



PL/M-386 Programmer's Guide

Order Number: 611052-002

In the United States, additional copies of this manual or other Intel literature may be obtained by writing:

Literature Distribution Center
Intel Corporation
P.O. Box 7641
Mt. Prospect, IL 60056-7641

Or you can call the following toll-free number: 1-800-548-4725

In locations outside the United States, obtain additional copies of Intel documentation by contacting your local Intel sales office. For your convenience, international sales office addresses are printed on the last page of this document. Contact your local sales office to obtain the latest specifications before placing your order.

Intel Corporation (Intel) makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Intel assumes no responsibility for any errors that may appear in this document. Intel makes no commitment to update nor to keep current the information contained in this document. No part of this document may be copied or reproduced in any form or by any means without prior written consent of Intel. Intel retains the right to make changes to these specifications at any time, without notice.

Intel software products are copyrighted by and shall remain the property of Intel Corporation. Use, duplication or disclosure is subject to restrictions stated in Intel's Software License Agreement.

U.S. GOVERNMENT RESTRICTED RIGHTS: These software products and documentation were developed at private expense and are provided with "RESTRICTED RIGHTS." Use, duplication, or disclosure by the Government is subject to restrictions as set forth in FAR 52.227-14 and DFAR 252.227-7013 et seq. or its successor.

The Intel logo, i960, Pentium, and iRMX are registered trademarks of Intel Corporation, registered in the United States of America and other countries. Above, i287, i386, i387, i486, Intel287, Intel386, Intel387, Intel486, Intel487 and EtherExpress are trademarks of Intel Corporation.

Adaptec is a registered trademark of Adaptec, Inc. AT, IBM and PS/2 are registered trademarks and PC/XT is a trademark of International Business Machines Corporation. All Borland products are trademarks or registered trademarks of Borland International, Inc. CodeView, Microsoft, MS, MS-DOS and XENIX are registered trademarks of Microsoft Corporation. Comtrol is a registered trademark and HOSTESS is a trademark of Comtrol Corporation. DT2806 is a trademark of Data Translation, Inc. Ethernet is a registered trademark of Xerox Corporation. Hayes is a registered trademark of Hayes Microcomputer Products. Hazeltine and Executive 80 are trademarks of Hazeltine Corporation. Hewlett-Packard is a registered trademark of Hewlett-Packard Co. MIX® is a registered trademark of MIX Software, Incorporated. MIX is an acronym for Modular Interface eXtension. MPI is a trademark of Centralp Automatismes (S.A.). NetWare and Novell are registered trademarks of Novell Corp. NFS is a trademark of Sun Microsystems, Inc. Phar Lap is a trademark of Phar Lap Software, Inc. Soft-Scope is a registered trademark of Concurrent Sciences, inc. TeleVideo is a trademark of TeleVideo Systems, Inc. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited. VAX is a registered trademark and VMS is a trademark of Digital Equipment Corporation. Visual Basic and Visual C++ are trademarks of Microsoft Corporation. All Watcom products are trademarks or registered trademarks of Watcom International Corp. Windows, Windows 95 and Windows for Workgroups are registered trademarks and Windows NT is a trademark of Microsoft in the U.S. and other countries. Wyse is a registered trademark of Wyse Technology. Zentec is a trademark of Zentec Corporation. Other trademarks and brands are the property of their respective owners.

Copyright © 1992 and 1995 Intel Corporation, All Rights Reserved

REVISION HISTORY

		DATE
-001	Original Issue	08/92
-002	Update for Release 2.2 of the OS	11/95

Quick Contents

- Chapter 1. Introduction
- Chapter 2. Language Elements
- Chapter 3. Data Declarations, Types, and Based Variables
- Chapter 4. Arrays and Structures
- Chapter 5. Expressions and Assignments
- Chapter 6. Flow Control Statements
- Chapter 7. Block Structure and Scope
- Chapter 8. Procedures
- Chapter 9. Built-In Procedures, Functions, and Variables
- Chapter 10. Features Involving the Target CPU and Numeric Coprocessor
- Chapter 11. Compiler Invocations and Controls
- Chapter 12. Sample Program
- Chapter 13. Extended Segmentation Models
- Chapter 14. Error and Warning Messages
- Appendix A. PL/M Reserved Words and Predeclared Identifiers
- Appendix B. PL/M Program Limits
- Appendix C. Grammar of the PL/M Language
- Appendix D. Differences Between PL/M Compilers
- Appendix E. Character Set
- Appendix F. Linking to Modules Written in Other Languages
- Appendix G. Run-time Interrupt Processing
- Appendix H. Run-time Support for PL/M Applications
- Index
- Service Information

Notational Conventions

The following notational conventions are used throughout this manual.

<i>Monospace Type</i>	indicates literal command syntax, and other actual input/output.
<i>italics</i>	indicate variable expressions and filenames. Substitute a value or a symbol.
<i>directory</i>	refers to a user-created directory. A forward slash (/) is used for iRMX directory paths. A backward slash (\) is used for DOS directory paths.
<i>pathname</i>	represents a fully-qualified reference to a file.

All numbers are decimal unless otherwise stated. Hexadecimal numbers include the H radix character (for example, 0FFH). Binary numbers include the B radix character (for example, 11011000B).

Contents

1 Introduction

Product Definition.....	19
Compatible Assemblers, Debuggers, and Utilities.....	20
Advantages of Using the PL/M Language.....	22
The Structure of a PL/M Program.....	23
Overview of PL/M Statements.....	24
Declaration Statements.....	24
Executable Statements.....	24
Built-in Procedures and Variables.....	25
Overview of PL/M Expressions.....	25
Input and Output.....	26

2 Language Elements

Character Set.....	27
Tokens, Separators, and the Use of Blanks.....	29
Identifiers and Reserved Words.....	30
Constants.....	30
Whole-number Constants.....	31
Floating-point Constants.....	32
Character Strings.....	33
Comments.....	33

3 Data Declarations, Types, and Based Variables

Variable Declaration Statements.....	36
Sample DECLARE Statements.....	36
Results of Variable Declarations.....	37
Combining DECLARE Statements.....	38
Initializations.....	39
The Implicit Dimension Specifier.....	42
Names for Execution Constants: the Use of DATA.....	43

Types of Declaration Statements	44
Compilation Constants (Text Substitution): The Use of LITERALLY	44
Declarations of Names for Labels	46
Results of Label Declarations	46
Declaration for Procedures	47
Data Types	48
Unsigned Binary Number Variables: Unsigned Arithmetic	50
INTEGER Variables: Signed Arithmetic	51
Signed Arithmetic	51
REAL Variables: Floating-point Arithmetic	51
Examples of Binary Scientific Notation	52
POINTER Variables and Location References	54
The @ Operator	55
Storing Strings and Constants via Location References	56
OFFSET Data Type and the Dot Operator	57
SELECTOR Variables	57
Based Variables	58
Location References and Based Variables	60
The AT Attribute	61
WORD32 WORD16 Type Mapping	64
Choosing WORD32 or WORD16	65

4 Arrays and Structures

Arrays	67
Subscripted Variables	68
Structures	69
Arrays of Structures	69
Arrays Within Structures	70
Arrays of Structures With Arrays Inside the Structures	70
Nested Structures	71
References to Arrays and Structures	72
Fully Qualified Variable References	72
Unqualified and Partially Qualified Variable References	73

5 Expressions and Assignments

Operands	75
Constants	76
Whole-number Constants in Unsigned Context	76
Whole-number Constants in Signed Context	76
String Constants	76
Variable and Location References	78

Subexpressions.....	78
Compound Operands.....	78
Arithmetic Operators.....	79
The +, -, *, and / Operators.....	79
The MOD Operator.....	82
Relational Operators.....	83
Logical Operators.....	85
Expression Evaluation.....	87
Precedence of Operators: Analyzing an Expression.....	87
Compound Operands Have Types.....	89
Relational Operators Are Restricted.....	90
Order of Evaluation of Operands.....	90
Choice of Arithmetic: Summary of Rules.....	91
Special Case: Constant Expressions.....	94
Assignment Statements.....	96
Implicit Type Conversions.....	96
Constant Expression.....	99
Multiple Assignment.....	100
Embedded Assignments.....	100

6 Flow Control Statements

DO and END Statements: DO Blocks.....	103
Simple DO Blocks.....	105
DO CASE Blocks.....	106
DO WHILE Blocks.....	108
Iterative DO Blocks.....	109
END Statement.....	112
IF Statement.....	112
Nested IF Statements.....	114
Sequential IF Statements.....	116
GOTO Statements.....	117
The CALL and RETURN Statements.....	118

7 Block Structure and Scope

Names Recognized Within Blocks.....	120
Restrictions on Multiple Declarations.....	122
Extended Scope: The PUBLIC and EXTERNAL Attributes.....	123
Scope of Labels and Restrictions on GOTOS.....	126

8	Procedures	
	Procedure Declarations.....	131
	Parameters.....	132
	Typed Versus Untyped Procedures	134
	Activating a Procedure: Function References and CALL Statements.....	135
	Indirect Procedure Activation	136
	Code Examples.....	138
	Exit from a Procedure: The RETURN Statement.....	139
	The Procedure Body	141
	Examples.....	141
	The Attributes: PUBLIC and EXTERNAL, INTERRUPT, REENTRANT	143
	Interrupts and the INTERRUPT Attribute.....	144
	Reentrancy and the REENTRANT Attribute	146

9	Built-in Procedures, Functions, and Variables	
	Obtaining Information About Variables.....	150
	The LENGTH Function.....	150
	The LAST Function.....	151
	The SIZE Function	151
	Explicit Type and Value Conversions.....	152
	The PL/M-386 LOW, HIGH, and DOUBLE Functions.....	159
	The FLOAT Function	160
	The FIX Function	160
	The INT Function.....	161
	The SIGNED Function	161
	The UNSIGN Function	162
	The Unsigned Binary Data Type Built-in Functions	163
	Signed Integer Data Type Built-in Function	163
	REAL Built-in Functions.....	164
	The SELECTOR Built-in Function.....	164
	The POINTER Built-in Function	164
	The OFFSET Built-in Function	165
	The ABS and IABS Functions	165
	Shift and Rotate Functions	166
	Rotation Functions	166
	Logical-shift Functions.....	167
	Algebraic-shift Functions	168
	Concatenate Functions.....	169

String Manipulation Procedures and Functions.....	170
The Copy String in Ascending Order Procedure	171
The Copy String in Descending Order Procedure	171
The Compare String Function.....	172
The Find Element Functions.....	173
The Find String Mismatch Function	174
The Translate String Procedure	175
The Set String to Value Procedure.....	176
PL/M-386 Bit Manipulation Built-ins.....	177
The Copy Bit String Procedure.....	177
The Find Set Bit Function	177
Miscellaneous Built-ins.....	179
The Move Bytes Procedure	179
The Time Delay Procedure.....	180
The Lock Set Function	180
The Lock Bit Functions.....	182
POINTER and SELECTOR-related Functions.....	183
The Return POINTER Value Function	183
The Return Segment Portion of POINTER Function.....	183
The Return Offset Portion of POINTER Function	183
The Set POINTER Bytes to Zero Variable	184
WORD16 Built-in Mapping	184

10 Features Involving the Target CPU and Numeric Coprocessor

Microprocessor Hardware-dependent Statements	185
The ENABLE and DISABLE Statements	185
The CAUSE\$INTERRUPT Statement.....	186
The HALT Statement	186
Microprocessor Hardware Flags	187
Optimization and the Hardware Flags.....	187
The CARRY, SIGN, ZERO, and PARITY Functions	188
The PLUS and MINUS Operators.....	188
Carry-rotation Functions.....	188
The Decimal Adjust Function.....	189
Microprocessor Hardware Registers	189
The Flags Register Access Variable	189
The STACKPTR and STACKBASE Variables.....	190

Microprocessor Hardware I/O	191
The Find Value in Input Port Function.....	191
The Access Output Port Array	191
The Read and Store String Procedure	192
The Write String Procedure	193
The Hardware Protection Model.....	193
The Task Register.....	193
The TASK\$REGISTER Variable.....	193
The Global Descriptor Table Register.....	194
The SAVE\$GLOBAL\$TABLE Procedure	195
The RESTORE\$GLOBAL\$TABLE Procedure.....	195
The Interrupt Descriptor Table Register.....	196
The SAVE\$INTERRUPT\$TABLE Procedure	196
The RESTORE\$INTERRUPT\$TABLE Procedure	197
The Local Descriptor Table Register	197
The LOCAL\$TABLE Variable.....	197
The Machine Status Register	198
The MACHINE\$STATUS Variable.....	198
The CONTROL\$REGISTER, DEBUG\$REGISTER, and TEST\$REGISTER Built-in Arrays.....	198
The CLEAR\$TASK\$SWITCHED\$FLAG Procedure.....	200
Segment Information	200
The GET\$ACCESS\$RIGHTS Function	200
The GET\$SEGMENT\$LIMIT Function	201
Segment Accessibility	202
The SEGMENT\$READABLE Function	202
The SEGMENT\$WRITABLE Function.....	202
Adjusting the Requested Privilege Level	203
The ADJUST\$RPL Function	203
The REAL Math Facility.....	204
Built-ins Supporting the REAL Math Unit.....	207
The INIT\$REAL\$MATH\$UNIT Procedure.....	207
The SET\$REAL\$MODE Procedure	207
The GET\$REAL\$ERROR Function	208
Saving and Restoring REAL Status	208
The SAVE\$REAL\$STATUS Procedure	209
The RESTORE\$REAL\$STATUS Procedure	210
Interrupt Processing.....	210
The WAIT\$FOR\$INTERRUPT Procedure	210
WORD16 Mapping for Built-ins	211
Intel486 Processor Built-ins.....	211

11	Compiler Invocation and Controls	
	Invocation Syntax on iRMX Systems	213
	Invocation Examples and Sign-on/Sign-off Messages under the iRMX OS	215
	Invocation Syntax on DOS Systems	216
	Invocation Examples and Sign-on/Sign-off Messages under DOS	216
	File Usage under DOS and the iRMX OS	217
	Input Files	217
	Work Files	217
	Print Files	218
	Object Files	218
	Executable Programs	220
	Introduction to Compiler Controls.....	221
	Input Format Control.....	226
	Code Generation and Object File Controls	226
	Segmentation Controls	226
	Listing Selection and Content Controls	227
	Listing Format Controls	227
	Source Inclusion Controls.....	227
	Conditional Compilation Controls.....	228
	Language Compatibility Control	231
	Predefined Switches	231
	Compiler Control Encyclopedia	232
	CODE NOCODE.....	232
	COND NOCOND	232
	DEBUG NODEBUG	233
	EJECT.....	233
	IF ELSE ELSEIF ENDIF	233
	INCLUDE.....	235
	INTERFACE.....	236
	LEFTMARGIN	241
	LIST NOLIST	241
	MOD486	242
	OBJECT NOOBJECT	242
	OPTIMIZE.....	243
	OVERFLOW NOOVERFLOW	258
	PAGELENGTH	258
	PAGewidth	259
	PAGING NOPAGING.....	259
	PRINT NOPRINT	259
	RAM ROM.....	260
	SAVE RESTORE	260

SET RESET.....	261
SMALL COMPACT MEDIUM LARGE FLAT.....	262
SMALL.....	262
COMPACT.....	263
MEDIUM.....	265
LARGE.....	265
FLAT.....	265
SUBTITLE.....	266
SYMBOLS NOSYMBOLS.....	266
TITLE.....	267
TYPE NOTYPE.....	267
WORD32 WORD16.....	267
XREF NOXREF.....	270
Program Listing.....	271
Sample Program Listing.....	271
Symbol and Cross-reference Listing.....	275
Compilation Summary.....	277

12 Sample Program

Introduction.....	279
FREQ Module.....	279
OPEN Module.....	283
PRINT Module.....	288
Include Files.....	294

13 Extended Segmentation Models

Overview.....	297
Introduction.....	298
Segmentation Controls Architecture Overview.....	299
Using Subsystems.....	303
Open Subsystems.....	308
Closed Subsystems.....	309
Communication Between Subsystems.....	310
Syntax.....	310
Placement of Segmentation Controls.....	313
Exporting Procedures.....	314
Large Matrix Example.....	316

14	Error and Warning Messages	
	PL/M Program Error and Warning Messages	321
	Fatal Command Tail and Control Error Messages	335
	Fatal Input/Output Error Messages	336
	Fatal Insufficient Memory Error Messages	336
	Fatal Compiler Failure Error Messages	337
	Insufficient Memory Warning Messages	337

A	PL/M Reserved Words and Predeclared Identifiers	
	Introduction	339

B	PL/M Program Limits	343
----------	----------------------------	-----

C	Grammar of the PL/M Language	
	Lexical Elements	346
	Character Sets	346
	Tokens	346
	Delimiters	346
	Identifiers	346
	Numeric Constants	347
	Strings	347
	PL/M Text Structure: Tokens, Blanks, and Comments	347
	Modules and the Main Program	348
	Declarations	349
	DECLARE Statement	349
	Variable Elements	349
	Label Element	350
	Literal Elements	350
	Factored Variable Element	350
	Factored Label Element	350
	The Structure Type	350
	Procedure Definition	351
	Attributes	351
	AT	351
	INTERRUPT	351
	Initialization	351

Units	352
Basic Statements	352
Assignment Statement	352
CALL Statement.....	352
GOTO Statement.....	352
Null Statement.....	352
RETURN Statement	352
Microprocessor-dependent Statements	353
Scoping Statements	354
Simple DO Statement	354
DO-CASE Statement	354
DO-WHILE Statement	354
Iterative DO Statement	354
END Statement.....	354
Procedure Statement.....	354
Conditional Clause	355
DO Blocks.....	355
Simple DO Blocks	355
DO-CASE Blocks.....	355
DO-WHILE Blocks	355
Iterative DO Blocks.....	355
Expressions	356
Primaries	356
Constants.....	356
Variable References.....	356
Location References	356
Operators.....	356
Structure of Expressions.....	357

D Differences Between PL/M Compilers

Differences between PL/M-86 and PL/M-80	359
Compatibility of PL/M-80 Programs and the PL/M-86 Compiler.....	360
Differences between PL/M-286 and PL/M-86	360
Compatibility of PL/M-86 Programs and the PL/M-286 Compiler.....	361
Differences between PL/M-386 and PL/M-286.....	361
Compatibility of PL/M-286 Programs and the PL/M-386 Compiler.....	362

E Character Set 363

F	Linking to Modules Written in Other Languages	
	Introduction	369
	Calling Sequence	371
	Procedure Prologue	373
	Procedure Epilogue	375
	Register Usage	376
	Segment Name Conventions	379
	C Language Compatibility	380
	Design Guidelines	381
	Code Example	381
	Compiling C and PL/M Modules.....	385

G	Run-time Interrupt Processing	
	General Information	387
	The Interrupt Descriptor Table	388
	Procedures and Tasks	388
	Interrupt Procedure Prologue and Epilogue	389
	Interrupt Tasks	392
	Exception Conditions in REAL Arithmetic	394
	Invalid Operation Exception.....	396
	Denormal Operand Exception	397
	Zero Divide Exception	397
	Overflow Exception	397
	Underflow Exception	398
	Precision Exception.....	398
	Writing a Procedure to Handle REAL Interrupts	399

H	Run-time Support for PL/M Applications	
	Numeric Coprocessor Support Libraries	405
	PL/M Support Libraries.....	406

Index		407
--------------	--	-----

Service Information	Inside Back Cover
----------------------------	-------	-------------------

Tables

1-1.	Assemblers, Debuggers, and Utilities	20
2-1.	PL/M Special Characters	28
3-1.	Declaration Elements	35
3-2.	Data Types	49
3-3.	WORD32 WORD16 Data Type Mapping	64
5-1.	Operator Precedence	79
5-2.	Summary of Expression Rules for PL/M-386	92
5-3.	Implicit Type Conversions in Assignment Statements for PL/M-386.....	97
9-1.	Value and Type Conversions for PL/M-386.....	153
11-1.	Compiler Controls.....	223
11-4.	WORD32 WORD16 Data Type Mapping	268
11-5.	WORD32 WORD16 Built-in Mapping	269
13-1.	Segmentation Controls and Memory Partitions.....	300
13-2.	Segmentation Controls, Register Addresses and Pointer Values.....	301
13-3.	Intel386 and Intel486 Microprocessor-specific ES Register Segmentation Controls, Register Addresses and Pointer Values	303
E-1.	Character Set.....	363
F-1.	Stack Representation for PL/M Parameters.....	371
F-2.	Summary of the Intel386 Microprocessor Register Usage.....	376
F-3.	Registers Used to Hold Simple Data Types	378
F-4.	Summary of PL/M-386 Segment Names.....	379
G-1.	Exception and Response Summary	395

Figures

1-1.	32-bit Protected Mode iRMX Application Development	21
3-1.	Successive Byte References of a Structure	62
7-1.	Inclusive Extent of Blocks	121
7-2.	Sample Program Modules Illustrating Valid GOTO Usage.....	128
7-3.	Sample Program Modules Illustrating Valid GOTO Transfers.....	129
10-1.	The Hardware Flags Register	190
10-2.	The REAL Error Byte	204
10-3.	The REAL Mode Word.....	205
11-1.	Sample Program Using Conditional Compilation (SET control).....	229
11-2.	Sample Program Showing the NOCOND Control.....	230
11-3.	Sample Program Showing the OPTIMIZE(0) Control	244
11-4.	Sample Program Showing the OPTIMIZE(1) Control	247
11-5.	Sample Program Showing the OPTIMIZE(2) Control.....	253
11-6.	Sample Program Showing the OPTIMIZE(3) Control	256
11-7.	Program Listing	272
11-9.	Cross-reference Listing	275
11-10.	Compilation Summary	277
12-1.	Source Code for FREQ Module.....	280
12-2.	Source Code for OPEN Module	284
12-3.	Source Code for PRINT Module	289
12-4.	Include File -- defns.inc	294
12-5.	Include File -- udi.inc.....	295
F-1.	Stack Layout at Point Where a Non-interrupt Procedure is Activated.....	372
F-2.	Stack Layout During Execution of a Non-interrupt Procedure Body.....	374
G-1.	Stack Layout at Point Where an Interrupt Procedure Gains Control.....	390
G-2.	Stack Layout during Execution of Interrupt Procedure Body	391
G-3.	Tag Word Format.....	401
G-4.	Memory Layout of the REAL Save Area in Protected Mode for the Intel386 Microprocessor	403

Introduction 1

This chapter introduces the PL/M-386 compiler and explains the process of developing software for execution by an Intel386™ or Intel486™ microprocessor-based system.

Product Definition

The PL/M-386 compiler is a software tool that translates PL/M source code into a relocatable object module. The object modules (in OMF386 format) are compatible with all other Intel OMF386-producing translators, such as ASM386, iC-386, and Fortran-386.

The PL/M-386 compiler translates a PL/M source text file into an object module and a listing file. Parts or all of a program can be compiled in a single compilation. Object modules can then be linked or bound in various combinations to form different applications. These applications can be run on a DOS host or as part of the iRMX® Operating System (OS) or its Human Interface layer.

The PL/M-386 compiler also provides a listing output, error messages, and a number of compiler controls that aid in developing and debugging programs.

⇒ **Note**

For information on invoking the compiler, see Chapter 11, Compiler Invocation and Controls. That chapter covers invocation on both iRMX and DOS-hosted systems.

Compatible Assemblers, Debuggers, and Utilities

Table 1-1 shows the compatible Intel assemblers, debuggers, and utilities. Figure 1-1 shows the role of these software tools in developing an iRMX application.

