed if the file did not have such a defaultable name. For example:

```
DIRECTORY
  Prog: FROM "OldProgFile";
```

could not be omitted if the component named *Prog* is contained in the file "OldProgFile.bcd," rather than in "Prog.bcd."

Fine point:

Please refer to subsection 7.4.4, concerning implicitly imported interfaces.

7.9.2 Explicit naming, IMPORTS, and EXPORTS *

In Mesa, names may be given to the interface records in an **IMPORTS** list (§ 7.4.1); the same is true in a configuration description. These names can then be used to supply the interfaces needed by component instances in the configuration. The notation for explicitly supplying interfaces to a component is similar to that for parameter lists in Mesa (except that there is no keyword notation for explicit imports parameter lists). For example, lines c2.1 through c2.5 above could have been written as

```

c2a.1: Config2A: CONFIGURATION
c2a.2:      IMPORTS alloc: Heap, io: IO, str: String
c2a.3:      CONTROL LexiconClient =
c2a.4:      BEGIN
c2a.5:      LexiconImpl[alloc, io, str];
...

```

The interfaces listed after *LexiconImpl* must correspond in order and (interface) type with the **IMPORTS** list for *LexiconImpl* (look at *LexiconImpl* in subsection 7.8.1 to check this).

A name may also be given to each component instance in a configuration by preceding the instance with "identifier:". This facility is necessary to distinguish multiple instances of the same prototype from one another. For example, we could name the *LexiconImpl* instance in line c2a.5 as follows:

```
alex: LexiconImpl[alloc, io, str];
```

LexiconImpl exports an interface whose type is *Lexicon*, and that interface record can also be named. The following further modification to line c2a.5 names it *lexRec*:

```
lexRec: Lexicon ← alex: LexiconImpl[alloc, io, str];
```

Here, as in Mesa, the type of *lexRec* follows the colon in the declaration, and *lexRec* is assigned the (single) interface exported by *LexiconImpl*. However, the type *Lexicon* is not actually necessary (it is inferred from *LexiconImpl*'s **EXPORTS** list), and the line could have been shortened to

```
lexRec ← alex: LexiconImpl[alloc, io, str];
```

Using all these explicit naming capabilities, we can now write a new version of the configuration in which none of the C/Mesa default naming is used:

```

c3.1:   Config3: CONFIGURATION
c3.2:   IMPORTS alloc: Heap, io: IO, str: String
c3.3:   CONTROL lexClient =
c3.4:   BEGIN
c3.5:   lexRec: Lexicon ← alex: LexiconImpl[alloc, io, str];
c3.6:   lexClient: LexiconClient[io, lexRec];
c3.7:   END.

```

An exported interface like *lexRec* need not always be set as the result of including a component instance like *alex* in the configuration. One can also assign interface records to one another as in the following two (equivalent) lines:

```
anotherLexRec: Lexicon ← lexRec;
anotherLexRec ← lexRec;
```

The form of **CRightSide** in these two statements only copies *lexRec*, whereas ones like line c3.5 above involve a "call" on a component prototype. The result of that "call" is an instance of the component and a set of results, the interface records exported by it.

7.9.3 Default names for interfaces and instances *

A component instance that is not explicitly given a name is given a default name equal to the name of the component prototype. Thus, the body of *Config2* is treated as if the programmer had written:

```
BEGIN
  LexiconImpl: LexiconImpl;
  LexiconClient: LexiconClient;
END.
```

Similarly, an unnamed interface is given a default name equal to the name of its interface type. So, another equivalent body for *Config2* is

```
BEGIN
  Lexicon: Lexicon ← LexiconImpl: LexiconImpl[];
  LexiconClient: LexiconClient;
END.
```

The empty imports parameter list in "*LexiconImpl*[]" specifies that a new instance of the prototype *LexiconImpl* is to be created. If the empty imports list were not there, the binder would interpret the appearance of *LexiconImpl* (the one after the colon) as the name of an already existing interface (*not* of an already existing module instance). When no assignment is specified, the empty imports parameter list is not necessary, as shown in the earlier examples.

Normally, omitting an imports parameter list (or, equivalently, specifying an empty list) means that the binder should use the default-named interfaces needed by that component instance. Thus, we could rewrite a completely explicit (and very wordy, but equivalent) version of *Config2*:

```
c2x.1: Config2X: CONFIGURATION
c2x.2:   IMPORTS Heap: Heap, IO: IO, String: String
c2x.3:   CONTROL LexiconClient =
c2x.4:   BEGIN
c2x.5:     Lexicon: Lexicon ← LexiconImpl: LexiconImpl[Heap, IO, String];
c2x.6:     LexiconClient: LexiconClient[IO, Lexicon];
c2x.7:   END.
```

Notice that the defaults greatly simplify a configuration, but that they also obscure a great deal of machinery concerned with naming things. It is important that the programmer not completely forget these details. Otherwise one could commit errors by not distinguishing between interface records and interface types, or between component instances and prototypes. For instance, this could be a problem if there are multiple component instances. Therefore, one is well advised to assign unique names to the instances.

7.9.4 Multiple exported interfaces from a single component *

A component can export more than a single interface. Assigning these exported interfaces to interface records is done using a Mesa-like extractor (§ 3.3.6). For example, if we had a

program module *StringsAndIOImpl* that exported both *String* and *IO*, we could use it in a modified *Config2* as follows:

```
c4.1:  Config4: CONFIGURATION
c4.2:      IMPORTS alloc: Heap
c4.3:      CONTROL LexiconClient =
c4.4:      BEGIN
c4.5:      [str: String, io: IO] ← StringsAndIOImpl[ ];
c4.6:      LexiconImpl[alloc, io, str];
c4.7:      LexiconClient[io, Lexicon];
c4.8:      END.
```

Line c4.5 assigns the exported interfaces obtained by instantiating *StringsAndIOImpl* (that is why it has an explicit, although empty imports parameter list following it) and declares their types as well. It would be equally correct to write

```
[str, io] ← StringsAndIOImpl[ ];
```

In this case the types for *io* and *str* would be inferred from the types of the interface records exported by *StringsAndIOImpl*. However, if the programmer had written instead,

```
[io, str] ← StringsAndIOImpl[ ];
```

with the positions of *io* and *str* reversed, that would have been accepted, but would have caused errors in both lines c4.6 and c4.7 because their inferred types would not match those explicit imports parameter lists. Be cautious when doing this.

Default names could also have been used for the exported interfaces in line c4.5, and *Config4* could simply have been written as

```
c4a.1: Config4A: CONFIGURATION
c4a.2:      IMPORTS Heap
c4a.3:      CONTROL LexiconClient =
c4a.4:      BEGIN
c4a.5:      StringsAndIOImpl;
c4a.6:      LexiconImpl;
c4a.7:      LexiconClient;
c4a.8:      END.
```

This would assign the exported interfaces to the default-named records *String* and *IO* and would use them in the defaulted import parameter lists for *LexiconImpl* and *LexiconClient*. Line c4a.5 could also show what *StringsAndIOImpl* exports using the default names for its exported records. This would give rise to the statement:

```
[String, IO] ← StringsAndIOImpl[ ];
```

Cases like this require that the user be aware of the distinction between interface records and interface types: *String* names an interface record here, but in line c4.5, it names an interface type.

7.9.5 Multiple components implementing a single interface *

An exported interface can be the result of contributions by a number of components. Think of the interface as a logical unit that may be implemented by a number of cooperating physical units (i.e., modules and configurations). For example, assume that *LexiconImpl* is divided into two modules *LexiconFAImpl* and *LexiconPIImpl*, with *LexiconFAImpl* providing the procedures *FindString* and *AddString*, and *LexiconPIImpl* providing *PrintLexicon*. Each exports *Lexicon*, but neither fully implements that interface. Still, *LexiconClient* will see a single interface in the following:

```
c5.1:  Config5: CONFIGURATION
c5.2:    IMPORTS Heap, IO, String
c5.3:    CONTROL LexiconClient =
c5.4:    BEGIN
c5.5:    lexRec: Lexicon ← LexiconFAImpl[ ];           -- use default imports
c5.6:    lexRec ← LexiconPIImpl[ ];                 -- merge interface contributions
c5.7:    LexiconClient[IO, lexRec];
c5.8:    END.
```

The two separate assignments to *lexRec* above actually merge the interface elements exported by the two modules. This merging does not allow any duplication of elements, and if both modules exported *PrintLexicon*, for example, an error would be generated during processing of *Config5* by the Binder.

The user may control the merging of interfaces himself using the **PLUS** operator. To obtain the same effect as above (but by explicit specification), one could write

```
lexRecFA ← LexiconFAImpl[ ];           -- one part
lexRecP ← LexiconPAImpl[ ];           -- the other part
lexRec ← lexRecFA PLUS lexRecP;       -- the merge
LexiconClient[IO, lexRec];            -- same as line c5.7
```

If the programmer wanted to use the original *LexiconFAImpl*, but use *LexiconPAImpl*'s *PrintLexicon* in the interface instead of *Lexicon*'s, he could use the **THEN** operator:

```
lexRec ← LexiconFAImpl[ ];           -- defines a complete interface
lexRecP ← LexiconPAImpl[ ];         -- defines one procedure
lexRecNew ← lexRecP THEN lexRec;    -- this order is important
LexiconClient[IO, lexRecNew];
```

The **THEN** operator makes an interface that includes all the elements defined by *lexRecP* (the left operand) together with those from *lexRec* (the right operand) that do not duplicate any in *lexRecP*. This could be useful if one simply wanted to test a new version of *PrintLexicon* procedure without altering *LexiconImpl* itself during the debugging period. Also, one could use **THEN** to provide a number of alternative *PrintLexicon* procedures, with the standard one incorporated in *LexiconImpl*.

7.9.6 Nested (local) configurations

Configurations may be defined within configurations, much like local procedures (§ 5.5) may be defined within other procedures in Mesa. They can then be instantiated and parametrized, and they can export interfaces (just like any configuration).

Nested configurations can be used to hide some of the interfaces exported by components in a configuration. For example, suppose that multiple instances of some component *ProgMod* were needed in a configuration, and further suppose that *ProgMod* exports the interface *ProgDefs*. Even if none of the exported *ProgDefs* interface records are needed in the configuration, they would each have to be given a unique name to avoid an interface merging error (§ 7.9.5).

This could be avoided by defining the following nested configuration:

```
NonexportingPM: CONFIGURATION = BEGIN ProgMod END.
```

Using *NonexportingPM* in place of *ProgMod* avoids the duplicate interface problem because the local configuration does not export the interface *ProgDefs* produced by instantiating *ProgMod* within it.

Nested configurations can also be used to avoid writing sequences of C/Mesa statements more than once. By collecting such a sequence in a nested configuration, one can get the effect of writing the whole sequence simply by instantiating the configuration.

The scope rules for names in C/Mesa allows a nested configuration to access interfaces and other (also nested) configurations outside it. So, one configuration can make instances of others. However, in its **IMPORTS** list, a nested configuration must name any interfaces that its components import but which are not satisfied within it. That is, interfaces are never automatically imported into a nested configuration..

7.9.7 Package creation: EXPORTS ALL

C/Mesa enables all of a configuration's interfaces to be exported easily using the **EXPORTS ALL** alternative of **CExports** in section 7.9.

All of a configuration's top-level interface records (but not module instances) can be exported by including **ALL** in the exports list. The top-level interface records are those that are exported by directly contained modules and subconfigurations. For example:

```
Exporter: CONFIGURATION
  IMPORTS . . .
  EXPORTS ALL =
  BEGIN
    SubConfiguration1:
    SubConfiguration2; . . .
    Module1; . . .
  END.
```

exports all top level interface records of *SubConfiguration1*, *SubConfiguration2*, *Module1*, etc. If a module instance must also be exported, just include it in the exports list; e.g.

EXPORTS ALL, *Module1*

A word of caution, however: **EXPORTS ALL** will export “hidden” interfaces (§ 7.9.6).

7.10 Loading and running modules and configurations

This section describes how configurations are loaded and run. Simple, atomic modules are discussed first, and then more general configurations.

Loading and running an atomic module is a sequence of five actions:

- (1) copying its executable code from the object file into memory,
- (2) allocating a global frame for its static variables,
- (3) filling in links for imported items,
- (4) filling in unsatisfied imported links in the running system that are satisfied by exports of the configuration being loaded.
- (5) initializing the module's variables and executing its main body code.

Actions (1), (2), (3), and (4) are accomplished by the loader. Action (5) can be accomplished by explicitly starting the instance or by means of a trap on the first call to any of its procedures. Both of these methods are described below.

7.10.1 Making copies of modules

A copy of a module may be made using the **NEW** operator. The syntax for **NEW** is

Expression ::= ... | **NEW** Variable

The **Variable** may be the name of a pointer to the frame of a program module or a **PROGRAM**. The **PROGRAM** may be imported or be the module containing the **NEW** statement. For example:

```
progInst1, progInst2: LONG POINTER TO FRAME [Prog];
...
progInst2 ← NEW ProgInst1;
```

The new instance has a copy of the bindings of the original, and shares its object code. However, it must be started to supply its program parameters (if any) and to initialize its global variables. If a module imports a program *ProgImpl* (§ 7.4.2), the operation “**NEW** *ProgImpl*” makes a copy of *ProgImpl*.

A **PROGRAM** is a transfer item (§ 7.3) and may be declared in a definitions module in the same way as a procedure is. Such a **PROGRAM** is part of the interface defined by that definitions module and may be imported by another module as part of that interface. Copies of that module can then be made using the **NEW** operation. For example, assume that the following declaration appears in the definitions module *Defs*:

```
ExportedProg: PROGRAM [i: INTEGER];
```

Any program that imports *Defs* will then have access to a value named *ExportedProg* which will have been bound (in step (3) of the loading process) to an instance of a program whose parameter types conform with those of *ExportedProg*. The only operation that a program can perform using this value is to **START** it, **RESTART** it, or make a copy of it using "**NEW *ExportedProg***." In summary, a program imported as part of an interface behaves like a value that is a pointer to a frame.

If a program *Prog* wishes to create a copy of itself, it can say:

```
copy: LONG POINTER TO FRAME [Prog];
...
copy ← NEW Prog;
```

7.10.2 How the loader binds interfaces

Each instance of a module or configuration may export some interfaces. To make these exported interfaces available for importation by other instances, the loader maintains a single global table of all exported interfaces. If any duplicate exported items are created as the result of loading, they supersede the existing items, as if a **THEN** (§ 7.9.5) had been done.

Complicated bindings done to hide interfaces, etc. must be done using the binder; only simple, straightforward bindings should be used at loading time.

7.10.3 **STARTing**, **STOPping**, and **RESTARTing** module instances

In a **PROGRAM** module, any non-constant initialization of global variables, and any statements not contained within procedures constitute the *mainline code* for the module. The process of executing the mainline code is referred to as *starting* the module. Mesa provides facilities for one program starting another explicitly, *but they are almost never used in real programs*. Modules almost never have parameters, they never **STOP**, and are either started explicitly by appearing in a **CONTROL** list of a configuration or implicitly by means of start traps. Modules that **EXPORT** variables (§ 7.4.3) need to be started before any other module attempts to reference those exported variables.

The remainder of this section may be considered as "starred."

The **START** operation suspends the execution of the program or procedure executing it and transfers control to a new, uninitialized instance of an atomic module. Additionally, if the program instance being started requires parameters, they are supplied as part of the **START**. Similarly, if the program being started is specified to return results (more details below), then the **START** operation must appear in a **RightSide** context, and the returned value is the value of the operation. Its syntax is

```
StartStmt      ::= START Call | ...
StartExpr      ::= START Call | ...
```

The variable following the word **START** must represent a global frame pointer or program variable; i.e., its type must conform to some **LONG POINTER TO FRAME** type or **PROGRAM** type. Here are some examples of its use:

```

START progInst;
START ExportedProg[5 + j];
x ← START progWithResult[firstArg: a, secondArg: b];  --keyword parameter list

```

When a program is started, it first executes code to initialize any static variables that were declared with initialization expressions. The initializations are done in the order in which the variables were declared in the program. Also, they may call both local and imported procedures (since descriptors for all imported procedures are filled in as part of loading and the **NEW** operation).

After all initialization expressions are complete, the *mainline* statements of the program are executed. Control can then return to the caller (the program or procedure that initiated the **START**) in one of two ways: the started program may **STOP**, or it may **RETURN** with results.

A program that executes a **STOP** can be **RESTARTED** later. **RESTART** is distinct from **START** primarily because it cannot pass parameters as **START** can. If a program does not return results, it may return either by an explicit use of **STOP** or by running off the end main body.

If a program declares (in its **ModuleHeader**) that it returns results, it uses **RETURN** statements just as does a procedure (and it cannot use **STOP**). A **RETURN** from a program does *not* deallocate its global frame. The syntax for **RESTART** and **STOP** is

```

RestartStmt      ::= RESTART Variable
StopStmt         ::= STOP

```

The **Variable** following **RESTART** must be a pointer to the frame for a program instance or a program variable, just as for **START**. A program that has not done a **STOP** cannot be **RESTARTED**. Attempting to do so will result in a run time error.

A module instance can also be **STARTED** "automatically." If a call is made on a procedure in a module that has not yet been started, a *start trap* occurs. The runtime system suspends the procedure call and **STARTS** the module with no parameters. When the *main body* of the module returns, the trap handler restarts the procedure call that was in progress when the trap occurred. If a start trapped module expects parameters, a *stack error* will occur at run time. If a start trapped module returns results, a different run time error will occur. (See the next section for further discussion of the start trap for configurations.)

Warning: A module must be **STARTED** either explicitly or implicitly before any attempt is made to access its variables through a **LONG POINTER TO FRAME**.

7.10.4 Loading and starting configurations, control modules

A configuration may contain a number of modules. A non-atomic configuration cannot be **STARTED** (what would it mean to start one?), but its **CONTROL** module can (if it has one). Basically, the **CONTROL** module acts as the representative for the whole configuration (since a C/Mesa configuration description does not contain executable Mesa statements). Thus, a program that **STARTS** the **CONTROL** module for a configuration has essentially **STARTED** the configuration. If the order of starting some of the instances in a configuration is important, they should all appear in the list of **CONTROL** modules in the desired order of

starting. The first control module of a configuration is obtained at run time as a result of the loading process.

A configuration can have a list of control modules. In this case, **STARTING** the control module returned by the loader will cause each module in the list to be **STARTED**, in the order given.

The start trap works for configurations as well as for atomic modules. If a start trap occurs for a module *M* in configuration *C* with control module *cM*, then the trap handler automatically starts *cM* before *M*. If the handler discovers, however, that *cM* has already been started, it will start *M* (since *cM* would have started *M* if it had intended to). In fact, if the handler starts *cM* but still finds *M* unstarted when *cM* and any other control modules **STOP**, it will start *M* itself before finally returning from the trap. Then the procedure call that caused the trap will be allowed to go through.

Fine points:

If an attempt is made to **RESTART** a program which has not been started, a **START** trap will occur and then the **RESTART** will proceed.

A **START** may have an optional catchphrase. This is discussed in chapter 8, but the form looks roughly as follows:

START *someInstance* [**ComponentList ! CatchPhrase**]

C/Mesa allows configurations to be named as control "modules." Control configurations simplify the use of packages in other programs. For example the configuration

```
CrossLister: CONFIGURATION
  IMPORTS ...
  EXPORTS ...
  CONTROL Lister, CrossStuffA =
  BEGIN
    Lister;
    CrossStuffA;
    CrossStuffB;
  END.
```

insures that all of *Lister*'s controls are started (§ 7.10.4) appropriately. Note that *CrossLister* can be written without any knowledge of *Lister*'s controls (which may be a fairly long list that changes over time), and that *Lister* does not have to export some control module just so that *CrossLister* can name it in a control list.



Signaling and signal data types

Signals are used to indicate when exceptional conditions arise during program execution, and to provide an orderly means of dealing with those conditions, at low cost if none are generated. (Unless otherwise stated, the term signal may stand for both **ERROR** and **SIGNAL**.) For example, it is common in most languages to write a storage allocator that returns a null (or otherwise invalid) pointer value if asked for a block whose size is too large. Any program that calls the allocator then embeds the call in an **IF** statement, and checks the return value to make sure that the request was satisfied. What that procedure then does is a very local decision.

In Mesa, one could write the allocator as if it *always* returned a valid pointer to an allocated block: calls to it would simply assign the returned value to a suitable pointer, without checking whether or not the allocation worked. If the caller needs to gain control when the allocator fails, the programmer attaches a **CatchPhrase** to the call; then if the allocator generates the signal *BlockTooLarge*, the caller will be able to catch the signal.

This way of handling exceptions has two important properties, one for the human reader of the program, and one for execution efficiency:

A reader of a program that calls the allocator can see immediately that an exceptional condition can arise (by the catch phrase on the call or the catch phrase appended to either the **BEGIN** of a block or the **DO** of a loop statement); he then knows that this is an unusual event and can read on with the normal program flow: **IF** statements do not have this characteristic of distinguishing one kind of branch from the other.

When the program is executing, the code to check the value returned by the allocator on every call is not present and therefore takes no space or execution time. If a signal is generated, there is more overhead to get to the catch phrase than a simple transfer; but since it happens infrequently, the overall efficiency is much higher than checking each call with an **IF** statement.

Signals work over many levels of procedure call, and it is possible for a signal to be generated by one procedure and be handled by another procedure much higher up in the call chain. We later discuss the mechanisms by which this is done; until then, examples show signals being caught by the caller of the procedure that generated the signal.

8.1 Declaring and generating SIGNALS and ERRORS

In its simplest form, a signal is just a name for some exceptional condition. Often, parameters are passed along with the signal to help a catch phrase that handles it to determine what went wrong. A catch phrase can return a result: the program that generated the signal receives this result as if it had called a normal procedure instead of a signal. Thus, it is sometimes possible to recover from a signal and allow the routine that generated it to continue on its merry way. Therefore, from the type viewpoint, signals correspond very closely to procedures; in fact, the type constructor for declaring signals is just a variation of the one for procedures:

```
SignalTC      ::= SignalOnError ParameterList RETURNS ResultList |
               SignalOnError ParameterList |
               SignalOnError RETURNS ResultList |
               SignalOnError

SignalOnError ::= SIGNAL | ERROR
```

For example, the signal *BlockTooLarge* might be defined to carry along with it two parameters, a *Zone* within which the allocator was trying to get a block, and the number of words needed to fill the current request. The catch phrase that handles the signal is expected to send back (i.e., return) an array descriptor for a block of storage to be added to the zone. The declaration of *BlockTooLarge* would look like

```
BlockTooLarge: SIGNAL [z: Zone, needed: CARDINAL]
                RETURNS [newStorage: DESCRIPTOR FOR ARRAY OF WORD];
```

A signal variable contains a unique name at run-time, which is a code identifying an *actual signal*, just as a procedure variable must be assigned an *actual procedure* before it can be used. If a procedure is imported from an interface (§ 7.4), any signals that it generates directly are probably contained in the same interface. Imported signals are bound by the same mechanisms as procedures. In addition, one may have signal variables that can be assigned any signal value of a compatible type.

The signal analog of an actual procedure is obtained by initializing a signal variable using the syntax "`= CODE`" in place of "`= BEGIN. .END`" for procedures. This causes the signal to be initialized to contain a unique value. The following syntax describes the initialization for an actual signal:

```
Initialization ::= = CODE | ...
```

A signal is generated by using it in a **SignalCall** as shown in the syntax below:

```
Statement      ::= SignalCall | ...

SignalCall     ::= SIGNAL Call | ErrorCall

ErrorCall      ::= RETURN WITH ERROR Call |
                 ERROR Call |
                 ERROR                                -- special error
```


Call is defined in section 5.2, and the called **Expression** must have some signal type in this case. A **SignalCall** can be used as an **Expression** as well as a **Statement**. For example,

```
newblock ← SIGNAL BlockTooLarge[zone, n];
```

Thus, generating a signal or error looks just like a procedure call, except for the additional word **ERROR** or **SIGNAL**.

Fine point:

Although it is not recommended, the keywords **SIGNAL** and **ERROR** may be omitted (except in the **RETURN WITH ERROR** construct). This makes the signal look exactly like a procedure call.

Initialization by **SIGNAL = CODE** produces a unique value that contains, in part, the global frame index of the module containing the initialization. If one creates a copy of the module with the **NEW** statement, signals raised by the two copies will be different. If the signal is declared and initialized in a procedure, recursive calls of the procedure will not generate different signal values.

If a signal is declared as an **ERROR**, it must be generated by an **ErrorCall**. If, however, it is declared as a **SIGNAL**, it can be generated by any **SignalCall**, including an **ErrorCall**. The difference between the two is that a catch phrase may *not* **RESUME** a signal generated by an **ErrorCall** (§ 8.2.5).

Except for a slight difference in the way the error is started (§ 8.2.3), the **RETURN WITH ERROR** construct behaves like the **ERROR** statement. Its primary use is in monitor **ENTRY** procedures (chapter 9).

The “special error” in the above **ErrorCall** syntax equation is used to indicate that something has gone wrong, without giving any indication of the cause; the statement

```
ERROR;
```

generates a system-defined error. Precisely because this construct gives no indication of the cause of the error, its use is discouraged. It is provided to cover those “impossible” cases that should never occur in correct programs but which it is always best to check for (such as falling out of a loop that should never terminate normally, or arriving at the **ENDCASE** of a **SELECT** statement that claims to handle all the cases). It can only be caught using the **ANY** option in a catch phrase (§ 8.2.3). It is customarily handled by the debugger.

8.1.1 ERROR in expressions

An expression with an **ERROR** type (or **SIGNAL** type raised as an error) conforms to any other type.

```
Color: TYPE = {red, orange, yellow, green, blue, violet};
c: Color;
button: [0..2];
invalidButtonColor: ERROR = CODE;
```

```
button ← SELECT c FROM
  red = > 0,
  yellow = > 1,
```

```

blue => 2,
ENDCASE => ERROR invalidButtonColor;

```

In the example, the only valid colors for *buttons* are *red*, *yellow*, and *blue*. Any other value results in an error. Such constructs allow an inexpensive way to get to the debugger in those “impossible” cases that arise from program errors.

Fine point:

In earlier versions of Mesa, the conformance of **ERROR** to any type was not implemented, so it was possible to declare **ERRORs** that “returned” values. While this is still legal, its use is unnecessary and highly discouraged.

8.2 Control of generated signals

Any program that needs to handle signals must anticipate that need by providing catch phrases for the various signals that might be generated. During execution, certain of these catch phrases will be *enabled* at different times to handle signals. Loosely speaking, when a signal *S* is generated, each of the procedures in the call hierarchy at that time will be given a chance to catch the signal, in a last-in-first-out order. Each such procedure *P*, if it has an enabled catch phrase, is given the signal *S* in turn, until one of them stops the signal from propagating any further (by mechanisms explained below). *P* may decide to reject *S* (in which case the next procedure in the call hierarchy will be considered), or *P* may decide to handle *S* by taking control and attempting to recover from the signal.

8.2.1 Preparing to catch signals: catch phrases

A catch phrase has the following form:

```

CatchTail      ::= Catch |
                ANY => Statement      -- ANY must come last

Catch          ::= ExpressionList => Statement

CatchSeries    ::= empty |
                CatchTail |
                Catch ; CatchSeries

```

The expressions in the **ExpressionList** must evaluate to signals. The special identifier **ANY** will match any signal (§ 8.2.3). Note that if **ANY** occurs, it must be last.

A catch phrase is written as part of an argument list, just after the last argument and before the right bracket. Catch phrases may appear in a procedure call, **SignalCall**, **NEW**, **START**, **RESTART**, **JOIN**, **FORK**, or **WAIT** (but not in a **RESUME** or **RETURN**). A catch phrase may also be appended to the **BEGIN** of a block or the **DO** of a loop statement by means of an **EnableClause**. The applicable syntax for a call and for a block or loop statement is

```

Call          ::= Variable [ ComponentList ! CatchSeries ] |
                Variable [ ! CatchSeries ] |
                ...

```

```

Block          ::= BEGIN-- (from section 4.4) -- or { ... } for BEGIN...END
                OpenClause
                EnableClause
                DeclarationSeries
                StatementSeries
                ExitsClause
                END

EnableClause   ::= empty |
                ENABLE CatchTail ; |
                ENABLE BEGIN CatchSeries END ; |
                ENABLE BEGIN CatchSeries ; END ; |
                ENABLE { CatchSeries } ; |
                ENABLE { CatchSeries } ; ;

```

Note that the **EnableClause** is always followed by a semi-colon, and **BEGIN...END** or **{...}** must be used if there is more than one **Catch** in an **EnableClause**.

The main difference between the two kinds of catch phrases (**ENABLE** and **!**) is the scope of their influence. A catch phrase on a **Call** is only enabled during that call. A catch phrase at the beginning of a compound or loop statement is enabled as long as control is in that block; it can catch a signal resulting from *any* call in the block (or generated in the block).

To clarify the scope of influence of **ENABLE** clauses, the following two diagrams are reproduced from subsection 4.4.2. The scope of each phrase extends over others with greater indentation.

```

BEGIN
  OpenClause
    EnableClause
      DeclarationSeries
      StatementSeries
    ExitsClause
  END

LoopControl
  DO
    OpenClause
      EnableClause
        DeclarationSeries
        StatementSeries
      LoopExitsClause
    ENDLOOP

```

Note that catch phrases enabled in the **EnableClause** of a **Block** or **LoopStmnt** are not in force in the **ExitsClause** or **LoopExitsClause**.

Procedures declared in the **DeclarationSeries** (of any enclosed **Block**) do not inherit the catch phrases in the **EnableClause** (this is not shown by the diagrams).

8.2.2 The scope of variables in catch phrases

Catch phrases are called to handle signals (the exact mechanisms are discussed in the next section). The naming environment that exists when a catch phrase is called (in order

of innermost to outermost scope) includes any parameters passed with that signal (these are declared as part of a signal's definition), and any variables to which the procedure or program activation *containing the catch phrase* has access.

If a Catch has more than one label (or the label ANY), where the types of those labels are not identical, then the signal's arguments are not accessible in the Statement chosen by that Catch.

If, however, there is exactly one type for the signals named in a **Catch's ExpressionList**, then the signal's arguments are accessible in the statement following "**= >**." The names used are the parameters given in the signal's declaration, just as for procedures. For example, a catch phrase for signal *BlockTooLarge* (defined earlier) might be used in a section of code such as:

```
-- in StorageDefs
BlockTooLarge: SIGNAL[z: Zone, needed: CARDINAL]
    RETURNS[newStorage: DESCRIPTOR FOR ARRAY OF CARDINAL];
GetMoreStorage: PROCEDURE[z: Zone, n: CARDINAL]
    RETURNS[DESCRIPTOR FOR ARRAY OF CARDINAL];

-- in a user program
p: POINTER TO Account;
...
p ← Allocate[SIZE[Account]!
    BlockTooLarge = > RESUME[GetMoreStorage[z, needed]]];
```

The names *z* and *needed* in the catch phrase refer to the parameters passed along with the signal from *Allocate* (see subsection 8.2.5 for a discussion of **RESUME**).

8.2.3 Catching signals

When a signal is generated, the signal code, and a descriptor for the actual arguments of the signal, are passed to a Mesa run-time procedure named *Signaller*. *Signaller's* definition is

```
Signaller: PROCEDURE[s: SignalCode, m: Message];
```

Here *s* identifies the signal being generated, and *m* contains its arguments. (Actually, different procedures are used to distinguish between **SIGNAL**, **ERROR**, and **RETURN WITH ERROR**.)

Signaller proceeds to pass the signal and its argument record from one enabled catch phrase to the next in an orderly fashion. Within each procedure invocation, catch phrases are visited in reverse order of the standard scope rules; inner blocks are visited, then outer blocks. The order, at the procedure level, follows the current call hierarchy, from the most recently called procedure to least recently called, *beginning with the procedure that generated the signal itself*. If the caller of a procedure is the outermost block of code for a program, the *Signaller* will follow its *return link* to continue propagating the signal (the return link points to the frame that last **STARTED** the module (§ 7.8.3)). Catch phrases are called by *Signaller* as if they were procedures of the following type:

```

CatchPhrase: PROCEDURE [s: SignalCode, m: Message]
    RETURNS [{Reject, Unwind, Resume}];

```

Fine point:

If a **RETURN WITH ERROR** is used in place of **SIGNAL** or **ERROR**, the procedure that generated the error is first deleted (after releasing the monitor lock, if it is an **ENTRY** procedure), and propagation of the error begins with its caller.

Because signals can be propagated right through the call hierarchy, the programmer must consider catching not only signals generated *directly* within any procedure that is called, but also any generated *indirectly* as a result of calling that procedure. Indirect signals are those generated by procedures called from within a procedure that you call, and that are not stopped before reaching you.

When a catch phrase is called, it behaves like a **SELECT** statement: it compares the signal code passed to it with each signal value in the **ExpressionList** of each **Catch** in the catch phrase. If the signal code matches one of the signal values, control enters the statement following the "**= >**" for that **Catch**. The **Catch** is said to have *caught* or *accepted* the signal. If no alternative in a catch phrase accepts the signal, there may be another enabled catch phrase in some surrounding block. If so, the first catch phrase sends control to the second one so that it can inspect the signal, and so on until the last enabled catch phrase in that routine has had a chance at the signal. If no catch phrase in the routine accepts the signal, control returns to *Signaller* with a value indicating that the signal was rejected, and *Signaller* propagates the signal to the next level in the call hierarchy.

Fine point:

A **Catch** consisting of "**ANY = > Statement**" automatically matches any signal code (and is the only way to catch the unnamed **ERROR** generated by the standalone **ERROR** statement discussed in section 8.1). The **ANY** catchall is intended primarily for use by the debugger, and should generally be avoided. It matches any signal, including **UNWIND** and all system-defined signals that might indicate some catastrophic condition (a double memory parity error, for example).

Within a catch phrase, you can use the statement **REJECT** to explicitly reject a signal, i.e., to terminate execution of that catch phrase and propagate the signal to the enclosing one.

Fine point:

If the same signal, *foo*, is enabled in several nested catch phrases in a procedure, each is given a chance to handle *foo* if the inner ones reject the signal.

Signaller continues propagating the signal up the call chain until it is exhausted, that is, until the root of the process has considered and rejected the signal. At that point, an *uncaught signal* has been generated, and drastic action must be taken.

Mesa guarantees that all signals will ultimately be caught and reported by the debugger to the user. This is helpful in debugging because the control context that existed when the signal was generated is still around and can be inspected to investigate the problem.

The declaration of *CatchPhrase* above indicates three reasons for returning to *Signaller*. The third reason, *Resume*, is discussed in subsection 8.2.5.

The first reason, *Reject*, tells *Signaller* to propagate the signal to the next visible catch phrase. There are three ways that a catch phrase can reject a signal: explicitly, with the **REJECT** statement, implicitly by not catching the signal, or by catching the signal, but having the catch phrase “fall off the end” without either a **RESUME** or an exit.

The second reason, *Unwind*, is used when a catch phrase has accepted a signal and is about to do some form of unconditional jump into the body of the routine containing it (this is the only form of “non-local goto” in Mesa). The jump may be generated by a **GOTO** statement (§ 4.4), an **EXIT** or **LOOP** (§ 4.5), or a **RETRY** or **CONTINUE** (see below). Immediately preceding such a jump, the catch phrase returns to *Signaller* with result *Unwind*; it also indicates the frame containing the catch phrase and the location for the jump. This causes *Signaller* to perform the following sequence of actions:

- (1) Starting at the point where the original signal was raised, *Signaller* passes the signal **UNWIND** to each catch phrase that rejected the original signal. (If a procedure says **RETURN WITH ERROR**, this process begins with its caller.) The propagation path for **UNWIND** is exactly the same as that of the original signal. Each catch phrase that rejected the original signal is given the opportunity to restore any invariants within its scope (by accepting **UNWIND**). After **UNWIND** is passed to the last catch phrase in a given procedure frame, *Signaller* deallocates the frame. When **UNWIND** reaches the catch phrase that accepted the original signal, it stops. At this point, all frames, beginning with the one containing the statement that raised the original signal through, but not including, the frame containing the statement that accepted the original signal, have been removed from the activation stack. Furthermore, if there are catch phrases nested within the catch phrase accepting the original signal, they have had a chance to restore any invariants within their scope.
- (2) *Signaller* then arranges for the jump to take place, and simply does a return to that frame, destroying itself in the process.

Fine point:

Inner blocks in the frame that caused the **UNWIND** are unwound so **INLINE ENTRY** procedures can release their **MONITOR** locks (see chapter 9).

Every Mesa program contains the pre-declared value

```
UNWIND: ERROR = CODE;
```

Fine points:

One cannot say **RETURN** in a catch phrase to return from the enclosing procedure. This is an implementation restriction that may be removed in the future, caused by the way in which a catch phrase is “called” like a procedure itself.

The **UNWIND** sequence gives each activation that is to lose control a chance to make consistent any data structures for which it is responsible. There are no constraints on the kinds of statements that it can use to do this: procedure calls, loops, or whatever are all legal. However, a catch for the **UNWIND** should never perform a control transfer that will also initiate an **UNWIND**:

```
START NextPhase[ ! UNWIND = > GOTO BailOut];
```

Because, in this case, the second **UNWIND** overrides the original, *Signaller* will stop propagating the **UNWIND** here instead of continuing up to the original catch phrase. Because **UNWIND** does not return to the original catch phrase, the control transfer contained there will not take place.