Table 1-1. Assemblers, Debuggers, and Utilities

Tool	Name for Intel386 and Intel486 Processors
assembler	ASM386
C compiler	iC-386
FORTRAN compiler	Fortran-386
debuggers	System Debugger (SDB) System Debug Monitor (SDM) Soft-Scope III Debugger
binder	BND386
librarian	LIB386
cross-reference	MAP386

These tools support modular application development. Refer to the following publications for further information:

- ASM386 - *ASM386 Macro Assembler Operating Instructions and ASM386 Assembly Language Reference*
- iC-386 - *iC-386 Compiler User's Guide*
- Fortran-386 - *Fortran 386 Compiler User's Guide*
- SDB and SDM - *iRMX System Debugger Reference*
- Soft-Scope III Debugger - *Soft-Scope® Debugger User's Guide*
- BND386 - *Intel386 Family Utilities User's Guide*
- LIB386 - *Intel386 Family Utilities User's Guide*
- MAP386 - *Intel386 Family Utilities User's Guide*

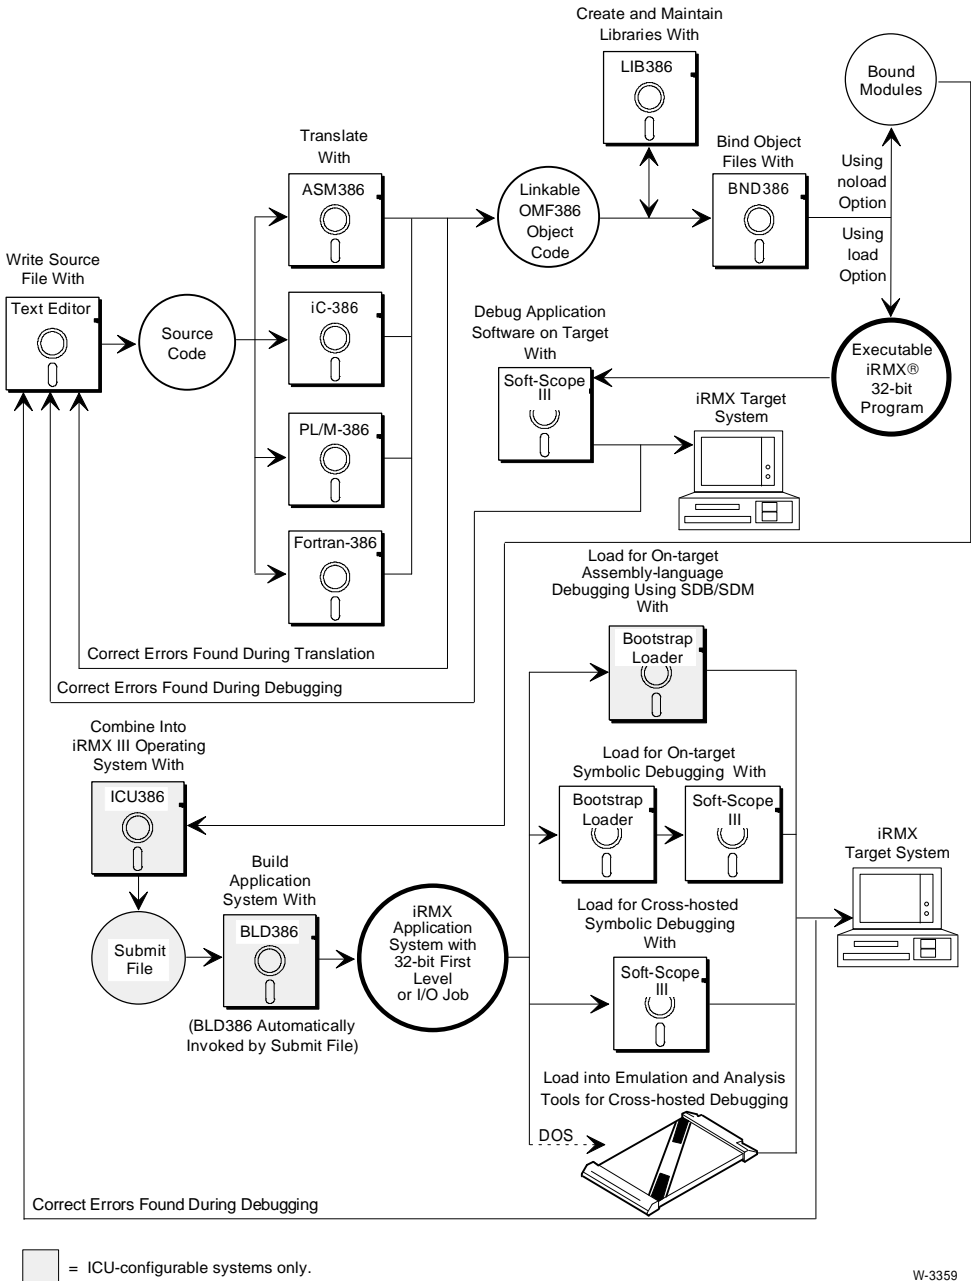


Figure 1-1. 32-bit Protected Mode iRMX Application Development

Advantages of Using the PL/M Language

PL/M programs are portable, which means that they are easily transferred from one microprocessor to another. When using PL/M, you need not be concerned with the instruction set of the target processor. Additionally, there is no need to be concerned with other details of the target processor, such as register allocation or assigning the proper number of bytes for each data item. The PL/M-386 compiler does these functions automatically. PL/M keywords and phrases are close to natural English, and many operations (including arithmetic and Boolean operations) can be combined into expressions. This enables the execution of a sequence of operations with just one program statement. Data types and data structures have functional attributes. For instance, in PL/M, the program can be written in terms of Boolean expressions, characters, and data structures, in addition to bytes, words, and integers.

Coding programs in a high-level language rather than assembly language involves thinking closer to the level used when planning the overall system design. Following is a list of the advantages of using PL/M, and the applications for which PL/M is best suited:

- PL/M block structure and control constructs aid and encourage structured programming.
- PL/M has facilities for data structures such as structured arrays and pointer-based dynamic variables.
- PL/M is a typed language. The compiler does data type compatibility checking during compilation to help detect logic errors in programs.
- PL/M data structuring facilities and control statements are designed in a logically consistent way. Thus, PL/M is a good language for expressing algorithms for systems programming.
- PL/M is a standard language used for application development on Intel systems. PL/M programs are compatible across the Intel386 and Intel486 family of microprocessors.
- PL/M was designed for programmers (generally systems programmers) who need access to the microprocessor's features such as indirect addressing and direct I/O for optimum use of all system resources.

In comparison with other languages, PL/M has more features than BASIC and is a more general-purpose language than either FORTRAN (best suited for scientific applications) or COBOL (designed for business data processing). PL/M accesses the microprocessor hardware features more easily than C. Additionally, in comparison to C, PL/M offers the ability to nest procedures and the program structure is easier to maintain.

The Structure of a PL/M Program

PL/M is a block-structured language; every statement in a program is part of at least one block. A block is a well-defined group of statements that begins with a `DO` statement or a procedure declaration and ends with an `END` statement.

A module is a labeled simple `DO`-block. A module must begin with a labeled `DO` statement and end with an `END` statement. Between the `DO` statement and the `END` statement other statements provide the definitions of data and processes that make up the program. These statements are said to be part of the block, contained within the block, or nested within the block. A module can contain other blocks but is never itself contained within another block. See Chapter 6 for a description of `DO`-blocks.

Every PL/M program consists of one or more modules, separately compiled, each consisting of one or more blocks. The two kinds of blocks are `DO`-blocks and procedure definition blocks.

A procedure definition block is a set of statements beginning with a procedure declaration and ending with an `END` statement. Other declarations and executable statements can be placed between these points, and are used later when the procedure is actually invoked or called into execution. The definition block is a further declaration of everything the procedure will use and do.

Overview of PL/M Statements

The two types of statements in PL/M are declarations and executable statements. All PL/M statements end with a semicolon (;).

Declaration Statements

The following is a simple example of a declaration statement:

```
DECLARE WIDTH BYTE;
```

This statement introduces the identifier `WIDTH` and associates it with the contents of 1 byte (8 bits) of memory. Now, rather than having to know the memory address of this byte, you can refer to it by the name `WIDTH`.

A group of statements intended to perform a function (i.e., a subprogram or subroutine) can be given a name by declaring them to be a procedure:

```
ADDER_UPPER: PROCEDURE (BETA) BYTE;
```

The statements that define the procedure follow the semicolon. This block of PL/M statements is invoked from other points in the program, and may involve passing parameters to the program. When a procedure has finished executing, control is returned immediately to the main program. This capability is the major feature enabling modular program construction.

Executable Statements

The following is an example of an executable statement:

```
CLEARANCE = WIDTH + 2;
```

The two identifiers, `CLEARANCE` and `WIDTH`, must be declared prior to this executable statement, which produces machine code to retrieve the `WIDTH` value from memory. Once the `WIDTH` value is obtained, 2 is added to it and the sum is stored in the memory location for `CLEARANCE`.

For most purposes, it is unnecessary to think in terms of memory locations when programming in PL/M. `CLEARANCE` and `WIDTH` are variables, and the assignment statement assigns the value of the expression `WIDTH + 2` to the variable `CLEARANCE`. The compiler automatically generates all the machine code necessary to retrieve data from memory, to evaluate the expression retrieved, and to store the result in the proper location.

Executable statements are discussed in the following chapters:

Assignment Statement	Chapter 5
CALL Statement	Chapter 8
CAUSE\$INTERRUPT Statement	Chapter 10
DISABLE Statement	Chapter 10
DO CASE Statement	Chapter 6
DO WHILE Statement	Chapter 6
ENABLE Statement	Chapter 10
END Statement	Chapter 6
Executable Functions	Chapter 9
GOTO Statement	Chapter 6
HALT Statement	Chapter 10
IF Statement	Chapter 6
Iterative DO Statement	Chapter 6
Nested IF Statement	Chapter 6
RETURN Statement	Chapter 8
Simple DO Statement	Chapter 6

Built-in Procedures and Variables

PL/M provides a variety of built-in procedures and variables. These include functions such as shifts and rotations, data type conversions, executable functions, block I/O, real math, and string manipulation (see Chapters 9 and 10).

Overview of PL/M Expressions

A PL/M expression is made up of operands and operators, and resembles a conventional algebraic expression.

Operands include numeric constants (such as 3.78 or 105) and variables (as well as other types discussed in Chapters 3 and 5). The operators include + and - for addition and subtraction, * and / for multiplication and division, and MOD for modular arithmetic.

As in an algebraic expression, elements of a PL/M expression can be grouped with parentheses.

An expression is evaluated using unsigned binary arithmetic, signed integer arithmetic, and/or floating-point arithmetic, depending on the types of operands in the expression (see Chapters 3 and 5).

Input and Output

PL/M does not provide formatted I/O capabilities like those of FORTRAN, BASIC, or COBOL. However, PL/M does provide built-in functions for direct I/O that do not require operating system run-time support. The PL/M-386 compiler has built-in functions which allow for single-byte, half-word or word I/O, as well as for block I/O (for strings of bytes, half-words, or single-words). For detailed information on these I/O functions, see Chapter 10.

□ □ □

Language Elements **2**

PL/M-386 programs are free-form, meaning there are no restrictions on where you place a statement on a line. You can use as many blanks (spaces) as necessary to format your program for readability.

Character Set

The PL/M-386 source program character set is the following subset of the ASCII character set:

A . . Z
a . . z
0 . . 9

and the following special characters:

= . / () + - ' * , < > : ; @ \$ _

and the blank (space), tab, carriage-return and line-feed characters. (Appendix E indicates if each ASCII character is a member of the PL/M character set and, if so, the hexadecimal value.)

PL/M does not distinguish between uppercase and lowercase letters, except in string constants. For example, the variable names `xyz` and `XYZ` are the same. (In this manual, all PL/M syntax is uppercase, by convention.)

Special characters have particular meaning in PL/M, as explained throughout this manual. Table 2-1 summarizes the meaning of special characters in PL/M.

Table 2-1. PL/M Special Characters

Symbol	Name	Use
=	equal sign	Two distinct uses: (1) assignment operator (2) relational test operator
:=	assign	embedded assignment operator
@	at-sign	location reference operator
.	dot	Three distinct uses: (1) decimal point (2) structure member qualification (3) address operator
/	slash	division operator
/* */		beginning-of-comment delimiter end-of-comment delimiter
()	left parenthesis right parenthesis	left delimiter of lists, subscripts, some expressions right delimiter of lists, subscripts, some expressions
+ -	plus minus	addition or unary plus operator subtraction or unary minus operator
'	apostrophe	string delimiter
*	asterisk	Two distinct uses: (1) multiplication operator (2) implicit dimension specifier
< > <= >= <>	less than greater than less or equal greater or equal not equal	relational test operator relational test operator relational test operator relational test operator relational test operator
: ;	colon semicolon	label terminator statement terminator
,	comma	list element delimiter
_	underscore	significant character in identifier
\$	dollar sign	Two distinct uses: (1) non-significant character embedded within number of identifier (2) significant as the first character on a control line in a source file

The PL/M compiler treats multiple contiguous blanks in PL/M source programs as single blanks, by ignoring all the blanks except the first one.

The compiler produces an error or warning message whenever it encounters a character other than those described above in a source program.

In addition to the source character set, PL/M allows the use of special character sets (such as Kanji characters), located from 0080H through 00FFH (excluding 0081H).

Tokens, Separators, and the Use of Blanks

The smallest meaningful unit of a PL/M statement is a token. Every token belongs to one of the following classes:

- Identifiers
- Reserved words
- Simple delimiters (all of the special characters, except the dollar sign, are simple delimiters)
- Compound delimiters (combinations of two special characters):

`<>, <=, >=, :=, /*, */`

- Numeric constants
- Character string constants

It is usually clear where one token ends and the next one begins. For example, in the assignment statement:

```
EXACT=APPROX*(HEIGHT-3)/SCALE;
```

EXACT, APPROX, HEIGHT, and SCALE are identifiers, 3 is a numeric constant, and all the other characters are simple delimiters.

If a delimiter (simple or compound) does not naturally occur between two tokens, you must separate them with one or more blank(s).

A comment can also be used as a separator.

Blanks can be inserted around any token without changing the meaning of the PL/M statement. Thus, the assignment statement:

```
EXACT = APPROX * ( HEIGHT - 3 ) / SCALE;
```

is equivalent to:

```
EXACT=APPROX*(HEIGHT-3)/SCALE;
```

Identifiers and Reserved Words

Identifiers name variables, procedures, symbolic constants, and statements. Statement identifiers are called labels. Identifiers can be up to 31 characters long. The first character must be alphabetic or the underscore (`_`), and the remaining characters may be alphabetic, numeric, or the underscore.

You can use the dollar sign character to improve the readability of an identifier or constant, but the dollar character is not meaningful to the compiler. An identifier or constant containing a dollar sign is equivalent to the same identifier without the dollar sign. Note that you must not use a dollar character in a procedure name within a subsystem definition. See Chapter 13.

Examples of valid identifiers are:

```
INPUT_COUNT
X
GAMM
LONGIDENTIFIERNUMBER3
LONG$$$IDENTIFIER$$NUMBER$$$3
_MAIN
INPUT$COUNT
INPUTCOUNT
```

The long identifiers are identical to the compiler. `INPUT$COUNT` and `INPUTCOUNT` are interchangeable, but are different from `INPUT_COUNT`.

Identifiers must be distinct from reserved words. If you want to use PL/M built-in procedures and variables, the identifiers in your source program must be distinct from the built-ins' predefined identifiers. Appendix A lists the reserved words and predefined identifiers.

Constants

A constant is a value that does not change during a program's execution. The three types of constants are whole-number constants, floating-point constants, and character strings.

Whole-number Constants

Whole-number constants can be binary, octal, decimal, or hexadecimal numbers. Specify the base of these constants by appending a B, Q, D, or H suffix. The compiler interprets numbers without a base suffix as decimal numbers. When they encounter characters that are invalid in the specified (or assumed) base, the compiler produces appropriate messages. If a constant contains characters invalid in the designated number base, it will be flagged as an error.

In PL/M-386, a whole-number constant can be an 8-bit, 16-bit or 32-bit value. It can also be a 64-bit value. The range of whole-number constants is non-negative. (The minus sign in front of a whole-number constant is not part of the constant.)

The first character of a hexadecimal number must be a numeric digit to avoid looking like an identifier. For example, write the hexadecimal form of the decimal value 163 as 0A3H (rather than A3H); otherwise the compiler will interpret it as an identifier.

Examples of valid whole-number constants are:

```
12AH 2 33Q 1010B 55D 0BF3H 65535 7770 3EACH 0F76C05H
```

Examples of invalid whole-number constants are:

- | | |
|-------|--|
| 12AF | Hexadecimal digits used without an H suffix, and invalid in the default decimal interpretation. |
| 12AD | The final D could be a suffix but the A is not a decimal digit. If hexadecimal is intended, a final H is needed. |
| 11A2B | A and 2 are not valid binary digits. If hexadecimal is intended, a final H is necessary. |
| 2ADGH | G is not a valid hexadecimal digit. |

For example, the maximum whole-number 16-bit constant is:

```
2**16-1 = 1111$1111$1111$1111B = 177777Q = 65535D = 0FFFFH
```

The maximum whole-number 32-bit constant is:

```
2**32-1 = 1111$1111$1111$1111$1111$1111$1111$1111B
        = 377777777777Q
        = 4294967295D
        = 0FFFFFFFFH
```

Floating-point Constants

The presence of a decimal point in a decimal constant creates a floating-point constant. Floating-point constants are represented in REAL precision (see Duty Types). Only decimal real constants are allowed.

At least one decimal digit (e.g., 0) must precede the decimal point. A fractional part is optional after the decimal point, as is the base-ten exponent, which is indicated by the letter E. This exponent must have at least one digit. Note that no fractional exponents are possible.

In PL/M-386, the range is $-2^{+(128)}$ to $-2^{(-126)}$, zero, $+2^{(-126)}$ to $+2^{+(128)}$. This range is approximately -3.4×10^{38} to $-8.4 \times 10^{(-37)}$, zero, and $8.4 \times 10^{(-37)}$ to 3.4×10^{38} .

The following are examples of valid floating-point constants:

```
5.30      176.0      1.88      3.14159      16.      222.2
53.0E-1   1.760E2     0.188E1   314159.E-5   1.6E+1   2.222E+2
```

Note that plus signs do not change the meaning of exponents.

The following are examples of invalid floating-point constants:

```
6          No decimal point
1.3AH      Hexadecimal not allowed in floating-point constants
10.011B    Binary not allowed
7.52Q      Octal not allowed
4.8E1AH/2  Only decimal constants in exponents; no hexadecimal, no expressions,
           no fractions
```


Character Strings

Character strings are printable ASCII characters enclosed within apostrophes. There are two types of character strings: 1) string constants and 2) character constants. A string constant is used to initialize variables or to pass a pointer. The maximum length of a string constant is 255. A character constant is used in expressions, and its value should fit into a double or machine word (32 bits). A string used as a character constant can contain from one to four characters.

To include an apostrophe in a string, write it as two apostrophes (e.g., the string `' 'Q'` comprises 2 characters, an apostrophe followed by a Q). Values 0080H through 00FFH (excluding 0081H) can be used in a quoted character string. Spaces are allowed but line-feeds are not. The compiler represents character strings in memory as ASCII codes, one 7-bit character code to each 8-bit byte, with a high-order zero bit. Strings of length 1 translate to single-byte values. Character constants of length 2 translate to 16-bit values, and those of length 3 or 4 translate to 32-bit values. For example:

```
'A' is equivalent to 41H
'AG' is equivalent to 4147H
'AGR' is equivalent to 414752H
'AGRX' is equivalent to 41475258H
```

Therefore, character constants can be used as 8-bit, 16-bit, or 32-bit values. Character constants longer than 4 characters exceed the 32-bit capacity. See also Appendix E, Character Set.

Comments

In PL/M, a comment is a sequence of characters delimited on the left by the character pair `/*` and on the right by the character pair `*/`. These delimiters instruct the compiler to ignore any text between them and to consider such text as not part of the program.

A comment can contain any printable ASCII or special character and can also include space, carriage-return, line-feed, and tab characters. If you embed a comment in a character string constant, it becomes part of the constant. A comment can appear anywhere that a blank character can appear except embedded within a token.

The following is an example of a PL/M comment:

```
/*This procedure copies one structure to another.*/
```

In this manual, comments are presented in lowercase to distinguish them visually from program code, which is presented in uppercase.



Data Declarations, Types, and Based Variables **3**

In PL/M-386, you can declare symbolic names for variables, constants, procedures and statements (labels). For each symbolic name, there must be one declaration at the beginning of the block containing the name, or in an outer, enclosing block. A declaration consists of an identifier, type, attributes and/or location. Multiple declarations of a name in a block are invalid. Required and optional declaration elements are shown in Table 3-1.

Table 3-1. Declaration Elements

Declaration	Must Use	Can Use
Variable Names	BYTE, INTEGER, CHARINT, SHORTINT, LONGINT, OFFSET, WORD, QWORD, HWORD, DWORD, REAL, STRUCTURE, ADDRESS*	linkage attributes:** PUBLIC or EXTERNAL; or location attributes: AT (location reference) variable initialization attribute: INITIAL (value-list)
Constant Names	type, as above, and constant initialization attribute: DATA (value-list)	linkage attributes as above
Label Names	LABEL	linkage attributes as above
Macros	LITERALLY 'string'	

* ADDRESS is equivalent to the OFFSET data type.

** Placement is important (see Variable Declaration Statements).

The declaration of a variable or constant identifier must precede use of the identifier in an executable statement. Although it is not good programming practice, you can call a reentrant procedure before defining it. You can either explicitly declare a statement label, or implicitly declare it by attaching it to an executable statement with a colon character.

Variable Declaration Statements

A `DECLARE` statement is a nonexecutable statement that introduces some object or collection of objects, associates names (and sometimes values) with them, and allocates storage if necessary. The most important use of `DECLARE` is for declaring variables.

A variable can be a scalar (i.e., a single quantity), an array, or a structure.

A scalar variable is a single object whose value may not be known at compile time and may change during the execution of the program.

An array is a list of scalars of the same data type, referred to by one identifier and distinguished by the subscript associated with each scalar.

A structure is an aggregate of scalars, arrays and/or structures with the same main identifier. The members of a structure are differentiated from each other by their own member-identifiers or field names. For example, `EMPLOYEES.NAME` would refer to the `NAME` field within the structure `EMPLOYEES`.

Sample DECLARE Statements

Note that when using linkage (`PUBLIC/EXTERNAL`) and initialization (`DATA/INITIAL`) attributes, the order of declaration is critical. Place linkage attributes before the initialization attribute, and after the type declaration.

For example:

```
DECLARE a$P BYTE PUBLIC INITIAL(4);
```

The following statements declare scalars:

```
DECLARE APPROX REAL;
DECLARE (OLD, NEW) BYTE;
DECLARE POINT WORD, VAL12 BYTE;
```

The first example declares a single scalar variable of type `REAL`, with the identifier `APPROX`.

The second example declares two scalars, `OLD` and `NEW`, both of type `BYTE`. This kind of statement is called a factored declaration, which is similar to the sequence:

```
DECLARE OLD BYTE;
DECLARE NEW BYTE;
```

A factored declaration (for structures and arrays) guarantees that the bytes will be contiguously located in memory, which may be useful in real-time applications (see also `Combining DECLARE Statements`). Separate declaration statements do not guarantee this.

The third example declares two scalars of different types: POINT is of type WORD, and VAL12 is of type BYTE.

The following statements declare arrays:

```
DECLARE DOMAIN (128) BYTE;  
DECLARE GAMMA (19) DWORD;
```

The first example declares the array DOMAIN, with 128 scalar elements of type BYTE. These elements are distinguishable by subscripting the name DOMAIN, using the range 0 to 127 for the subscripts. For example, the third element of DOMAIN can be referred to as DOMAIN(2). The first element of every array has subscript 0.

The second example declares the array GAMMA, with 19 scalar elements of type DWORD. The subscripts for this array can range from 0 to 18.

The third example declares a structure with two scalar members:

```
DECLARE RECORD STRUCTURE (KEY BYTE, INFO WORD);
```

The two members are a BYTE scalar that can be referred to as RECORD.KEY and a WORD scalar that can be referred to as RECORD.INFO. The word named by RECORD.INFO is the second and third bytes of this structure.

Structures are discussed in further detail in Chapter 4.

Results of Variable Declarations

Valid variable declarations result in the following:

- The name is given a unique address.
- The variable is considered to have the attributes declared.

All subsequent uses of the variable in the block where it is declared refer to the same address (except for based variables, discussed in Based Variables).

A valid variable declaration also requires all references to the variable to conform to the rules for the current attributes (i.e., those attributes having priority in the current block). Thus, the compiler can flag a large variety of errors caused by incompatible references within the current block. The variable reference must be consistent with the variable declaration.

Combining DECLARE Statements

A separate DECLARE statement is not required for each declaration. For example, instead of writing the two DECLARE statements:

```
DECLARE CHR BYTE INITIAL ( 'A' );  
DECLARE COUNT INTEGER;
```

Both declarations can be written in a single DECLARE statement, as follows:

```
DECLARE CHR BYTE INITIAL ( 'A' ), COUNT INTEGER;
```

This declare statement contains two declaration elements, separated by a comma. A declaration element is the text for declaring one identifier (or one factored list of identifiers). Every DECLARE statement contains at least one declaration element. If a DECLARE statement contains more than one declaration element, they are separated by commas.

Most of the examples shown previously have only one declaration element in each DECLARE statement. In the preceding example, the text CHR BYTE INITIAL ('A') is one declaration element; the text COUNT INTEGER is another.

Another way of combining declaration elements is called a factored declaration as indicated above in this section. For example, the non-factored declarations:

```
DECLARE A BYTE, B BYTE;  
DECLARE C WORD, D WORD;  
DECLARE E DWORD, F DWORD;
```

can be combined as:

```
DECLARE (A,B) BYTE, (C,D) WORD, (E,F) DWORD;
```

In each factored declaration, the allocated locations are contiguous. Elements declared in a nonfactored declaration statement are not necessarily contiguous.

Use factored declarations if the order in which variables are allocated is important.

Variables declared in a factored declaration (i.e., variables within a parenthesized list that are not based, are not used as parameters, or are not EXTERNAL), are stored contiguously in the order specified. (If a based variable occurs in a parenthesized list, the variable is ignored when storage is allocated.)

The declaration elements in a single DECLARE statement are independent of each other, as if they were declared in separate DECLARE statements.

Initializations

Initialization guarantees that the variables being initialized have a particular value before program execution begins. Every constant should be initialized. Variables can also be initialized. There are no default values for constants or variables. Of course, variables can be initialized by an assignment statement such as the following:

```
PI = 3.1415927;      /* PI must first be declared REAL */
VAR13 = 10;         /* VAR13 must be declared earlier */
```

However, in PL/M-386, the compiler can set up these values during the compilation rather than using both instruction space and execution time to initialize variables in the program.

There are two kinds of compile-time initializations: `INITIAL`, used with variables, and `DATA`, used for constants. (`DATA` is explained in greater detail later in this section.) In both initializations, the initialization attribute is placed after the type in the declaration. For example:

```
DECLARE FAMILY WORD INITIAL (2);
```

Additionally, when using a linkage attribute (`PUBLIC/EXTERNAL`), place the linkage attribute after the type declaration and before the initialization attribute.

`INITIAL` causes initialization to occur during program loading for variables that have storage allocated for them. Such variables can subsequently be changed during execution (just as any other variable). These variables will not be reinitialized on a program restart.

The following rules apply to both `INITIAL` and `DATA`:

- `INITIAL` and `DATA` cannot be used together in the same declaration.
- `INITIAL` can occur only in declarations at the outer level of a module. `DATA`, however, can occur in declarations at any level.
- No initializations are permitted with based variables, formal parameters (see Chapter 8), or with the `EXTERNAL` attribute (see Chapter 7).
- Either `INITIAL` or `DATA` can follow use of the `AT` attribute. However, if this use of `INITIAL` or `DATA` causes multiple initializations, the result cannot be predicted.

- The initializing value should fit into the space allocated by the data type. The only exception is initialization of `HWORD` when the offset is derived with a dot operator. For example:

```
DECLARE HH HWORD INITIAL (.B)
```

In this case, the real offset is truncated to give the lower 16 bits. A warning message is issued when an `OFFSET` value is truncated.

The general form of the `INITIAL` attribute is as follows:

```
INITIAL (value-list)
```

Where:

value-list is a sequence of values separated by commas.

Values are taken one at a time from the value list and used to initialize the individual scalars being declared. The initialization is performed in the same manner as an assignment. Initial values for members of an array or structure must be specified explicitly. For character string constants, the characters are taken one at a time to initialize an 8-bit scalar, two at a time to initialize a 16-bit scalar, four at a time to initialize a 32-bit scalar, and eight at a time to initialize a 64-bit scalar.

The expressions used with the `INITIAL` attribute have the following restrictions:

- For real variables only: An expression, which can contain a unary `+` or `-` operator, can only be a single floating-point constant which can be used to initialize a `REAL` scalar only.
- For `POINTER` variables only: A restricted expression can be a location reference formed with the `@` operator, which must refer to a variable already declared or to a constant list.
- For all other types (except `SELECTOR`): A restricted expression can be a constant expression containing no operators except `+` or `-`. A constant expression has only whole-number constants as operands (e.g., `2048, 256+5`), as explained in Chapter 5. The constant expression is evaluated as if it were being assigned to the scalar being initialized, using the rules described in Chapter 5.
- For `OFFSET` or `WORD` variables only: A constant expression containing only the `+` and `-` operators, and operands that can be whole-number constants and/or `"."` location references. If the expression contains a `"."` location reference, only the `+` operator can precede it. Any combination of `+` and `-` operators can follow the `"."` location reference. For example: `5+ .xyz-10`.

The declaration:

```
DECLARE THRESHOLD BYTE INITIAL (48);
```

declares the BYTE scalar THRESHOLD and initializes the scalar to a value of 48.

The declaration:

```
DECLARE EVEN (5) BYTE INITIAL (2, 4, 6, 8, 10);
```

declares the BYTE array EVEN and initializes its five scalar elements to 2, 4, 6, 8, and 10, respectively.

The declaration:

```
DECLARE COORD STRUCTURE (HIGH$BOUND WORD,  
    VALUE (3) BYTE,  
    LOW$BOUND BYTE) INITIAL (302, 3, 6, 12, 0);
```

declares the structure COORD and initializes it as follows:

```
COORD.HIGH$BOUND to 302  
COORD.VALUE(0) to 3  
COORD.VALUE(1) to 6  
COORD.VALUE(2) to 12  
COORD.LOW$BOUND to 0
```

If a string occurs in the value list, it is taken apart from left to right and the pieces are stored in the scalars being initialized. One character is stored in each BYTE scalar, two characters in each WORD scalar, and four in each DWORD scalar. For example:

```
DECLARE GREETING (5) BYTE AT (@HI) INITIAL ('HELLO');
```

causes GREETING(0) to be initialized with the ASCII code for H, GREETING(1) with the ASCII code for E, and so on.

All the examples shown previously have had value lists that match up one-for-one with the scalars being declared. The value list can have fewer elements than are being declared. Thus:

```
DECLARE DATUM (100) BYTE INITIAL (3, 5, 7, 8);
```

will work. The first four elements of the array DATUM are initialized with the four elements in the value list, and the remainder of the array is left uninitialized. However, the value list cannot have more elements than are being declared.

The Implicit Dimension Specifier

Often, when initializing an array, you want the array to have the same number of elements as the value list. This can be done conveniently by using the implicit dimension specifier in place of an ordinary dimension specifier (a parenthesized constant). The implicit dimension specifier has the form:

```
( *)
```

Also use the implicit dimension specifier to define an external or based array whose precise number of elements is either unknown or insignificant. Thus the declaration:

```
DECLARE FAREWELL(*) BYTE PUBLIC INITIAL ('GOODBYE, NOW');
```

declares a public BYTE array, FAREWELL, with enough elements to contain the string 'GOODBYE, NOW' (namely 12), and initializes the array elements with the characters of the string. To reference this array in another program module, declare it as follows:

```
DECLARE FAREWELL(*) BYTE EXTERNAL;
```

See Chapter 7 for more information about PUBLIC and EXTERNAL attributes.

Note that the INITIAL and DATA *value-lists* must not be present when the implicit dimension specifier is used with an external array; otherwise, INITIAL and DATA *value-lists* are required. Also, the LENGTH, LAST, and SIZE built-ins cannot be used on an external array that was declared with the implicit dimension specifier.

The following is an example of an implicit dimension in a based declaration:

```
DECLARE X BASED P(*) BYTE;
```

The implicit dimension specifier cannot be used after the parenthesized list of identifiers in a factored declaration (unless it is declared EXTERNAL). Additionally, an implicit dimension specifier cannot be used to specify an array that is a member of a structure.

The implicit dimension specifier can be used with any value list; it is not restricted to strings.

Names for Execution Constants: the Use of DATA

A variable is the name of a single data item intended to be used and altered by a program. If the variable is not altered during execution, it is a constant.

For example, the formula for the circumference of a circle ($R \times 2 \times \pi$) or (radius $\times 2 \times \pi$) could be written in PL/M as:

```
C = R * 2.0 * 3.14159;
```

in which *C* and *R* would be variables. The declarations for *C* and *R* would have to precede the executable statement, and could appear as:

```
DECLARE (C, R) REAL;
```

If *pi* is used often enough, simplify writing of statements by using *PI* to declare a symbolic name with that value as follows:

```
DECLARE PI REAL DATA (3.1415927);
```

An array of constants requires a list of values. For example:

```
DECLARE FIBONACCI(9) BYTE DATA (0,1,1,2,3,5,8,13,21);
```

The form and use of the *DATA* initialization is identical to that of *INITIAL* except for the following differences:

- *DATA* causes storage to be allocated in the program's constant data segment. The content and meaning of the name cannot be changed during execution. The name should never appear on the left-hand side of an assignment statement. This is not the case with *INITIAL*.
- *DATA* initializations can be used in declarations at any block level in the program. *INITIAL* can occur only at the module level, that is, inside the *DO*-block that is the module itself, and outside any sub-blocks that the module may contain.
- If the keyword *DATA* is used in a *PUBLIC* declaration when compiling with the *ROM* option, *DATA* must also be used in the *EXTERNAL* declaration of program modules that reference it. However, no *value-list* can be used since the data is defined elsewhere. *INITIAL* cannot be combined with *EXTERNAL*.
- Use of the *AT* attribute forces a name to be associated with a specific memory location, which can defeat the purpose of the *DATA* initialization. This will not happen with *INITIAL* unless the variables and locations are explicitly redefined using multiple *AT*s.
- If the first declaration has a data initialization, then the variable that is *AT* that location is also referred to as *DATA*, i.e., cannot have a value assigned into it.

Types of Declaration Statements

Compilation Constants (Text Substitution): The Use of LITERALLY

If the program is large enough to have many declarations, declaring a compilation constant will save time at the keyboard, as follows:

```
DECLARE DCL LITERALLY 'DECLARE' ;
```

Thereafter, during compilation, every time `DCL` appears alone (not as part of a word), the full string `DECLARE` will be substituted by the compiler. Subsequent declarations can be written as follows:

```
DCL AREA REAL ;  
DCL SIZE WORD ;
```

A declaration using the reserved word `LITERALLY` defines a parameterless macro for expansion at compile-time. Declare an identifier to represent a character string, which will then be substituted for each occurrence of the identifier in subsequent text. This expansion will not take place in strings or constants. The form of the declaration is:

```
DECLARE identifier LITERALLY 'string' ;
```

Where:

identifier is any valid PL/M identifier.

string is a sequence of arbitrary characters (limited by the size of the symbol table) from the PL/M set (except an apostrophe).

An apostrophe can be included in a string by writing it as two consecutive apostrophes.

The following example illustrates another use of LITERALLY:

```
DECLARE TRUE LITERALLY 'OFFH', FALSE LITERALLY '0';

DECLARE ROUGH BYTE;
DECLARE (X, Y, DELTA, FINAL) REAL;
. . .
ROUGH = TRUE;
DO WHILE ROUGH;
    X = SMOOTH (X, Y, DELTA);
        /* SMOOTH is a procedure declared elsewhere. */
    IF (X-FINAL) < DELTA THEN
        ROUGH = FALSE;
END;
. . .
```

This example of a LITERALLY declaration defines the Boolean values TRUE and FALSE in a manner consistent with the way PL/M handles relational operators (see Chapter 5). Literal substitution for fixed values makes a program more readable.

LITERALLYs can also be used to declare quantities that are fixed for one compilation, but are subject to change from one compilation to the next. Consider the following example:

```
DECLARE BUFFER$SIZE LITERALLY '32';
DECLARE PRINT$BUFFER(BUFFER$SIZE) WORD;
. . .
PRINT$BUFFER(BUFFER$SIZE - 10) = 'G';
. . .
```

A future change to BUFFER\$SIZE can be made in one place, at the first declaration, and the compiler will propagate the change throughout the program during compilation. This eliminates the need to search the program for the occurrences of 32 that are BUFFER\$SIZE references and not some other reference to 32.

Declarations of Names for Labels

A label marks the location of an instruction. Labels are permitted only on executable statements, not on declarations.

A name can be declared as a label both explicitly and implicitly. Explicit label declarations are used mainly to enable module-to-module references (see Chapter 7). The three explicit label declarations have the following formats:

```
DECLARE PART3 LABEL;  
DECLARE START1 LABEL PUBLIC; /* for intermodule reference */  
DECLARE PHASE2 LABEL EXTERNAL; /* for intermodule reference */
```

The rules for explicit label declarations are discussed in detail in Chapter 7.

In implicit label declarations (used more commonly than explicit label declarations), the name is placed at the very beginning of the executable statement to which the name is supposed to point. For example:

```
START2: ALPHA = 127;
```

This statement defines the label `START2` as pointing to the location of the PL/M instruction shown. If this block has no explicit declaration of `START2`, such as the following:

```
DECLARE START2 LABEL;
```

then the compiler takes the definition of `START2` as an implicit declaration as well as a definition, as if the declaration had occurred at the start of the last simple `DO` or procedure statement. If there is an explicit declaration, then the actual placement of the label remains simply a definition.

Labels are used to indicate significant instructions or the starting point of instruction sequences. Labels can be useful reference points for understanding the parts of a program, or targets for the transfer of control during execution (as discussed under `GOTO` and `CALL` in Chapter 6).

Results of Label Declarations

Valid label declarations result in the following:

- The declared name can be used to point to an executable instruction.
- The use of the declared name as a variable in its block is disallowed.
- If the label is also defined in its block by appearing in an executable statement, the address of that statement will be assigned as the value of the label.

Declaration for Procedures

To declare a procedure, give its name with a statement of the form:

```
name: PROCEDURE
```

followed optionally by parameters, type and/or attributes. The definition of the procedure then follows. The procedure definition is the set of statements declaring items used in the procedure (including any parameters) and the executable statements of the procedure itself. The definition ends with an `END` statement, optionally including the procedure name.

The complete declaration of a procedure includes all the statements from the `PROCEDURE` statement through the `END` statement. This definition/declaration must appear before the procedure name is used in an executable statement, just as variable and constant names must be declared before their use.

The only exceptions arise when the full definition may appear in another separately compiled module where it is declared `PUBLIC`, or when a procedure has been declared `REENTRANT`. A `PUBLIC` procedure can be used (called) only if the calling module meets the following requirements:

1. The procedure has been declared with the `EXTERNAL` attribute (so the linker or binder will search for it).
2. Each formal parameter the procedure uses has been declared so the compiler can verify correct usage when this module invokes the procedure. End this local declaration with an `END` statement.

For example:

```
SUMMER: PROCEDURE (A, B) EXTERNAL;  
        DECLARE A WORD, B BYTE;  
END SUMMER;
```

See Chapter 7 for details on intermodule references. See Chapter 8 for details on procedure definition and use.

Data Types

Data types apply not only to variables, but to every value processed by a PL/M program. This includes values returned by procedures as well as values calculated by processing expressions. Data type specifications determine the value an object can have, how this value is stored in memory, and the operations that can be used on the value.

The PL/M-386 compiler recognizes five classes of data, each of which has one or more data types.

There are several unsigned binary number types: `BYTE` (8-bit number), `HWORD` (16-bit number), `WORD` (32-bit number), and `DWORD` (64-bit number). The `OFFSET` type is a 32-bit number that represents the offset portion of a pointer, which has its own type: `POINTER`. (The `POINTER` type itself is also recognized.) Note that the compiler controls `WORD32` and `WORD16` automate mapping 32- and 16-bit types. These controls are discussed in Chapter 11.

There are four signed integer data types: `INTEGER` (32-bit number); `CHARINT` (8-bit number); `SHORTINT` (16-bit number).

PL/M-386 recognizes the floating-point data type `REAL`, for signed 32-bit numbers.

Throughout this manual, the data types are referenced according to the data type class. Table 3-2 summarizes the data type classes for the Intel386 and Intel486 microprocessors. See the sections at the end of this chapter for a discussion on the PL/M-386 compiler's `WORD32` | `WORD16` mapping.

⇒ **Note**

Although the PL/M-386 compiler assumes a 32-bit word it also accepts PL/M-286 code as input. PL/M-286 code can take advantage of the 32-bit data type provided by the Intel386 and Intel486 microprocessors when compiled with the PL/M-386 compiler.

Table 3-2. Data Types

Data Type and Value	
Unsigned Binary Number	Description
BYTE	8-bit number ranging from 0 to 255. Occupies one byte of memory.
WORD	Occupies two contiguous bytes of memory. The least significant 8 bits are stored in the lower address.
WORD	32-bit number ranging from 0 to 4,294,967,295. Occupies two contiguous WORDs of memory. The least significant 16 bits are stored in the lower address.
DWORD	64-bit number ranging from 0 to (2**64) -1. Occupies two contiguous WORDs of memory. The least significant 32 bits are stored in the lower address.
OFFSET	32-bit number that represents the offset portion of a POINTER. (ADDRESS supported by PL/M-80 and PL/M-86/286, is equivalent to OFFSET.)
SHORTINT	16-bit number from -32768 to +32767 occupies contiguous bytes of memory. The least significant 8 bits are stored in the low address. Internally stored in two's complement notation.
INTEGER	32-bit number ranging from -2,147,483,648 to +2,147,483,647. Occupies four contiguous bytes of memory. The least significant 16 bits are stored in the low address. Internally stored in two's complement notation. WORD32's LONGINT is equivalent to INTEGER.
CHARINT	8-bit number ranging from -128 to +127. Occupies one byte of memory. Internally stored in two's complement notation.
Real Numbers	Description
REAL	Signed, floating-point number. Occupies four contiguous bytes of memory.
Pointers	Description
POINTER	The value is the address of the memory storage location. Consists of a segment selector portion and an offset portion.
Selectors	Description
SELECTOR	The value is equivalent to the segment selector portion of a POINTER. Can be used as the base of a based variable.

Unsigned Binary Number Variables: Unsigned Arithmetic

Unsigned arithmetic is used to perform any arithmetic operation on unsigned binary number variables. All of the PL/M operators can be used with these data types. Arithmetic and logical operations on such variables yield a result of one of the unsigned binary number types, depending on the operation and the operands. Relational operations always yield a true or false result of type `BYTE`.

With unsigned arithmetic, if a large value is subtracted from a smaller one, the result is the two's complement of the absolute difference between the two values. For example, if a `BYTE` value of 1 (00000001 binary) is subtracted from a `BYTE` value of 0 (00000000 binary), the result is a `BYTE` value of 255 (11111111 binary).

Also, the result of a division operation is always truncated (rounded down) to a whole number. For example, if an `HWORD` value of 7 (0000000000000111 binary) is divided by a `BYTE` value of 2 (00000010 binary), the result is an `HWORD` value of 3 (0000000000000011 binary).

When declaring a variable that may be used to hold or produce a negative result, it is advisable to make the variable either a signed integer or real. If the variable is supposed to hold or produce a non-integer, it must be declared as `REAL`. Use of the appropriate data types will reduce the occurrences of incorrect results from arithmetic operations (see Chapter 5).

INTEGER Variables: Signed Arithmetic

The sign bit is 0 if the `INTEGER` value is positive or zero, and 1 if the value is negative. The magnitude is given in two's complement notation.

Signed Arithmetic


For the Intel386 and Intel486 microprocessors, arithmetic operations on signed variables use 32-bit signed arithmetic to hold signed intermediate or final results. Thus, addition and subtraction always produce mathematically correct results if overflow does not occur. (See also the `OVERFLOW` control in Chapter 11.) Relational operations are signed arithmetic comparisons that yield a true or false result of type `BYTE`.

However, as with unsigned binary number operands, division produces only an `INTEGER` result. The result is rounded toward zero (i.e., down if the result is positive, up if the result is negative).

Only the arithmetic and relational operators can be used with signed operands. Logical operators are not allowed except for constant expressions within cast parentheses (see Chapter 5).

REAL Variables: Floating-point Arithmetic

The value of a `REAL` variable is a signed floating-point number that occupies four contiguous bytes of memory, which may be viewed as 32 contiguous bits in the single precision format. The bits are divided into fields as follows:



OSD567

The byte with the lowest address contains the least significant 8 bits of the significand, and the byte with the highest address contains the sign bit and the most significant 7 bits of the exponent field.

The sign bit is 0 if the `REAL` value is positive or zero, and 1 if the `REAL` value is negative.

The exponent field contains a value offset by 127. In other words, the actual exponent can be obtained from the exponent field value by subtracting 127. This field is all 0s if the `REAL` value is zero.

The significand contains the binary digits of the fractional part of the REAL value when this part is represented in binary scientific notation. This field is all 0s if the REAL value is zero.

Operations on REAL operands use signed floating-point arithmetic to yield a result of type REAL. The implementation guarantees that the result of each operation is the closest floating-point number to the mathematical real-number result (if overflow or underflow does not occur). The relational operators and the arithmetic operators +, -, *, and / can be used with REAL operands: the MOD operator and the logical operators are not allowed. Arithmetic operations yield a result of type REAL and relational operations yield a true or false result of type BYTE.

The PL/M compiler extends the utility of the REAL data types by holding intermediate results in the numeric coprocessor's temporary-real format (80-bit). This format preserves 64 bits of precision and the full range of representable numbers. The exponent in this format is 15 bits instead of 8 in the single precision format.

The increased exponent range greatly reduces the likelihood of underflow and overflow, and eliminates roundoff as a source of error until the final assignment of the result is performed. Underflow, overflow, and roundoff errors are probable for intermediate computations as well as in the final result. For example, an intermediate underflow result might later be multiplied by a very large factor, providing a final result of acceptable magnitude.

Examples of Binary Scientific Notation

1. Consider the following binary number (which is equivalent to the decimal value 10.25):

```
1010.01B
```

The dot (.) in this number is a binary point. The same number can be represented as:

```
1.01001B * 2**3
```

This is binary scientific notation, with the binary point immediately to the right of the most significant digit. The digits 01001 are the fractional part, and 3 is the exponent. This value would be represented in the single precision format as follows:

- The sign bit would be 0, because the value is positive.
- The exponent field would contain the binary equivalent of $127+3=130$.
- The leftmost digits of the fraction field would be 01001, and the remainder of this field would be all 0s.

The complete 32-bit representation would be:

0 1000010 010010 0000000000000000

and the contents of the four contiguous memory bytes would be as follows:

highest address:	01000001
	00100100
	00000000
lowest address:	00000000

Note that the most significant digit is not actually represented, because by definition it is a 1 unless the REAL value is zero. If the REAL value is zero, the entire 32-bit representation is all 0s.

2. Consider the fraction 1/16, or 0.0625. In binary, it is:

$$1.0000\text{B} * 2^{**}(-4)$$

In single precision format, the actual exponent -4 would be represented as 123 (127-4), and the fraction field would contain all 0s.

In the single precision format, the largest possible value for a valid exponent field is 254, which corresponds to an actual exponent of 127. Therefore, the largest possible absolute value for a positive or negative REAL value is:

$$1.111111111111111111111111111111\text{B} * 2^{**}127$$

or approximately $3.37 * 10^{**}38$.

The lowest permissible exponent field value for a non-zero REAL value is 1, which corresponds to an actual exponent of -126. Therefore, the smallest possible absolute value for a positive or negative REAL value is:


$$1.0\text{B} * 2^{**}(-126)$$

or approximately $8.43 * 10^{**}(-37)$.

POINTER Variables and Location References

The value of a `POINTER` variable is the address of the microprocessor's storage location and consists of a segment selector portion (see Chapter 9) and an offset portion.

The bits are divided as follows:



OSD577

`POINTER` variables are important as bases for based variables.

Only the relational operators for equality and inequality (`=` or `<` `>`) can be used with `POINTER` operands, yielding a true or false result of type `BYTE`. No arithmetic or logical operations are allowed (see Chapter 5).

A `POINTER` can be viewed as a structure of `SELECTOR` and `OFFSET` rather than a scalar. Therefore, arithmetic with `POINTERS` (e.g., `PTR+1`) is illegal.

The value of a `POINTER` variable can be created or changed in the following ways:

- The variable can be initialized when declared, using `INITIAL` or `DATA` with an address created with `.`
- The variable can be assigned an address created via the `@` operator (described in the following section). This is the most commonly used method.
- The variable can be assigned the value of a `POINTER` variable or function (including `NIL`, described in Chapter 9).
- The variable can be assigned a value generated by the `BUILD$PTR` function (also described in Chapter 9).
- `POINTER` type conversion (cast). Changing from one value to another is different from the `POINTER` built-in function (see Chapter 9).
- In `SMALL RAM` model, the `POINTER` is actually the offset portion only. In this case, all operations on the `PL/M-386 OFFSET` data type can be used, including arithmetic.

The @ Operator

A location reference is formed with the @ operator. A location reference has a value of type `POINTER`, that is, a location address. An important use of location references is to supply values for `POINTER` variables.

The basic form of a location reference is as follows:

```
@ variable-ref
```

Where:

variable-ref is the name of a variable.

The value of this location reference is the actual run-time location of the variable.

The *variable-ref* may also refer to an unqualified array or structure name. The pointer value is the location of the first element or member of the array or structure.

Consider the following declarations:

```
DECLARE RESULT REAL;  
DECLARE XNUM(100) BYTE;  
DECLARE RECORD STRUCTURE (KEY BYTE,  
    INFO(25) BYTE,  
    HEAD POINTER);  
DECLARE LIST(128) STRUCTURE (KEY BYTE,  
    INFO(25) BYTE,  
    HEAD POINTER);
```

The `@RESULT` is the location of the `REAL` scalar `RESULT`, and `@XNUM(5)` is the location of the 6th element of the array `XNUM`. `@XNUM` is the location of the beginning of the array, that is, the location of the first element (element 0).

The `RECORD STRUCTURE` declares a byte called `KEY` followed by 25 bytes called `INFO(0)`, `INFO(1)`, and so on, followed by the `POINTER` variable named `HEAD`. Because `KEY`, `INFO`, and `HEAD` are all declared part of the `RECORD` structure, their contents must be referred to as `RECORD.KEY`, `RECORD.INFO(0)`, . . . , `RECORD.INFO(24)`, and `RECORD.HEAD`.

Refer to the addresses of these elements of the `RECORD` structure by using the @ operator. `@RECORD.HEAD` is the location of the `POINTER` scalar `RECORD.HEAD` and `@RECORD` is the location of the structure, which is the same as that of the `BYTE` scalar `RECORD.KEY`. `@RECORD.INFO` is the location of the first element of the 25-byte array `RECORD.INFO`, whereas `@RECORD.INFO(7)` is the location of the 8th element of the same array.

LIST is an array of structures. The location reference @LIST(5).KEY is the location of the scalar LIST(5).KEY. Note that @LIST.KEY is illegal because it does not identify a unique location (i.e., the KEY of which LIST).

The location reference @LIST(0).INFO(6) is the location of the scalar LIST(0).INFO(6). Also, @LIST(0).INFO is the location of the first element of the same array (i.e., the location of the array itself).

A special case exists when the identifier used as *variable-ref* is the name of a procedure. This procedure must be declared at the outer level of the program module. No actual parameters can be given (even if the procedure declaration includes formal parameters). The value of the location reference in this case is the location of the entry point of the procedure. (See Chapter 8 and Appendices F and G.)

Storing Strings and Constants via Location References

Another form of location reference is the following:

```
@(constant list)
```

Where:

```
constant list
```

is a sequence of one or more BYTE constants separated by commas and enclosed by parentheses.

When this type of location reference is made, space is allocated for the constants. The constants are stored in this space (contiguously, in the order given by the list), and the value of the location reference is the location of the first constant. If RAM is specified on the compiler invocation command, constants are stored in the DATA segment. If ROM is specified on the compiler invocation command, constants are placed in the CODE segment (see Chapters 11 and 13).

Values in the constant list are treated as if they were in a BYTE array initialization list.

Strings can be included in the list. For example, if the operand:

```
@('NEXT VALUE')
```

appears in an expression, it causes the string 'NEXT VALUE' to be stored in memory (one character per byte, thus occupying 10 contiguous bytes of storage). The value of the operand is the location of the first of these bytes; in other words, it is a pointer to the string.

OFFSET Data Type and the Dot Operator

A dot operator is provided for compatibility with PL/M-80 programs. The dot operator (.) is similar to the @ operator, but produces an address of type WORD. This address represents an offset in the current data segment (for variables) or in the current code segment (for procedures). Use this address with caution, because it can produce unexpected results in a PL/M program that contains more than one data segment or more than one code segment.


In a PL/M-386 program, wherever WORD can be used, OFFSET can also be used. The main difference between the two types is in casting.

To create or change the value of an OFFSET variable, it can be assigned an OFFSET variable or function, or assigned the result of the built-in function OFFSET\$OF, or OFFSET type conversion, or the dot operator (see Chapter 9).

SELECTOR Variables

The value of a SELECTOR variable is equivalent to the segment selector portion of a POINTER, and can also be used as the base of a based variable.

In PL/M-386, the bits of the SELECTOR portion of a POINTER are shown below:



OSD578

The sections of this diagram are discussed in detail in Chapter 10.

Only the logical and relational operators for equality and inequality (=, <, > and <>) can be used with SELECTOR operands, yielding a true or false result of type BYTE. No arithmetic operations are allowed (see Chapter 5).

To create or change the value of a SELECTOR variable it can be assigned a SELECTOR variable or function, or assigned the result of the built-in function SELECTOR\$OF or SELECTOR type conversion (see Chapter 9).

The results of the @ and dot operators cannot be assigned directly to SELECTOR variables. They must first be converted to the SELECTOR type with the built-in functions SELECTOR\$OF and SELECTOR.

Based Variables

Sometimes, the address of a variable is not known until the program is actually run. For instance, if a procedure is written to swap two bytes and this procedure is called from various places in the code, the addresses of the two bytes are not known when writing the procedure definition.