Consider the following example :

```

Sig1: ERROR = CODE;
Sig2: ERROR = CODE;

Proc1: PROC [x: CARDINAL] =
  BEGIN                                     -- block 1
  ENABLE {                                  -- catch phrase 1
    Sig1 => GOTO punt;                    -- catch case 1
    Sig2 => {...};                        -- catch case 2
    UNWIND => {...} }                    -- catch case 3
  Statement1;
  Statement2;
  BEGIN                                     -- block 2
  ENABLE                                  -- catch phrase 2
    Sig1 => {...};                        -- catch case 4
  Statement3;
  Statement4;
  OtherProc[x!                            -- catch phrase 3
    Sig2 => {...};                        -- catch case 5
    UNWIND => {...} ];                    -- catch case 6
  END;                                     -- block 2, scope catch phrase 2
  Statement5;
  EXITS                                    -- end scope catch phrase 1
    punt => Statement6;
  END;                                     -- Proc1, block 1

  ...

OtherProc: PROC [x: CARDINAL] = {
  StillOtherProc[x!                        -- catch phrase 4
    Sig1 => {...};                        -- catch case 7
    Sig2 => {...};                        -- catch case 8
    UNWIND => {...} ] };                  -- catch case 9

StillOtherProc: PROC [x: CARDINAL] = {
  IF x = 0 THEN ERROR Sig1 ELSE ERROR Sig2};

```

Suppose that *Sig1* is raised by *StillOtherProc*. Catch phrase 4 will get the signal first. If catch case 7 rejects the signal, catch phrase 3 sees it next and rejects it implicitly. Catch phrase 2 will see the signal next. Suppose that catch case 4 also rejects. This brings us to catch phrase 1, which exits by going to the label "punt." This will cause *Signaller* to raise the signal **UNWIND** in *StillOtherProc* at the point where *Sig1* was raised. Consequently, the catch phrases that previously rejected *Sig1* get a chance to see **UNWIND**: first catch phrase 4 then catch phrase 3 then catch phrase 2. When *Signaller* gets back to catch phrase 1, it realizes that this the catch phrase requesting the exit, so it stops propagating **UNWIND** and

allows the branch to happen. The frames are freed for any procedures whose invocations are completely exited, in this case *StillOtherProc* and *OtherProc*.

8.2.4 RETRY and CONTINUE in catch phrases

Besides **GOTO**, **EXIT**, and **LOOP**, there are two other statements, **RETRY** and **CONTINUE**, which initiate an **UNWIND**. These can only be used within catch phrases.

RETRY means “go back to the *beginning* of the statement to which this catch phrase belongs”; **CONTINUE** means “go to the statement *following* the one to which this catch phrase belongs” (what is called *Next-Statement* in chapter 4).

For a catch phrase in a **Call**, the catch phrase “belongs” to the statement containing that **Call**. Thus, if the signal *NoAnswer* is generated for the call below, the assignment statement is retried:

```
answer ← GetReply[Send[“What next?”]! NoAnswer => RETRY];
```

On the other hand, if **CONTINUE** had been used instead, the statement after the assignment would be executed next (and the assignment would *not* be performed). For example, suppose the procedure *ReadLine* reads characters from a file up to a carriage return and appends them onto the string *buffer*. If reading beyond the end of file raises the signal *StreamError*, the call

```
ReadLine[! StreamError => IF buffer.length > 0 THEN CONTINUE];
```

deals with the case of no carriage return after that last line in the file. If there is no such final line, other catch phrases higher on the call chain are given a chance to catch the signal.

For a catch phrase after **ENABLE**, there are two cases to consider, blocks and loops. In a block, the catch phrase “belongs” to the block of smallest scope that contains the **ENABLE**; the next section shows an example. In a loop, the catch phrase “belongs” to the *body* of the loop, and **CONTINUE** really means “go around the loop again.” The following two examples are equivalent:

```

UNTIL p = NIL                                -- iteration test
DO
  BEGIN                                       -- containing block of ENABLE
  ENABLE
    TryList2 => BEGIN
      p ← list2;
      CONTINUE;                               -- GOTO statement following end of containing
                                              -- block. Effectively a GOTO the iteration test
                                              -- without any further loop processing.
    END                                       -- CatchSeries
  --start loop processing
  LoopStatement1;
  ...
  IF boolCondition THEN SIGNAL TryList2;
  ...

```



```

        END                                -- containing block of ENABLE
    ENDLOOP;

    UNTIL p = NIL                          -- iteration test
    DO                                       -- ENABLE belongs to body of UNTIL loop
        ENABLE
        TryList2 = > BEGIN
            p ← list2;
            CONTINUE:                       -- same effect as LOOP, proceed to iteration test
                                           -- without any further loop processing.
            END                             -- CatchSeries
        --beginning of loop body
        LoopStatement1;
        ...
        IF boolCondition THEN SIGNAL TryList2;
        ...
    ENDLOOP;

```

If we replaced `CONTINUE` with `RETRY` in the above two examples, they would still be equivalent but the point at which control is resumed would be different. Since `RETRY` means “go back to the beginning of the statement to which this catch phrase belongs,” control would resume with *LoopStatement1* without testing the loop condition (or altering the value of any loop control variable).

In either case, recall that an *Unwind* is initiated prior to completion of a `RETRY` or `CONTINUE`.

If a procedure call in the **Initialization** clause of a declaration contains a catch phrase, this catch phrase cannot contain `RETRY` or `CONTINUE` since it is in no well defined statement.

8.2.5 Resuming from a catch phrase: RESUME

The third alternative available to a catch phrase, after *Reject* and *Unwind*, is *Resume*. This option is invoked by using the `RESUME` statement to return values (or perhaps just control) from a catch phrase to the routine that generated the signal. To that routine, it appears as if the signal call were a procedure call that returns some results. The syntax for `RESUME` is just like that for `RETURN`:

```

Statement      ::= ResumeStmt | RETRY | CONTINUE | ...

ResumeStmt     ::= RESUME |
                RESUME [ ComponentList ]

```

When *Signaller* receives a *Resume* from a catch phrase, it simply returns and passes the accompanying results to the routine that originally called it (i.e., that generated the signal). If the signal was generated by an `ErrorCall` and a catch phrase requests a *Resume*, *Signaller* simply generates a signal itself (which results in a recursive call on *Signaller*); its declaration is

```

ResumeError: PUBLIC ERROR;

```

Since it is an `ERROR`, one cannot legally `RESUME` it.

The ability to **RESUME** and return values gives the ability to deal with exceptional conditions in a way that is quite inexpensive in the non-exceptional case. For example, consider the declaration

```
StringBoundsFault: SIGNAL [s: STRING] RETURNS [ns: STRING];
```

This signal allows the user to deal with the situation where characters are to be added to a string that is already "full." Thus, the call

```
AppendChar[str, c ! StringBoundsFault = >
  BEGIN
    ns ← AllocateString[s.maxlength + 10];
    AppendString[ns, s ];
    FreeString[s ];
    RESUME [str ← ns ];
  END;
```

allocates a larger string and updates the local variable whenever the string is about to overflow. Of course, the procedure *AppendChar* has to be written in such a way as to deal with the signal being resumed with a new string value. This application of signals can cause errors if there are any procedures between the signaller and the catcher that have their own idea about the location of the string. One possible fix (if such situations are possible) is to have a second signal

```
StringMoved: SIGNAL [old, new: STRING] = CODE;
```

that is raised by *AppendChar* after *StringBoundsFault* is resumed.

The presence or absence of the **ComponentList** depends on whether the signal caught is declared to return values. In a **Catch** whose **ExpressionList** contains more than one signal, one can **RESUME** only if all signals have equivalent types. For example:

```
ASig: TYPE = SIGNAL RETURNS [CARDINAL];
sig1: ASig;
sig2: ASig;
sig3: SIGNAL RETURNS [CARDINAL];
sig4: SIGNAL;
...
ENABLE
  BEGIN
    sig1, sig2 = > RESUME [3];           -- legal
    sig1, sig3 = > RESUME [0];           -- legal
    sig1, sig4 = > RESUME [1];           -- illegal
  END;
```

8.3 Signals within signals

What happens if, in the course of handling a signal, *firstSignal*, a catch phrase (or some procedure called by it) generates another signal, *secondSignal*? Handling nested signal generation is almost exactly like non-nested signal propagation. Generating the signal will call *Signaller* (recursively, since the instance of *Signaller* responsible for the first signal is still around), and it propagates the new signal back through the call hierarchy by

calling a second activation of *Signaller*, say "*Signaller2*." When in the course of doing this it encounters the previous activation of *Signaller* ("*Signaller1*"), then something different must be done.

If *firstSignal* is not the same as *secondSignal*, *Signaller2* propagates it right through *Signaller1*, and all the activations beyond it are also given a chance to catch *secondSignal*.

On the other hand, if *secondSignal* = *firstSignal*, then all of the routines whose frames lie beyond *Signaller1*, up to the frame containing the catch phrase called by *Signaller1*, have already had a chance to handle *firstSignal*, so they are not given it again. In order to skip around that section of the call hierarchy, *Signaller2* simply copies the appropriate state variables from *Signaller1*. Next, *Signaller2* skips over the frame containing the catch phrase (by following its return link), and continues propagating *secondSignal* normally.

The main import of nested signals is to consider not only what signals can be generated, directly or indirectly, by called procedures, but also those that can be generated by catch phrases in that procedure or even the catch phrases of any calling procedures, also either directly or indirectly.



Processes and concurrency

Mesa provides language support for concurrent execution of multiple processes. This allows programs that are inherently parallel in nature to be clearly expressed. The language also provides facilities for synchronizing such processes by means of entry to monitors and waiting on condition variables.

The next section discusses the forking and joining of concurrent processes. Later sections deal with monitors, how their locks are specified, and how they are entered and exited. Condition variables are discussed, along with their associated operations.

9.1 Concurrent execution, FORK and JOIN

The FORK and JOIN statements allow parallel execution of two procedures. Their use also requires the new data type PROCESS. Since the Mesa process facilities provide considerable flexibility, it is easiest to understand them by first looking at a simple example.

9.1.1 A process example

Consider an application with a front-end routine providing interactive composition and editing of input lines:

```
ReadLine: PROCEDURE [s: STRING] RETURNS [CARDINAL] =  
  BEGIN  
    c: CHARACTER;  
    s.length ← 0;  
    DO  
      c ← ReadChar[];  
      IF ControlCharacter[c] THEN DoAction[c]  
      ELSE AppendChar[s,c];  
      IF c = CR THEN RETURN [s.length];  
    ENDLOOP;  
  END;
```

The call

```
n ← ReadLine[buffer];
```

will collect a line of user type-in up to a *CR* and put it in some string named *buffer*. Of course, the caller cannot get anything else accomplished during the type-in of the line. If there is anything else that needs doing, it can be done concurrently with the type-in by *forking* to *ReadLine* instead of calling it:

```
p: PROCESS RETURNS [CARDINAL];
```

```
p ← FORK ReadLine[buffer];
```

```
...
```

```
<concurrent computation>
```

```
...
```

```
n ← JOIN p;
```

This allows the statements labeled <concurrent computation> to proceed in parallel with user typing (clearly, the concurrent computation should not reference the string *buffer*). The **FORK** construct spawns a new process whose result type matches that of *ReadLine*. (*ReadLine* is referred to as the "root procedure" of the new process.)

Later, the results are retrieved by the **JOIN** statement, which also deletes the spawned process. Obviously, this must not occur until both processes are ready (i.e., have reached the **JOIN** and the **RETURN**, respectively); this rendezvous is synchronized automatically by the process facility.

Note that the types of the arguments and results of *ReadLine* are always checked at compile time, whether it is called or forked.

The one major difference between calling a procedure and forking to it is in the handling of signals; see subsection 9.5.1 for details.

9.1.2 Process language constructs

The declaration of a **PROCESS** is similar to the declaration of a **PROCEDURE**, except that only the return record is specified. The syntax is formally specified as follows:

```
TypeConstructor    ::= ... | ProcessTC
ProcessTC          ::= PROCESS ReturnsClause
ReturnsClause      ::= empty | RETURNS ResultList    -- from section 5.1
ResultList         ::= FieldList                    -- from section 5.1
```

Suppose that *f* is a procedure and *p* a process. In order to fork *f* and assign the resulting process to *p*, the **ReturnClause** of *f* and that of *p* must be compatible, as described in section 5.2.

The syntax for the **FORK** and **JOIN** statements is straightforward:

Statement	:: = ... JoinCall
Expression	:: = ... ForkCall JoinCall
ForkCall	:: = FORK Call
JoinCall	:: = JOIN Call
Call	:: = (§ 5.2 and § 8.2.1)

The **ForkCall** always returns a value (of type **PROCESS**) and thus a **FORK** cannot stand alone as a statement. Unlike a procedure call, which returns a **RECORD**, the value of the **FORK** cannot be discarded by writing an empty extractor. The action specified by the **FORK** is to spawn a process parallel to the current one, and to begin it executing the named procedure.

The **JoinCall** appears as either a statement or an expression, depending upon whether or not the process being joined has an empty **ReturnsClause**. It has the following meaning: When the forked procedure has executed a **RETURN** and the **JOIN** is executed (in either order),

the returning process is deleted, and

the joining process receives the results, and continues execution.

A catch phrase can be attached to either a **FORK** or **JOIN** by specifying it in the **Call**. Note, however, that such a catch phrase does not catch signals incurred during the execution of the procedure; see subsection 9.5.1 for further details.

There are several other important similarities with normal procedure calls which are worth noting:

The types of all arguments and results are checked at compile-time.

There is no *intrinsic* rule against multiple activations (calls and/or forks) of the same procedure coexisting at once. Of course, it is always possible to write procedures which will work incorrectly if used in this way, but the mechanism itself does not prohibit such use.

One expected pattern of usage of the above mechanism is to place a matching **FORK/JOIN** pair at the beginning and end of a single textual unit (i.e., procedure, compound statement, etc.) so that the computation within the textual unit occurs in parallel with that of the spawned process. This style is encouraged, but is *not* mandatory; in fact, the matching **FORK** and **JOIN** need not even be done by the same process. Care must be taken, of course, to insure that each spawned process is joined only once, since the result of joining an already deleted process is undefined. Note that the spawned process always begins and ends its life in the same textual unit (i.e., the target procedure of the **FORK**).

While many processes will tend to follow the **FORK/JOIN** paradigm, there will be others whose role is better cast as continuing provision of services, rather than one-time calculation of results. Such a "detached" process is never joined. If its lifetime is bounded at all, its deletion is a private matter, since it involves neither synchronization nor delivery of results. No language features are required for this operation; see the *Pilot*

Programmer's Manual for the description of the system procedure provided for detaching a process.

9.2 Monitors

Generally, when two or more processes are cooperating, they need to interact in more complicated ways than simply forking and joining. Some more general mechanism is needed to allow orderly, synchronized interaction among processes. The interprocess synchronization mechanism provided in Mesa is a variant of *monitors* adapted from the work of Hoare, Brinch Hansen, and Dijkstra. The underlying view is that interaction among processes always reduces to carefully synchronized access to shared data, and that a proper vehicle for this interaction is one which unifies:

- the synchronization
- the shared data
- the body of code which performs the accesses

The Mesa monitor facility allows considerable flexibility in its use. Before getting into the details, let us first look at a slightly over-simplified description of the mechanism and a simple example. The remainder of this section deals with the basics of monitors (more complex uses are described in section 9.4); **WAIT** and **NOTIFY** are described in section 9.3.

9.2.1 An overview of monitors

A monitor is a module instance. It thus has its own data in its global frame, and its own procedures for accessing that data. Some of the procedures are public, allowing calls into the monitor from outside. Obviously, conflicts could arise if two processes were executing in the same monitor at the same time. To prevent this, a monitor lock is used for mutual exclusion (i.e., to insure that only one process may be in each monitor at any one time). A call into a monitor (to an entry procedure) implicitly acquires its lock (waiting if necessary); and returning from the monitor releases it. The monitor lock serves to guarantee the integrity of the global data, which is expressed as the monitor invariant -- i.e. an assertion defining what constitutes a "good state" of the data for that particular monitor. It is the responsibility of every entry procedure to restore the monitor invariant before returning, for the benefit of the next process entering the monitor.

Things are complicated slightly by the possibility that one process may enter the monitor and find that the monitor data, while in a good state, nevertheless indicates that that process cannot continue until some other process enters the monitor and improves the situation. The **WAIT** operation allows the first process to release the monitor lock and await the desired condition. The **WAIT** is performed on a condition variable, which is associated by agreement with the actual condition needed. When another process makes that condition true, it will perform a **NOTIFY** on the condition variable, and the waiting process will continue from where it left off (after reacquiring the lock, of course.)

For example, consider a fixed block storage allocator providing two entry procedures: *Allocate* and *Free*. A caller of *Allocate* may find the free storage exhausted and be obliged to wait until some caller of *Free* returns a block of storage.

```

StorageAllocator: MONITOR =
  BEGIN
    StorageAvailable: CONDITION;
    Block: TYPE = RECORD [ . . . ];           -- or some other data type
    ListPtr: TYPE = LONG POINTER TO ListElmt;
    ListElmt: TYPE = RECORD [ block: Block, next: ListPtr ];
    FreeList: ListPtr;

    Allocate: ENTRY PROCEDURE RETURNS [p: ListPtr] =
      BEGIN
        WHILE FreeList = NIL DO
          WAIT StorageAvailable
        ENDLOOP;
        p ← FreeList; FreeList ← p.next;
      END;

    Free: ENTRY PROCEDURE [p: ListPtr] =
      BEGIN
        p.next ← FreeList; FreeList ← p;
        NOTIFY StorageAvailable
      END;
  END.

```

Note that it is clearly undesirable for two asynchronous processes to be executing in the *StorageAllocator* at the same time. The use of entry procedures for *Allocate* and *Free* assures mutual exclusion. The monitor lock is released while **WAITING** in *Allocate* in order to allow *Free* to be called (this also allows other processes to call *Allocate* as well, leading to several processes waiting on the queue for *StorageAvailable*).

9.2.2 Monitor locks

The most basic component of a monitor is its *monitor lock*. A monitor lock is a predefined type, which can be thought of as a small record:

```
MONITORLOCK: TYPE = PRIVATE RECORD [locked: BOOLEAN, queue: Queue];
```

The monitor lock is private; its fields are defined with default initialization and are never accessed explicitly by the Mesa programmer. Instead, it is used implicitly to synchronize entry into the monitor code, thereby authorizing access to the monitor data (and in some cases, other resources, such as I/O devices, etc.) The next section describes several kinds of monitors which can be constructed from this basic mechanism. In all of these, the idea is the same: during entry to a monitor, it is necessary to acquire the monitor lock by:

1. waiting (in the queue) until: *locked* = FALSE (default),
2. setting: *locked* ← TRUE.

9.2.3 Declaring monitor modules, ENTRY and INTERNAL procedures

In addition to a collection of data and an associated lock, a monitor contains a set of procedures that do operations on the data. *Monitor modules* are declared much like program or definitions modules; for example:

```
M: MONITOR [arguments] =
  BEGIN
  ...
  END.
```

The procedures in a monitor module are of three kinds:

```
Entry procedures
Internal procedures
External procedures
```

Every monitor has one or more *entry* procedures; these acquire the monitor lock when called, and are declared as:

```
P: ENTRY PROCEDURE [arguments] RETURNS [results] = ...
```

The entry procedures will usually comprise the set of public procedures visible to clients of the monitor module. (There are some situations in which this is not the case; see external procedures, below). The usual Mesa default rules for **PUBLIC** and **PRIVATE** procedures apply.

Many monitors will also have *internal* procedures: common routines shared among the several entry procedures. These execute with the monitor lock held, and may thus freely access the monitor data (including condition variables) as necessary. Internal procedures should be private, since direct calls to them from outside the monitor would bypass the acquisition of the lock (for monitors implemented as multiple modules, this is not quite right; see section 9.4, below). Internal procedures can be called only from an entry procedure or another internal procedure. They are declared as follows:

```
Q: INTERNAL PROCEDURE [arguments] RETURNS [results] = ...
```

The attributes **ENTRY** or **INTERNAL** may be specified only on a procedure in a **MONITOR** module (or on an **INLINE** procedure in a definitions module). Subsection 9.2.4 describes how one declares an interface for a monitor.

Some monitor modules may wish to have *external* procedures. These are declared as normal non-monitor procedures:

```
R: PROCEDURE [arguments] RETURNS [results] = ...
```


Such procedures are logically outside the monitor, but are declared within the same module for reasons of logical packaging. For example, a public external procedure might do some preliminary processing and then make repeated calls into the monitor proper (via a private entry procedure) before returning to its client. Being outside the monitor, an external procedure must *not* reference any monitor data (including condition variables), nor call any internal procedures. The compiler checks for calls to internal procedures and

usage of the condition variable operations (**WAIT**, **NOTIFY**, etc.) within external procedures, but does not check for accesses to monitor data.

Fine point:

Actually, non-changing read-only global variables *may* be accessed by external procedures; it is changeable monitor data that is strictly off-limits.

Generally speaking, a chain of procedure calls involving a monitor module has the general form:



Any deviation from this pattern is likely to be a mistake. A useful technique to avoid bugs and increase the readability of a monitor module is to structure the source text in the corresponding order:

```

M: MONITOR =
  BEGIN
    < External procedures >
    < Entry procedures >
    < Internal procedures >
    < Initialization (main-body) code >
  END.

```

9.2.4 Interfaces to monitors

In Mesa, the attributes **ENTRY** and **INTERNAL** are associated with a procedure's body, not with its type. Consequently, these attributes are associated with procedure bodies of **MONITOR** modules or **INLINE** procedures of definitions modules. Typically, internal procedures of a monitor module are not exported anyway, although they may be for a multi-module monitor (§ 9.4.4).

From the client side of an interface, a monitor appears to be a normal program module, hence the keywords **MONITOR** and **ENTRY** do not appear. For example, a monitor *M* with entry procedures *P* and *Q* might appear as:

```

MDefs: DEFINITIONS =
  BEGIN
    M: PROGRAM [arguments];
    P, Q: PROCEDURE [arguments] RETURNS [results];
    .
    .
  END.

```

9.2.5 Interactions of processes and monitors

One interaction should be noted between the process spawning and monitor mechanisms as defined so far. If a process executing within a monitor forked to an internal procedure of the same monitor, the result would be two processes inside the monitor at the same time, which is the exact situation that monitors are supposed to avoid. The following rule is therefore enforced:

A **FORK** may have as its target any procedure *except an internal procedure of a monitor*.

Fine points:

In the case of a multi-module monitor (§ 9.4.4) calls to other monitor procedures through an interface cannot be checked for the **INTERNAL** attribute, since this information is not available in the interface (§ 9.2.4).

The **JOIN** mechanism does not release any monitor locks its process holds while it waits for its argument process to complete. Executing a **JOIN** from within an **ENTRY** procedure of a monitor will result in deadlock if the process being joined calls **ENTRY** procedures of the same monitor. Consequently, **JOIN** statements within **ENTRY** procedures require careful consideration by the programmer and are not, in general, recommended.

9.3 Condition variables

Condition variables are declared as:

c: **CONDITION**;

All the fields of a condition variable are private to the process mechanism; condition variables may be accessed only via the operations defined below. It is important to note that it is the condition *variable* which is the basic construct; a condition (i.e., the contents of a condition variable) should *not* itself be thought of as a meaningful object; it may *not* be assigned to a condition variable, passed as a parameter, etc.

9.3.1 Wait, notify, and broadcast

A process executing in a monitor may find some condition of the monitor data which forces it to wait until another process enters the monitor and improves the situation. This can be accomplished using a condition variable, and the three basic operations: **WAIT**, **NOTIFY**, and **BROADCAST**, defined by the following syntax:

Statement ::= ... | **WaitStmt** | **NotifyStmt**

WaitStmt ::= **WAIT** Variable OptCatchPhrase

NotifyStmt ::= **NOTIFY** Variable | **BROADCAST** Variable

A condition variable *c* is always associated with some Boolean expression describing a desired state of the monitor data, yielding the general pattern:

Process waiting for condition:

```

WHILE ~BooleanExpression DO
  WAIT c
ENDLOOP;

```

Process making condition true:

```

  make BooleanExpression TRUE;           -- i.e., as side effect of modifying global
                                           -- data
  NOTIFY c;

```

Consider the storage allocator example from subsection 9.2.1. In this case, the desired **BooleanExpression** is "*FreeList # NIL*." There are several important points regarding **WAIT** and **NOTIFY**, some of which are illustrated by that example:

WAIT always releases the lock while waiting, in order to allow entry by other processes, including the process which will do the **NOTIFY** (e.g., *Allocate* must not lock out the caller of *Free* while waiting, or a deadlock will result). Thus, the programmer is always obliged to restore the monitor invariant (return the monitor data to a "good state") before doing a **WAIT**.

NOTIFY, on the other hand, retains the lock, and may thus be invoked *without* restoring the invariant; the monitor data may be left in an arbitrary state, so long as the invariant is restored before the next time the lock is released (by exiting an entry procedure, for example).

A **NOTIFY** directed to a condition variable on which no one is waiting is simply discarded. Moreover, the built-in test for this case is more efficient than any explicit test that the programmer could make to avoid doing the extra **NOTIFY**. (Thus, in the example above, *Free* always does a **NOTIFY**, without attempting to determine if it was actually needed.)

Each **WAIT** *must* be embedded in a loop checking the corresponding condition (e.g., *Allocate*, upon being notified of the *StorageAvailable* condition, still loops back and tests again to insure that the freelist is actually non-empty). This rechecking is necessary because the condition, even if true when the **NOTIFY** is done, may become false again by the time the awakened process gets to run. (Even though the freelist is always non-empty when *Free* does its **NOTIFY**, a third process could have called *Allocate* and emptied the freelist before the waiting process got a chance to inspect it.)

Given that a process awakening from a **WAIT** must be careful to recheck its desired condition, the process doing the **NOTIFY** can be somewhat more casual about insuring that the condition is actually true when it does the **NOTIFY**. This leads to the notion of a *covering condition variable*, which is notified whenever the condition desired by the waiting process is *likely* to be true; this approach is useful if the expected cost of false alarms (i.e., extra wakeups that test the condition and wait again) is lower than the cost of having the notifier always know precisely what the waiter is waiting for.

The last two points are somewhat subtle, but quite important; condition variables in Mesa act as *suggestions* that their associated Boolean expressions are likely to be true and should therefore be rechecked. They do *not* guarantee that a process, upon awakening

from a **WAIT**, will necessarily find the condition it expects. The programmer should never write code which implicitly assumes the truth of some condition simply because a **NOTIFY** has occurred.

It is often the case that the user will wish to notify *all* processes waiting on a condition variable. This can be done using:

```
BROADCAST c;
```

This operation can be used when several of the waiting processes should run, or when *some* waiting process should run, but not necessarily the head of the queue.

Consider a variation of the *StorageAllocator* example:

```
StorageAllocator: MONITOR =
  BEGIN
  StorageAvailable: CONDITION;
  ...

  Allocate: ENTRY PROCEDURE [size: CARDINAL] RETURNS [p:ListPtr] =
  BEGIN
  UNTIL <storage chunk of size words is available> DO
    WAIT StorageAvailable
  ENDLOOP;
  p ← <remove chunk of size words>;
  END;

  Free: ENTRY PROCEDURE [p:ListPtr, size: CARDINAL] =
  BEGIN
  ...
  <put back storage chunk of size words>
  ...
  BROADCAST StorageAvailable
  END;
END.
```

In this example, there may be several processes waiting on the queue of *StorageAvailable*, each with a different *size* requirement. It is not sufficient to simply **NOTIFY** the head of the queue, since that process may not be satisfied with the newly available storage while another waiting process might be. This is a case in which **BROADCAST** is needed instead of **NOTIFY**.

An important rule of thumb: *it is always correct to use a **BROADCAST**. **NOTIFY** should be used instead of **BROADCAST** if both of the following conditions hold:*

It is expected that there will typically be several processes waiting in the condition variable queue (making it expensive to notify all of them with a **BROADCAST**), and

It is known that the process at the head of the condition variable queue will always be the right one to respond to the situation (making the multiple notification unnecessary);

If both of these conditions are met, a **NOTIFY** is sufficient, and may represent a significant efficiency improvement over a **BROADCAST**. The allocator example in subsection 9.2.1, in which fixed length blocks are allocated, is a situation in which **NOTIFY** is preferable to **BROADCAST**.

As described above, the condition variable mechanism, and the programs using it, are intended to be robust in the face of "extra" **NOTIFY**s. The next section explores the opposite problem: "missing" **NOTIFY**s.

Fine point:

When a program **WAIT**s, it releases the monitor lock. When it returns from the **WAIT**, it reacquires the lock. The address of the condition variable has to be calculated twice. If this address is obtained by a complicated expression, there is a subtle restriction. The address calculation cannot do a **WAIT** in the same process. In other words, consider the procedure

CondProc: PROCEDURE RETURNS [POINTER TO CONDITION];

If a program contains the statement

WAIT *CondProc*[] ↑

then the execution of *CondProc* cannot **WAIT**.

9.3.2 Condition variable timeouts

One potential problem with waiting on a condition variable is the possibility that one may wait "too long." There are several ways this could happen, including:

- Hardware error (e.g., "lost interrupt")
- Software error (e.g., failure to do a **NOTIFY**)
- Communication error (e.g., lost packet)

To handle such situations, waits on condition variables are allowed to *time out*. This is done by associating a *timeout interval* with each condition variable, which limits the delay that a process can experience on a given **WAIT** operation. If no **NOTIFY** has arrived within this time interval, one will be generated automatically.

The timeout of a condition variable is not client-visible, but it may be changed by calling procedures in the **Process** interface. When a **CONDITION** variable comes into existence, its timeout is initialized to "infinity," which disables timeouts. Programs that rely on condition variables timing out could call the procedure **Process.SetTimeout** (see the *Pilot Programmer's Manual* for details). To change the timeout back to "infinity," a program should call **Process.DisableTimeouts**.

The waiting process will perceive a timeout as a normal **NOTIFY**. (Some programs may wish to distinguish timeouts from normal **NOTIFY**s; this requires checking the time as well as the desired condition on each iteration of the loop.)

No facility is provided to time out waits for monitor locks. This is because there would be, in general, no way to recover from such a timeout.

9.4 More about monitors

The next few sections deal with the full generality of monitor locks and monitors.

9.4.1 The `LOCKS` clause

Normally, a monitor's data comprises its global variables, protected by the special global variable `LOCK`:

```
LOCK: MONITORLOCK;
```

This implicit variable is declared automatically in the global frame of any module whose heading is of the form:

```
M: MONITOR [arguments]
  IMPORTS ...
  EXPORTS ... =
```

In such a monitor it is generally not necessary to mention `LOCK` explicitly at all. For more general use of the monitor mechanism, it is necessary to declare at the beginning of the monitor module exactly which `MONITORLOCK` is to be acquired by entry procedures. This declaration appears as part of the program type constructor that is at the head of the module. The syntax is as follows:

```
ProgramTC      ::= ... | MONITOR ParameterList ReturnsClause LocksClause

LocksClause    ::= empty | LOCKS Expression |
                 LOCKS Expression USING identifier : TypeSpecification
```

If the `LocksClause` is empty, entry to the monitor is controlled by the distinguished variable `LOCK` (automatically supplied by the compiler). Otherwise, the `LocksClause` must designate a variable of type `MONITORLOCK`, a record containing a distinguished lock field (§ 9.4.2), or a pointer that can be dereferenced (perhaps several times) to yield one of the preceding. If a `LocksClause` is present, the compiler does not generate the variable `LOCK`.

If the `USING` clause is absent, the lock is located by evaluating the `LOCKS` expression in the context of the monitor's main body; i.e., the monitor's parameters, imports, and global variables are visible, as are any identifiers made accessible by a global `OPEN`. Evaluation occurs upon entry to, and again upon exit from, the entry procedures (and for any `WAITS` in entry or internal procedures). The location of the designated lock can thus be affected by assignments within the procedure to variables in the `LOCKS` expression. To avoid disaster, it is essential that each reevaluation yield a designator of the same `MONITORLOCK`. This case is described further in subsection 9.4.4.

If the `USING` clause is present, the lock is located in the following way: every entry or internal procedure must have a parameter with the same identifier and a compatible type as that specified in the `USING` clause. The occurrences of that identifier in the `LOCKS` clause are bound to that procedure parameter in every entry procedure (and internal procedure

doing a **WAIT**). The same care is necessary with respect to reevaluation; to emphasize this, the distinguished argument is treated as a read-only value within the body of the procedure. See subsection 9.4.5 for further details.

9.4.2 Monitored records

For situations in which the monitor data cannot simply be the global variables of the monitor module, a *monitored record* can be used:

```
r: MONITORED RECORD [x: INTEGER, . . . ];
```

A monitored record is a normal Mesa record, except that it contains an automatically declared field of type **MONITORLOCK**. As usual, the monitor lock is used implicitly to synchronize entry into the monitor code, which may then access the other fields in the monitored record. The fields of the monitored record must *not* be accessed except from within a monitor which first acquires its lock. In analogy with the global variable case, the monitor lock field in a monitored record is given the special name *LOCK*; generally, it need not be referred to explicitly (except during initialization; § 9.6).

Fine point:

A more general form of monitor lock declaration is discussed in subsection 9.4.6

CAUTION: If a monitored record is to be passed around (e.g., as an argument to a procedure) this should always be done by reference using a **LONG POINTER TO MONITORED RECORD**. Copying a monitored record (e.g., passing it by value) will generally lead to chaos.

The assignment operation is not available for updating objects containing **MONITORLOCK** or **CONDITION** values: updating of such objects must be done component-by-component.

9.4.3 Monitors and module instances

Even when all the procedures of a monitor are in one module, it is not quite correct to think of the module and the monitor as identical. For one thing, a monitor module, like an ordinary program module, may have several instances. In the most straightforward case, each instance constitutes a separate monitor. More generally, through the use of monitored records, the number of monitors may be larger or smaller than the number of instances of the corresponding module(s). The crucial observation is that in all cases:

There is a one-to-one correspondence between monitors and monitor locks.

The generalization of monitors through the use of monitored records tends to follow one of two patterns:

Multi-module monitors, in which several module instances implement a single monitor.

Object monitors, in which a single module instance implements several monitors.

Fine point:

These two patterns are *not* mutually exclusive; multi-module object monitors are possible, and may occasionally prove necessary.

9.4.4 Multi-module monitors

In implementing a monitor, the most obvious approach is to package all the data and procedures of the monitor within a single module instance (if there are multiple instances of such a module, they constitute separate monitors and share nothing except code). While this will doubtless be the most common technique, the monitor may grow too large to be treated as a single module.

Typically, this leads to multiple modules. In this case the mechanics of constructing the monitor are changed somewhat. There must be a central location that contains the monitor lock for the monitor implemented by the multiple modules. This can be done either by using a **MONITORED RECORD** or by choosing one of the modules to be the “root” of the monitor. Consider the following example:

```

BigMonRoot: MONITOR IMPORTS ... EXPORTS ... =
  BEGIN
    monitorDatum1: ...
    monitorDatum2: ...
    ...
    p1: PUBLIC ENTRY PROCEDURE ...
    ...
  END.

BigMonA: MONITOR
  LOCKS root -- could equivalently say root.LOCK
  IMPORTS root: BigMonRoot ... EXPORTS ... SHARES BigMonRoot =
  BEGIN
    ...
    p2: PUBLIC ENTRY PROCEDURE ...
      x ← root.monitorDatum1; -- access the protected data of the monitor
    ...
  END.

BigMonB: MONITOR
  LOCKS root
  IMPORTS root: BigMonRoot ... EXPORTS ... SHARES BigMonRoot =
  BEGIN OPEN root;
  ...
    p3: PUBLIC ENTRY PROCEDURE ...
      monitorDatum2 ← ...; -- access the protected data via an OPEN
  ...
  END.

```

The monitor *BigMon* is implemented by three modules. The modules *BigMonA* and *BigMonB* have a **LOCKS** clause to specify the location of the monitor lock: in this case, the distinguished variable *LOCK* in *BigMonRoot*. When any of the entry procedures *p1*, *p2*, or *p3* is called, this lock is acquired (waiting if necessary), and is released upon returning.

The reader can verify that no two independent processes can be in the monitor at the same time.

Note that since the *LOCK* field is private in *BigMonRoot*, the modules *BigMonA* and *BigMonB* must **SHARE** *BigMonRoot*. Another way to accomplish access to the lock would be to specify a **PUBLIC GlobalAccess** (§ 7.5) for *BigMonRoot*.

Fine point:

Using this form of **LOCKS** clause for a multi-module monitor, introduces undesirable compilation dependencies. In particular, whenever *BigMonRoot* is recompiled, so must the rest of the modules implementing the monitor. This problem can be avoided by exporting an explicitly declared **MONITORLOCK** (§ 9.4.6) to an interface shared by the implementing modules, and locking that **MONITORLOCK**.

Another means of implementing multi-module monitors is by means of a **MONITORED RECORD**. Use of **OPEN** allows the fields of the record to be referenced without qualification. Such a monitor is written as:

```

MonitorData: TYPE = MONITORED RECORD [x: INTEGER, ... ];

MonA: MONITOR [pm: LONG POINTER TO MonitorData]

    LOCKS pm
    IMPORTS ...
    EXPORTS ... =
BEGIN OPEN pm;
P: ENTRY PROCEDURE [...] =
    BEGIN
    ...
    x ← x + 1;           -- access to a monitor variable
    ...
    END;
...
END.

```

The **LOCKS** clause in the heading of this module (and each other module of this monitor) leads to a **MONITORED RECORD**. Of course, in all such multi-module monitors, the **LOCKS** clause will involve one or more levels of indirection (**POINTER TO MONITORED RECORD**, etc.) since passing a monitor lock by value is not meaningful. As usual, Mesa will provide one or more levels of automatic dereferencing as needed.

More generally, the target of the **LOCKS** clause can evaluate to a **MONITORLOCK** (i.e., the example above is equivalent to writing "**LOCKS pm.LOCK**").

CAUTION: The meaning of the target expression of the **LOCKS** clause *must not change* between the call to the entry procedure and the subsequent return (i.e., in the above example, changing *pm* would invariably be an error) since this would lead to a different monitor lock being released than was acquired, resulting in total chaos.