For this type of manipulation, PL/M uses based variables. A based variable is one that is pointed to by another variable called its base. This means the base contains the address of the desired (based) variable. A variable is made **BASED** by inserting in its declaration the word **BASED** and the identifier of the base (which must already have been declared).

A based variable is not allocated storage by the compiler. At different times during program execution the based variable may actually refer to different places in memory, because the variable's base may be changed by the program.

To declare an address based variable, first declare its base, which must be of type **POINTER**, **SELECTOR**, **WORD**, or **OFFSET**. Next, declare the based variable itself as follows:

```
DECLARE I BYTE;  
DECLARE ITEM$PTR POINTER;  
DECLARE ITEM BASED ITEM$PTR BYTE;
```

In these declarations, a reference to **ITEM** is, in effect, a reference to the **BYTE** value pointed to by the current value of **ITEM\$PTR**. Thus, the sequence:

```
ITEM$PTR = @I;  
ITEM = 77H;
```

loads the **BYTE** value of 77 (hex) into the variable **I**.

PL/M supports more than one level of based variable, so variables can be based on based variables.

For example, the following declarations are valid:

```
DECLARE PTR1 POINTER;  
DECLARE PTR2 BASED PTR1 POINTER;  
DECLARE STR1 BASED PTR2 STRUCTURE (  
    X REAL,  
    Y REAL);
```

The following restrictions apply to bases:

- No initializations are permitted with based variables.
- The base must be of type `POINTER`, `SELECTOR`, `WORD`, or `OFFSET`. However, use a base of type `OFFSET` or `WORD` with caution because it does not contain a full microprocessor address. `OFFSET`- or `WORD`-based variables are addressed relative to the current DS register.
- The base cannot be subscripted. That is, it cannot be an array element.

The word `BASED` must immediately follow the name of the based variable in its declaration, as in the following examples:

```
DECLARE (AGE$PTR, INCOME$PTR, RATING$PTR, CATEGORY$PTR) POINTER;  
DECLARE AGE BASED AGE$PTR BYTE;  
DECLARE (INCOME BASED INCOME$PTR, RATING BASED RATING$PTR) WORD;  
DECLARE (CATEGORY BASED CATEGORY$PTR)(100) WORD;
```

In the first `DECLARE` statement, the `POINTER` variables `AGE$PTR`, `INCOME$PTR`, `RATING$PTR`, and `CATEGORY$PTR` are declared. They are used as bases in the next three `DECLARE` statements.

In the second `DECLARE` statement, a `BYTE` variable called `AGE` is declared. The declaration implies that whenever `AGE` is referenced by the running program, its value will be found at the location given by the current value of the `POINTER` variable `AGE$PTR`.

The third `DECLARE` statement declares two based variables, both of type `WORD`.

The fourth `DECLARE` statement defines a 100-element `WORD` array called `CATEGORY`, based on `CATEGORY$PTR`. When any element of `CATEGORY` is referenced at run time, the current value of `CATEGORY$PTR` is the location of the array `CATEGORY` (i.e., its first element).

The other elements follow contiguously. The parentheses around the tokens `CATEGORY BASED CATEGORY$PTR` make the statement more readable, but are not required.

⇒ **Note**

Debug information is available for only the first level of indirection when using variables based on `BASED` variables.

Location References and Based Variables

An important use of location references is to supply values for bases. Thus, the @ operator, together with the based variable concept, gives PL/M a very powerful facility for manipulating pointers.

For example, to refer to the three different REAL variables NORTH\$ERROR, EAST\$ERROR, and HEIGHT\$ERROR at different times with the single identifier ERROR, write:

```
DECLARE (NORTH$ERROR, EAST$ERROR, HEIGHT$ERROR) REAL;  
DECLARE ERROR$PTR POINTER;  
DECLARE ERROR BASED ERROR$PTR REAL;  
.  
.  
.  
ERROR$PTR = @NORTH$ERROR;
```

The value of ERROR\$PTR is the location of NORTH\$ERROR. A reference to ERROR is, in effect, a reference to NORTH\$ERROR. Later in the program, write:

```
ERROR$PTR = @HEIGHT$ERROR;
```

Now a reference to ERROR is, in effect, a reference to HEIGHT\$ERROR. In the same way, the value of the pointer can be made the location of EAST\$ERROR, and a reference to ERROR can be made a reference to EAST\$ERROR.

This technique is useful for manipulating complicated data structures and for passing locations to procedures as parameters. Examples are given in Chapter 8.

The AT Attribute

The AT attribute causes the address of a variable to be the specified location. The AT attribute has the form:

```
AT ( location )
```

Where:

location must be a location reference formed with the @ operator.

AT must refer to a nonbased variable that has already been declared. If there is a subscript expression, it must be a constant expression containing no operators except + and -.

The following are examples of valid AT attributes:

```
AT (@BUFFER)
AT (@BUFFER(128) )
AT (@NAMES( INDEX + 1) )
```

In the last example, INDEX represents a whole-number constant that has been previously declared with a LITERALLY declaration. The compiler replaces this name with the declared whole-number constant, thus satisfying the restrictions previously mentioned.

The first nonbased variable in a factored declaration containing the AT attribute will have the address specified by *location*. Other variables in the same declaration will, in sequence, refer to successive locations thereafter.

For example, the declaration:


```
DECLARE (CHAR$A, CHAR$B, CHAR$C) BYTE AT (@BUFFER);
```

causes the BYTE variable CHAR\$A to refer to the location of BUFFER. The variables CHAR\$B and CHAR\$C are located in the next two bytes after CHAR\$A.

The declaration:

```
DECLARE T(10) STRUCTURE (X(3) BYTE,
                          Y(3) BYTE,
                          Z(3) BYTE) AT (@DATA$BUFFER);
```

sets up structure references to 90 bytes. They are organized so that each of the 10 members of T refers to nine bytes. The first three use the name X, the second three Y, and the last three Z. Figure 3-1 illustrates this structure.



OSD533

Figure 3-1. Successive Byte References of a Structure

The preceding declaration, using the AT attribute, causes the beginning of the structure T, namely the scalar T(0).X(0), to be located at the same location as a previously declared variable called DATA\$BUFFER. The other scalars making up the structure will follow this location in logical order: T(0).X(1), T(0).X(2), and so on up to T(9).Z(2), which is the last scalar, located in the 89th byte after the location of DATA\$BUFFER.

However, no memory locations for these 90 scalars are allocated by this declaration. You determine the contents of the memory space beginning at @DATA\$BUFFER.

The following rules apply to the AT attribute:

- AT cannot be used with variables that are based, EXTERNAL, or parameters.
- AT can be used with the PUBLIC attribute, if it immediately follows the word PUBLIC. However, the location cannot be a location reference to a variable that is EXTERNAL.

The AT attribute can be used to make variables equivalent, providing more than one way of referring to the same information. For example:

```
DECLARE DATUM HWORD;  
DECLARE ITEM BYTE AT (@DATUM);
```

causes ITEM to be declared a BYTE variable at the same location that has just been allocated for the HWORD variable DATUM. Thus, any reference to ITEM is, in effect, a reference to the low-order byte of DATUM (because HWORD values are stored with the low-order 8 bits preceding the high-order 8 bits).

The following is another example using the AT attribute:

```
DECLARE VECTOR (6) BYTE;  
DECLARE SHORT$VECTOR STRUCTURE (FIRST (3) BYTE,  
                                SECOND (3) BYTE)  
                                AT (@VECTOR);
```

In this example, a six-element BYTE array called VECTOR is declared. Additionally, a structure of two three-byte arrays, SHORT\$VECTOR.FIRST and SHORT\$VECTOR.SECOND, is declared.

The first scalar of this structure, SHORT\$VECTOR.FIRST(0), is located at the same location as the first element of the array VECTOR.

Thus, there are two ways to refer to the same six bytes. For example, the fifth byte in the group can be referenced as either VECTOR(4) or SHORT\$VECTOR.SECOND(1).

When a variable is declared with the AT attribute, the compiler does not optimize the machine code generated to access that variable.

WORD32 | WORD16 Type Mapping

The PL/M-386 compiler supports two primary controls, WORD32 and WORD16, for unsigned binary number and signed integer data types, which provide some basic data type and language semantics compatibility for the Intel386 and Intel486 family of microprocessors. These controls specify the basic WORD size and thus affect the representation of certain data types. The default for PL/M-386 is WORD32. The WORD16 control does not specify 16-bit code (a parameter pushed on the stack is still four bytes), but maps the names of some data types into others. Internally, all processing is the same (e.g., signed arithmetic is 32-bit for both WORD16 and WORD32). To accommodate existing 16-bit code where data type representation is critical, WORD16 can be used to map word size to the convention used in earlier versions of the PL/M compiler. Table 3-3 lists the data type representation for WORD32 and WORD16.

Table 3-3. WORD32 | WORD16 Data Type Mapping

Unsigned Binary Number Data Types	WORD32 (default)	WORD16
BYTE	8-bit	8-bit
HWORD	16-bit	8-bit
WORD	32-bit	16-bit
DWORD	64-bit	32-bit
QWORD	64-bit	64-bit
Signed Integer Data Types	WORD32	WORD16
CHARINT	8-bit	8-bit
SHORTINT	16-bit	8-bit
INTEGER	32-bit	16-bit
LONGINT	32-bit	32-bit



Note

In PL/M-386, ADDRESS is equivalent to the OFFSET data type. OFFSET is a 32-bit data type that represents the offset portion of a POINTER. The size of OFFSET is not affected by the WORD32|WORD16 compiler control.

When writing new PL/M code, or when updating existing PL/M code, it is best to declare variables used for local addressing (i.e., those that are assigned from or initialized to the dot operator location references, assigned from the `OFFSET$OF` function, or used with the `BUILD$PTR` function or the `STACK$PTR` built-in) as `OFFSET` (or `ADDRESS`).

In PL/M-386, `WORD` is the natural 32-bit data type of the language on which all operations are available. However, in ASM386 a `WORD` is 16 bits and a `DWORD` is 32 bits.

Choosing `WORD32` or `WORD16`

The `WORD32|WORD16` compiler control determines how the data types in the source code are interpreted by the PL/M-386 compiler. See Chapter 11 for a description of the `WORD32|WORD16` control and syntax.

When compiling new PL/M-386 source code, use `WORD32` to take full advantage of the Intel386 or Intel486 microprocessors' features.

When recompiling existing PL/M-86 or PL/M-286 code, consider the source code to determine which compiler control to use. `WORD32` is usually preferable. Use `WORD16` if one of the conditions listed below applies to the source code. Note that the `WORD16` control does not have any effect on the `CMPB` instruction. This always remains as a 32-bit instruction.

- Scalar types are mapped to external data, such as `STRUCTURES` defined to represent data records read from a peripheral device. The format of the data from the peripheral device will not change regardless of the microprocessor processing it.
- Data is overlaid, for example:

```
DECLARE W HWORD (B1,B2) BYTE AT (@W);
DECLARE P POINTER, B BASED P (2) BYTE, WW BASED P WORD;
```

In this example, code may depend on the fact that two `BYTES` overlaying the `HWORD` constitute both halves of the `WORD` completely. Similarly, code can depend on the fact that the `LOW` or `HIGH` of an `HWORD` returns 8 bits.

- Loops depend on the size of a `WORD` type. Operations dependent on a variable overflow could produce unexpected results.



Arrays and Structures 4

Arrays

For increased efficiency, it is often desirable to use a single identifier to refer to a whole group of scalars, and to distinguish the individual scalars by means of a subscript (i.e., a value enclosed in parentheses). Such a list, in which the scalars are all the same type, is called an array.

An array is declared by using a dimension specifier. The dimension specifier is a nonzero whole-number constant enclosed in parentheses. The value of the constant specifies the number of array elements (individual scalar variables) making up the array. For example:

```
DECLARE ITEMS (100) BYTE;
```

causes the identifier `ITEMS` to be associated with 100 array elements, each of type `BYTE`. One byte of storage is allocated for each of these scalars.

The elements of an array are stored contiguously, with the first element in the lowest location and the last element in the highest location. No storage is allocated for a based array, but the elements are considered to be contiguous in memory.

The declaration:

```
DECLARE (WIDTH, LENGTH, HEIGHT) (100) REAL;
```

is similar to the following sequence:

```
DECLARE WIDTH (100) REAL;  
DECLARE LENGTH (100) REAL;  
DECLARE HEIGHT (100) REAL;
```

The difference between the two declarations is that contiguous storage is guaranteed for variables declared in a single parenthesized list, whereas variables declared in consecutive declarations are not necessarily stored contiguously.

This causes each of the three identifiers, `WIDTH`, `LENGTH`, and `HEIGHT`, to be associated with 100 array elements of type `REAL`, so that 300 elements of type `REAL` have been declared in all. For each of these scalars, four contiguous bytes of storage are allocated.

Subscripted Variables

To refer to a single element of a previously declared array, use the array name followed by a subscript enclosed in parentheses. This construct is called a subscripted variable.

For example, as a result of the following DECLARE statement:

```
DECLARE ITEMS (100) BYTE;
```

each byte can be referenced as an individual item using `ITEMS(0)`, `ITEMS(1)`, `ITEMS(2)`, and so on up to `ITEMS(99)`.

Notice that the first element of an array has subscript 0, not 1. Thus, the subscript of the last element is 1 less than the dimension specifier.

To add the third element of the array `ITEMS` to the fourth, and store the result in the fifth, write the PL/M assignment statement as follows:

```
ITEMS(4) = ITEMS(2) + ITEMS(3);
```

The subscript of a subscripted variable need not be a whole-number constant. It can be another variable, or any PL/M expression that yields a `BYTE`, `HWORD`, `WORD`, `OFFSET`, `SHORTINT`, `CHARINT`, or `INTEGER` value.

Thus, the construction:

```
VECTOR(ITEMS(3) + 2)
```

refers to some element of the array `VECTOR`. Which element this construction refers to depends on the expression `ITEMS(3) + 2`. This value, in turn, depends on the value stored in `ITEMS(3)`, the fourth element of array `ITEMS`, at the time when the reference is processed by the running program. If `ITEMS(3)` contains the value 5, then `ITEMS(3) + 2` is equal to 7 and the reference is to `VECTOR(7)`, the eighth element of the array `VECTOR`.

The following sequence of statements will sum the elements of the 10-element array `NUMBERS` by using an index variable named `I`, which takes values from 0 to 9:

```
DECLARE SUM BYTE;                /* To avoid overflow, */
DECLARE NUMBERS(10) BYTE;        /* SUM should add up  */
DECLARE I BYTE;                  /* to less than 255  */

SUM = 0;
DO I = 0 TO 9;
    SUM = SUM + NUMBERS(I);
END;
```

Subscripted array variables can be used anywhere a variable can be used, including the left side of an assignment statement if the array elements are of a scalar type.

Structures

Just as an array enables one identifier to refer to a collection of elements of the same type, a structure enables one identifier to refer to a collection of structure members that may have different data types. Each member of a structure has a member identifier.

A structure member can be another structure; these nested structures are described in the section titled, Nested Structures.

The following is an example of a structure declaration:

```
DECLARE AIRPLANE STRUCTURE (  
    SPEED REAL,  
    ALTITUDE REAL);
```

This statement declares two `REAL` scalars, both associated with the identifier `AIRPLANE`. Once this declaration has been made, the first scalar can be referred to as `AIRPLANE.SPEED` and the second as `AIRPLANE.ALTITUDE`. These names are also called the members of this structure.

A structure can have many members (see Appendix B for the correct limit). The members of a structure are stored contiguously in the order in which they are specified. (No storage is allocated for a based structure, but the members are considered to be contiguous in memory.)

Individual structure members cannot be based and cannot have any attributes (see Chapter 3).

Arrays of Structures

With PL/M, arrays of structures can be created. The following `DECLARE` statement creates an array of structures that can be used to store `SPEED` and `ALTITUDE` for 20 `AIRPLANES` instead of one:

```
DECLARE AIRPLANE (20) STRUCTURE (  
    SPEED REAL,  
    ALTITUDE REAL);
```

This statement declares 20 structures associated with the array identifier `AIRPLANE`, each distinguished by subscripts from 0 to 19. Each of these structures consists of two `REAL` scalar members. Thus, storage is allocated for 40 `REAL` scalars.

To refer to the `ALTITUDE` of the 17th `AIRPLANE`, write `AIRPLANE(16).ALTITUDE`.

Arrays Within Structures

An array can be used as a member of a structure, as follows:

```
DECLARE PAYCHECK STRUCTURE (  
    LAST$NAME(15) BYTE,  
    FIRST$NAME(15) BYTE,  
    MI BYTE,  
    AMOUNT REAL);
```

This structure consists of two 15-element `BYTE` arrays, `PAYCHECK.LAST$NAME` and `PAYCHECK.FIRST$NAME`, the `BYTE` scalar `PAYCHECK.MI`, and the `REAL` scalar `PAYCHECK.AMOUNT`.

To refer to the fourth element of the array `PAYCHECK.LAST$NAME`, write `PAYCHECK.LAST$NAME(3)`.

Arrays of Structures With Arrays Inside the Structures

Given that an array can be made up of structures, and a structure can have arrays as members, the two constructions can be combined to write:

```
DECLARE FLOOR (30) STRUCTURE (  
    OFFICE (55) BYTE);
```

The identifier `FLOOR` refers to an array of 30 structures, each of which contains one array of 55 `BYTE` scalars. This could be thought of as a 30-by-55 matrix of `BYTE` scalars. To reference a particular scalar value (for example, element 46 of structure 25) write `FLOOR(24).OFFICE(45)`. Note that the scalar elements of each `OFFICE` array are stored contiguously, and the `OFFICE` arrays are elements of the `FLOOR` array and are stored contiguously.

Alter the preceding `PAYCHECK` structure declaration to make it an array of structures, as follows:

```
DECLARE PAYROLL (100) STRUCTURE (  
    LAST$NAME(15) BYTE,  
    FIRST$NAME(15) BYTE,  
    MI BYTE,  
    AMOUNT REAL);
```

This is an array of 100 structures, each of which can be used during program execution to store the last name, first name, middle initial, and amount of pay for one employee. `LAST$NAME` and `FIRST$NAME` in each structure are 15-byte arrays for storing the names as character strings.

To refer to the *K*th character of the first name of the *N*th employee, write:

```
PAYROLL(N-1).FIRST$NAME(K-1)
```

where *N* and *K* are previously declared variables to which appropriate values have been assigned. This might be convenient in a routine for printing out payroll information.

Nested Structures

A member of a structure can also be another structure; this is called a nested structure.

Nested structures are subject to the same rules as all structures. They can contain their own member identifiers, whether these are scalars, arrays, or structures.

The following example shows nested structures:

```
DECLARE EMPLOYEE (100) STRUCTURE (  
    ID WORD,  
    NAME STRUCTURE (  
        LAST$NAME (15) BYTE,  
        FIRST$NAME (15) BYTE,  
        MI BYTE),  
    AGE BYTE,  
    JOB WORD,  
    PAY STRUCTURE (  
        RATE REAL,  
        OTRATE REAL,  
        BENEFITS STRUCTURE (  
            OPTIONS REAL,  
            CHOSEN BYTE)  
        )  
    ) ;
```

The preceding declaration statement is for an array (named `EMPLOYEE`) of 100 structures. Each of the 100 elements of `EMPLOYEE` is a structure with the following members: a `WORD` scalar named `ID`, a nested structure called `NAME`, a `BYTE` scalar named `AGE`, a `WORD` scalar named `JOB`, and a nested structure named `PAY`.

The `NAME` structure has two arrays (`LAST$NAME` and `FIRST$NAME`) of 15 bytes each for members, as well as a `BYTE` scalar named `MI`.

The `PAY` structure has two `REAL` scalars (`RATE` and `OTRATE`) for members, as well as a nested structure named `BENEFITS`. `BENEFITS` has the `REAL` scalar `OPTIONS` and the `BYTE` scalar `CHOSEN` as members.

The preceding example contains two levels of nested structures. The structures `NAME` and `PAY` are at the first level of nesting; the structure `BENEFITS` is at the second level of nesting. See Appendix B for the maximum limit on nested structures.

References to Arrays and Structures

A variable reference is the use, in program text, of the identifier of a variable that has been declared. A variable reference can be fully qualified, partially qualified, or unqualified.

Fully Qualified Variable References

A fully qualified variable reference specifies a single scalar. For example, given the following declarations:

```
DECLARE AVERAGE REAL;
DECLARE ITEMS (100) BYTE;
.
.
.
DECLARE RECORD STRUCTURE (
    KEY BYTE,
    INFO WORD);
DECLARE NODE (25) STRUCTURE (
    SUBLIST (100) BYTE,
    RANK BYTE);
```

then `AVERAGE`, `ITEMS(5)`, `RECORD.INFO`, and `NODE(21).SUBLIST(32)` are all fully qualified variable references. Each refers unambiguously to a single scalar.

Note that qualification can only be applied to variables that have been appropriately declared. A subscript can only be applied to an identifier that has been declared with a dimension specifier. A member-identifier can be applied only to an identifier declared as a structure identifier. The compiler flags violations of these rules as errors.

Unqualified and Partially Qualified Variable References

Unqualified and partially qualified variable references can be used only in location references (see Chapter 3) and in the built-in procedures `LENGTH`, `LAST`, and `SIZE` (see Chapter 9).


An unqualified variable reference is the identifier of a structure or an array, without a member-identifier or subscript. For example, with the declarations in the previous section, `ITEMS` and `RECORD` are unqualified variable references. An unqualified variable reference is a reference to the entire array or structure. `@ITEMS` is the location of the entire array `ITEMS` (the location of its first byte). Similarly, `@RECORD` is the location of the first byte of the structure `RECORD`.

A partially qualified variable reference does not refer to a single scalar even using a subscript and/or member-identifier with an identifier.

For example, in the declaration in the previous section, `NODE(15)` and `NODE(12).SUBLIST` are partially qualified variable references.

When used with the `@` operator, partially qualified variable references are taken to mean the first byte that fits the description. Thus, `@NODE(15)` is the location of the first byte of the structure `NODE(15)`, which is an element of the array `NODE`. Similarly, `@NODE(12).SUBLIST` is the location of the first byte of the array `NODE(12).SUBLIST`, which is a member of the structure `NODE(12)`, which is an element of the array `NODE`.

Because it is ambiguous, `@NODE.SUBLIST` cannot be used. In a location reference referring to an array consisting of structures, a subscript must be given before a member-identifier can be added to the reference. The rule is different for partially qualified variable references in connection with the built-in procedures `LENGTH`, `LAST`, and `SIZE`, as explained in Chapter 9.



Expressions and Assignments

5

A PL/M expression consists of scalar operands (values) combined by arithmetic, logical, and relational operators. For example:

```
A + B
A + B - C
A*B + C/D
A*(B + C) - (D - E)/F
```

where +, -, *, and / are arithmetic operators for addition, subtraction, multiplication, and division, and A, B, C, D, E, and F represent operands. The parentheses group operands and operators to control the order of evaluation.

This chapter describes the rules governing PL/M expressions. Although these rules may appear complex, most of the expressions used in actual programs are simple. In particular, when the operands of arithmetic and relational operators are all of the same type, the resulting expression is easy to understand.

Operands

Operands are the building blocks of expressions. An operand is a quantity with a value at run time on which an arithmetic, logical, or relational operation is performed by an operator. In the preceding examples, A, B, C, etc., are identifiers of scalar variables that have values at run time.

Operands in expressions can also be numeric constants and fully qualified variable references. The following sections describe all of the types of operands that are permitted.

Constants

A numeric constant can be an operand in an expression. However, its type must be appropriate, as discussed in the following paragraphs.

A numeric constant that contains a decimal point is of type `REAL`. A numeric constant that does not contain a decimal point is a whole-number constant.

You can use a whole-number constant in either signed context or unsigned context. In unsigned context, a whole-number constant is treated as an unsigned binary number data type. In signed context, a whole-number constant is treated as a signed integer data type (see Table 3-2).

Whole-number Constants in Unsigned Context

In PL/M-386, if the `WORD32` control is in effect, a whole number constant in unsigned context is treated as follows:

- As a `BYTE` value if it ranges from 0 to 255
- As a `HWORD` value if it ranges from 256 to 65,535
- As a `WORD` value if it ranges from 65,536 to 4,294,967,295 (i.e., $2^{32}-1$)
- As a `DWORD` value if it ranges from 2^{32} to $2^{64}-1$

Whole-number Constants in Signed Context

In signed context, a whole-number constant is always treated as an `INTEGER` value. In PL/M-386, the range is -2,147,483,648 to 2,147,483,647. Additionally, small integer values are extended into 32-bit values with no change to the arithmetic value.

String Constants

A string constant containing not more than four characters can also be used as an operand. If a string constant has only one character, it is treated as a `BYTE` constant whose value is the 8-bit `ASCII` code for the character. If a string constant is a two-character string, it is treated as an `HWORD` constant in PL/M-386. The value of the two-character string is formed by stringing together the `ASCII` codes for the two characters, with the code for the first character forming the most significant 8 bits of the 16-bit number.

In PL/M-386, if the `WORD32` control is in effect, a three- or four-character string constant is treated as a `WORD` constant whose value is formed by stringing together the ASCII codes for all of the characters. The first character represents the high 8 bits, the second character represents the second most significant 8 bits, and so on. If the string has three characters, the ASCII NUL character is inserted in front of the first character to form a four-character string.

Strings of more than four characters are illegal as operands in expressions, and can be used in only two contexts: as initialization values for an array or as part of a location reference that points to the location at which the string constant is stored (see Chapter 3).

Variable and Location References

As described in Chapter 4, fully qualified variable references uniquely specify a single scalar value. (Partially qualified references, also discussed in Chapter 4, have very restricted uses.) Any fully qualified variable reference can be used as an operand in an expression. When the expression is evaluated, the reference is replaced by the value of the scalar.

A function reference is the name of a typed procedure that has been declared previously, along with any parameters required by the procedure declaration. The value of a function reference is the value returned by the procedure.

For example, in the statement:

$$I = J + \text{ABS}(L);$$

the absolute value of L will be returned by the function ABS and then added to the value of J before being stored in I . If L is -27 , the result will be the same as writing:

$$I = J + 27;$$

For a complete discussion of procedure and function references, see Chapter 8. Location references are described in Chapter 3.

Subexpressions

A subexpression is an expression enclosed in parentheses, which can be used as an operand in an expression. A subexpression can be used to group portions of an expression together, just as in ordinary algebraic notation.

Compound Operands

All the operand types previously described are primary operands. An operand can also be a value calculated by evaluating some portion of the total expression. For example, in the expression:

$$A + B * C$$

(where A , B , and C are variable references), the operands of the $*$ operator are B and C . The operands of the $+$ operator are A , and the result of the compound operand $B * C$. Notice that this expression is evaluated as if it had been written as follows:

$$A + (B * C)$$

This analysis of an expression to determine which operands belong to which operators, and which groups of operators and operands form compound operands, is discussed in Expression Evaluation. Table 5-1 lists operator precedence.

Table 5-1. Operator Precedence

Operator Class	Operator	Interpretation
Precedence	()	Controls order of evaluation: expressions within parentheses are evaluated before the action of any outside operator on the parenthesized items
Unary	+, -	Single positive operator, single negative operator
Arithmetic	*, /, MOD +, -	Multiplication, division, modulo (remainder) division, addition, subtraction
Relational	<, <=, <> >=, >	Less than, less than or equal to, not equal to, =, equals, greater than or equal to, greater than
Logical	NOT AND OR, XOR	Logical negation Logical conjunction Logical inclusion disjunction, logical exclusive disjunction

Arithmetic Operators

PL/M has the following five principal arithmetic operators:

+ - * / MOD

These operators are used as in ordinary algebra to combine two operands. Each operand can have an unsigned binary number data type value; a signed integer data type value; or a REAL number data type value (except that REAL operands cannot be used with the MOD operator).

Arithmetic operations cannot be used with POINTER and SELECTOR variables.

The +, -, *, and / Operators

The operators +, -, *, and / perform addition, subtraction, multiplication, and division on operands of any data type except the POINTER and SELECTOR data types. The following rules govern these operations.

- Both operands must be of the same class (i.e., both operands must be unsigned, signed, or real). Mixing operands of different classes is illegal. However, an operand of one class can be converted, in an expression, to another class with the use of a built-in conversion function (see Chapter 9).

- Unsigned Arithmetic

- Unsigned Addition and Subtraction

If both operands are of the same data type, the result is of the same data type (e.g., `BYTE + or - BYTE` produces a `BYTE` result).

If the operands are of different data types, the smaller operand is extended with high-order 0 bits to the size of the larger operand; the addition or subtraction is then performed as though both operands are of the same type. For example, for `BYTE + or - WORD`, the `BYTE` operand is zero-extended by 8 bits to `WORD` size; then the operation is performed with the `WORD` operands to produce a `WORD` result. For a `BYTE + or - DWORD` the `BYTE` operand is zero-extended by 24 bits to `DWORD` size; then the operation is performed with the `DWORD` operands to produce a `DWORD` result. For `WORD + or - DWORD`, the `WORD` operand is zero-extended by 16 bits to `DWORD` size; then the operation is performed with two `DWORD` operands producing a `DWORD` result.

- Unsigned Multiplication and Division

Assuming `WORD32`, if both operands are of type `BYTE`, the `*` and `/` operations produce an `WORD` result; if both operands are of type `WORD`, the `*` and `/` operations produce a `WORD` result. If both operands are of type `WORD`, the `*` and `/` operations produce a `WORD` result. If both operands are of type `OFFSET`, the `*` and `/` operations produce an `OFFSET` result. If both operands are of type `DWORD`, the `*` and `/` operations produce a `DWORD` result.

For mixed unsigned operands, the same rules as for addition and subtraction apply. The smaller operand is zero-extended to the size of the larger operand, then the multiplication or division is performed as though both operands are of the same type. The results are as described in the preceding paragraph.

If one operand is a whole-number constant or a string constant, it is treated as a `WORD` or `WORD` depending on its value (see `Whole-number Constants` in `Unsigned Context` and `String Constants`).

All arithmetic for signed operands is 32-bit signed integer arithmetic. The names of the storage type (e.g., `CHARINT`) do not imply what type of arithmetic is performed, only the size of storage assigned for the variable.

During signed arithmetic an expression can overflow only if it overflows the machine word (32 bits). However, overflow is possible when the value is assigned to a variable with `SHORTINT` or `CHARINT` data type. If the value is assigned to `CHARINT`, 24 high-order bits are truncated to form the `CHARINT` value. If the value is assigned to `SHORTINT`, 16 high-order bits are truncated to form the `SHORTINT` value. If the value is assigned to `INTEGER`, it is not changed.

Assignment overflow is detected using the `OVERFLOW` control (see Chapter 11).

Constants are always represented as integer constants, regardless of their value.

- Real arithmetic

Both operands are always of type `REAL`. Thus, the `+`, `-`, `*`, and `/` operations produce a result of type `REAL`.

If one operand is a constant, it must be typed as a floating-point constant, that is, it must have a decimal point. Mixing `REAL` operands with whole-number constants is not allowed. For example, if `R` is a `REAL` variable, `R+1.0` is a legal expression, but `R+1` is illegal. Also, `1.0+1` is illegal, because it mixes a `REAL` constant with a whole-number constant.

- Arithmetic expressions containing operands of type `SELECTOR` or `POINTER` are illegal.
- If both operands are whole-number constants, the operation depends on the context in which it occurs, as explained in Special Case: Constant Expressions.
- The result of division by 0 is undefined, except for `REAL` values (see Appendix G).

A unary `-` operator is also defined in PL/M. It takes a single operand, to which it is prefixed. A minus sign that has no operand to the left of it is taken to be a unary minus.

A unary `-` operator makes `(-A)` equivalent to `(0-A)`, where `A` is any operand. The 0 is a `BYTE` value if `A` is an unsigned binary number data type. The 0 is an `INTEGER` value if `A` is a signed integer data type; or a `REAL` value if `A` is a real number data type. If `A` is a whole-number constant, its type and the unary `-` operation depend on the context as explained in Special Case: Constant Expressions. In unsigned context, `(-1)` is assigned a `BYTE` value `(0-1)` which is equivalent to `OFFH`. In signed context, `(-1)` is assigned an `INTEGER` value `(0-1)` which is equivalent to `OFFFFFFFFH` for PL/M-386.

Finally, a unary `+` has no effect; `(+A)` is equivalent to `(A)`.

The MOD Operator

MOD performs division, except the result is not the quotient, but rather the remainder left after integer division. The result has the same sign as the operand on the left side of the MOD operator.

REAL operands cannot be used with the MOD operator; only unsigned or signed operands can be used.

For example, if A and B are INTEGER variables with values of 35 and 16, respectively, then $A \text{ MOD } B$ yields an INTEGER result of 3, and $-A \text{ MOD } B$ yields -3.

Unlike the / operator, the MOD operator must be separated from surrounding letters and digits by blanks or other separators.

Relational Operators

The following relational operators are used to compare operands of the same type:

<	less than
>	greater than
<=	less than or equal to
>=	greater than or equal to
<>	not equal to
=	equal to

Relational operators are always binary operators, taking two operands, to yield a `BYTE` result. Relational operators can be used with all types.

If both operands are unsigned, then unsigned arithmetic will be used to compare the two values. As with the arithmetic operators, mixing unsigned data types are allowed, with the smaller operand being zero-extended to the size of the larger operand.

Whole-number and string constant operands are treated as `BYTE`, `HWORD`, `WORD`, or `OFFSET`. Unsigned data types are primarily used to represent positive values. Negative numbers are represented by two's complement in the smallest unsigned data type that can hold the value. For example, -2 is represented as the `BYTE` value of `0FEH`. If `B` is a `BYTE` variable, then the relational expression `B>=-2` is `TRUE` only if `B` has the value of 254 or 255, because the expression -2 (when evaluated unsigned) has a `BYTE` value of 254.

In `PL/M-386`, if both operands are signed, then signed 32-bit integer arithmetic is used to compare the two values. `CHARINT` or `SHORTINT` are sign extended to `INTEGER` values. The calculated value is then assigned to the specified data type.

If both operands are real, floating point arithmetic will be used to compare the two values. Only floating-point constants (i.e., constants containing a decimal point) can be mixed with `REAL` operands.

Two `POINTER` operands can be compared for equality but greater than, less than, and inequality operations cannot be used. In `PL/M-386` only, two `POINTERS` are equal only if they are bitwise equal (i.e., if both segment selector portions are equal and both offset portions are equal).

Two `SELECTOR` operands can be compared for equality, inequality, less than, and greater than.

Since constants cannot be typed as `POINTER` or `SELECTOR`, comparison between `POINTER` or `SELECTOR` operands and constants is illegal.

As with arithmetic operations, operands of different classes cannot be mixed together in relational operations. An operand of one class can be converted, in an expression, to another class using a built-in conversion function (see Chapter 9).

If the specified relation between the operands is true, a `BYTE` value of `0FFH` (or `1111$1111B`) is returned. Otherwise, the result is a `BYTE` value of `00H` (or `0000$0000B`). Thus, in all cases, the result is of type `BYTE`, with all 8 bits set to 1 for a true condition, or to 0 for a false condition. For example:

<code>(6>5)</code>	result is <code>0FFH</code> (true)
<code>(64<=4)</code>	result is <code>00H</code> (false)

Values of true and false resulting from relational operations are useful in conjunction with `DO WHILE` statements and `IF` statements, as described in Chapter 6. In the context of a `DO WHILE` statement or `IF` statement, only the least significant bit of a true or false value is used. Thus, each value with the least significant bit set (including `0FFH`) is considered true and each value with the least significant bit 0 is considered false. A `BYTE` value is returned.

Logical Operators

PL/M has the following four logical (Boolean) operators:

NOT AND OR XOR

These operators are used with the unsigned binary number data type, or whole-number or string constant operands to perform logical operations on 8, 16, or 32 bits.

NOT is a unary operator, taking one operand only. It produces a result of the same type as its operand: each bit of the result is the one's complement of the corresponding bit of the original value.

The remaining operators (AND, OR, XOR) each take two operands, and perform bitwise and, or, and exclusive or, respectively. The bits of an AND result are 1 only when the corresponding bit in each operand is 1. The bits of an OR result are 1 when the corresponding bit of either operand is 1, and 0 only when both operands are 0. The bits of an XOR result are 0 only when the corresponding bits of the operand are the same (i.e., both 1 or both 0); the result has a 1 when one operand is 1 and the corresponding bit of the other operand is 0.

When both operands are of the same type, the result is the same type as the operands.

As with the arithmetic and relational operators, unsigned data types can be mixed in any combination for logical operations. Whole-number operands are treated as BYTE, HWORD, or WORD values in PL/M-386. The only exception is an expression composed only of whole numbers within the cast parentheses; then the constants have integer context and the numbers are extended to the 16-bit signed value. The usual bitwise logical operation then takes place (as explained above for 16-bit numbers for bitwise operations). Mixing OFFSET with WORD produces an OFFSET result.

The following are examples of logical operations:

NOT 11001100B	result is 00110011B
10101010B AND 11001100B	result is 10001000B
10101010B OR 11001100B	result is 11101110B
10101010B XOR 11001100B	result is 01100110B

Note that true and false values resulting from relational operations can be combined with logical operators:

NOT (6>5)	result is 00H (false)
(6>5) AND (1>2)	result is 00H (false)
(6>5) OR (1>2)	result is 0FFH (true)
(LIM = Y) XOR (Z<2)	result is 0FFH (true) if LIM = Y and Z<2 or if LIM<>Y and Z<2, but result is 00H (false) if both relations are true or both false

Note that in the statement:

$$A = (\text{NOT } B)$$

parentheses must be used as indicated. Failure to do so will result in a syntax error because relational operators (=) have higher precedence than logical operators (NOT).

The following are examples of whole numbers. In this example the parentheses enclose items to be converted (casted).

`OFFSET(10101010B AND 11001100B)` gives `OFFSET(010001000B)`. This is the 16-bit result obtained with the simple logical operation written above, except that the offset type is returned. PL/M-386 extends the result to 32 bits.

`WORD (-4 AND 7)` gives `WORD (0FFFFFFFCH AND 00000007H)` which gives `WORD(4)` or the unsigned 16-bit value of 4.

Expression Evaluation

Precedence of Operators: Analyzing an Expression

In PL/M, operators have an implied order that determines how operands and operators are grouped and analyzed during compilation.

The PL/M operators are listed in Table 5-1 (page 5-79) from highest to lowest precedence; those that take effect first are listed first. Operators in the same line are of equal precedence and are evaluated as they are encountered in a left to right reading of an expression.

The order of evaluation in an expression is controlled first by parentheses, then by operator precedence, and finally by left to right order.

The compiler first evaluates operands and operators enclosed in paired parentheses as subexpressions, working from the innermost to the outermost pairs of parentheses. The value of the subexpression is then used as an operand in the remainder of the expression.

Parentheses are also used around both subscripts and the parameters of function or procedure references. These are not subexpressions, but they too must be evaluated before the remainder of the expressions or references can be evaluated at a higher level.

When there is more than one operator in an expression, evaluate the results by beginning with the operator with the highest precedence. If the operators are of equal precedence, evaluate them left to right, as follows:

Example

$(A + B) * C$ is not the same as $A + B * C$
 $A + B * C$ means the same as $A + (B * C)$
 $A/B * C$ means the same as $(A/B) * C$

Reason

Parentheses form subexpressions
Operator precedence
Left to right, equal precedence

The following are examples of precedence ranking:

$A + B * C$	is equivalent to $A + (B * C)$
$A + B - C * D$	is equivalent to $(A + B) - (C * D)$
$A + B + C + D$	is equivalent to $((A + B) + C) + D$
$A / B * C / D$	is equivalent to $((A / B) * C) / D$
$A > B \text{ AND NOT } B < C - 1$	is equivalent to $(A > B) \text{ AND } (\text{NOT } (B < (C - 1)))$

In the last four examples, the application of the left-to-right rule for operators with the same precedence is shown. In the second, third, and fifth examples, the left-to-right rule for operators of equal precedence makes no difference in the value of the expression. But in the fourth example, the left-to-right rule is critical.

The following example shows the action of the rules of precedence on a longer expression:

$$(-B + \text{SQRT}(B * B - 4.0 * A * C)) / (2.0 * A)$$

Assume A , B , and C are variables of type `REAL`, and `SQRT` is a procedure of type `REAL` which returns the square root of the value passed to it as a parameter. In this case, the parameter is the expression $B * B - 4.0 * A * C$. Floating point constants (4.0, 2.0) are used rather than whole-number constants (4, 2) because it is invalid to combine whole-number constants with `REAL` variables.

The compiler first analyzes the portions of the expressions within the innermost parentheses, then the procedure parameter and the subexpression $2.0 * A$. (The subexpression is also called a compound operand because its result is used in evaluating the whole expression.)

In a left-to-right scan, the two operands of the first `*` operator are both equal to the value of B . The operands of the second `*` operator are 4.0 and the value of A . The operands of the third `*` operator are the results of the second evaluation (i.e., the compound operand $4.0 * A$) and the value of C . The operands of the fourth `*` operator are 2.0 and the value of A .

The subexpression $2.0 * A$ is now completely analyzed, but the parameter expression still contains a minus (`-`) operator that has not been analyzed. The operands of this operator are the result of evaluating $B * B$ and the result of evaluating $4.0 * A * C$. Once the evaluations are done, the parameter expression is analyzed and its value can be calculated.

This value does not become an operand in the overall expression. It is passed to the procedure `SQRT`, which returns the square root of the parameter. This returned value then becomes an operand in the remainder of the full original expression:

$$(-B + \textit{returned value}) / (2.0 * A)$$

Now that the innermost subexpressions have been analyzed and evaluated, a division operator whose left operand must be evaluated further remains. This outer subexpression is $-B +$ the returned square root: there are two operators. The first is a unary minus ($-$) and its operand is the value of B . The second is the binary plus ($+$) operator, with two operands: the value of $-B$ and the value of $\text{SQRT}(B * B - 4.0 * A * C)$. $-B$ has the same meaning as $0-B$, which is to be added to the known value of the square root indicated. The final operator is division ($/$), whose two operands are known: the value of $(-B + \text{SQRT}(B * B - 4.0 * A * C))$ and the value of $(2.0 * A)$.

Three important points must be emphasized about expression evaluation, as discussed in the next three sections.

Compound Operands Have Types

Compound operands have types as do primary operands. All of the primary operands used in the preceding example were of type `REAL`, which results in compound operands of type `REAL`. It is always valid for all the operands in an arithmetic expression to be of the same type, and the result will be that type also. Combining `BYTE` values can validly create a `WORD` or `WORD` value. Combining a signed integer data type value always creates an `INTEGER` value.

In an expression containing mixed data types, any combinations can be used as long as the types belong to the same class (i.e., unsigned binary number, signed integer, real, pointer, or selector). Data types (of the same class) can be mixed as operands in expressions, whether they are constants or variables.

Mixing types of different classes in arithmetic, logical, or relational expressions is invalid. For example, if F and G are `INTEGER` variables and H and K are `REAL` variables, then the expressions $F > K$ and $H + G$ are invalid.

Due to operator precedence, some combinations can occur validly in the same expression without being directly combined. In the following logical expression:

$(F > G \text{ AND } H < K)$

the subexpression $F > G$ yields a `BYTE` value, as does the subexpression $H < K$. Then the `BYTE` values are `AND`ed together. This expression is legal despite an apparent mixing of types. G and H could not be the operands for two reasons:

1. The relational operators are of higher precedence than the `AND` operator.
2. Only unsigned operands are legal with logical operators.

Relational Operators Are Restricted

In the absence of parentheses denoting a subexpression, the result of a relational operation (comparison) cannot become an operand in another relational operation. Thus, the expression:

$$A \leq X \leq B$$

is invalid in PL/M because the second \leq operator would have to use the result of the first \leq operator as one of its operands.

In PL/M the valid expression is as follows:

$$A \leq X \text{ AND } X \leq B$$

Parentheses also could have created a valid expression; for example:

$$(A \leq X) \leq B$$

However, in this expression the result does not have the desired meaning: $A \leq X$ becomes a byte of value 0 if A is greater than X, 0FFH if A is not greater than X. Thus, if A is 0, X is 1, and B is 2:

$$(0 \leq 1) \leq 2$$

evaluates to:

$$(0\text{FFH}) \leq 2$$

and yields a FALSE value. This is contrary to the original intention.

Order of Evaluation of Operands

Operators and operands are not bound in the same order as the order in which operands are evaluated.

The rules of analysis specify which operands are bound to each operator. The following example shows how operands are bound to operators:

$$A + B * C$$

B and C are the operands of the $*$ operator, and A and the value of $B * C$ are the operands of the $+$ operator. B and C must be evaluated before the $*$ operation can be performed, and the compound operand $B * C$ must be evaluated before the $+$ operation is performed.

However, it is not obvious whether B will be evaluated before C or vice versa. A could be evaluated before either B or C, and its value stored until the $+$ operation is performed.

The rules of PL/M do not specify the order in which subexpressions or operands are evaluated in each statement. This flexibility enables the compiler to optimize the object code it produces, as described in Chapter 11. In most cases, the order of evaluation makes no difference. However, certain embedded assignments (see Assignment Statements) or function references (see Chapter 8) change the value of an operand in the same expression.

Choice of Arithmetic: Summary of Rules

As described in Chapter 3, PL/M uses three distinct kinds of arithmetic: unsigned, signed, and floating-point. Whenever an arithmetic or relational operation is carried out, PL/M uses one of these types of arithmetic, depending on the types of the operands.

Table 5-2 is a summary of the rules that determine which type of arithmetic is used in each case. The table also lists the data type of the result for each kind of arithmetic operation. The notes following the table provide additional information. (see Relational Operators and Logical Operators for rules governing relational and logical operations.)

In PL/M-386, `OFFSET` operands are always 32-bit unsigned operands.

In expressions, whole-number constants are always converted to the value of the equivalent data type.

Table 5-2. Summary of Expression Rules for PL/M-386

Variable Type	Kind of Arithmetic	Operand Type	Operation	Result	Notes
BYTE HWORD	Unsigned	BYTE w/BYTE	+ or - * / or MOD	BYTE HWORD	range: 0 to 255 0 to 65,535
WORD DWORD		HWORD w/HWORD	+ or - * / or MOD	HWORD WORD	range: 0-65,535 0 to 2**32-1
		BYTE w/HWORD becomes HWORD w/HWORD	+ or - * / or MOD	HWORD WORD	BYTE is extended with 8 high-order zeros to an HWORD value
		WORD w/WORD	any arithmetic	WORD	range: 0 to 2**32-1
		BYTE w/WORD becomes WORD w/WORD	any arithmetic	WORD	BYTE is extended with 24 high-order zeros to a WORD value
		HWORD w/WORD becomes WORD w/WORD	any arithmetic	WORD	HWORD is extended with 16 high-order zeros to a WORD value
		DWORD w/DWORD	any arithmetic	DWORD	range: 0-2**63-1
		BYTE w/DWORD becomes DWORD w/DWORD	any arithmetic	DWORD	BYTE is extended with 56 high-order zeros to a DWORD value
		HWORD w/DWORD becomes DWORD w/DWORD	any arithmetic	DWORD	HWORD is extended with 48 high-order zeros to a DWORD value
	WORD w/DWORD becomes DWORD w/DWORD	any arithmetic	DWORD	WORD is extended with 32 high-order zeros to a DWORD value	

continued

Table 5-2. Summary of Expression Rules for PL/M-386 (continued)

OFFSET	Unsigned	OFFSET w/OFFSET	any arithmetic	OFFSET	range: 0 to 2**32-1
		BYTE w/OFFSET becomes OFFSET w/OFFSET	any arithmetic	OFFSET	BYTE is extended with 24 high-order zeros to an OFFSET value
		WORD w/OFFSET becomes OFFSET w/OFFSET	any arithmetic	OFFSET	WORD is extended with 16 high-order zeros to an OFFSET value
		DWORD w/OFFSET becomes OFFSET w/OFFSET	any arithmetic	OFFSET	range: 0 to 2**32-1
		OFFSET w/DWORD becomes DWORD w/DWORD	any arithmetic	DWORD	OFFSET is extended with 32 high-order zeros to a DWORD value
CHARINT SHORTINT INTEGER	Signed	INTEGER w/INTEGER	+ or - * / or MOD	INTEGER	-2**31 to +2**31-1
REAL	Floating Point	REAL w/REAL	+ - * or /	REAL	
POINTER		POINTER w/POINTER	=	BYTE	0 or 0FFH
SELECTOR	Unsigned	SELECTOR w/SELECTOR	=, <>, <, or >	BYTE	0 or 0FFH

Note: CHARINT and SHORTINT are sign extended to INTEGER before expression evaluation.

The combinations of operands shown in Table 5-2 are the only usable combinations of arithmetic operations and operands. For example, an operand of the signed integer data type cannot be combined with an operand of the unsigned binary number data type. However, explicit conversion can be coded in-line using the PL/M built-ins described in Chapter 9.

Special Case: Constant Expressions

The rules already given explain expressions like:

$$A + 3 * B$$

where there is a single whole-number constant. However, if there is an expression like:

$$3 - 5 + A$$

then the kind of arithmetic that will be used to evaluate $3 - 5$ must be considered, because both operands are whole-number constants.

The answer, in this case, depends on the type of operand A . If A is an unsigned binary number, then $3 - 5$ is considered to be in unsigned context. Unsigned arithmetic is used to evaluate $3 - 5$, giving a `BYTE` result of 254. Unsigned arithmetic is then used to add this result to A .

For PL/M-386, if A is a signed integer, then $3 - 5$ is in signed context. Signed 32-bit arithmetic is used to evaluate $3 - 5$. Signed 32-bit arithmetic is then used to add this result to A .

If A is of type `REAL`, `POINTER`, or `SELECTOR`, the expression is illegal.

Any compound operand, subexpression, or expression that contains only whole-number constants as primary operands is called a constant expression. Floating-point constants are of type `REAL` and are treated as the values of `REAL` variables.

In this expression:

$$3 - 5 + 500 + A$$

$3 - 5$ is a constant expression that forms part of the larger constant expression $3 - 5 + 500$.

If the constant expression is not the entire expression, its value is an operand in the expression. The context is created by the other operand of the same operator.

In the preceding example, suppose the operand A has a `BYTE` value. Then the constant expression $3 - 5 + 500$ is in unsigned context. The constants 3 and 5 are treated as `BYTE` values, and 500 is treated as a `WORD` or `HWORD` value. The operation $3 - 5$ gives a `BYTE` result of 254, and this is extended to a `WORD` or `HWORD` value of 254 before adding 500. This results in a `WORD` or `HWORD` value of 754. It is exactly as if the expression had been written as follows:

$$754 + A$$

If `A` had a `SHORTINT` value, the constant `3 + 5 - 500` would be in signed context; signed 32-bit arithmetic is used for the operation `3 - 5 + 500`. The result (498) is added to the value of `A` to form a 32-bit signed temporary result.

In summary, if the context is created by an unsigned binary number data type operand, the constant expression is in unsigned context. If the context is created by a signed integer data type operand, the constant expression is in signed context. Note that if the context is created by a real number, pointer or selector data type operand, the constant expression is illegal.

If the constant expression is the entire expression, then it belongs in one of the categories listed below. For additional information, see *Assignment Operators*.

- Constant expression as right-hand part of an assignment statement: context is created by the variable to which the expression is being assigned.
- Constant expression as subscript of an array variable: evaluated as if being assigned to an `INTEGER` variable.
- Constant expression in the `IF` part of an `IF` statement: evaluated as if being assigned to a `BYTE` variable.
- Constant expression in a `DO WHILE` statement: evaluated as if being assigned to a `BYTE` variable.
- Constant expression as start, step, or limit expression in an iterative `DO` statement: evaluated as if being assigned to a variable of the same type as the index variable in the same iterative `DO` statement.
- Constant expression in a `DO CASE` statement: evaluated as if being assigned to a `WORD` variable.
- Constant expression as an actual parameter in a `CALL` statement or function reference: evaluated as if being assigned to the corresponding formal parameter in the procedure declaration.
- Constant expression in a `RETURN` statement: evaluated as if being assigned to a variable of the same type as the (typed) procedure that contains the `RETURN` statement.
- Constant expression inside an explicit type conversion (cast built-ins); evaluated as if being assigned to an `INTEGER` variable, shorter values are extended to 16 bits or 32 bits. The only exception is that relational operators can be used and are performed bitwise on 16-bit or 32-bit constant values.