There are a few other issues regarding multi-module monitors which arise any time a tightly coupled piece of Mesa code must be split into multiple module instances and then spliced back together. For example:

If the lock is in a **MONITORED RECORD**, the monitor data will probably need to be in the record also. While the global variables of such a multi-module monitor are covered by the monitor lock, they do *not* constitute monitor data in the normal sense of the term, since they are not uniformly visible to all the module instances.

Making the internal procedures of a multi-module monitor **PRIVATE** will not work if one module wishes to call an internal procedure in another module. (Such a call is perfectly acceptable so long as the caller already holds the monitor lock.) Instead, a second interface (hidden from the clients) is needed as part of the "glue" holding the monitor together. Note however, that Mesa cannot currently check that the procedure being called through the interface is an internal one (§ 9.2.4).

9.4.5 Object monitors

Some applications deal with *objects*, implemented, say, as records named by pointers. Often it is necessary to insure that operations on these objects are *atomic*, i.e., once the operation has begun, the object will not be otherwise referenced until the operation is finished. If a module instance provides operations on some class of objects, the simplest way of guaranteeing such atomicity is to make the module instance a monitor. This is logically correct, but if a high degree of concurrency is expected, it may create a bottleneck; it will serialize the operations on *all* objects in the class, rather than on *each* of them individually. If this problem is deemed serious, it can be solved by implementing the objects as monitored records, thus effectively creating a separate monitor for each object. A single module instance can implement the operations on all the objects as entry procedures, each taking as a parameter the object to be locked. The locking of the parameter is specified in the module heading via a **LocksClause** with a **USING** clause. For example:

```

ObjectRecord: TYPE = MONITORED RECORD [ ... ];

ObjectHandle: TYPE = LONG POINTER TO ObjectRecord;

ObjectManager: MONITOR [arguments]
  LOCKS object USING object: ObjectHandle
  IMPORTS ...
  EXPORTS ... =
  BEGIN
    Operation: PUBLIC ENTRY PROCEDURE [object: ObjectHandle, ... ] =
      BEGIN
        ...
      END;
    ...
  END.

```

Note that the argument of **USING** is evaluated in the scope of the arguments to the entry procedures, rather than the global scope of the module. In order for this to make sense, each entry procedure, and each internal procedure that does a **WAIT**, must have an argument which matches exactly the name and type specified in the **USING** subclause. All other components of the argument of **LOCKS** are evaluated in the global scope, as usual.

As with the simpler form of `LOCKS` clause, the target may be a more complicated expression and/or may evaluate to a monitor lock rather than a monitored record. For example:

```
LOCKS p.q.LOCK USING p: LONG POINTER TO ComplexRecord . . .
```

CAUTION: Again, the meaning of the target expression of the `LOCKS` clause *must not change* between the call to the entry procedure and the subsequent return (i.e., in the above example, changing `p` or `p.q` would almost surely be an error).

CAUTION: It is important to note that global variables of object monitors are very dangerous; they are *not* covered by a monitor lock, and thus do *not* constitute monitor data. If used at all, they must be set only at module initialization time and must be read-only thereafter.

9.4.6 Explicit declaration of monitor locks

It is possible to declare monitor locks explicitly:

```
myLock: MONITORLOCK;
```

The normal cases of monitors and monitored records are essentially stylized uses of this facility via the automatic declaration of `LOCK`, and should cover all but the most obscure situations. For example, explicit declarations are useful in defining `MACHINE DEPENDENT` monitored records. (Note that the `LOCKS` clause becomes mandatory when an explicitly declared monitor lock is used.) More generally, explicit declarations allow the programmer to declare records with several monitor locks, declare locks in local frames, and so on; this flexibility can lead to a wide variety of subtle bugs, hence use of the standard constructs whenever possible is strongly advised.

9.4.7 Inline `ENTRY` procedures

The syntax for definitions modules allows the specification of a `LOCKS` clause. This is to allow inline `ENTRY PROCEDURES` to be declared in the interface. In order for this to make sense, the monitor lock must be an interface variable, or the procedures must deal with an object style monitor. No special restrictions (other than those that apply to all `INLINE` bodies) need be met when declaring inline `ENTRY PROCEDURES` within the program module of a monitor.

9.5 Signals

9.5.1 Signals and processes

Each process has its own call stack, down which signals propagate. If the signaller scans to the bottom of the stack and finds no catch phrase, the signal is propagated to the debugger. The important point to note is that forking to a procedure is different from calling it, in that the forking creates a gap across which signals cannot propagate. This implies that in practice, one cannot casually fork to any arbitrary procedure. The only suitable targets for forks are procedures which catch any signals they incur, and which never generate any signals of their own.

9.5.2 Signals and monitors

Signals require special attention within the body of an entry procedure. A signal raised with the monitor lock held will propagate without releasing the lock and possibly invoke arbitrary computations. For errors, this can be avoided by using the `RETURN WITH ERROR` construct.

```
RETURN WITH ERROR NoSuchObject;
```

Recall from chapter 8 that this statement has the effect of removing the currently executing frame from the call chain before issuing the `ERROR`. If the statement appears within an entry procedure, the monitor lock is released before the error is started as well. Naturally, the monitor invariant must be restored before this operation is performed.

For example, consider the following program segment:

```
Failure: ERROR [kind: CARDINAL] = CODE;

Proc: ENTRY PROCEDURE [...] RETURNS [c1, c2: CHARACTER] =
  BEGIN
    ENABLE UNWIND = > ...
    ...
    IF cond1 THEN ERROR Failure[1];
    IF cond2 THEN RETURN WITH ERROR Failure[2];
    ...
  END;
```

Execution of the construct `ERROR Failure[1]` raises a signal that propagates until some catch phrase specifies an exit. At that time, unwinding begins; the catch phrase for `UNWIND` in *Proc* is executed and then *Proc*'s frame is destroyed. Within an entry procedure such as *Proc*, the lock is held until the unwind (and thus through unpredictable computation performed by catch phrases).

Fine point:

If the body of the `ENTRY` procedure has an `EXITS` clause, the scope of the `ENABLE` on the body does not include the `EXITS` clause (§ 8.2.1). Consequently, if an error is raised in the `EXITS` clause of an `ENTRY` procedure's outer block, an `ENABLE UNWIND` clause on the body of this procedure will not be invoked and the monitor lock will not be released.

Execution of the construct `RETURN WITH ERROR Failure[2]` releases the monitor lock and destroys the frame of *Proc* before propagation of the signal begins. Note that the argument list in this construct is determined by the declaration of *Failure* (not by *Proc*'s `RETURNS` clause). The catch phrase for `UNWIND` is not executed in this case. The signal *Failure* is actually raised by the system, after which *Failure* propagates as an ordinary error (beginning with *Proc*'s caller).

When the `RETURN WITH ERROR` construct is used from within an internal procedure, the monitor lock is *not* released; `RETURN WITH ERROR` will release the monitor lock in precisely those cases that `RETURN` will.

Another important issue regarding signals is the handling of **UNWIND**s; any entry procedure that may experience an **UNWIND** must catch it and clean up the monitor data (restore the monitor invariant):

```
P: ENTRY PROCEDURE [ . . . ] =
  BEGIN ENABLE UNWIND = > BEGIN <restore invariant> END;
  .
  .
  .
  END;
```

At the end of the outermost **UNWIND** catch phrase, the compiler will append code to release the monitor lock before the frame is unwound. It is important to note that a monitor always has at least one cleanup task to perform when catching an **UNWIND** signal: the monitor lock must be released. To this end, the programmer should be sure to place an enable-clause on the body of every entry procedure that might evoke an **UNWIND** (directly or indirectly). If the monitor invariant is already satisfied, no further cleanup need be specified, but *the null catch-phrase must be written* so that the compiler will generate the code to unlock the monitor:

```
BEGIN ENABLE UNWIND = > NULL;
```

This should be omitted *only* when it is certain that no **UNWIND**s can occur.

Another point is that signals caught by the **OptCatchPhrase** of a **WAIT** operation should be thought of as occurring after reacquisition of the monitor lock. Thus, like all other monitor code, catch phrases within a monitor are always executed with the monitor lock held.

9.6 Initialization

When a new monitor comes into existence, its monitor data must be set to some appropriate initial values; in particular, the monitor lock and any condition variables must be initialized. Mesa takes responsibility for initializing them by the defaulting mechanism (§ 3.7).

Monitor data in global variables can be initialized using the normal Mesa initial value constructs in declarations. Monitor locks and condition variables in the global frame will also be initialized automatically (although in this case, the programmer does not write any explicit initial value in the declaration).

Records allocated from an **UNCOUNTED_ZONE** using **z.NEW** will have any fields of type **MONITORLOCK** or **CONDITION** properly initialized. Any other fields can be set at the time of creation by either field defaults or by having an initialization constructor within the application of **NEW**. It is illegal to set the value of such a record with a constructor after initialization since the running system may have queues threaded through the private fields of the **MINOTORLOCK** or **CONDITION**. After initialization, such records must be updated field-by-field rather than with constructors.

Fine point:

If a record containing fields of type **MONITORLOCK** or **CONDITION** is allocated from some free storage managed by means other than the **NEW** construct, its fields must be explicitly set by calls on operating system supplied interfaces. See the *Pilot Programmer's Manual* for descriptions of the appropriate procedures.

Since initialization code modifies the monitor data, it must have exclusive access to it. The programmer should insure this by arranging that the monitor not be called by its client processes until it is ready for use.



A

Pronouncing Mesa

The following suggestions may be helpful in reading Mesa programs:

For	Read
= >	chooses
←	gets
$n: T$	n is a T
$m.field$	m 's <i>field</i>
$p \uparrow$	p 's referent
@ x	address of x
[$a..b$]	(the interval) a through b
[$a..b$]	(the interval) a up to b
($a..b$)	(the interval) above a through b
($a..b$)	(the interval) above a up to b
FOR $i \leftarrow j, k \dots$	for i getting first j , thereafter $k \dots$
$f [x, y, z]$	f of x, y and z
!	enabling

We leave as an exercise for the reader the following statement, attributed to Oscar Hammerstein II.

$i \leftarrow weary$ AND Sick [*trying*];



B

Programming conventions

This appendix proposes some style conventions for Mesa programs. Style conventions are valuable since they make it easier for individuals to understand other programmers' code. The conventions described here have been in general use within the Systems Development Department of Xerox for a number of years and have served very well. The conventions cannot be hard and fast rules, and there can be compelling reasons for not following some convention in a particular instance. Nevertheless, to the extent that we all usually follow them, it will help us in reading (and therefore in modifying) one another's programs. The recommended conventions are summarized below. Each convention is given as a short rule which may be followed by some examples of its application.

The Mesa compiler only uses blanks, TABs and carriage returns as separators for basic lexical units such as identifiers; extra ones do not hurt. Furthermore, it allows you to write identifiers in any combination of upper and lower case letters: the identifiers *Alpha*, *ALPHA*, *alpha* and *Alpha* are all legal (but different) identifiers.

B.1 Names

B.1.1 Capitalization

Capitalize the first letter of a name if it identifies a module (either interface or implementation), procedure, signal, type, or label, otherwise the first letter is lower case. Variables, record components, enumerated type literals, and constants should all have a lower case first letter:

```
Space: DEFINITIONS = ...  
Factorial: PROC [i: LONG INTEGER] RETURNS [LONG INTEGER];  
Complex: TYPE = RECORD [real, imag: REAL];  
EndOfStream: SIGNAL [nextIndex: CARDINAL];  
Handle: TYPE = LONG POINTER TO Object;  
nullHandle: Handle = NIL;  
Color: TYPE = {red, yellow, blue};
```

Capitalize the first letter of each embedded word of a multi-word name:

```
EndOfStream: SIGNAL [nextIndex: CARDINAL];
nullHandle: Handle = NIL;
```

Case shift alone should not be used to distinguish identifiers; in general different identifiers should differ in at least two characters. There are some exceptions: if one has a type, *Foo*, it is acceptable to declare a variable of that type as *foo*.

```
badID, badid, bADid, BADid: LONG INTEGER;           -- bad!

handle: Handle;                                     -- acceptable

h: Handle;                                          -- acceptable
```

B.1.2 Qualification

Identifiers from interfaces should be qualified by their interface names or abbreviations for them:

```
DIRECTORY
  Exec, Format, TTY;
SimpleExample: PROGRAM
  IMPORTS Exec, Format, TTY =
  ...
  ReverseName: Exec.ExecProc = BEGIN ... END;
  execOut: Format.StringProc;
```

Renaming an interface in an `OPEN` or `IMPORTS` clause is acceptable; if you do so, use the assigned name throughout the module.

```
DIRECTORY
  Exec, Format, TTY;
SimpleExample: PROGRAM
  IMPORTS E: Exec, Format, TTY =
  ...
  ReverseName: E.ExecProc = BEGIN ... END;
```

The renaming form of `OPEN` should be used instead of the unqualified form. There are virtually no valid reasons for using an unqualified `OPEN` and many compelling reasons not to. In the past, programmers would `OPEN` an interface over a limited scope (e.g., a procedure or block) within which identifiers from that interface are used heavily, but this practice is now discouraged in favor of using a renaming open, possibly with a very brief new name.

```
DIRECTORY
  Ascii, TTY;
SimpleExample: PROGRAM
  IMPORTS TTY =
  BEGIN
  OPEN A: Ascii;
  h: TTY.Handle;
  ...
```

```

TTY.PutChar[h, A.ControlG];  -- ring the bell
...
END.

```

B.1.3 Module naming

All module source file names have the extension ".mesa".

```

ObjectSupport.mesa
ObjectSupportImpl.mesa
ListSortRef.mesa
SimpleExample.mesa

```

A **DEFINITIONS** module does not need any standard suffix on its name.

```

NSString: DEFINITIONS = ...
TIP: DEFINITIONS = ...
Space: DEFINITIONS = ...
MFile: DEFINITIONS = ...

```

A **DEFINITIONS** module that contains an interface and types used internally in a package but not exported to the world-at-large often has a name that ends in "Ops."

```

WindowOps: DEFINITIONS = ...
TajoOps: DEFINITIONS = ...

```

A **PROGRAM** or **MONITOR** module name should have the suffix "Impl" or "Pack" if its name without the suffix would conflict with the name of a **DEFINITIONS** module.

```

ObjectSupport: DEFINITIONS = ...
ObjectSupportImpl: PROGRAM      -- name would conflict with
  EXPORTS ObjectSupport = ...   -- ObjectSupport without Impl

ObjectMachinery: PROGRAM        -- name doesn't conflict with ObjectSupport
Fun: PROGRAM =                  -- name conflicts with nothing

```

A **CONFIGURATION** file name should have the extension ".config".

```

FileSystem.config
Spy.config
Compiler.config

```

A **CONFIGURATION** module does not need any standard suffix on its name.

```

FileSystem: CONFIGURATION
  EXPORTS MFile = ...
Spy: CONFIGURATION
  EXPORTS SpyOps = ...
Compiler: CONFIGURATION
  EXPORTS CompilerOps = ...

```

B.2 Types

Name types, subranges, enumerated types, and records rather than defining their anonymous counterparts. This is especially true in definitions modules.

```
DeckIndex: TYPE = {0..52};                -- yes
CardDeck: TYPE = ARRAY DeckIndex OF Card;
CardDeck: TYPE = ARRAY {0..52} OF Card;    -- no
```

Use closed lower bounds starting at 0 and open upper bounds for interval types.

```
[0..NumberOfItems)                        -- preferred form
FOR i IN [0..n) DO IF a[i] = a[n] THEN ... ENDLOOP;
```

Use positional notation for single-component argument lists, record constructors and extractors *provided the arguments have appropriate names and incompatible types*.

```
MFile.Release [file];
execOut.Number [column, columnFormat];
```

Use a keyword form if the arguments are inappropriately named (and you can't change the names), or if there are manifest constants which are not self explanatory (numbers, TRUE, FALSE):

```
String.StringToNumber[s: Name, radix: 10]; -- yes
String.StringToNumber[name, 10];           -- no
Time.AppendCurrent[s: string, zone: TRUE]; -- yes
Time.AppendCurrent[string, TRUE];          -- no
```

Use the keyword form when there are two or more constituents if some constituents have equivalent types. Keyword form should be used when most components are being defaulted and only a few given values:

```
stream.Put{
  block: [blockPointer: p, startIndex:0, stopIndexPlusOne: nBytes],
  endRecord: FALSE};
[parent: parent, mapped: mapped] ← Space.GetAttributes[space]; -- yes
[parent,,,,,mapped] ← Space.GetAttributes[space];              -- no
```

Name procedure parameters in the definitions of procedure types in DEFINITIONS modules:

```
ExecProc: TYPE = PROC[h:Handle, cleintData: LONG POINTER ← NIL]
  RETURNS [outcome: Outcome ← normal];
```

Name procedure result fields, especially if there is more than one field or one of the fields has boolean type (naming can define the sense of the boolean).

```
GetTimes: PROC[h: Handle] RETURNS [create, write, read: Time.packed];
FindItem: PROC[k: Key] RETURNS [v: Value, found: BOOLEAN];
```

It is acceptable to omit the name of a single result when its meaning is perfectly clear:

```
IsReady: PROC [h: Handle] RETURNS [BOOLEAN];    -- PROC name sets sense of BOOLEAN
```

Procedures with a single argument returning a single result may be named using the following convention:

```
<ResultTypeName>From<ArgumentTypeName>: PROC [a: ArgumentTypeName]
                                         RETURNS [r: ResultTypeName]; e.g.,
PageFromLongPointer: PROC [ptr: LONG POINTER] RETURNS [page: Page];
```

B.3 Exceptions: SIGNALS and ERRORS

B.3.1 General

Use **ENDCASE** only to generate an **ERROR** or to treat "none of the above," but not to handle specific cases.

```
Color: TYPE = {red, orange, yellow, green, blue, violet};
c: Color;
button: [0..2];
invalidButtonColor: ERROR = CODE;
```

```
button ← SELECT c FROM
    red => 0,
    yellow => 1,
    blue => 2,
    ENDCASE => ERROR invalidButtonColor;
```

Use **SIGNAL** or **ERROR** when generating a **SIGNAL** or **ERROR**;

```
SIGNAL foo;                                -- yes
ERROR gliche;                              -- yes
foo;                                       -- no
```

Avoid using the anonymous **ERROR**.

Signals that return results should pass enough information such that any catcher on the stack can react to the signal. If a catcher wishes to resolve a problem it should have enough information to do so.

```
BufferTooSmall: SIGNAL [offender: Buffer, lengthNeeded: CARDINAL]
    RETURNS [newBuffer: Buffer];
```

Care should be taken with signals that move data structures to be sure that intermediate procedure instances don't hold onto the address of the old one. See the discussion in subsection 8.2.5.

B.3.2 In DEFINITIONS modules

Declare **ERRORS** as **ERRORS** and **SIGNALS** as **SIGNALS**. Don't declare a **SIGNAL** and then generate an **ERROR** with it at run time.