Assignment Statements

Results of computations can be stored as values of scalar variables. At any given moment, a scalar variable has only one value; however, this value can change with program execution. The PL/M assignment statement changes the value of a variable. Its simplest form is:

```
variable =expression;
```

where *expression* is any PL/M expression, as described in the preceding sections. This expression is evaluated, and the resulting value is assigned to (that is, stored in) the variable. This variable can be any fully qualified variable reference except a function reference. The old value of the variable is lost.

For example, following execution of the statement:

```
RESULT = A + B;
```

the variable `RESULT` will have a new value, calculated by evaluating the expression `A + B`.

Implicit Type Conversions

In an assignment statement, if the type of the value of the right-hand expression is not the same as the type of the variable on the left side of the equal sign, then either the assignment is illegal or an implicit type conversion occurs. For PL/M-386, all unsigned binary number, signed integer and real data type values are converted automatically. Chapter 9 includes a description of built-in functions that, when invoked, perform explicit conversions for use in expressions or assignments.

For implicit type conversions, the data type of the value on the right-hand side of the assignment statement is always forced to equal the data type of the value on the left-hand side of the assignment statement. This is done either by extending the value of the expression, or by truncating the value of the expression by the appropriate number of high-order bits so that the data types of both sides of the assignment statement are the same.

The implicit type conversions that occur for assignment statements are summarized in Table 5-3.

Table 5-3. Implicit Type Conversions in Assignment Statements for PL/M-386*

Expression Result Type	Variable on Left of Assignment Statement	Conversion
BYTE	HWORD	BYTE value is extended by 8 high-order 0 bits to HWORD value
	WORD	BYTE value is extended by 24 high-order 0 bits to WORD value
	DWORD	BYTE value is extended by 56 high-order 0 bits to DWORD value
	OFFSET	BYTE value is extended by 24 high-order 0 bits to OFFSET value
HWORD*	BYTE	8 high-order bits of HWORD value are truncated to convert it to a BYTE value
	WORD	HWORD value is extended by 16 high-order 0 bits to convert it to a WORD value
	DWORD	HWORD value is extended by 48 high-order 0 bits to convert it to a DWORD value
	OFFSET	HWORD value is extended by 16 high-order 0 bits to convert it to an OFFSET value
WORD*	BYTE	24 high-order bits of WORD value are truncated to convert it to a BYTE value
	HWORD	16 high-order bits of WORD value are truncated to convert it to a HWORD value
	DWORD	WORD value is extended by 32 high-order 0 bits to convert it to a DWORD value
	OFFSET	No conversion (both WORD and OFFSET are 32-bits)
DWORD*	BYTE	56 high-order bits of DWORD value are truncated to convert it to BYTE value
	HWORD	48 high-order bits of DWORD value are truncated to convert it to HWORD value
	WORD	32 high-order bits of DWORD value are truncated to convert it to WORD value
	OFFSET	32 high-order bits of DWORD value are truncated to convert it to OFFSET value

* Assuming WORD32.

continued

**Table 5-3. Implicit Type Conversions in Assignment Statements for PL/M-386*
(continued)**

Expression Result Type	Variable on Left of Assignment Statement	Conversion
OFFSET**	BYTE	24 high-order bits of OFFSET value are truncated to convert it to a BYTE value
	HWORD	16 high-order bits of OFFSET value are truncated to convert it to a HWORD value
	WORD	No conversion is necessary (both WORD and OFFSET are 32 bits)
	DWORD	OFFSET value is extended by 32-high-order 0 bits to convert it to a DWORD value
INTEGER*	CHARINT	24 high-order bits of INTEGER value are truncated to convert it to CHARINT value
	SHORTINT	16 high-order bits of INTEGER value are truncated to convert it to SHORTINT value
REAL	REAL	Automatically converted to 32-bit value

* Assuming WORD32.

** A warning message is issued if OFFSET values are truncated.

Note that implicit conversion is not performed for POINTER or SELECTOR values. For assignment statements with POINTER or SELECTOR expressions, the left side of the assignment statement would be of the same type as the expression.

Constant Expression

BYTE variable on the left: The constant expression is evaluated in unsigned context. If the resulting value is equal to or greater than 0 and equal to or less than 255, it is treated as a **BYTE** value and no conversion is necessary. If the resulting value is greater than 255, it is truncated to type **BYTE** by dropping all except its 8 low-order bits.

INTEGER variable on the left: The constant expression is evaluated in signed context. No conversion is necessary.

REAL variable on the left: The assignment is illegal unless all values on the right are floating-point constants. If the value of the constant expression is out of the range for **REAL** variables, an overflow exception occurs (see Chapter 10 and Appendix G).

WORD variable on the left: The constant expression is evaluated in unsigned context. If the resulting value is equal to or greater than 0 and equal to or less than 65,535, it is treated as an **WORD** value, and no conversion is necessary. If the resulting value is greater than 65,535, it is truncated to type **WORD** by dropping all except its 16 low-order bits.

WORD variable on the left: The constant expression is evaluated in unsigned context. No conversion is necessary.

DWORD variable on the left: The constant expression is evaluated in unsigned context and is zero-extended to a **DWORD** value.

OFFSET variable on the left: The constant expression is evaluated in unsigned context. No conversion is necessary.

CHARINT variable on the left: The constant expression is evaluated in 32-bit **INTEGER** arithmetic. If the value is less than -128 or greater than +127, it is truncated to 8 bits.

SHORTINT variable on the left: The constant expression is evaluated in 32-bit **INTEGER** arithmetic. If the value is outside the given range for **SHORTINT** (-32,768 to +32,767), it is truncated to 16 bits.

Constants cannot be assigned to **POINTER** or **SELECTOR** variables.

Type conversion built-ins can be used to change the type of a constant expression to the type required for assignment. The entire expression within the type conversion is evaluated in signed context.

Multiple Assignment

It is often convenient to assign the same value to several variables at the same time. This is accomplished in PL/M by listing all the variables to the left of the equal sign, separated by commas. The variables `LEFT`, `CENTER`, and `RIGHT` can all be set to the value of the expression `INIT + CORR` with the single assignment statement:

```
LEFT, CENTER, RIGHT = INIT + CORR;
```

The variables on the left-hand side of a multiple assignment must be all of the same class, that is, all unsigned, all signed, all `POINTER`, all `SELECTOR`, or all `REAL`. Then the conversion rules described previously in this chapter are applied separately to each assignment.

⇒ Note

The order in which the assignments are performed is not guaranteed. Therefore, if a variable on the left side of a multiple assignment also appears in the expression on the right side, the results are undefined.

Embedded Assignments

A special form of assignment can be used within PL/M expressions. The form of this embedded assignment is:

```
variable := expression
```

and can appear anywhere an expression is allowed. The expression (everything to the right of the `:=` assignment symbol) is evaluated and stored in the variable on the left. Parentheses are used to specify the limits of an embedded assignment within an assignment statement. The value of the embedded assignment is the same as that of its right half. For example, the expression:

```
ALT + (CORR := TCORR + PCORR) - (ELEV := HT/SCALE)
```

results in exactly the same value as:

```
ALT + (TCORR + PCORR) - (HT/SCALE)
```

except that the intermediate results `TCORR + PCORR` and `HT/SCALE` are stored in `CORR` and `ELEV`, respectively. These names for intermediate results can then be used at a later point in the program without recalculating their values. The names must have been declared earlier.

The rules of PL/M do not specify the order in which subexpressions or operands are evaluated. When an embedded assignment changes the value of a variable that also appears elsewhere in the same expression, the results cannot be guaranteed.

For example, the following expression:

$$A = (X := X + 4) + Y * Y + X;$$

could mean either of the following interpretations:

$$A1 = (X + 4) + Y * Y + (X + 4);$$
$$A2 = (X + 4) + Y * Y + X;$$

Avoid this ambiguity by removing the embedded assignment from the expression and using a separate assignment statement to achieve the desired effect as follows:

$$X = X + 4;$$
$$A1 = X + Y * Y + X;$$
$$X = X + 4;$$
$$A2 = X + Y * Y + X - 4;$$
$$A3 = X + 4 + Y * Y + X;$$
$$X = X + 4;$$

□ □ □

Flow Control Statements 6

This chapter describes statements that alter the sequence of PL/M statement execution and that group statements into blocks.

DO and END Statements: DO Blocks

Procedures and DO blocks are the basic building units of modular programming in PL/M. (Procedures are discussed in Chapter 8.)

This chapter discusses all four kinds of DO-blocks. Each DO block begins with a DO statement and includes all subsequent statements through the closing END statement. The four kinds of DO blocks are as follows:

- Simple DO block

```
DO;          /* all statements executed, each in order */
  statement-0;
  statement-1;
  statement-2;
  .
  .
END;
```

- DO CASE block

```
DO CASE select_expression; /* one statement executed */
  case-0-statement; /* executed if select_expression = 0 */
  case-1-statement; /* executed if select_expression = 1 */
  .
  .
END;
```

- DO WHILE block

```
DO WHILE expression_true;
  statement-0; /* all executed repeatedly if expression */
  statement-1; /* true, none executed if false. */
  .
  .
END;
```

- Iterative DO block

```

DO counter = start-expr TO limit-expr BY step-expr;
  statement-0; /* all statements executed a number */
  statement-1; /* of times depending on comparison */
                /* of counter with limit expression */
  .
  .
END;
```

The last two blocks are also referred to as DO-loops because the executable statements within them can be executed repeatedly (in sequence) depending on the expressions in the DO statement.

Any DO statement can have multiple labels on it, and only the last of these can appear between the word END and the next semicolon. For example:

```

A: B: C: D: EM: DO;
    .
    .
    .
    END EM ;          /* end of block EM; */
                    /* A, B, C, D also end here. */
```

As mentioned in Chapter 3, the placement of declarations is restricted. Except for use in procedures, declarations are permitted only at the top of a simple DO block, before any executable statements of the block. (This DO can, of course, be nested within other DOs or procedures. Chapter 7 discusses the scope of declared names.)

Each DO block can contain any sequence of executable statements, including other DO blocks. Each block is considered by the compiler as a unit, as if it were a single executable statement. This fact is particularly useful in the DO CASE block and the IF statement, both discussed in this chapter.

The discussions that follow describe the normal flow of control within each kind of DO block. The normal exit from the block passes through the END statement to the statement immediately following. These discussions assume that none of the statements in the block causes control to bypass that process. A GOTO statement with the target outside the block would be one such bypass. (GOTOS are discussed later in this chapter.)

Simple DO Blocks

A simple DO block merely groups, as a unit, a set of statements that will be executed sequentially (except for the effect of GOTOS or CALLS):

```
DO;  
  statement-0;  
  statement-1;  
  . . .  
  statement-n;  
END;
```

For example:

```
DO;  
  NEW$VALUE = OLD$VALUE + TEMP;  
  COUNT = COUNT + 1;  
END;
```

This simple DO block adds the value of TEMP to the value of OLD\$VALUE and stores it in NEW\$VALUE. It then increments the value of COUNT by one.

DO blocks can be nested within each other as shown in the following example:

```
ABLE: DO;  
      statement-0;  
      statement-1;  
BAKER: DO;  
      statement-a;  
      statement-b;  
      statement-c;  
      END BAKER;  
      statement-2;  
      statement-3;  
      END ABLE;
```

The first DO statement and the second END statement bracket one simple DO block. The second DO statement and the first END statement bracket a different DO block inside the first one. Notice how indentation (using tabs or spaces) can be used to make the sequence more readable, so that it can be seen at a glance that one DO block is nested inside another. It is recommended that this practice be followed in writing PL/M programs. See Appendix B for the number of DO blocks that can be nested.

A simple DO block can delimit the scope of variables, as discussed in Chapter 7.

DO CASE Blocks

A DO CASE block begins with a DO CASE statement, and selectively executes one of the statements in the block. The statement is selected by the value of an expression. The maximum number of cases is given in Appendix B. The form of the DO CASE block is:

```
DO CASE select_expression;  
    statement-0;  
    statement-1;  
    . . .  
    statement-n;  
END;
```

In the DO CASE statement, `select_expression` must yield an unsigned binary number (excluding `DWORD`) or a signed integer value. If the expression is a constant expression, it is evaluated as if it were being assigned to a `WORD` variable. The value (call the value K) must be between 0 and n , inclusive. K is used to select one of the statements in the DO CASE block, which is then executed. The first case (`statement-0`) corresponds to $K = 0$; the second (`statement-1`) corresponds to $K = 1$, and so forth. Only one statement from the block is selected. This statement is then executed (only once). Control then passes to the statement following the `END` statement of the DO CASE block.



Note

If the run-time value of the expression in the DO CASE statement is less than 0 or greater than n (where $n + 1$ is the number of statements in the DO CASE block), the effect of the DO CASE statement is undefined. This may have disastrous effects on program execution. Therefore, if there is any possibility that this out-of-range condition may occur, the DO CASE block should be contained within an IF statement that tests the expression to make sure that it has a value that will produce meaningful results.

An example of a DO CASE block is:

```
DO CASE SCORE;
; /* case 0 */
CONVERSIONS=CONVERSIONS + 1; /* case 1 */
SAFETIES = SAFETIES + 1; /* case 2 */
FIELDGOALS = FIELDGOALS + 1; /* case 3 */
; /* case 4 */
; /* case 5 */
TOUCHDOWNS=TOUCHDOWNS + 1; /* case 6 */
END;
```

When execution of this CASE statement begins, the variable SCORE must be in the range 0 to 6. If SCORE is 0, 4, or 5 then a null statement (consisting of only a semicolon, and having no effect) is executed; otherwise the appropriate statement is executed, causing the corresponding variable to be incremented.

A more complex DO CASE block is the following:

```
SELECT = COUNT - 5;
IF SELECT <= 2 AND SELECT >= 0 THEN
  DO CASE SELECT;

    X = X + 1; /* Case 0 */

  DO; /* Begin Case 1 */
    X = Y + 10;
    Y = Y + 1;
  END; /* End Case 1 */

  DO I = LAST$HI + 1 TO TOP - 6; /* Begin Case 2 */
    Z(I) = X * Y + 1;
    W(I) = Z(I) * Z(I);
    V(I) = W(I) - Z(I);
  END; /* End Case 2 */

  END; /* End DO CASE block */
ELSE CALL ERROR;
```

If SELECT and COUNT are INTEGER variables, negative values could occur. The DO CASE block is placed within an IF statement to guarantee that execution of the DO CASE block will not be attempted if the value of SELECT is less than 0 or greater than 2. Instead, a procedure called ERROR (declared previously) will be activated.

The preceding example illustrates the use of a simple DO block as a single PL/M statement. The DO CASE statement can select Case 1 or Case 2 and cause multiple statements to be executed. This is only possible because they are grouped as a simple DO block, which acts as a single statement.

DO WHILE Blocks

DO WHILE and IF statements examine only the least significant bit of the value of the expression. If the value is an odd number (least significant bit = 1), it will be considered true. If it is even (least significant bit = 0), it will be considered false. If the expression is relational, e.g., A<B, the result will have a value of 00H or 0FFH, but this is incidental; it may have any unsigned value.

A DO WHILE block begins with a DO WHILE statement, and has the following form:

```
DO WHILE expression;          /* expression must yield */
    statement-0;              /* an unsigned value */
    statement-1;
    . . .
    statement-n;
END;
```

The effect of this statement is as follows:

1. First the unsigned expression following the reserved word WHILE is evaluated. If the rightmost bit of the result is 1, then the sequence of statements up to the END is executed.
2. When the END is reached, the expression is evaluated again, and again the sequence of statements is executed only if the value of the expression has a rightmost bit of 1.
3. The block is executed over and over until the expression has a value whose rightmost bit is 0. Execution then skips the statements in the block and passes to the statements following the END statement.

Consider the following example:

```
AMOUNT = 1;
DO WHILE AMOUNT <= 3;
    AMOUNT = AMOUNT + 1;
END;
```

The statement AMOUNT = AMOUNT + 1 is executed exactly 3 times. The value of AMOUNT when program control passes out of the block is 4.

Iterative DO Blocks

An iterative DO block begins with an iteration statement and executes each statement in the block, in order, repeating the entire sequence. The form of the iterative DO block is:

```
DO counter = start-expr TO limit-expr BY step-expr ;
  statement-0 ;
  statement-1 ;
  .
  .
  .
END ;
```

The BY `step-expr` phrase is optional; if omitted, a step of 1 is the default.

For PL/M-386, the counter must be a non-subscripted variable of unsigned type: BYTE, HWORD, WORD, or OFFSET, or a signed integer data type: INTEGER, CHARINT, or SHORTINT.

An example of an iterative DO block is:

```
DO I = 1 TO 10;
  CALL BELL;
END;
```

where BELL is the name of a procedure that causes a bell to ring. The bell will ring ten times.

Another example shows how the index-variable can be used within the block:

```
AMOUNT = 0;
DO I = 1 TO 10;
  AMOUNT = AMOUNT + I;
END;
```

The assignment statement is executed 10 times, each time with a new value for I. The result is to sum the numbers from 1 to 10 (inclusive) and leave the sum (namely, 55) as the value of AMOUNT.

The next example uses `step-expr`:

```
/* Compute the product of the first N odd integers */
PROD = 1;
DO I = 1 TO (2*N-1) BY 2;
  PROD = PROD*I;
END;
```

The type of counter (signed or unsigned) affects the following factors in the execution flow of iterative DOs:

- When `step-expr` is evaluated.
- What causes execution to exit the DO block.

The following steps constitute the general execution sequence of an iterative DO block, with both signed and unsigned variables and expressions in the DO itself. Type is mentioned only for steps in which actions or consequences vary according to type. Where the signed case is different, it is described in parentheses. The discussion following this description summarizes the rules and their results for signed and unsigned data types.

1. The `start-expr` is evaluated and assigned to counter.
2. The `limit-expr` is evaluated and compared with counter. (If counter and `limit-expr` are of signed type, then `step-expr` is also newly evaluated at this time.)
 - a. If counter is greater than `limit-expr`, execution exits the DO and passes to the statement following the next END (unless `step-expr` is a negative signed value; if so, the exit occurs only if counter is less than `limit-expr`).
 - b. Otherwise, the statements within the DO block are executed in order until the END statement is reached.
 - c. At the END, a `step-expr` of unsigned type (BYTE, HWORD, or WORD for PL/M-386) is newly evaluated.
3. The counter is incremented by the value of `step-expr`. For unsigned counters, if the new value is less than the old value (due to modulo arithmetic as explained next), the loop is exited immediately. Otherwise, control returns to step 2.

An 8-bit BYTE can represent numbers no larger than 11111111B (255 decimal). The largest number a 16-bit WORD (or HWORD) can represent is 1111111111111111B, which is 65535 decimal. The largest number a 32-bit WORD can represent is 0FFFFFFFH, which is 4,294,967,295 decimal. Adding 1 to these values gives a result of 0. Thus, the new counter can be less than the old.

These rules and their consequences can be summarized in two broad cases:

1. Starting with a non-negative `step-expr`, the loop is exited as soon as any one of the following conditions become true:
 - a. The new counter is greater the new `limit-expr`.
 - b. A signed `step-expr` becomes negative and the new counter is still less than the new `limit-expr`.
 - c. An unsigned `step-expr` causes a lower counter than the one just used.
2. When starting with a negative and signed `step-expr`, then the loop is exited as soon as either of the following two conditions occurs:
 - a. The new counter is less than the new `limit-expr`.
 - b. The new `step-expr` becomes non-negative and the new counter is greater than the new `limit-expr`.

Upon exit from the iterative `DO` block:

1. In all cases `step-expr` has been reevaluated.
2. In all but one case `limit-expr` has been reevaluated. When an unsigned counter has just gone over and become smaller, `limit-expr` is unchanged from its value during the last loop.
3. In all cases `counter` has been changed, but the step value that was added to it varies. If signed, `counter` has been incremented by the former step value before it was reevaluated. For unsigned counters, the newer step has been used.

The following distinctions are important:

- In every case, `start-expr` is evaluated only once and `limit-expr` is evaluated before any execution.
- A signed `step-expr` is evaluated in step 2; other `step-exprs` are evaluated in step 3.
- With an unsigned counter, there cannot be a negative step. Furthermore, stepping down to a `limit-expr` that is less than `start-expr` is not possible because the loop will be exited immediately.

END Statement

An END statement must terminate all DO blocks. An END statement has the following syntax:

```
END [name];
```

Where:

name is the optional name that (if present) should match the label of the corresponding DO statement.

IF Statement

The IF statement provides conditional execution of statements. It takes the form:

```
IF expression THEN statement-a;  
ELSE statement-b; /*optional*/
```

The reserved word THEN and the statement following it are required. The reserved word ELSE and the statement following it are optional.

The IF statement has the following effect: first *expression* is evaluated as if it were being assigned to a variable of type BYTE. If the result is true (rightmost bit is 1) then *statement-a* is executed. If the result is false (rightmost bit is 0), then *statement-b* is executed. Following execution of the chosen alternative, control passes to the next statement following the IF statement. Thus, of the two statements (*statement-a* and *statement-b*) only one is executed.

Consider the following program fragment:

```
IF NEW > OLD THEN RESULT = NEW;  
ELSE RESULT = OLD;
```

Here, RESULT is assigned the value of NEW or the value of OLD, whichever is greater. This code causes exactly one of the two assignment statements to be executed. RESULT always gets assigned some value, but only one assignment to RESULT is executed.

In the event that *statement-b* is not needed, the ELSE part may be omitted entirely. Such an IF statement takes the form:

```
IF expression THEN statement-a;
```


Here, `statement-a` is executed if the value of `expression` has a rightmost bit of 1. Otherwise, nothing happens, and control immediately passes on to the next statement following the `IF` statement.

For example, the following sequence of PL/M statements will assign to `INDEX` either the number 5, or the value of `THRESHOLD`, whichever is larger. The value of `INIT` will change during execution of the `IF` statement only if `THRESHOLD` is greater than 5. The final value of `INIT` is copied to `INDEX` in any case:

```
INIT = 5;
IF THRESHOLD > INIT THEN INIT = THRESHOLD;
INDEX = INIT;
```

The power of the `IF` statement is enhanced by using `DO` blocks in the `THEN` and `ELSE` parts. Since a `DO` block can be used wherever a single statement can be used, each of the two statements in an `IF` statement may be a `DO` block. For example:

```
IF A = B THEN
    DO;
        EQUAL$EVENTS = EQUAL$EVENTS + 1;
        PAIR$VALUE = A;
        BOTTOM = B;
    END;
ELSE
    DO;
        UNEQUAL$EVENTS = UNEQUAL$EVENTS + 1;
        TOP = A;
        BOTTOM = B;
    END;
```

`DO` blocks nested within an `IF` statement can contain further nested `DO` blocks, `IF` statements, variable and procedure declarations, and so on.

Nested IF Statements

Any IF statement (including the ELSE part, if any) can be considered a single PL/M statement (although it is not a block). Thus, the statement to be executed in a THEN or an ELSE clause may in fact be another IF statement.

An IF statement inside a THEN clause is called a nested IF. Nesting may be carried to several levels without needing to enclose any of the nested IF statements in DO blocks, as in the following construction:

```
IF expression-1 THEN
    IF expression-2 THEN
        IF expression-3 THEN statement-a;
```

Here are three levels of nesting. Note that `statement-a` will be executed only if the values of all three expressions are true. Thus, the preceding example is equivalent to:

```
IF expression-1 AND expression-2 AND expression-3
THEN statement-a;
```

Notice that the preceding example of nesting does not have an ELSE part. When using nested IF statements, it is important to understand the following rule of PL/M:

- A set of nested IF statements can have only one ELSE part, and it belongs to the innermost (that is, the last) of the nested IF statements.

This rule could also be restated as follows:

- When an IF statement is nested within the THEN part of an outer IF statement, the outer IF statement may not have an ELSE part.

For example, the construction:

```
IF expression-1 THEN
    IF expression-2 THEN statement-a
    ELSE statement-b;
```

is legal and means that if the values of both `expression-1` and `expression-2` are true, then `statement-a` will be executed. If the value of `expression-1` is true and the value of `expression-2` is false, then `statement-b` will be executed. If the value of `expression-1` is false, neither `statement-a` nor `statement-b` will be executed, regardless of the value of `expression-2`.

The preceding construction is equivalent to:

```
IF expression-1 THEN
    DO;
        IF expression-2 THEN statement-a;
        ELSE statement-b;
    END;
```

This construction is much more readable and offers less opportunity for error.

If the intention is for the ELSE part to belong to the outer IF statement, then the nesting must be done by means of a DO block:

```
IF expression-1 THEN
    DO;
        IF expression-2 THEN statement-a;
    END;
ELSE statement-b;
```

Note that the meaning of this construction differs completely from the previous one.

Finally, consider the following:

```
IF expression-1 THEN
    IF expression-2 THEN
        IF expression-3 THEN statement-a;
        ELSE statement-b;
        ELSE statement-c; /* illegal statement */
    ELSE statement-d; /* illegal statement */
```

This construction is illegal because only one ELSE part is allowed. If the intention is for the ELSE parts to match the IF parts as indicated by the indenting, the nesting must be done with DO blocks, as follows:

```
IF expression-1 THEN
    DO;
        IF expression-2 THEN
            DO;
                IF expression-3 THEN statement-a;
                ELSE statement-b;
            END;
        ELSE statement-c;
    END;
ELSE statement-d;
```

Sequential IF Statements

Consider the following example. An ASCII-coded character is stored in a `BYTE` variable named `CHAR`. If the character is an `A`, `statement-a` should be executed. If the character is a `B`, `statement-b` should be executed. If the character is a `C`, `statement-c` should be executed. If the character is not `A`, `B`, or `C`, `statement-x` should be executed. The code for doing this could be written as follows, using `IF` statements that are completely independent of one another:

```
IF CHAR = 'A' THEN statement-a;
IF CHAR = 'B' THEN statement-b;
IF CHAR = 'C' THEN statement-c;
IF CHAR <> 'A' AND CHAR <> 'B' and CHAR <> 'C'
    THEN statement-x;
```

This sequence is inefficient because all four `IF` statements (six tests in all) will be carried out in every case, which is wasteful when one of the earlier tests succeeds.