Try to define a small number of **ERRORS** and **SIGNALS** in an interface. Normally define a single **ERROR** or **SIGNAL** with a parameter to distinguish the exact reason for it. Only use separate signals when they must have different types for reasons of resumption.

```
Error: ERROR [stream: Handle, code: ErrorCode];
ErrorCode: TYPE = {invalidHandle, indexOutOfRange, invalidOperation,
  fileTooLong, fileNotAvailable, invalid, other};
IOSignal: SIGNAL [sc: SignalCode];
SignalCode: TYPE = {emptyBuffer, unmatchedLeftParen, unmatchedStringDelim,
  rubout, unprintableValue, typeMismatch, unknownFormat};
StringBoundsFault: SIGNAL [text: LONG STRING] RETURNS [LONG STRING];
```

B.3.3 In PROGRAM modules

Only exceptions that are part of the abstraction should emanate from a module. Don't let exceptions that are "part of" the interfaces of invoked routines escape: handle them completely within your abstraction, or transform them into exceptions that are part of your interface (i.e., insulate your clients from details of your implementation).

```
... inputNumber ← SelectNumber[
  !String.InvalidNumber => [inputNumber ← 0; CONTINUE]];
```

Private signals may be used for error conditions which can not be handled by the client:

```
LogicError: PRIVATE ERROR = CODE;
...
IF Factorial[4] # 24 THEN ERROR LogicError;
```

B.4 Module histories

Introduce the module with one or more comment lines giving the name of the module, the name of the person who last edited it, and the time when the edit was made.

Very Brief:

```
-- Foo.Mesa                                21-Apr-83 By DKnutsen
```

Brief:

```
-- FILE: OnlineMergSortRefImpl.mesa
-- Last Edited by MBrown on March 9, 82 5:02 pm
```

Keep a **CHANGE LOG** at the end of each source module for keeping track of who first created the module, who has changed it, why it was changed, and when.

```
...
END.                                -- space.mesa
```

CHANGE LOG

```
17-Mar-82 17:11:48  DKnutsen   Created file from Space.mesa of 8-Aug-81
23-Aug-82 11:43:17  Luniewski  Mods for long page numbers and counts
31-Mar-83  9:50:11  DKnutsen   Result of Map and MapAt was named wrong. Get Page< = >LongPointer from Environment.
6-Jul-83   17:2:42  DKnutsen   ActivateProc arg is PROCEDURE, not UNSPECIFIED.
```

B.5 Documentation of definitions modules

If a definitions module has no separate documentation, the documentation should be in the module itself immediately following each item defined. This is typical and convenient for private definitions files, where the module may be recompiled if it is necessary to update the documentation. Such a self-documenting definitions file should have complete descriptions of procedures and signals, with attention to special cases and error conditions and actions. The module should also contain a conceptual overview, as necessary.

```
Sort: PROC [
    itemList: ItemList,
    compare: Proc[Item, Item] RETURNS [Comparison]] RETURNS [ItemList];

-- Destructive sort of itemList; returns sorted list containing the same
-- items. The order of equal items is not preserved.
```

If a definitions module has separate documentation, the declarations in the module may have either no embedded documentation or perhaps terse notes and reminders. This is typical and convenient for public definitions files, where recompiling the interface to correct or improve the documentation would be unacceptable.

B.6 Module organization

Significant parameters in the implementation should be named and declared first in the module.

In general, organize procedures alphabetically. If there are important reasons for grouping procedures, alphabetize within groups.

B.7 Layout

Write statements one per line, unless several simple statements *which together perform a single function* will fit on one line.

Indent the labels of a **SELECT** (including the **ENDCASE**) one level, and the statements a second level (unless a statement will fit on the same line with the label).

```

SELECT predecessor FROM
  NIL => {listElt.next ← list.first; list.first ← listElt};
  list.last =>
    BEGIN
      list.last.next ← listElt;
      list.last ← listElt;
      listElt.next ← NIL;
    END;
  ENDCASE =>
    BEGIN
      listElt.next ← predecessor.next;
      predecessor.next ← listElt;
    END;

```

Indent one level for the statement following a **THEN** or **ELSE** (unless it fits on the same line). Put **THEN** on the same line with **IF**, and don't indent **ELSE** with respect to **IF**. If the statement following **ELSE** is another **IF**, write both on the same line.

```

IF (order = userSupplied) THEN
  BEGIN
    IF (orderPredicate = NIL) OR (matchPredicate = NIL) THEN RETURN[NIL];
  END
ELSE IF matchPredicate = NIL THEN matchPredicate ← Equal;

```

Indent one level for each compound **BEGIN-END**, **DO-ENDLOOP**, or bracket pair in a record declaration.

When the rules for **IF** and **SELECT** call for indenting a statement, do not indent an extra level for a **BEGIN**.

It is fine to put a compound statement or loop on a single line if it will fit.

If a statement won't fit on a single line, indent the continuation line(s) by one level.

Among other things, these rules have the property that they allow a program to be easily converted to a form in which the bracketing is implied by the indentation.

Running a source file through the Formatter will give it a "standard" indentation.

B.8 Spaces

The following rules for spaces should be broken when necessary, but are a good general guide:

A space after a comma, semicolon, or colon, and none before

No spaces immediately inside brackets or parentheses

No spaces around single-character operations: *****, **-**, etc., except for **←**.


```

SELECT i*j + k FROM
  1, IN [7..10] => {var1 ← 10, var2 ← 20};
  2, 5, > 10 => Statement1;
ENDCASE;

```

B.9 Miscellaneous

You may use { . . . } for blocks of up to three lines; they typically should only be used on innermost blocks. Use **BEGIN END** for longer and enclosing blocks.

Brackets should always be used when calling procedures with no parameters.

```

procWithNoParameters;           -- no
procWithNoParameters[];       -- yes

```

Do not use ↑ with the "." or array indexing operators. Let the compiler do the dereferencing for you.

```

Person: TYPE = RECORD
  [
    age:{0..200},
    sex: {male, female},
    party: {Democratic, Republican}
  ];
winner: POINTER TO Person;

winner.party ← Democratic;    -- winner ↑ .party ← Democratic;

actualArray: ARRAY [0..20] OF INTEGER;
arrayPtr: POINTER TO ARRAY [0..20] OF INTEGER ← @actualArray;
arrayFinger: POINTER TO POINTER TO ARRAY [0..20] OF INTEGER ← @arrayPtr;
actualArray[1] ← 3;           -- two equivalent statements follow
arrayPtr[1] ← 3;             -- arrayPtr ↑ [1] ← 3
arrayFinger[1] ← 3;         -- arrayFinger ↑ ↑ [1] ← 3

```

Use **SELECT TRUE** instead of a long chain of **IF...ELSE IF... ELSE IF. . .**;

```

SELECT TRUE FROM
  condition1 => {lengthy consequence1};
  condition2 => {lengthy consequence2};
  . . .
ENDCASE => BEGIN ... END;

```

Always use **USING** in the **DIRECTORY** clause. It is sometimes convenient to omit **USING** clauses when the program is undergoing massive development, and then use a program called the Lister to generate them when the program stabilizes. See the *Mesa User's Guide* for instructions on how to run the Lister.



C

Mesa machine dependencies

This appendix contains a number of machine-dependent constants and definitions for the implementation of Mesa.

C.1 Numeric limits

The numeric limits are the following:

FIRST [INTEGER] = -32768 = -2^{15} and has internal representation 100000B
LAST [INTEGER] = 32767 = $2^{15}-1$ and has internal representation 077777B
LAST [CARDINAL] = 65535 = $2^{16}-1$ and has internal representation 177777B
FIRST [LONG INTEGER] = -2147483648 = -2^{31}
LAST [LONG INTEGER] = 2147483647 = $2^{31}-1$
LAST [LONG CARDINAL] = 4294967295 = $2^{32}-1$

C.2 ASCII character set and ordering of character values

The following list gives the characters of the ASCII character set in increasing order, accompanied by their literal representations. Control characters are represented as $\uparrow a$. In addition, a number of special characters such as SP (space), DEL (rubout) are denoted by their generally accepted names.

Octal Value	Character Name(s)	Octal Value	Character Name(s)
000C	NUL	100C	'@
001C	\uparrow A	101C	'A
002C	\uparrow B	102C	'B
003C	\uparrow C	103C	'C
004C	\uparrow D	104C	'D
005C	\uparrow E	105C	'E
006C	\uparrow F	106C	'F
007C	\uparrow G, BELL	107C	'G
010C	\uparrow H, BS	110C	'H
011C	\uparrow I	111C	'I
012C	\uparrow J, LF	112C	'J
013C	\uparrow K	113C	'K
014C	\uparrow L	114C	'L
015C	\uparrow M, CR	115C	'M
016C	\uparrow N	116C	'N
017C	\uparrow O	117C	'O

C

Mesa Machine Dependencies

020C	↑P	120C	'P
021C	↑Q	121C	'Q
022C	↑R	122C	'R
023C	↑S	123C	'S
024C	↑T	124C	'T
025C	↑U	125C	'U
026C	↑V	126C	'V
027C	↑W	127C	'W
030C	↑X	130C	'X
031C	↑Y	131C	'Y
032C	↑Z	132C	'Z
033C	ESC	133C	'[
034C		134C	'\
035C		135C	']
036C		136C	'↑
037C		137C	'←
040C	' , SSpace	140C	
041C	'!	141C	'a
042C	'"	142C	'b
043C	'#	143C	'c
044C	'\$	144C	'd
045C	'%	145C	'e
046C	'&	146C	'f
047C	' , a single quote	147C	'g
050C	'('	150C	'h
051C	')	151C	'i
052C	'*	152C	'j
053C	'+	153C	'k
054C	'.	154C	'l
055C	'.	155C	'm
056C	'.	156C	'n
057C	'/'	157C	'o
060C	'0	160C	'p
061C	'1	161C	'q
062C	'2	162C	'r
063C	'3	163C	's
064C	'4	164C	't
065C	'5	165C	'u
066C	'6	166C	'v
067C	'7	167C	'w
070C	'8	170C	'x
071C	'9	171C	'y
072C	'.'	172C	'z
073C	'.'	173C	'{'
074C	'<	174C	'
075C	'='	175C	'}'
076C	'>	176C	'~
077C	'?	177C	DEL



D

Binder Extensions

The implementation of the Mesa binder provides two extensions for controlling the space occupied by Mesa programs at runtime. These are specified with the **CPacking** and **CLinks** clauses (section 7.7).

D.1 Code packing

It is possible to pack together the code for several modules into a single segment. This is useful for two reasons:

Since the code is allocated an integral number of pages, there is some wasted space in the last page ("breakage"). If several modules are combined into a single segment, the breakage is amortized over all the modules, and there is less waste on the average.

All the modules will be brought into and out of memory together, as a unit; a reference to any module in the pack will cause all the code to be brought in. Modules which are tightly coupled dynamically are good candidates for packing (for example, resident code should probably always be packed).

Of course, it is possible to "over pack" a configuration; the segments might become so large that there will never be room in memory for more than one of them at a time (this should remind you of an overlay system). *Packing is a tradeoff, and should be used with caution.*

D.1.1 Syntax

The segments are specified at the beginning of the configuration by giving a list of the modules which comprise each one. Any number of **PACK** statements may appear. The scope of the packing specification is the whole configuration, and not subconfigurations or individual module instances, because there is at most one copy of a module's code in any configuration.

```
ConfigDescription ::= Directory CPacking Configuration .
CPacking          ::= empty | CPackSeries ;
CPackList         ::= PACK IdList
```

CPackSeries ::= **CPackList** | **CPackSeries** ; **CPackList**

Each **PackList** defines a single segment; the code for all the modules in the **IdList** will be packed into it. The identifiers in the **IdList** must refer to modules in the configuration, and not to module instances; it is the code and not the global frames that are being packed (the frames are always packed when they are allocated by the loader).

It is illegal to specify the same module in more than one **PackList**. Even though there may be multiple instances of the module (i.e., multiple global frames) in the configuration, the code is shared by all of them, and therefore can only appear in one pack.

Finally, it is perfectly fine to reach inside a previously bound configuration that is being instantiated and single out some or all of its modules for packing. Of course, you must know something about the structure of that configuration in order to do this.

D.1.2 Restrictions

Obviously, the **PACK** statements apply only if the code is being moved to the output file; otherwise, the pack lists are ignored (and no warning message is given). This allows the programmer to debug the configuration without shuffling the code from file to file, thereby saving time. When making the final version, the packing can be effected with a binder switch, without having to modify the source of the configuration description.

Once some modules have been packed together, they cannot be taken apart and repacked with other modules later on, when they are bound into some other configuration.

Fine point:

If a previously bound configuration contains a pack, referencing any module of the pack gets the whole thing. So it is possible to pack a module and a pack together, or even to pack two packs. It is never possible to unpack a pack.

In general, code packing should be specified only to the extent that no unpacking will ever be desired. Once the packing is done, it can't be undone, unless you start over with the individual modules.

D.2 External links

In previous Mesa systems, links to the externals referenced by a program (imported procedures, signals, errors, frames, and programs) were always stored in the module's global frame. This allows each instance of a module to be bound differently, and it allows binding to be done at runtime without modification of the module's code segment. However, it has two drawbacks:

The links are only referenced by the module's code, and are therefore not needed when the code is swapped out. Hence, the links logically belong in the code segment.

If two instances of a module are bound identically (the usual case), the links must be stored twice.

Fine Point:

To determine the amount of space required for external links, see the compiler's typescript file. Each link occupies two words.

The Mesa binder optionally places links in the code segment. This option is enabled by constructs in the configuration language, and is further controlled by binder and loader switches.

D.2.1 Syntax

For each component of a configuration, the link location is specified using the `LINKS` construct defined below. The default is frame links.

CLinks ::= `empty` | `LINKS : CODE` | `LINKS : FRAME`

A link specification can optionally be attached to each instantiation of a module, overriding the current default, so that the link location can be different for each instance.

CRightSide ::= `Item` | `Item [] CLinks` | `Item [IdList] CLinks`

Alternately, the link option can be specified in the configuration header. This merely changes the default option for the configuration; it will apply to all components (including nested configurations) unless it is explicitly overridden.

CHead ::= `CONFIGURATION CLinks Imports CExports ControlClause`

This construction works much like the `PUBLIC / PRIVATE` options in Mesa, and it nests in the same way. A link option attached to a configuration changes the default for all components within it, but that default can be overridden for a particular module (or nested configuration) by specifying a different link option.

D.2.2 Restrictions

This scheme has the consequence that, *if a module with code links has multiple instances, each instance must be bound the same.*

As with code packing, the code links option takes effect only when the code is being moved to the output file. At this point, the binder will make room for the links as it copies the code if any module sharing that code has requested code links. Again, this allows a programmer to debug without the expense of moving the code (using frame links), and then to effect the code links option with a binder switch, without changing the source of the configuration description.

Fine point:

Once space for code links has been added to a configuration, it cannot be undone by a later binding. On the other hand, space for code links can always be added to a (previously bound) configuration, even if it did not specify code links in its description.

Using code links has one drawback: it slows down the binding and loading process, as the code must be swapped in and rewritten. The binder must make room in the code segment for the links, as described above. And because the loader resolves imports of previously

loaded modules, as well as the imports of the module being loaded, it may have to swap in (and perhaps update and swapout) the code segment for every module in the system.

Fine point:

If a module that is part of the boot file has unresolved code links, they get filled in when the exporting module is loaded. Unfortunately, they stay filled in when the system is booted again, even though the implementation has not been loaded yet. Modules built into the boot file should not have unresolved code links.

Because of the overhead involved, the loader will not automatically attempt to use code links, even if the space is available in the code segment. A loader switch must be used to effect this option.

Documentation of binder and loader switches is in the *Executive* section of the *Mesa User's Guide*.



E

Mesa reserved words

Listed below are all of the Mesa reserved words. Words marked with an astrisk are *predeclared* rather than *reserved*. Predeclared identifiers can be redefined (but seldom should be).

ABORTED	EXPORTS
ABS	FALSE
ALL	FINISHED
AND	FIRST
ANY	FOR
APPLY	FORK
ARRAY	FRAME
BASE	FREE
BEGIN	FROM
BOOL	GO
BOOLEAN	GOTO
BROADCAST	IF
CARDINAL	IMPORTS
CHAR	IN
CHARACTER	INLINE
CODE	INT
COMPUTED	INTEGER
CONDITION*	INTERNAL
CONTINUE	ISTYPE
DECREASING	JOIN
DEFINITIONS	LAST
DEPENDENT	LENGTH
DESCRIPTOR	LOCKS
DIRECTORY	LONG
DO	LOOP
ELSE	LOOPHOLE
ENABLE	MACHINE
END	MAX
ENDCASE	MDSZone
ENDLOOP	MIN
ENTRY	MOD
ERROR	MONITORED
EXIT	MONITOR
EXITS	MONITORLOCK*

E**Mesa Reserve Words**

NARROW	RETRY
NAT	RETURN
NATURAL	RETURNS
NEW	SELECT
NIL*	SEQUENCE
NOT	SHARES
NOTIFY	SIGNAL
NULL	SIZE
OF	START
OPEN	STATE
OR	STOP
ORD	STRING
ORDERED	StringBody*
OVERLAID	SUCC
PACKED	THEN
POINTER	THROUGH
PORT	TRANSFER
PRED	TRASH
PRIVATE	TRUE*
PROC	TYPE
PROCEDURE	UNCOUNTED
PROCESS	UNSPECIFIED*
PROGRAM	UNTIL
PUBLIC	UNWIND*
READONLY	USING
REAL*	VAL
RECORD	VAR
REJECT	WAIT
RELATIVE	WHILE
REPEAT	WITH
RESIDENT	WORD*
RESTART	ZONE
RESUME	