`A` must be tested for in all cases. However, `B` needs to be tested only if the test for `A` fails and `C` needs to be tested only if both previous tests fail. Finally, if the tests for `A`, `B`, and `C` all fail, no further tests are needed and `statement-x` must be executed. To improve the code, rewrite it as follows:

```
IF CHAR = 'A' THEN statement-a;
ELSE IF CHAR = 'B' THEN statement-b;
ELSE IF CHAR = 'C' THEN statement-c;
ELSE statement-x;
```

Notice that this sequence is not a case of nested `IF` statements as described in the preceding section. `IF` statements are nested only when one `IF` statement is inside the `THEN` part of another. In the next example, `IF` statements are inside the `ELSE` parts of other `IF` statements. This construction is called sequential `IF` statements. It is equivalent to the following:

```
IF CHAR = 'A' THEN statement-a;
ELSE DO;
    IF CHAR = 'B' THEN statement-b;
    ELSE DO;
        IF CHAR = 'C' THEN statement-c;
        ELSE statement-x;
    END;
END;
```

Sequential `IF` statements are useful whenever a set of tests is to be made, but the remaining tests should be skipped whenever one of the tests succeeds. This construction works in such cases because all the remaining tests are in the `ELSE` part of the current test.

GOTO Statements

A GOTO statement alters the sequential order of program execution by transferring control directly to a labeled statement. Sequential execution then resumes, beginning with the target statement. The GOTO statement has the following form:

```
GOTO label
```

For example:

```
GOTO ABORT;
```

The appearance of *label* in a GOTO statement is called a label reference, not a label definition.

The reserved word GOTO can also be written GO TO, with an embedded blank.

For reasons discussed in Chapter 7, GOTO statements are restricted. The only possible GOTO transfers are the following:

- From a GOTO statement in the outer level of some block to a labeled statement in the outer level of the same block.
- From a GOTO statement in an inner block to a labeled statement in the outer level of an enclosing block (not necessarily the smallest enclosing block). However, if the inner block is a procedure block, the transfer can only be to a statement in the outer level of the main program module.
- From any point in one program module to a labeled statement in the outer level of the main program module. To jump to such a label, the label must be declared to have extended scope, (i.e., declare it PUBLIC in the main module and EXTERNAL in the module containing the GOTO).


The use of GOTOS is necessary in some situations. However, in most situations where control transfers are desired, the use of an iterative DO, DO WHILE, DO CASE, IF, or a procedure activation (see Chapter 8) is preferable. Indiscriminate use of GOTOS will result in a program that is difficult to understand, correct, and maintain.

The CALL and RETURN Statements

The CALL and RETURN statements are mentioned here only for completeness, since they control the flow of a program. However, they are discussed in detail in Chapter 8.

The CALL statement is used to activate an untyped procedure (one that does not return a value).

The RETURN statement is used within a procedure body to cause a return of control from the procedure to the point from which it was activated.



Block Structure and Scope 7

This chapter explains the meaning of outer level and the concept of scope, including the use of the linkage attributes, `PUBLIC` and `EXTERNAL`.

The outer level of a block means statements (or labels) contained in the block but not contained in any nested blocks. The term exclusive extent also has this meaning. The inner level, or inclusive extent, includes this outer level and all nested blocks as well.

A block at the same level as another block means that both blocks are contained by exactly the same outer blocks.

The scope of an object means those parts of a program where its name, type, and attributes are recognized (i.e., handled according to a given declaration). An object means a variable, label, procedure, or symbolic (named) constant (i.e., a compilation constant or execution constant as discussed in Chapter 3). A program is the complete set of modules that are ultimately executed as a unit.

Names Recognized Within Blocks

As shown throughout this manual, PL/M is a block-structured language that enables design implementation for problem solving, data processing, and hardware control.

PL/M is used to create blocks of code containing declarations followed by executable statements. These blocks are ordered and nested in such a way as to simplify and clarify the flow of data and control. (See Appendix B for maximum block nesting.) A collection of these blocks that performs a single function, or a small set of related functions, is usually compiled as one module, as discussed in Chapter 1.

Beyond the advantages of modularity, simplicity, and clarity, the nesting of blocks serves another very basic purpose: names declared at an outer level are known to all statements of all nested blocks as well.

A new meaning can be declared for any such name within a nested simple DO or procedure block, thereby cutting off its earlier meaning for this block. But if this option is not chosen, its meaning is established by a single declaration at an outer level. (The only objects that do not require declarations prior to use are labels and reentrant procedures.)

In Figure 7-1, everything inside the figure (except the title) constitutes the inclusive extent of block MMM (in this case, module MMM). KK is known throughout this block, including all nested blocks.

Everything inside the large box constitutes the inclusive extent of block SORT. JJ and II are known throughout this block, but not outside it. JJ and II are not known before the label SORT or after the END SORT statement.

Everything inside the small box constitutes the inclusive extent of block FIND. Since this is not a simple DO or procedure block, declarations are not allowed. All prior declarations shown are available for use within FIND.


```

MMM: DO;      /* Beginning of module */
      DECLARE RECORD (128) STRUCTURE
          (KEY BYTE,
           INFO WORD);
      DECLARE CURRENT STRUCTURE
          (KEY BYTE,
           INFO WORD);
      DECLARE KK BYTE;
      KK = 127;
          /* Instructions here would read in data. */

SORT:
  DO;
  DECLARE (JJ,ii) INTEGER;
  DO JJ = 1 TO 127;
  CURRENT.KEY = RECORD(JJ).KEY;
  CURRENT.INFO = RECORD(JJ).INFO;
  II = JJ;

  FIND:
    DO WHILE II > 0 AND
      RECORD(II-1).KEY > CURRENT.KEY;
      RECORD(II).KEY = RECORD(II-1).KEY;
      RECORD(II).INFO = RECORD(II-1).INFO;
      II = II-1;
    END FIND;

  RECORD(II).KEY = CURRENT.KEY;
  RECORD(II).INFO = CURRENT.INFO;
  END;
  END SORT;

          /* Instructions here would write out data from the records.
*/
END MMM;      /* End of module */

```

Figure 7-1. Inclusive Extent of Blocks

In Figure 7-1, the area within the large box and outside the small box is the exclusive extent (the outer level) of block `SORT`. The area within the small box is the exclusive (and inclusive) extent of block `FIND`. To the instructions within the `FIND` block, `SORT`'s exclusive extent is an outer level. The outermost level (or module level) is the area outside the large box enclosing the `SORT` block.

Restrictions on Multiple Declarations

In any given block, a known name cannot be redeclared at the same level as its original declaration. A new declaration is permitted inside a nested simple `DO` or procedure block, where it automatically identifies a new object despite the existence of the same name at a higher level. The new object will be the only one known by this name within its block, and it will be unknown outside its block, where the prior name maintains its meaning. These observations also apply when a name is redeclared in another block at the same level as the block containing the original declaration.

When a name is declared only in a separate block at the same level, there is no way to access it except in that block where it is declared. The definition is not at an outer level to the current block. Any local declaration that is supplied establishes a new separate object whose values bear no relation to those of the other.

The reason for these rules, as for many in programming, is that there must be no ambiguity about what address/location is meant by each name in the program. The preceding declaration rules give freedom to choose names appropriate to a given block, without interfering with exterior uses of them. But when a name is redeclared, its outer-level meaning is inaccessible until execution exits the block containing the new declaration. For example:

```
A: DO;
    DECLARE X, Y, Z BYTE;
L1: X = 2;
    Y = X;
    Z = X;
B: DO;
    DECLARE X, Y BYTE;
    X = 3;
    Y = X;
L2: Z = X;
END B;
L3: /* At this point, X=2, Y=2, Z=3, because */
    /* the value of the redeclared X was used */
    /* to fill Z. If statement L2 were outside */
    /* the loop labeled B, then Z would be 2 */
    /* because the outer X value would be used */
```

Extended Scope: The PUBLIC and EXTERNAL Attributes

The PUBLIC and EXTERNAL attributes permit the scope of names to be extended for all objects except modules; a module name cannot be declared with either attribute.

To extend the scope means to make the names available for use in modules other than the one where they are defined. (The names are already available to nested blocks in this module.) Extended scope includes names for variables, labels, procedures, and execution constants.

For example, the statement:

```
DECLARE FLAG BYTE PUBLIC;
```

causes a byte named FLAG to be allocated, and its address made known to any other module where the following declaration occurs:

```
DECLARE FLAG BYTE EXTERNAL;
```

Similarly, if one module has a procedure declaration block that begins:

```
SUMMER: PROCEDURE (A,B) WORD PUBLIC;
        DECLARE (A,B) BYTE;
        .      /* other declarations can go here */
        .      /* executable statements go here, */
        .      /* defining the procedure          */
        .
END SUMMER;
```

then any other module may invoke SUMMER if it first declares:

```
SUMMER: PROCEDURE (A,B) WORD EXTERNAL;      /* A,B can be any */
        DECLARE (A,B) BYTE;      /* names but these names must */
                                   /* match them and each type must */
END SUMMER;      /* match its public definition */
```

The use of `PUBLIC` and `EXTERNAL` must follow a strict set of rules to prevent ambiguity of location or definition. These rules are as follows:

1. These attributes can be used only in a declaration at the outermost level of a module (i.e., never in a nested block).
2. Only one can appear in any declaration, no more than once. Thus:

```
DECLARE ZETA BYTE PUBLIC EXTERNAL;           /* error */
DECLARE RHO WORD PUBLIC PUBLIC;              /* error */
```

and similar constructs are all invalid.

3. Names can be declared `PUBLIC` no more than once. The `PUBLIC` declaration is the defining declaration: the address it creates is used in each procedure or module where the same name is declared `EXTERNAL`. Do not create more than one `PUBLIC` address for any name.
4. Names can be declared `EXTERNAL` only if they are also declared `PUBLIC` in a different module of the program. The `EXTERNAL` attribute is essentially a request to use a `PUBLIC` address. An `EXTERNAL` without a `PUBLIC` is a dead letter. Lack of a definition elsewhere will result in a link-time error.
5. Where the name is declared `EXTERNAL`, it must be given the same type as where it is declared `PUBLIC`. Any contradiction of type would violate the intention to use the location(s) and content(s) defined elsewhere. If the name is declared `PUBLIC` and has the `DATA` attribute, all `EXTERNAL` declarations must also use `DATA`, but cannot assign a value to the constant being declared.
6. Similarly, names declared `EXTERNAL` must not be given a location (using the `AT` clause), or an initialization (using `DATA` or `INITIAL`). Such usage would contradict the fact that names are being defined in another module. However, in the module where this name is declared `PUBLIC`, the use of `AT`, `DATA` (with initialization values present), or `INITIAL` is allowed.
7. Neither `PUBLIC` nor `EXTERNAL` can be applied to a name that is based. For example:

```
DECLARE PTR1 POINTER;
DECLARE V1 BASED PTR1 PUBLIC;
```

is invalid. The reason: by definition, `V1` has no home of its own; its location is always determined by `PTR1`. Thus, to declare `V1` `PUBLIC` or `EXTERNAL` does not permit the correct assignment of addresses. `PTR1`, on the other hand, always contains the current address of `V1`. Declaring the base, in this case `PTR1`, to be `PUBLIC` or `EXTERNAL` is always permissible since it permits valid results.

⇒ **Note**

The PL/M compiler will generate external records only for items that are actually referenced in the program.

8. When extending the scope of a name with the `PUBLIC` attribute and `DATA` or `INITIAL`, the placement in the `DECLARE` statement is critical. `PUBLIC` must be placed after the type declaration and before the `DATA` or `INITIAL` attribute. For example:

```
DECLARE a$P BYTE PUBLIC INITIAL(4);
```

(Additional restrictions on the use of `PUBLIC` and `EXTERNAL` procedures are described in Chapter 8.)

Following these rules will enable consistent and reliable execution of programs using names with extended scope. A `PUBLIC` definition occurring in one module will then be used by all related references to that name in separate modules; that is, references which declare the name `EXTERNAL`. The following diagram illustrates this:

```
MOD1: DO;
    DECLARE V1 BYTE PUBLIC;
    .
    .
    .
    END MOD1;
MOD2: DO;
    DECLARE V1 BYTE EXTERNAL;
    QQ4: PROCEDURE PUBLIC;
    .
    .
    .
    END QQ4;
    END MOD2;
```

Both references to `V1` will use the same definition (location) for `V1`, namely, the definition in module `MOD1`. Similarly, if any module needed to call procedure `QQ4`, it would first need a declaration like this:

```
QQ4:PROCEDURE EXTERNAL;  
END QQ4;
```

so that a subsequent `CALL QQ4` would correctly pass control to that procedure in `MOD2`.

Scope of Labels and Restrictions on GOTOS

Labels are subject to exactly the same rules of scope previously discussed.

A label is unknown outside the block where it is declared. As discussed in Chapter 1, a label is either declared explicitly at the beginning of a simple `DO` or procedure block, or the compiler considers it to be declared there as soon as it is defined by use anywhere in the block. Therefore, the discussion of what names are known in which blocks applies directly to labels as well as to other names.

The label on a block is not part of the block it names. For example, the name on the `DO` enclosing the module itself is not part of the `DO`; it merely names it. For nested blocks, a label is again not part of the block it names, but belongs instead to the outer level as part of that first enclosing block.

If a name used as a label on a block is defined inside that block, it will name a new item, be it label, variable, or constant. There will be no confusion with the outer label name. This fact leads to important restrictions on the use of the `GOTO` statement:

1. It is impossible for a `GOTO` to transfer control from an outer block to a labeled statement inside a nested block.
2. Moreover, a `GOTO` can transfer control from one block to another in the same module only if the target block encloses the one containing the `GOTO` (and only if the name of that target label is not declared in the nested block).

Furthermore, a label with the `PUBLIC` attribute is permitted only in the main module. This has the interesting consequence of forcing all other transfers of control (i.e., those not involving a return to the main module) to use procedure calls. This favors the development of orderly, modularized, traceable programs.

Only four GOTO transfers are possible; these are as follows:

1. From one point in a block to another statement also in the same level of the same block.
2. From an inner, nested DO block (not a nested procedure) to a statement in the outer level of any enclosing block.
3. From a procedure to a statement in the outer level of the main program in the same module.
4. To a main-program label that is declared PUBLIC, from any point in any module that declares that label EXTERNAL.

Recall that only labels at the outer level of a main program can be declared PUBLIC.

Program structure and declarations are shown in Figure 7-2. Figure 7-3 illustrates the only legal GOTO transfers that are permitted among the given labels in Figure 7-2. A single-headed arrow means the transfer is valid only in the direction shown. A double-headed arrow means that a GOTO can be used in either direction.

```

MAIN: DO;
    DECLARE (LAB33, LAB77) LABEL PUBLIC;
    DECLARE IT BYTE;
    . . .
LAB33: . . . ;
        DO;
        . . .
        END;
    . . .
LAB77: . . . ;
        DO WHILE IT > 0;
        . . .
        END;
    . . .
END MAIN;
MOD1: DO;
    DECLARE (LAB33,LAB77) LABEL EXTERNAL;
    P1: PROCEDURE;
        L1: . . . ;
            DO;
            DECLARE KO BYTE;
            P2: PROCEDURE;
                . . .
                L2: . . . ;
                . . .
            END P2;
            END;
        L3: . . . ;
        . . .
    END P1;
END MOD1;
MOD2: DO;
    . . .
    DECLARE (LAB33,LAB77) LABEL EXTERNAL;
    P4: PROCEDURE;
        . . .
        L4: . . . ;
        . . .
        L5: . . . ;
            DO;
                L6: . . . ;
                . . .
            END;
        . . .
        L7: . . . ;
        . . .
    END P4;
    LB: . . . ;
END MOD2;

```

Figure 7-2. Sample Program Modules Illustrating Valid GOTO Usage

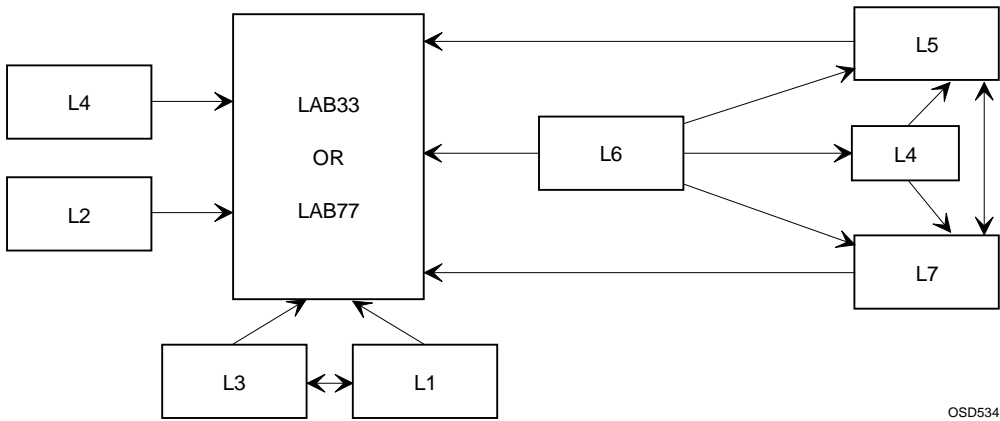


Figure 7-3. Sample Program Modules Illustrating Valid GOTO Transfers



A procedure is a section of PL/M code that is declared without being executed, and then activated from other parts of the program. A function reference or `CALL` statement activates the procedure, even if it is physically located elsewhere. Program control is transferred from the point of activation to the beginning of the procedure code, and the code is executed. Upon exit from the procedure code, program control is passed back to the statement immediately after the point of activation.

The use of procedures forms the basis of modular programming. It facilitates making and using program libraries, eases programming and documentation, and reduces the amount of object code generated by a program. The following sections review how to declare and activate procedures.

Procedure Declarations

A procedure must be declared, just as variables must be declared. Thereafter, any reference to a procedure must occur within the scope defined by the procedure declaration. Also, a procedure cannot be used (called, or invoked in an expression) until after the `END` statement of the procedure declaration unless it is reentrant.

A procedure declaration consists of three parts: a `PROCEDURE` statement, a sequence of statements forming the procedure body, and an `END` statement.

The following is a simple example:

```
DOOR$CHECK: PROCEDURE;  
    IF FRONT$DOOR$LOCKED AND SIDE$DOOR$LOCKED THEN  
        CALL POWER$ON;  
    ELSE CALL DOOR$ALARM;  
END DOOR$CHECK;
```

where `POWER$ON` and `DOOR$ALARM` are procedures declared previously in the same program.



Note

The name `DOOR$CHECK` in a `PROCEDURE` statement has the same appearance as a label definition, but it is not considered a label definition, and a procedure name is not a label. `PROCEDURE` statements cannot be labeled.

The name `DOOR$CHECK` is a PL/M identifier, which is associated with this procedure. The scope of a procedure is governed by the placement of its declaration in the program text, just as the scope of a variable is governed by the placement of its `DECLARE` statement (see Chapter 7 for a detailed description). Within this scope, the procedure can be activated by the name used in the `PROCEDURE` statement.

A procedure declaration, like a `DO` block, controls the scope of variables as described in Chapter 7. Also, like a simple `DO` block, a procedure declaration can contain `DECLARE` statements, which must precede the first executable statement in the procedure body.

As in a `DO` block, the identifier in the `END` statement has no effect on the program, but helps legibility and debugging. If used, it should be the same as the procedure name.

The parameter list and the type are discussed in the following two sections.

Parameters

Formal parameters are non-based scalar variables declared within a procedure declaration; their identifiers appear in the parameter list in the `PROCEDURE` statement. The identifiers in the list are separated by commas and the list is enclosed in parentheses. No subscripts or member-identifiers can be used in the parameter list.

If the procedure has no formal parameters, the parameter list (including the parentheses) is omitted from the `PROCEDURE` statement.

Each formal parameter must be declared as a non-based scalar variable in a `DECLARE` statement preceding the first executable statement in the procedure body. However, procedure parameters are not stored according to the same rules as other declared variables. In particular, do not assume that a parameter is stored contiguously with other variables declared in the same factored variable declaration.

When a procedure that has formal parameters is activated, the `CALL` statement or function reference contains a list of actual parameters. Each actual parameter is an expression whose value is assigned to the corresponding formal parameter in the procedure before the procedure begins to execute.

For example, the following procedure takes four parameters, called PTR, N, LOWER, and UPPER. It examines N contiguously stored BYTE variables. The parameter PTR is the location of the first of these variables. If any of these variables is less than the parameter LOWER or greater than the parameter UPPER, the ERRORSET procedure (declared previously in the program) is activated:

```
RANGE$CHECK: PROCEDURE(PTR, N, LOWER, UPPER);
    DECLARE PTR POINTER;
    DECLARE (N, LOWER, UPPER, I) BYTE;
    DECLARE ITEM BASED PTR(1) BYTE;

    DO I = 0 TO N - 1;
        IF (ITEM(I) < LOWER) OR (ITEM(I) > UPPER)
            THEN CALL ERRORSET;
        /* ERRORSET is a procedure declared previously */
    END;
END RANGE$CHECK;
```

Notice that the array ITEM is declared to have only one element. Since it is a based array, a reference to any element of ITEM is really a reference to some location relative to the location represented by PTR. In writing the procedure RANGE\$CHECK, a dimension specifier that is any arbitrary number greater than zero must be supplied for ITEM so that references to ITEM can be subscripted. But it does not matter what the dimension specifier is (1 is arbitrarily used here).

Having made this declaration, suppose that 25 variables are stored contiguously in an array called QUANTS. To check that all of these variables have values within the range defined by the values of two other BYTE variables, SMALL and LARGE, write:

```
CALL RANGE$CHECK (@QUANTS, 25, SMALL, LARGE);
```

When this CALL statement is processed, the following sequence occurs:

- The four actual parameters in the CALL statement (@QUANTS, 25, SMALL, and LARGE) are assigned to the formal parameters PTR, N, LOWER, and UPPER, which were declared within the procedure RANGE\$CHECK. Since ITEM is based on PTR and the value of PTR is @QUANTS, every reference to an element of ITEM becomes a reference to the corresponding element of QUANTS.
- The executable statements of the procedure RANGE\$CHECK are executed. If any of the values are less than the value of SMALL or greater than the value of LARGE, the procedure ERRORSET is activated.
- Finally, control returns to the statement following the CALL statement.

Notice how the use of a based variable, with the base passed as a parameter, allows the procedure to have its own unchanging name (ITEM) for a set of variables which may be a different set each time the procedure is activated.

Parameters are placed on the stack in left-to-right order. The stack grows from higher locations to lower locations, so the first parameter occupies the highest position on the stack, and the last parameter occupies the lowest position. For more information, see Appendix F.



Note

PL/M does not guarantee the order in which multiple actual parameters will be evaluated when the procedure is activated. If one actual parameter changes another actual parameter, the results are undefined. This can occur if an expression used as an actual parameter contains an embedded assignment or function reference that changes another actual parameter for the same procedure.

Typed Versus Untyped Procedures

The preceding `RANGE$CHECK` procedure is an untyped procedure. No type is given in the `PROCEDURE` statement, and it does not return a value. An untyped procedure is activated by using its name in a `CALL` statement.

A typed procedure, also called a function, has a type in its `PROCEDURE` statement: an unsigned binary number, signed `INTEGER`, `REAL` number, `POINTER` or `SELECTOR` data type. Such a procedure returns a value of this type, which is used in an expression or stored as the value of a variable. The procedure is activated by using its name as an operand in an expression; this special type of variable reference is called a function reference.

When the expression is processed at run time, the function reference causes the procedure to be executed. The function reference itself is then replaced by the value returned by the procedure. The expression containing the function reference is then evaluated, and program execution continues in normal sequence.

Like an untyped procedure, a typed procedure can have parameters. They are handled as described in the previous section.

The body of a typed procedure can contain a `RETURN` statement with an expression, as explained later in this chapter.



Note

The body of a typed procedure can contain code (such as an assignment statement) that changes the value of some variable declared outside the procedure. This is called a side effect.

Recall that PL/M does not guarantee the order in which operands in an expression are evaluated. Therefore, if a function used in an expression changes the value of another variable in the same expression, the value of the expression depends on whether the function reference or the variable is evaluated first.

If the analysis of the expression does not force one of these operands to be evaluated before the other, then the value of the expression is undefined.

This situation can be avoided by using parentheses to segregate any typed procedure that has a side effect, or by using this procedure in an assignment statement first to create an unambiguous sequence.

Activating a Procedure: Function References and CALL Statements

The two forms of procedure activation depend on whether the procedure is typed or untyped. An untyped procedure is activated by means of a CALL statement, which has the form:

```
CALL name;
```

or

```
CALL name (parameter list);
```

For example:

```
CALL REORDER (@RANK$TABLE, 3);
```

(An alternate form of the CALL statement is discussed later.)

A typed procedure is activated by means of a function reference, which is an operand in an expression. A function reference has the form:

```
name
```

or

```
name (parameter list)
```

This occurs as an operand in an expression, as in the following example:

```
TOTAL = SUBTOTAL + SUM$ARRAY (@ITEMS, COUNT);
```

where `SUM$ARRAY` is a previously declared typed procedure. The value added to `SUBTOTAL` will be the value returned by `SUM$ARRAY` using the actual parameters (`@ITEMS, COUNT`).

In both forms of procedure activation, the elements of the parameter list are called actual parameters to distinguish them from the formal parameters of the procedure declaration. At the time of activation, each actual parameter is evaluated and the result is assigned to the corresponding formal parameter in the procedure declaration. Then, the procedure body is executed. Any PL/M expression may be an actual parameter if its type is the same as that of the corresponding formal parameter.

The actual parameter list in a procedure activation must also match the formal parameter list in the procedure declaration. That is, it must contain the same number of parameters of the same type (except as described in the next paragraph) in the same order. If the procedure is declared without a formal parameter list, then no actual parameter list can be used in the activation.

As in expression evaluation and assignment statements (see Chapter 5), a few type conversions are performed automatically when necessary in activating and returning from a procedure. The built-in explicit type conversion procedures described in Chapter 9 can also be used to force the value of an expression to a desired type.