Collected grammar

Mesa is a living language which has undergone many changes since its initial implementations. Extensions and refinements continue to be made. Consequently, the BNF which is found in the chapters of this manual may not reflect the current state of Mesa syntax. The included BNF, and accompanying explanations, is our strategy for describing Mesa language concepts accurately. However, the "official" grammar used by the parser of the Mesa 11.0 compiler has been reproduced in its entirety below. This grammar accepts a superset of valid Mesa programs. A subsequent pass of the compiler corrects this deficiency.

```
compilationUnit ::= directory identlist proghead block .  
                  | directory identlist defhead defbody .  
  
defbody         ::= BEGIN open declist END  
                  | BEGIN open declist ; END  
                  | { open declist }  
                  | { open declist ; }  
  
defhead        ::= DEFINITIONS locks imports shares = public  
  
directory      ::= empty  
                  | DIRECTORY ;  
                  | DIRECTORY includelist ;  
  
exports       ::= empty  
                  | EXPORTS  
                  | EXPORTS modulelist  
  
idlist        ::= id  
                  | id , idlist  
  
imports       ::= empty  
                  | IMPORTS  
                  | IMPORTS modulelist  
  
includeitem   ::= id : FROM string using  
                  | id : TYPE using
```

		id using
		id : TYPE id using
includelist	:: =	includeitem includelist , includeitem
interface	:: =	imports exports shares
locks	:: =	empty LOCKS primary lambda
lambda	:: =	empty USING ident typeexp
moduleitem	:: =	id id : id
modulelist	:: =	moduleitem modulelist , moduleitem
proghead	:: =	PROGRAM arguments interface = public MONITOR arguments locks interface = public
shares	:: =	empty SHARES idlist
using	:: =	empty USING [] USING [idlist]
typeexp	:: =	id predefinedtype typeid typecons
arglist	::	empty fieldlist
arguments	:: =	arglist returnlist
base	:: =	empty BASE
bounds	:: =	exp .. exp
default	:: =	empty ← defaultopt
defaultopt	:: =	empty TRASH NULL

	exp ' TRASH	-- second <i>embedded punctuation</i>
	exp ' NULL	-- second <i>embedded punctuation</i>
	exp	
dependent	:: = empty	
	MACHINE DEPENDENT	
element	:: = id (exp)	
	(exp)	
	id	
elementlist	:: = empty	
	elementlist'	
elementlist'	:: = element	
	elementlist' , element	
fieldlist	:: = []	
	[pairlist]	
	[typelist]	
identlist	:: = id :	
	id position :	
	id , identlist	
	id position , identlist	
indextype	:: = empty	
	typeexp	
interval	:: = [bounds]	
	[bounds)	
	(bounds]	
	(bounds)	
length	:: = [exp]	
monitored	:: = empty	
	MONITORED	
optbits	:: = empty	
	: bounds	
ordered	:: = empty	
	ORDERED	
packed	:: = empty	
	PACKED	
pairitem	:: = identlist public typeexp default	
pairlist	:: = pairitem	
	pairlist , pairitem	

predefinedtype	:: =	BOOL BOOLEAN CARDINAL CHAR CHARACTER INT INTEGER NAT NATURAL REAL CONDITION MDSZone MONITORLOCK STRING StringBody UNSPECIFIED WORD	-- equivalent to BOOLEAN -- equivalent to CHARACTER -- equivalent to LONG INTEGER -- equivalent to NATURAL
pointerprefix	:: =	POINTER POINTER interval	
pointertype	:: =	pointerprefix pointerprefix TO readonly typeexp	
position	:: =	(exp optbits)	
public	:: =	empty PUBLIC PRIVATE	
readonly	:: =	empty READONLY	
reclist	:: =	[] NULL [pairlist] [typelist] [pairlist , variantpair] [variantpart default] [variantpair]	
returnlist	:: =	empty RETURNS fieldlist	
tagtype	:: =	* typeexp	
transfermode	:: =	PROCEDURE PROC PORT SIGNAL ERROR	

		PROCESS PROGRAM
typeappl	:: =	typeappl . id id length typeid length typeappl length
typecons	:: =	interval id interval typeid interval dependent { elementlist } dependent monitored RECORD reclist ordered base pointertype VAR typeexp packed ARRAY indextype OF typeexp DESCRIPTOR FOR readonly typeexp transfermode arguments id RELATIVE typeexp typeid RELATIVE typeexp UNCOUNTED ZONE LONG typeexp FRAME [id] typeappl
typeid	:: =	id id id typeid typeid'
typeid'	::	id . id typeid' . id
typelist	:: =	typecons default typeid default id id ← defaultopt typecons default , typelist typeid default , typelist id , typelist id ← defaultopt , typelist
variantitem	:: =	idlist = > reclist
variantlist	:: =	variantitem variantlist , variantitem
variantpair	:: =	identlist public variantpart default
variantpart	:: =	SELECT vcasehead FROM variantlist ENDCASE SELECT vcasehead FROM variantlist , ENDCASE packed SEQUENCE vcasehead OF typeexp

vcasehead	:: =	ident public tagtype COMPUTED tagtype OVERLAID tagtype
statement	:: =	IF exp THEN statement IF exp THEN balstmt ELSE statement casehead casestmtlist ENDCASE = > statement basicstmt
balstmt	:: =	IF exp THEN balstmt ELSE balstmt casehead casestmtlist ENDCASE = > balstmt basicstmt
basicstmt	:: =	lhs lhs ← exp [explist] ← exp block casehead casestmtlist ENDCASE forclause dotest DO scope doexit ENDLOOP EXIT LOOP GOTO id GO TO id RETURN optargs transfer lhs lhs . FREE [exp optcatch] WAIT lhs ERROR STOP NULL RESUME optargs REJECT CONTINUE RETRY lhs ← STATE STATE ← lhs
binditem	:: =	exp id : exp
bindlist	:: =	binditem bindlist , binditem
block	:: =	BEGIN scope exits END { scope exits }
casehead	:: =	SELECT exp FROM WITH binditem SELECT optexp FROM

caselabel	:: = ident typeexp caselabel'
caselabel'	:: = casetest caselabel' , casetest
casestmtitem	:: = caselabel = > statement
casestmtlist	:: = empty casestmtlist' casestmtlist' ;
casestmtlist'	:: = casestmtitem casestmtlist' ; casestmtitem
casetest	:: = optrelation exp
catchany	:: = ANY = > statement
catchcase	:: = lhslst = > statement
catchhead	:: = empty catchhead catchcase ;
catchlist	:: = catchhead catchhead catchcase catchhead catchany catchhead catchany ;
odelist	:: = orderlist oodelist ; orderlist
controlid	:: = ident typeexp id
declaration	:: = identlist public entry readonly typeexp initialization identlist public TYPE = public typeexp default identlist public TYPE optsize
declist	:: = declaration declist ; declaration
direction	:: = empty DECREASING
doexit	:: = empty REPEAT exitlist REPEAT exitlist FINISHED = > statement REPEAT exitlist FINISHED = > statement ;

dotest	:: = empty UNTIL exp WHILE exp
enables	:: = empty ENABLE catchcase ; ENABLE catchany ; ENABLE BEGIN catchlist END ; ENABLE { catchlist } ;
entry	:: = empty ENTRY INTERNAL
exititem	:: = idlist = > statement
exitlist	:: = empty exitlist' exitlist' ;
exitlist'	:: = exititem exitlist' ; exititem
exits	:: = empty EXITS exitlist
forclause	:: = empty FOR controlid direction IN range FOR controlid ← exp , exp THROUGH range
initialization	:: = empty ← initvalue = initvalue
initvalue	:: = inline block CODE MACHINE CODE BEGIN codelist END MACHINE CODE { codelist } TRASH NULL exp
inline	:: = empty INLINE
lhslist	:: = lhs lhslist , lhs
open	:: = empty OPEN bindlist ;

optargs	:: = empty [explist] lhs
optsize	:: = empty [exp]
scope	:: = open enables statementlist open enables declist ; statementlist
statementlist	:: = empty statementlist' statementlist' ;
statementlist'	:: = statement statementlist' ; statement
transfer	:: = SIGNAL ERROR RETURN WITH ERROR START RESTART JOIN NOTIFY BROADCAST TRANSFER WITH RETURN WITH
exp	:: = transferop lhs IF exp THEN exp ELSE exp casehead caseexplist ENDCASE = > exp lhs ← exp [explist] ← exp ERROR disjunct
addop	:: = + -
caseexpitem	:: = caselabel = > exp
caseexplist	:: = empty caseexplist' caseexplist' ,
caseexplist'	:: = caseexpitem caseexplist' , caseexpitem
conjunct	:: = conjunct AND negation negation

desclist	:: =	exp , exp opttype exp
disjunct	:: =	disjunct OR conjunct conjunct
explist	:: =	orderlist keylist
factor	:: =	addop primary primary
ident	::	id position : id :
keyitem	:: =	id : optexp
keylist	:: =	keyitem keylist , keyitem
lhs	:: =	id char NARROW [exp opttype optcatch] LOOPHOLE [exp opttype] APPLY [exp , exp optcatch] (exp) lhs qualifier
multop	:: =	* / MOD
negation	:: =	not relation relation
not	:: =	~ NOT
optcatch	:: =	empty ! catchlist
optexp	:: =	empty TRASH NULL exp
optrelation	:: =	not relationtail relationtail
opttype	:: =	empty , typeexp

orderlist	:: =	optexp orderlist , optexp
prefixop	:: =	LONG ABS PRED SUCC ORD MIN MAX BASE LENGTH
primary	:: =	num lnum flnum string lstring NIL [explist] prefixop [orderlist] VAL [orderlist] ALL [orderlist] lhs . NEW [typeexp initialization optcatch] typeop [typeexp] SIZE [typeexp] SIZE [typeexp , exp] ISTYPE [exp , typeexp] @ lhs DESCRIPTOR [desclist] lhs
product	:: =	product multop factor factor
qualifier	:: =	. prefixop . typeop . SIZE [explist optcatch] . id ↑
range	:: =	id empty id interval typeid interval interval typeid
relation	:: =	sum optrelation sum

relationtail ::= **IN range**
 | **relop sum**

relop ::= **empty**
 | **#**
 | **<**
 | **< =**
 | **>**
 | **> =**

sum ::= **sum addop product**
 | **product**

transferop ::= **SIGNAL**
 | **ERROR**
 | **START**
 | **JOIN**
 | **NEW**
 | **FORK**

typeop ::= **FIRST**
 | **LAST**
 | **NIL**



Index

- !, 8-4
- +, 2-13
- ;; see *semicolon*
- , 2-13
- , 2-3
- ., 3-21
- ::=, 1-2
- =>, 4-5, 8-4
- @, 3-32
- ↑, 3-30
- ←, 2-2
- <<...>>, 2-3
- abbreviated identifier, 3-1
- ABS, 2-2
- ABS
 - with unsigned operands, 3-44
- Access, 7-22
- access control, 7-22
 - global default, 7-25
 - in type definitions, 7-24
- activation (procedure), 5-7
 - of nested procedures, 5-17
- actual procedure, 5-9
- actual tag, 6-20
- AddingOperator, 2-13
- address-containing type, 3-47
- aggregate data type, 3-3
- aggregate variable, 3-11
- alignment
 - of arrays, 3-14
 - of parameter/results records, 5-4
 - of pointer referents, 3-34
- ALL, 3-46
 - in array constructors, 3-12, 3-15
 - in exports list, 7-41
- allocation
 - of string literal, 6-3
 - of storage, 6-31
- AlternateName, 4-12
- AND, 2-17
- ANY, 8-3
 - in catch phrase, 8-7
- APPLY, 5-9
- argument record, 5-1
- arguments, 5-8
 - relationship to parameters, 5-8
- ARRAY, 6-6, 3-11
 - as argument, 6-7
 - as compile-time constant, 3-16
 - assignment, 6-6
 - component type, 3-11
 - index type, 3-11
 - length of, 3-14
 - READONLY, 6-6
 - restrictions on assignment, 4-3
- array constructor, 3-12
 - keyword, 3-16
- array descriptor, 6-6
- ArrayTC, 3-13
- Assignment, 4-16
- assignment
 - restrictions on, 4-3
- assignment expression, 4-2, 2-18
 - order of evaluation, 4-3
- assignment statement, 2-4, 4-2
- AssignmentExpr, 4-2, 2-18
- AssignmentStmt, 2-4, 4-2, 3-27
- automatic dereferencing, 3-36
- B, 2-6
- Backus-Naur Form, 1-2
- balancing, 3-37
- BASE, 6-9, 3-15
 - and automatic dereferencing, 3-37
 - as pointer attribute, 3-33
- BASE pointer, 3-35, 6-1
- base type, of subrange, 3-9
- BaseOption, 6-9
- BCD, 7-3

- BEGIN**, 2-2, 4-8
- binary configuration description, 7-3
- BinaryTree example, 5-13
- binder, 7-33
- binding, 7-1
 - done by loader, 7-43
 - of imports and exports, 7-17
- blank space
 - as separator, 2-3
 - in **STRING**, 2-3
 - in formatting, B-8
- Block**, 4-8, 4-9, 8-5
- BNF**, 1-2
- BOOLEAN**, 2-6, 3-6
 - as element type, 3-4
- BOOLEAN** expression, 4-4
- BOOLEAN** operators, 2-17
- bound variant, 6-20
- bound variant types, 6-14
- bounds checking, 2-10
- BoundsType**, 6-26
- bracket pair, 4-9
- BROADCAST**, 9-8
 - when to use, 9-10
- BuiltinCall**, 2-12
- C**, 2-7, 7-34
- Call**, 5-6
- call by value, 5-2
- CallStmt**, 5-6
- capitalization of identifiers, B-1
- CARDINAL**, 2-5
 - as element type, 3-4
- Catch**, 8-4
 - catch phrase, 8-4
 - location of, 8-4
 - scope of variables in, 8-5
 - with **FORK** or **JOIN**, 9-3
 - within a monitor, 9-18
- CatchPhrase**, 8-1
- CatchSeries**, 8-4
- CHARACTER**, 2-7
 - as element type, 3-4
 - operations on, 2-16
- characterLiteral**, 2-7
- client, 7-29
- client module, 7-5
- C/Mesa**, 7-3
 - semantics of, 7-34
 - syntax of, 7-34
- CODE**, 8-2
- coercion, 3-38
- comma, 4-8
- comment, 2-3
- CommonPart**, 3-28
- compatibility
 - of procedure types, 5-5
- compilation dependencies, 7-26
- compilation order, 7-10
- compile-time constant, 2-21
 - array as, 3-16
- Component**, 3-15
 - component type, 3-11
- ComponentList**, 3-15
- ComponentType**, 3-13
- concurrent processes, 9-1
- CONDITION**, 9-8
 - condition variable, 9-4
 - covering, 9-9
 - updating timeout field, 4-3
 - timeouts, 9-11
- ConditionTest**, 4-15
- CONFIGURATION**, 7-34
- configuration
 - control, 7-45
 - default names in, 7-38
 - instance of, 7-33
 - nested, 7-41
 - self-contained, 7-34
- configuration description, 7-33
- configuration language, 7-29
- configuration prototype, 7-35
- conformance
 - of array types, 3-13
 - of pointers, 3-33
 - of procedure arguments, 5-8
 - of procedure types, 5-5
 - of subranges, 3-9
 - of types, 2-4
- conformance relations, 2-10
- Conjunction**, 2-17
- constant, 2-4
 - compile-time, 2-21
 - identifier, 3-5
 - run time, 7-12
- Constructor**, 3-15
- constructor
 - keyword, 5-10
- CONTINUE**, 8-10
- CONTROL**, 7-34
- control configuration, 7-45
- control module, 7-3, 7-44
- control variable, 4-17
- ControlVariable**, 4-16
- D**, 2-6
- dangling reference, 3-32
 - in binding, 7-2
- data type
 - aggregate, 3-3
 - constructed, 3-1
- deadlock, 9-9
- Declaration**, 3-1, 2-22
- DeclarationSeries**, 4-9

- DECREASING**, 4-16
- defaults
 - extended, 3-45
 - for global access, 7-25
 - for procedure arguments, 7-14
 - in an elided/omitted field, 3-24
 - in argument/result records, 5-12
 - in parameter list, 5-4
 - of procedure arguments, 5-6
 - in record type, 3-17
 - in record field, 3-24
 - in result list, 5-4
 - inheritance of, 3-45
 - of elided/omitted component, 3-46
 - of built-in types, 3-45
- DefaultOption**, 3-18
- DefaultSpecification**, 3-25
- defining occurrences, 7-9
- DEFINITIONS**, 5-16
- DEFINITIONS** module, 7-1, 7-5, 7-6
- dereferencing, 3-30
- DESCRIPTOR**, 6-6, 3-15
- DescriptorTC**, 6-6
- determining representation, 3-42
- Direction**, 4-16
- DIRECTORY**, 7-4, 5-16
 - CMesa, 7-35
- Disjunction**, 2-17
- DO**, 2-2, 4-15
- dot notation, 2-19, 3-7
 - and interface items, 7-28
 - discriminated variant types, 6-23
 - with unary operators, 2-20
- dynamic array, 6-6
- Element**, 3-7
- element type, 3-3
- ElementList**, 3-7
- ElementType**, 3-4
- elision
 - in result record, 5-12
 - of array component, 3-15
 - of element in returns list, 5-10
 - of record component, 3-23
- ELSE**, 2-2, 4-3
- ElseClause**, 4-3
- empty, 1-2, 3-15
- empty constructor, 5-7
- empty extractor, 3-27
- empty interval, 3-14
- ENABLE**, 8-4, 8-5
 - scope of, 8-5
- EnableClause**, 8-5
- END**, 2-2, 4-8
- ENDCASE**, 4-7
 - guidelines for use, B-5
- ENDLOOP**, 2-2, 4-15
- ENTRY** attribute, 9-7
- entry procedure
 - declaration of, 9-6
 - INLINE**, 9-17
 - JOIN** statement in, 9-8
- enumerated type, 3-5
 - ordered, 3-6
- enumeration, 3-3
 - machine dependent, 3-7
 - sparse, 3-8
- EnumerationTC**, 3-5
- equivalence
 - of array types, 3-13
 - of pointers, 3-33
 - of procedure types, 5-5
 - of record types, 7-11
 - of subranges, 3-9
 - of types, 2-9
- ERROR**, 8-1
 - omission of keyword, 8-3
 - unnamed, 8-3
- ErrorCall**, 8-2
- EXIT**, 4-18, 4-19
 - in catch phrase, 8-8
- EXITS**, 4-10
- ExitsClause**, 8-5, 4-10
- ExitSeries**, 4-10
- ExitStmt**, 4-19
- expansion
 - of **INLINE** procedure, 5-18
- exporting, 7-2
 - explicit renaming, 7-36
 - of interface, 7-20
 - of program module, 7-20
 - of type, 7-25
- export record, 7-2
- EXPORTS**, 7-4
 - C/Mesa, 7-35
- EXPORTS ALL**, 7-41
- ExprChoiceList**, 4-8
- Expression**, 4-1, 2-12
- expression
 - BOOLEAN**, 4-4
 - assignment, 4-2
 - SELECT**, 4-8
- extended defaults, 3-45
- ExtractItem**, 3-27
- Extractor**, 3-26
- extractor, 3-26
 - type of, 3-38
- Factor**, 2-13
- field list, 3-17
 - named, 3-17, 5-10
 - unnamed, 3-17
- field name, 3-17
- FieldDescription**, 3-18, 7-23

- FieldId, 3-28
- FieldIdList, 3-28
- FieldList, 3-18
- FieldPosition, 3-28
- FinalStmtChoice, 4-6
- FINISHED, 4-18
- FinishedExit, 4-18
- FIRST, 3-6, 2-20
- floating point, 2-8
- FOR, 4-16
- FORK, 9-2
 - target of, 9-8
- ForkCall, 9-3
- FRAME, 7-5
- frame, 4-9
- FREE, 6-1, 6-32
- free conformance, 3-41
 - of arrays, 3-13
 - of pointers, 3-33
 - of procedure types, 5-5
- FROM, 4-6, 7-4
- FunctionCall, 2-12
- fundamental operation, 2-5
 - on procedure types, 5-4
 - on strings, 6-3
 - standard implementation of, 7-26
- GCD, 2-1
- global frame, 7-5
 - allocation of, 7-33
 - allocation of literals from, 6-3
- GlobalAccess, 7-22, 7-25
- GOTO, 4-18
 - in catch phrase, 8-8
- GOTO statement, 4-9
- GotoStmt, 4-10
- H, 2-6
- home module, 7-22
- identifier, 1-3, 3-3
 - abbreviated, 3-1
 - capitalization of, B-1
 - case of letters, 2-3
 - predeclared, 2-12
- identifier, 2-3
- identifier constant, 3-5
- identifier list, 2-4
- IdList, 2-4-3
- IF, 2-2, 4-3
- IF expression, 4-4
 - target type of, 3-38
- IF statement, 4-3
 - order of evaluation, 4-4
- IfExpr, 4-4
- IfStmt, 4-3
- implementor, 7-29
- importing, 7-1
 - explicit, 7-21
 - implicit, 7-21
 - in DEFINITIONS modules, 7-21
 - explicit renaming, 7-36
- IMPORTS, 7-4
 - C/Mesa, 7-35
- ImportsList, 7-21
- IN, 2-16, 4-16
 - balancing and, 3-41
- IncludeItem, 7-6
- indentation
 - guidelines for, B-7
- index type, 3-11
- indexed reference, 3-12
- IndexedReference, 3-14, 3-33
- IndexType, 3-13
- inherent type, 3-37
- InitialExpr, 4-16
- Initialization, 5-9, 2-22
- initialization, 2-20
 - of monitor, 9-19
 - of monitored record, 9-19
 - of relative pointers, 6-12
 - of signal, 8-2
- INLINE, 5-9
- inline procedures, 5-3, 5-18
 - efficiency of, 7-16
 - expansion of, 5-18
 - in interface, 7-14
- InlineOption, 5-9
- instance
 - of configuration, 7-33
- INTEGER, 2-2
 - as element type, 3-4
 - declaration of, 2-4
- interface, 7-1
 - exporting, 7-20
 - system, 7-30
- interface item, 7-12
- interface record, 7-2
 - vs. interface type, 7-17
- interface type, 7-17
- INTERNAL, 9-7
- internal procedure
 - declaration of, 9-6
- Interval, 2-16
- interval
 - empty, 2-17
- ISTYPE, 3-39
- Iteration, 4-16
- IterativeControl, 4-15
- JOIN, 9-2
- JOIN statement
 - in entry procedure, 9-8
- JoinCall, 9-3
- jump table, 4-7
- keyword array constructor, 3-16

- keyword constructor, 3-22, 5-10
- keyword extractors, 5-10
- keyword name, 3-23
- KeywordComponent**, 3-22
- KeywordComponentList**, 3-22
- KeywordExtract**, 3-27
- KeywordExtractList**, 3-27
- Label**, 4-10
- LAST**, 3-6, 2-20
- LeftItem**, 4-6
- LeftSide**, 2-4-3
- LENGTH**, 3-15
- length, 6-1
- lexical unit, 2-3
- link, 7-33
- LINKS**, 7-34
- Literal**, 2-12
- loading, 7-33
- local frame
 - allocation of literals from, 6-3
- local procedures, 5-17
- local variables, 5-2
- lock, monitor, 9-4
- LOCKS**, 9-12
- LOCKS** clause, 9-14
- LocksClause**, 9-12
- LONG**, 2-14, 3-35
- LONG CARDINAL**, 2-7
- LONG DESCRIPTOR**, 6-8
- LONG INTEGER**, 2-7
- LONG POINTER**, 3-30
- LONG STRING**, 6-5
- LongTC**, 6-8, 3-35
- LOOP**, 4-19
 - in catch phrase, 8-8
- loop body, 4-15
- loop control, 4-15
 - direction, 4-17
- LOOP** statement, 4-19
- loop termination
 - conditional, 4-17
 - forcible, 4-17
 - normal, 4-17
- LoopCloseStmt**, 4-19
- LoopControl**, 4-15
- LoopExits**, 4-18
- LoopExitsClause**, 4-18
- LOOPHOLE**, 3-39
- LoopRange**, 4-16
 - evaluation of bounds, 4-17
- LoopStmt**, 4-15
- MACHINE DEPENDENT** enumeration, 3-7
- MACHINE DEPENDENT** record, 3-20
 - positional form, 3-29
- MACHINE DEPENDENT SEQUENCE**, 6-27
- MachineDependent**, 3-7
- main data space, 3-30
- mainline code, 7-43
- MAX**, 2-13
 - balancing, 3-41
- maxlength
 - of string, 6-1
- MDS**, 3-30
- MDSZone**, 6-31
- MIN**, 2-13
 - balancing, 3-41
- MOD**, 2-2, 2-13
- module, 7-1
 - accessing items in, 7-7
 - atomic, 7-33
 - change log for, B-6
 - copy of, 7-42
 - DEFINITIONS**, 7-5, 7-6
 - format of, B-7
 - indentation of, B-7
 - monitor, 9-6
 - naming conventions, B-2
 - OPENING**, 4-13
 - organization of, B-7
 - privileged, 7-22
 - program, 7-1
 - RESTARTING**, 7-44
 - STARTING**, 7-43
 - STOPPING**, 7-44
- MONITOR**, 7-5, 9-6
- monitor, 9-4
 - deadlocked, 9-8
 - entry procedure, 9-6
 - external procedure, 9-6
 - initialization of, 9-19
 - internal procedure, 9-6
- module
 - declaration of, 9-6
 - structure of, 9-7
- multi-module, 9-13
- object, 9-13
- monitor lock, 9-4
 - and **UNWIND**, 9-18
 - exported, 9-15
 - restrictions on assignment, 4-3
 - updating of, 9-13
- MONITORLOCK**, 9-5, 2-22, 9-12
- MONITORED RECORD**, 9-12, 9-14
 - initialization of, 9-19
- multiple implementors, 7-40
- MultiplyingOperator**, 2-13
- name scope, 7-9, 3-21
- named field list, 5-10
 - as result list, 5-10
- NamedField**, 3-28
- NamedFieldList**, 3-18, 7-23
- naming

- in configurations, 7-36
 - of modules, B-2
 - of types, B-3
- NARROW**, 3-39, 7-28
- NATURAL**, 2-5
 - as element type, 3-4
- Negation**, 2-17
- nesting
 - of configurations, 7-41
 - of procedures, 5-16
 - of signals, 8-12
- NEW**, 6-1
 - of module instance, 7-19
- Next-Statement**, 4-1
- NextExpr**, 4-16
- NIL**, 2-20, 3-32
 - lengthening of, 3-35
- non-privileged module, 7-22
- non-interface item, 7-12
- non-local variable, 5-3
- non-transfer variable, 7-12
 - exported, 7-20
- NOT**, 2-16
- NOTIFY**, 9-4
 - when to use, 9-10
- NULL**, 4-7, 3-24
- NULL** statement, 4-7
- NullStmt**, 4-7
- number**, 2-5
- number, floating point, 2-8
- numeric literal, 2-5
- numeric type, 2-5
 - long, 2-8
 - short, 2-8
 - subrange of, 3-10
- object file, 7-33
- object module, 7-1
- object notation, 7-28
- object-oriented programming, 7-28
- omission
 - of element in returns list, 5-10
- opaque type, 7-25
- OPEN**, 4-14
 - guidelines for use of, B-2
- OPEN** clause, 4-12, 7-8
 - scope of, 4-12
- OpenClause**, 4-12
- OpenItem**, 4-14
- operations
 - on characters, 2-16
- operators
 - BOOLEAN**, 2-17
 - precedence of, 2-19
 - numeric, 2-12
 - relational, 2-16
- OptionalTyped**, 3-15
- OR**, 2-17
- ORD**, 2-19
- ORDERED**, 3-32
- ORDERED POINTER**, 3-32
- ordered type, 2-16, 3-3
- ordinary statement, 4-1
- PACK**, 7-34
- PACKED**, 6-6, 3-14
- packing
 - during binding, 7-35
 - in parameter or result record, 5-4
- PackingOption**, 3-13
- parameter record, 5-1
 - alignment of, 5-4
- ParameterList**, 5-3
- parameters, 5-8
 - of procedure, 5-1
 - of signals, 8-2
 - relationship to arguments, 5-8
- phrase class, 1-2
- PLUS**, 7-40
- POINTER**, 6-9, 3-32
- pointer, 3-3
 - arithmetic on, 3-35
 - BASE**, 3-35, 6-1
 - dereferencing, 6-12
 - lengthening, 3-35
 - ORDERED**, 3-32
 - READONLY**, 3-31
 - reference type, 3-30
 - referent, 3-30
 - relational operations on, 3-34
 - RELATIVE**, 6-9, 3-35, 6-10
 - subrange, 3-37
- pointer subrange type, 3-33
- POINTER TO UNSPECIFIED**, 3-35
- PointerTC**, 6-9, 3-32
- positional constructor, 3-22
- positional notation
 - of imports, 7-22
- PositionalComponentList**, 3-15
- PositionalExtractList**, 3-27
- precedence, 2-13
 - of **IF** operator, 4-5
 - of assignment operator, 4-2
- PRED**, 2-19
- PredefinedType**, 2-22
- Predicate**, 4-3
- Primary**, 3-5, 2-12-3
- principal instance
 - of imported interface, 7-29
 - of interface, explicit, 7-21
- PRIVATE**, 7-4
 - interface components, 7-15
 - declaration of, 7-22
 - in type definitions, 7-24

- PRIVATE** items
 - accessing, 7-25
- privileged module, 7-22
- PROC**, 5-4
- PROCEDURE**, 5-3
- procedure, 5-1
 - activation of, 5-7
 - actual, 5-2
 - arguments of, 5-1
 - default arguments for, 7-14
 - entry, 9-12
 - forked, 9-2
 - inline, 5-3
 - in monitor module, 9-6
 - internal, 9-6
 - "local," 5-12
 - nested, 5-16
 - results of, 5-1
 - termination of, 5-8
- procedure call, 5-1, 5-6
- procedure types, 5-3
 - equivalence of, 5-5
 - free conformance of, 5-5
- procedure variable, 5-4, 5-9
- ProcedureBody**, 5-9
- ProcedureTC**, 5-3
- PROCESS**, 9-1
- processes, 9-1
 - concurrent, 9-2
 - detached, 9-3
 - synchronized, 9-4
 - waiting, 9-8
- Product**, 2-13
- PROGRAM**, 7-5, 5-14, 7-42
- program module, 7-1
- program variable, 7-18
- prototype
 - configuration, 7-35
 - program, 7-35
- PUBLIC**, 7-4, 5-17, 7-13
 - declaration of, 7-22
 - in type definitions, 7-24
- qualification, 3-17, 3-21
 - of identifier constant, 3-5
 - of identifiers, 7-7
 - of items from **OPEN** module, 4-13
 - of pointer referents, 3-31
- QualifiedReference**, 3-21
- range assertion, 3-11
- range error, 3-9
- READONLY**, 3-33, 7-12
- READONLY POINTER**, 3-31
- REAL**, 2-8
- realLiteral**, 2-8
- RECORD**, 3-19
- record, 3-16
 - basic operations on, 3-17
 - declaration of, 3-19
 - initialization of, 3-20
 - MACHINE DEPENDENT**, 3-28
 - monitored, 9-12
 - OVERLAID**, 6-23
 - packing of, 3-20
 - parameter/return, 5-12
 - restrictions on assignment, 4-3
 - single-component, 5-2
- record constructor, 3-17
- record types
 - equivalence of, 7-11
- RecordTC**, 3-19, 6-15
- recursion, 5-2
- reentrant procedures, 5-2, 5-7
- reference type of pointer, 3-30
- REJECT**, 8-7
- relational operators, 2-16
- RelationalOperator**, 2-16
- RELATIVE**, 6-9
- relative array descriptors, 6-13
- relative pointer, 3-35, 6-1
 - creating with **FIRST**, 6-12
- relative pointer types, 6-11
- RelativeTC**, 6-9
- REPEAT**, 4-18
- Repetition**, 4-16
- reserved words, 2-3
- RESTART**, 7-19
- result record, 5-2
 - alignment of, 5-4
- ResultList**, 5-3
- results of procedure, 5-1
- RESUME**, 8-11
- ResumeStmt**, 8-11
- RETRY**, 8-10
- RETURN**, 5-2
 - default values for, 5-12
- return link, 8-6
- RETURN WITH ERROR**, 8-2, 9-18
- RETURNS**, 5-3
- ReturnsClause**, 5-3
 - absence of, 5-10
- ReturnStmt**, 5-10
- RightSide**, 2-4
- root procedure, 9-2
- running, 7-42
- run-time error, 7-2
- scalar type, 3-3
- scaleFactor**, 2-9
- scope
 - defined by procedures, 5-18
 - of **ENABLE** clause, 8-5
 - of identifiers, 7-9
 - of local variables, 5-17

- of OPEN clause, 7-8
 - of procedure, 5-10
 - of qualified reference, 7-9
 - of variables in catch phrases, 8-5
- scope guidelines, 7-10
- SELECT**, 4-5
- SELECT** expression, 4-8
- SELECT** statement, 4-5
 - order of evaluation, 4-6
- SelectExpr**, 4-8
- SelectStmt**, 4-6
- self-contained (configuration), 7-34
- semicolon, 2-2
 - as separator, 4-9
 - following **ThenClause**, 4-4
- SeqTag**, 6-26
- SEQUENCE**, 6-26
- sequence, 6-1
 - MACHINE DEPENDENT**, 6-27
 - allocation of, 6-28
 - declaration of, 6-25
 - initialization of, 6-28
 - limitations of, 6-25
 - operations on, 6-29
 - tags, 6-26
- SHARE**, 7-22
- SHARES**, 7-5
- side effects, 5-7
- SIGNAL**, 8-1
 - actual, 8-2
 - catching, 8-6
 - declaration of, 8-2
 - guidelines for use, B-5
 - indirect, 8-7
 - in monitored procedures, 9-17
 - in forked procedures, 9-17
 - initialization of, 8-2
 - nested, 8-12
 - omission of keyword, 8-3
 - parameters of, 8-2
 - rejection of, 8-8
 - uncaught, 8-7
- SignalCall**, 8-2
- Signaller**, 8-6
- SignalOnError**, 8-2
- SignalTC**, 8-2
- signed number, 2-5
- single-component record, 5-2
- SIZE**, 2-20
 - of array, 3-14
 - of record, 3-20
 - of type, 7-27
- space, see *blankspace*
- sparse enumeration, 3-8
- START**, 7-19
- start trap, 7-21
- Statement**, 4-1
- statement
 - assignment, 4-2
 - EXIT**, 4-19
 - GOTO**, 4-9
 - LOOP**, 4-19
 - NULL**, 4-7
 - OPEN**, 4-12
 - ordinary, 4-1
 - SELECT**, 4-5
- StatementSeries**, 4-9
- static variables, 7-33
- STOP**, 7-44
- STRING**, 6-1, 2-22
- string
 - allocation, 6-32
 - declaration of, 6-4
 - initialization of, 6-4
- string expression, 6-3
- string literal, 6-3
 - allocation, 6-3
- string variable, 6-4
- StringBody**, 6-2
 - initialization, 6-2
- Subrange**, 4-6
- subrange, 3-3
- subrange type, 3-8
- subrange variable, 3-8
- SubrangeTC**, 3-8
- succ**, 2-19
- Sum**, 2-13
- syntax
 - of C/Mesa, 7-34
 - of module, 7-4
- TAB**, 2-3
- Tag**, 3-28
- target type, 3-37
- TC**, 3-3
- Test**, 4-6
- TestList**, 4-6
- TEXT**, 6-30
- THEN**, 2-2, 4-3, 4-4, 7-40
- ThenClause**, 4-3
- THROUGH**, 4-16
- time out of condition variable, 9-11
- transfer constant, 7-12
 - exported, 7-20
- transfer item, 7-12
- transfer variable, 7-12
 - exported, 7-20
- trap, 7-33
- TRASH**, 3-24
- TYPE**, 3-1
- type
 - access control of, 7-24
 - address-containing, 3-47

- concrete, 7-25
- equivalent, 2-9
- exported, 7-25
- of procedure, 5-4
- opaque, 7-25
- ordered, 2-16, 3-3
- naming of, B-3
- procedure, 5-3
- type conformance, 2-9
- type constructor, 2-5, 3-2
- type conversion, 3-38
 - automatic, 3-38
- type definitions
 - non-interface, 7-1
- type determination, 3-3
- type identifier, 3-2
- type rule, 3-37
- type specification, 2-4
- TypeConstructor**, 3-2
- TypeDeclaration**, 3-1
- TypeIdentifier**, 3-2, 6-19, 7-8
- TypeSpecification**, 3-2, 2-22, 6-26
- uncaught signal, 8-7
- UNCOUNTED_ZONE**, 6-5
- unnamed field list
 - as **ResultList**, 5-10
- UnnamedFieldList**, 3-18
- unsigned number, 2-5, 3-35
- UNSPECIFIED**, 2-22
- UNTIL**, 2-2, 4-15
- UNWIND**, 8-8, 9-19
 - and monitor locks, 9-18
- USING**, 7-5, 5-16
 - in monitor lock clause, 9-12
 - in object monitors, 9-16
- USING** clause, 7-10
- VAL**, 2-19
- VAR**, 7-12
- Variable**, 2-12
- variable
 - READONLY**, 7-13
 - condition, 9-4
 - local, 5-2
 - nonlocal, 5-3
 - procedure, 5-4, 5-9
 - program, 7-18
 - static, 7-33
 - subrange, 3-8
- variant record, 6-1
 - accessing components of, 6-19
 - access control, 7-23
 - actual tag, 6-17
 - allocation from zones, 6-32
 - allocation of, 6-25
 - assignment, 6-19
 - bound, 6-18
 - computed tag, 6-17
 - declaration of, 6-15
 - defaults, 6-25
 - discrimination, 6-20
 - initialization, 6-15, 6-19
 - notation, 6-23
 - ragged, 6-18
 - tag, 6-14
- VariantFieldList**, 3-28
- VariantPart**, 6-26
- voiding
 - in result record, 5-12
 - of array component, 3-15
 - of component, 3-23
 - of element in returns list, 5-10
 - of record field, 3-24
- WAIT**, 9-4
- WHILE**, 4-15
- WHILE** expression
 - target type of, 3-38
- WORD**, 2-22
- zone, 6-1
 - allocation of, 6-32
 - deallocation, 6-32

Reader's Feedback

Xerox's Technical Publications Departments want to provide documents that meet the needs of all our product users. Your comments help us correct and improve our publications. Please take a few minutes to respond. If you have comments on the product this document describes, contact your Xerox representative.

1. Did you find any errors in this publication? What were they? On which pages?

2. Were there any areas that were hard to understand because of descriptions or wording? What were they? Where?

3. Did this publication give you all the information you needed? If not, what was missing?

4. Was this manual at the right level for your needs? If not, what other types of publications do you need?

5. What *one thing* could we do to improve this manual for you?

NAME _____ DATE _____

TITLE _____ COMPANY _____

ADDRESS _____

CITY _____ STATE _____ ZIP _____