Indirect Procedure Activation

The `CALL` statement, in the form shown in the preceding section, activates an untyped procedure by its name. It is also possible to activate an untyped procedure by its location. This is done by means of a `CALL` statement with the form:

```
CALL identifier[.member-identifier] [(parameter list)];
```

The identifier cannot be subscripted; however it can be a structure reference. The identifier must be a fully qualified `POINTER` or `WORD` type variable reference for PL/M-86 and PL/M-286, and a fully qualified `POINTER`, `OFFSET`, or `WORD` type variable reference for PL/M-386. Its value is assumed to be the location of the entry-point of the procedure being activated.

**Note**

Calls through 48-bit `POINTERS` will be translated into long calls whereas calls through 32-bit `OFFSETS`, `WORDS`, or `POINTERS` (in the `SMALL` case) will be translated into short calls (relative to the current code segment).

The identifier for the indirect procedure activation cannot be an `HWORD`. Therefore, all variables used for indirect calling in programs that are recompiled from PL/M-286 and use the `WORD16` control should have `DWORD`, `OFFSET` (or `ADDRESS`) data types.

A normal `CALL` uses the name of the procedure; the compiler checks to make sure that the correct number of parameters is supplied and performs automatic type conversion on the actual parameters.

When the `CALL` statement uses a location, the compiler does not check the number of parameters or perform type conversion. However, type conversion is performed if the actual argument is a constant expression. The constant expression is evaluated in unsigned context, as described in Chapter 5. If the number of parameters is wrong or if an actual parameter is not of the same type as the corresponding formal parameter, the results are unpredictable.

Code Examples

The following code examples illustrate an indirect call for the COMPACT model. The first example is a procedure which, when compiled, generates warnings.

```
1      $COMPACT
2
3      CALLF:DO;
4      DECLARE dummy word,
5      inner_p pointer,
6      main_p pointer;
7
8      funct1:PROCEDURE;
9          DECLARE i WORD;
10         i = 0;
11         RETURN;
12     END funct1;
13     funct:PROCEDURE;
14         DECLARE i WORD;
15         i = 0;
16         inner_p = @funct1;
17         call inner_p;
18         RETURN;
19     END funct;
20
21     dummy = .funct;
22     CALL dummy;
23     main_p = @funct;
24     CALL main_p;
25     END callf;
```

Warnings are generated at lines 16 and 23. The warnings occur because of conflicts in FAR and NEAR calls. In most cases of using the COMPACT segmentation model, indirect function calls are NEAR calls. The "@" operator causes FAR function calls. Therefore, indirectly activating a function using the "@" operator in a COMPACT model causes a FAR call, however, the function will execute a NEAR return. This causes the compiler to generate a warning.

The warning is based on stack corruption. A long call pushes the segment selector and offset addresses onto the stack. COMPACT functions do a NEAR RETURN (unless they are on the EXPORT list). Therefore, only the OFFSET for the RETURN address is popped. This leaves the previously pushed segment selector on the stack.

The following example properly demonstrates indirect procedure calls in the COMPACT model. This method uses the "." operator to generate a NEAR call. This operator is similar to the "@" operator except it generates an address of the type WORD.

```
1      $COMPACT
2
3      CALLF: DO;
4      DECLARE dummy WORD;
5
6      funct:PROCEDURE;
7          DECLARE i WORD;
8          i=0;
9      RETURN;
10     END funct;
11
12     dummy=funct;
13     CALL dummy;
14     END CALLF;
```



Note

Do not use the "." operator when using a pointer to a function as required by certain iRMX system calls. These calls, such as **rq_create_task** and **rq_create_job**, expect a pointer to a task address, not just the offset. The interface to iRMX system libraries requires a 32-bit pointer as a parameter. The "@" operator must be used when the pointer to the start address of the task is passed to the iRMX system call. No compiler warning is generated because the task never returns, causing no stack corruption.

Exit from a Procedure: The RETURN Statement

The execution of a procedure is terminated in one of three ways:

- By execution of a RETURN statement within the procedure body. A typed procedure must terminate with a RETURN statement that has an expression.
- By executing a GOTO to a statement outside the procedure body. The target of the GOTO must be at the outer level of the main program (see Chapter 7).
- By reaching the END statement that terminates the procedure declaration.

The RETURN statement takes one of two forms:

```
RETURN;
```

or

```
RETURN expression;
```

The first form is used in an untyped procedure. The second form is used in a typed procedure. The value of the expression becomes the value returned by the procedure. It is evaluated as if it were being assigned to a variable of the same type as used on the PROCEDURE statement.

The Procedure Body

The statements within the procedure body can be any valid PL/M statements, including CALL statements as well as nested procedure declarations.

Examples

1. The following is a typed procedure declaration:

```
AVG: PROCEDURE (X, Y) REAL;  
      DECLARE (X, Y) REAL;  
      RETURN (X + Y)/2.0;  
END AVG;
```

This procedure could be used as follows:

```
SMALL = 3.0;  
LARGE = 4.0;  
MEAN = AVG (SMALL, LARGE);
```

The effect would be to assign the value 3.5 to MEAN.

2. The following is an untyped procedure:

```
AOUT: PROCEDURE (ITEM);  
      DECLARE ITEM WORD;  
      IF ITEM >= OFFH THEN COUNTER = COUNTER + 1;  
      RETURN;  
END AOUT;
```

Here COUNTER is some variable declared outside the procedure (i.e., it is a global variable). This procedure could be activated as follows:

```
CALL AOUT (UNKNOWN);
```

If the value of the variable UNKNOWN is greater than or equal to OFFH, the value of COUNTER will be incremented.

3. This example demonstrates an important use of based variables:

```
SUM$ARRAY: PROCEDURE (PTR,N) BYTE;  
            DECLARE PTR POINTER,  
                    ARRAY BASED PTR(1) BYTE,  
(N,SUM,I)BYTE;  
            SUM = 0;  
            DO I = 0 TO N;  
                SUM = SUM + ARRAY(I);  
            END;  
            RETURN SUM;  
        END SUM$ARRAY;
```

This procedure returns the sum of the first $N + 1$ elements (from the zeroth to the N th) of a `BYTE` array pointed to by `PTR`. Notice that `ARRAY` is declared to have 1 element. Since it is a based variable, no space is allocated for it. It must be declared as an array (with a non-zero dimension) so that it can be subscripted in the iterative `DO` block. The choice of 1 as the constant in the dimension specifier is arbitrary and does not restrict the value of N that may be supplied when the procedure is activated.

The procedure could be used as follows to sum the elements of a 100-element `BYTE` array named `PRICE`, and to assign the sum to the variable `TOTAL`:

```
TOTAL = SUM$ARRAY(@PRICE,99);
```

The Attributes: PUBLIC and EXTERNAL, INTERRUPT, REENTRANT

The PUBLIC and EXTERNAL attributes can be included in PROCEDURE statements to give procedures extended scope. Extended scope is discussed in Chapter 7.

A procedure declaration with the PUBLIC attribute is called a defining declaration. A procedure declaration with the EXTERNAL attribute is called a usage declaration. Most of the rules for PUBLIC and EXTERNAL appear in Chapter 7. The following additional rules apply to the use of the EXTERNAL attribute in a procedure declaration:

1. The EXTERNAL attribute cannot be used in the same PROCEDURE statement as a PUBLIC or REENTRANT attribute. Note, however, that the defining declaration of a procedure may have the REENTRANT attribute.
2. A usage (EXTERNAL) declaration of a procedure should have the same number of parameters as the defining (PUBLIC) declaration. Variable types and dimension specifiers should match up in the same sequence in both declarations. The names of the parameters need not be the same. Note that a discrepancy between the parameter lists in the defining declaration and in a usage declaration will not be automatically detected (see Chapter 11 for a description of the TYPE control to detect such an error at module linkage time).
3. The procedure body of a usage declaration cannot contain anything except the declarations of the formal parameters. The formal parameters must be declared with the same types as in the defining declaration.
4. No labels can appear in a usage declaration.

⇒ Note

The PL/M compiler will generate external records only for items that are actually referenced in the program.

For example, the procedure AVG (from example 1 in "The Procedure Body") can be altered by giving it the PUBLIC attribute:

```
AVG: PROCEDURE (X,Y) REAL PUBLIC;  
      DECLARE (X,Y) REAL;  
      RETURN (X + Y)/2.0;  
END AVG;
```

Another module would have a usage declaration, as follows:

```
AVG: PROCEDURE (X,Y) REAL EXTERNAL;  
      DECLARE (X,Y) REAL;  
END AVG;
```

Now, in the module with the usage declaration, `AVG` can be referenced in an executable statement:

```
MIDDLE = AVG (FIRST, LATEST);
```

thereby activating the procedure `AVG` as declared in the first module.

Interrupts and the INTERRUPT Attribute

The `INTERRUPT` attribute enables definition of a procedure to handle some condition signaled by a microprocessor interrupt (e.g., from a peripheral device). A procedure with this attribute is activated when the corresponding interrupt signal is received in the target system. The PL/M statement `CAUSE$INTERRUPT` (*constant*) can also be used to initiate an interrupt signal (see Chapter 10).

Note that the following discussion applies only to interrupt procedures; interrupt tasks are discussed in Appendix G.

The `INTERRUPT` attribute can be used only in declaring an untyped procedure with no parameters at the outermost level of a program module. It must be declared `PUBLIC` or `EXTERNAL` (and optionally `REENTRANT`). The form is:

```
INTERRUPT
```

At build time, an interrupt vector is assigned to each interrupt procedure.

The following discussion of the microprocessor interrupt mechanism clarifies how interrupt procedures work. Additional information can be found in Appendix G.

The microprocessor interrupt mechanism has two states: enabled or disabled. With the `ENABLE` statement, interrupts can take effect. The `DISABLE` statement prevents interrupts from having any effect. The `HALT` statement also enables interrupts. (The state of the microprocessor interrupt mechanism upon initialization is determined by the operating system.)

When some peripheral device sends an interrupt to the CPU, it is ignored if the interrupt mechanism is disabled. If interrupts are enabled, the interrupt is processed as follows:

1. The CPU completes any instruction currently in execution.
2. The CPU sends an acknowledge interrupt signal, then the interrupting device sends its interrupt number.
3. The interrupt mechanism is disabled. This prevents any other device from interfering.
4. Control passes to the interrupt procedure whose number matches the number sent by the peripheral device. If no such procedure has been established, the results are undefined (since the vector that transfers control may be uninitialized).
5. When the procedure is through (by executing a RETURN or reaching the END of the procedure), the interrupt mechanism is enabled so other devices can be serviced, and control returns to the point where the interrupt occurred.

It is possible (as with other untyped procedures) for the procedure to terminate by executing a GOTO with a target outside the procedure in the outer level of the main program module. In this case, control will never be returned to the point where the program was interrupted, and interrupts will not be enabled automatically.

The following is an example of an interrupt procedure for a system where a peripheral device initiates an interrupt whenever the temperature of a device exceeds a certain threshold. The interrupt procedure turns on the annunciator light, updates a status word, and returns control to the program:

```
HITEMP: PROCEDURE INTERRUPT 100 PUBLIC;
        CALL ANNUNCIATOR(1);
        /* This will result in an output from the microprocessor
        to turn on annunciator light number 1, the
        high-temperature warning. */

        ALERT = ALERT OR 00000010B;
        /* This puts a 1 in one of the bit positions
        of ALERT, which contains a bit pattern
        representing current alerts. */
END HITEMP;
```

Reentrancy and the REENTRANT Attribute

With the `REENTRANT` attribute, a procedure can suspend execution temporarily, restart with new parameters, and then later complete the original execution successfully as if there had been no interruption.

This ability is desirable in two circumstances: (1) if the procedure (`PROC1`) activates itself (called direct recursion), or (2) if the procedure activates another procedure (`PROC2`) that will reactivate `PROC1` before `PROC1` has finished its original processing (called indirect recursion).

Without the `REENTRANT` attribute, storage for procedure variables is allocated statically, in fixed locations within the data segment of the object module. Re-entering such a procedure would write over the earlier contents of such locations making it impossible to complete the original suspended execution.

When the attribute `REENTRANT` is used in declaring a procedure, its variables are not stored with other variables in the data section, but are stored on the stack. Thus preserved, each set can be used independently by each invocation of the procedure.

Hence, multiple sets of variables might need to be stored on the stack during recursive use of such procedures. A stack size must be specified (when binding the program module) that is large enough for all such storage needed by all multiple invocations that may be active at one time.

A procedure with the `REENTRANT` attribute may be activated before it is declared. This permits direct recursion, where the procedure activates itself and permits indirect recursion, where the procedure activates a second procedure and the second procedure activates the first, or activates a third procedure, which activates a fourth, and so forth, with the result that the first procedure is activated before it terminates.


The following rules summarize the use of the `REENTRANT` attribute:

- Any procedure that can be interrupted and is also activated from within an interrupt procedure should have the `REENTRANT` attribute.

Note that this may apply to an interrupt procedure that runs with interrupts enabled because it contains an `ENABLE` statement. If there is any possibility that it will be interrupted by its own interrupt, it should have the `REENTRANT` attribute. This situation is equivalent to recursion.

- Any procedure that is directly recursive (activates itself) should have the `REENTRANT` attribute.
- Any procedure that is indirectly recursive (activates another procedure and is activated itself as a result) should have the `REENTRANT` attribute.

- Any procedure that is activated by a reentrant procedure should also have the `REENTRANT` attribute. In other words, if there is any possibility that a procedure can be activated while it is already running, it should be `REENTRANT`.
- The `REENTRANT` attribute cannot be used in the same declaration as the `EXTERNAL` attribute. (It may be used with the `PUBLIC` attribute.)
- The `REENTRANT` attribute can only be used in a `PROCEDURE` statement at the outer level of a module.
- A procedure declaration with the `REENTRANT` attribute cannot have a nested procedure declaration.



Built-in Procedures, Functions, and Variables **9**

Built-in procedures, functions, and variables are already declared in the PL/M code. This makes it unnecessary to write code to perform the particular functions that built-ins are designed to perform. The following built-in procedures, functions, and variables are discussed in this chapter:

- `LENGTH`, `LAST`, and `SIZE` functions – these functions return information concerning variables. For example, the `SIZE` function returns the number of bytes occupied by a scalar, array, or structure.
- Explicit type and value conversion functions – these functions provide explicit conversion for types and values.
- Shift and rotate functions – these functions move bits using a pattern of 8, 16, or 32 bits.
- String manipulation procedures and functions – these procedures and functions move strings, compare strings, search strings for a match or a mismatch, translate strings, and set strings to a specified value.
- Bit manipulation procedures – these functions copy (and move) a bit string and search bit strings for a set bit.
- `MOVE` bytes – this procedure moves a specified number of bytes from one location to another.
- Time delay – this procedure causes a time delay.
- Lock set – this function enables a software synchronization lock.
- Lock bit – this function enables a memory location lock.
- `POINTER` and `SELECTOR` functions – these functions enable the manipulation of location addresses in the microprocessor's memory.

The identifiers for these built-ins are subject to the rules of scope (described in Chapter 7). This means that the name of a built-in procedure or variable can be declared to have a local meaning (scope) within the program. Within the scope of such a declaration, the built-in is unavailable. This distinguishes these identifiers from reserved words (listed in Appendix A), which cannot be used as identifiers in declarations.

No built-in procedure can be used within a location reference (e.g., @LENGTH(LIST)). No built-in variable can be used within a location reference, except as specifically noted in the following sections.

Obtaining Information About Variables

PL/M has three built-in procedures that take variable names as actual parameters and return information based on the declarations of the variables: LENGTH, LAST, and SIZE.

The LENGTH Function

LENGTH is a built-in WORD function that returns the number of elements in an array; it is activated by a function reference with the form:

```
LENGTH ( variable-ref )
```

Where:

variable-ref must be a non-subscripted reference to an array.

The array can be a member of a structure; it cannot be an EXTERNAL array using the implicit dimension specifier (see Chapter 3).

The value returned is the number of elements assigned to the array in the declaration statement (i.e., the value of the dimension specifier).

If the array is not a structure member, then the reference must be an unqualified variable reference. If the array is a structure member, then the reference is a partially qualified variable reference. For example, given the declaration:

```
DECLARE RECORD STRUCTURE (KEY BYTE,  
                           INFO(3) WORD);
```

LENGTH(RECORD.INFO) is a valid function reference and returns a WORD value of 3.

If the array is a member of a structure, and that structure is an element of an array, a special case arises. Given the declaration:

```
DECLARE LIST (4) STRUCTURE (KEY BYTE,  
                           INFO (3) WORD);
```

then all of the following function references are correct and return the value 3:

```
LENGTH(LIST(0).INFO)  
LENGTH(LIST(1).INFO)  
LENGTH(LIST(2).INFO)  
LENGTH(LIST(3).INFO)
```

In other words, the subscript for the array `LIST` is irrelevant when a member-identifier is supplied, because the arrays within the structures are all the same length. PL/M has a shorthand form of partially qualified variable reference in the `LENGTH`, `LAST`, and `SIZE` function references. For example:

```
LENGTH(LIST.INFO)
```

is a valid function and returns the value 3.

The LAST Function

`LAST` is a built-in `WORD` function that returns the subscript of the last element in an array. It is activated by a function reference with the form:

```
LAST (variable)
```

Where:

variable must be a non-subscripted reference to an array.

The array can be a member of a structure; it cannot be an `EXTERNAL` array using the implicit dimension specifier (see Chapter 3).

The value returned is the subscript of the last element of the array. For a given array, `LAST` will always be one less than `LENGTH`. When used with a based variable, `LAST` returns the value assigned in the declaration statement. This is not necessarily the actual value.

As in the `LENGTH` function, a shorthand form of partially qualified variable reference is allowed in the case where the array is a member of a structure that is also an array element.

The SIZE Function

`SIZE` is a built-in `WORD` function that returns the number of bytes occupied by a scalar, array or structure. It is activated by a function reference with the form:

```
SIZE (variable)
```

Where:

variable is a fully qualified, partially qualified, or unqualified reference to any scalar, array, or structure. The variable cannot be an `EXTERNAL` declaration that uses the implicit dimension specifier (see Chapter 3).

The value returned is the number of bytes required by the variable referenced. When used with a based variable, `SIZE` returns the value assigned in the declaration statement. This is not necessarily the actual (current) value.

If the reference is fully qualified, it refers to a scalar, and the value is the number of bytes required for the scalar. If the reference is unqualified, it refers to an entire structure or array, and the value is the total number of bytes required for the structure or array.

If the reference is partially qualified, it refers either to a structure member that is an array or nested structure, or to an array element that is a structure. The value is the number of bytes required for the array or structure.

As in the `LENGTH` function, a shorthand form of partially qualified variable reference is allowed in the case where the array or scalar is a member of a structure and the structure is an array element.

Explicit Type and Value Conversions

The functions in this section provide explicit conversion from one data type to another and from signed values to or from absolute magnitudes.

Explicit type and value conversion functions are invoked as:

```
function-name (expression)
```

In Tables 9-1 and 9-2, each function name is followed by the expression type expected, the purpose of the function, and the nature of the value it returns to the expression that invoked it. For each function there is only one possible class of expressions (e.g., `HIGH` accepts only unsigned values) that can be converted. For the type conversions (`BYTE`, `WORD`, `DWORD`, `INTEGER`, `REAL`, `POINTER`, and `SELECTOR`, `OFFSET`, `HWORD`, `CHARINT`, and `SHORTINT`), the context of the entire expression is always a signed integer value. Table 9-1 gives the value and type conversions for PL/M-386 when the `WORD32` control is in effect.

Table 9-1. Value and Type Conversions for PL/M-386

Procedure Name	Parameter Type	Function	Result Returned
LOW	BYTE		BYTE value unchanged
	WORD	Converts WORD value to BYTE value	Low-order BYTE of WORD
	DWORD or OFFSET	Converts DWORD or OFFSET value to WORD value	Low-order DWORD of DWORD or OFFSET
	QWORD	Converts QWORD value to WORD value	Low-order WORD of QWORD
HIGH	BYTE		zero
	WORD	Converts WORD value to BYTE value	High-order BYTE of WORD
	DWORD or OFFSET	Converts DWORD or OFFSET value to WORD value	High-order DWORD of DWORD or OFFSET
	QWORD	Converts QWORD value to WORD value	High-order WORD of QWORD
DOUBLE	BYTE	Converts BYTE value to WORD value	WORD, by appending 8 high-order zero bits
	WORD	Converts WORD value to DWORD value	DWORD, by appending 16 high-order zero bits
	QWORD or OFFSET	Converts QWORD or OFFSET value to QWORD value	QWORD, by appending 32 high-order zero bits
	QWORD	QWORD value unchanged	
FLOAT	CHARINT SHORTINT INTEGER	Converts signed integer value to REAL value	Same value of type REAL
FIX	REAL	Converts REAL value to INTEGER value	Same value of type INTEGER if within range -2^{31} to $+2^{31}-1$ otherwise undefined
INT	BYTE WORD DWORD	Converts unsigned binary value to INTEGER value, interprets parameter as positive	Same value of type INTEGER if within range -2^{31} to $+2^{31}-1$ otherwise
SIGNED	BYTE	Converts unsigned integer value to INTEGER value	BYTE value is extended with 24 high-order zeros
	WORD		WORD value is extended with 16 high-order zeros
	DWORD		DWORD value unchanged

continued

Table 9-1. Value and Type Conversions for PL/M-386 (continued)

Procedure Name	Parameter Type	Function	Result Returned
UNSIGN	CHARINT SHORTINT INTEGER	Converts INTEGER value to WORD value	Signed INTEGER value is interpreted as unsigned WORD value
ABS	REAL	Converts negative real value to positive real value	Absolute value of parameter: value unchanged if positive -(value) if negative. Result type is same as parameter type.
IABS	CHARINT SHORTINT INTEGER	Converts negative integer to positive integer	Absolute value of parameter: value unchanged if positive -(value) if negative. If -(value) is out of range, result is undefined. Result type is same as parameter type.
BYTE	any unsigned type	Converts any unsigned type to BYTE	BYTE value, by truncation
	any signed type	Converts any signed type to BYTE	BYTE value, by truncation
	REAL	Converts any REAL type to BYTE	BYTE (CHARINT (real))
	SELECTOR	Converts SELECTOR to BYTE	BYTE value, by truncation
	POINTER	Converts offset portion of POINTER to BYTE	BYTE (OFFSET\$OF (pointer))

continued

Table 9-1. Value and Type Conversions for PL/M-386 (continued)

Procedure Name	Parameter Type	Function	Result Returned
HWORD	any unsigned	Converts any unsigned type to HWORD	HWORD value, by truncation or zero extension
	any signed type	Converts any signed type to HWORD	HWORD value, by truncation or sign extension
	REAL	Converts any real type to HWORD	HWORD (SHORTINT (real))
	SELECTOR	Converts SELECTOR to HWORD	HWORD type, value unchanged
	POINTER	Converts offset portion of POINTER to HWORD	HWORD (OFFSET\$OF (pointer))
WORD	any unsigned type	Converts any unsigned type to WORD	WORD value, by truncation or zero extension
	any signed type	Converts any signed type to WORD	WORD value, by sign extension
	REAL	Converts any real type to WORD	WORD (INTEGER (real))
	SELECTOR	Converts SELECTOR to WORD	WORD value, by zero extension
	POINTER	Converts offset portion of POINTER to WORD	WORD (OFFSET\$OF (pointer))

continued

Table 9-1. Value and Type Conversions for PL/M-386 (continued)

Procedure Name	Parameter Type	Function	Result Returned
DWORD	any unsigned type	Converts any unsigned type to DWORD	DWORD value, by zero extension
	any signed type	Converts any signed type to DWORD	DWORD value, by sign extension
	REAL	Converts any real type to DWORD	DWORD (INTEGER (real))
	SELECTOR	Converts SELECTOR to DWORD	DWORD value, by zero extension
	POINTER	Converts offset portion of POINTER to DWORD	DWORD (OFFSET\$OF (pointer))
CHARINT	any unsigned	Converts any unsigned type to CHARINT	CHARINT value, by truncation
	any signed type	Converts any signed type to CHARINT	CHARINT value, by sign-extension
	REAL	Converts any real type to CHARINT	CHARINT (FIX(real))
	SELECTOR	Converts SELECTOR to CHARINT	CHARINT value, by truncation
	POINTER	Converts offset portion of POINTER to CHARINT	CHARINT (OFFSET\$OF (pointer))

continued

Table 9-1. Value and Type Conversions for PL/M-386 (continued)

Procedure Name	Parameter Type	Function	Result Returned
SHORTINT	any unsigned type	Converts any unsigned type to SHORTINT	SHORTINT value, by zero extension or truncation
	any signed type	Converts any signed type to SHORTINT	SHORTINT value, by sign extension
	REAL	Converts any real type to SHORTINT	SHORTINT (FIX (real))
	SELECTOR	Converts SELECTOR to SHORTINT	SHORTINT value
	POINTER	Converts offset portion of POINTER to SHORTINT	SHORTINT (OFFSET\$OF (pointer))
INTEGER	any unsigned type	Converts any unsigned type to INTEGER	INTEGER value, by zero extension or truncation
	any signed type	Converts any signed type to INTEGER	INTEGER value, by sign extension
	REAL	Converts any real type to INTEGER	INTEGER (FIX (real))
	SELECTOR	Converts SELECTOR to INTEGER	INTEGER value, by zero extension
	POINTER	Converts offset portion of POINTER to INTEGER	INTEGER (OFFSET\$OF (pointer))
REAL	any unsigned type (except OFFSET)	Converts any unsigned type to REAL	REAL (SIGNED (unsigned))
	any signed type	Converts any signed type to REAL	FLOAT (signed)
	REAL		value unchanged

continued

Table 9-1. Value and Type Conversions for PL/M-386 (continued)

Procedure Name	Parameter Type	Function	Result Returned
SELECTOR	any unsigned binary type	Converts any unsigned binary type to SELECTOR	SELECTOR value, by zero extension or truncation
	OFFSET		Current data segment selector
	any unsigned integer data type	Converts any signed integer data type to SELECTOR	SELECTOR value by sign extension or truncation
	POINTER		Selector portion of the POINTER
	REAL		Cannot be used
OFFSET	any unsigned type	Converts any unsigned type to OFFSET	OFFSET, by zero extension or truncation
	any signed type	Converts any signed type to OFFSET	OFFSET, by sign extension
	SELECTOR		zero (0)
	POINTER		OFFSET\$OF (pointer)
POINTER	any unsigned type	Converts value of any unsigned type to POINTER	BUILD\$PTR (DS, OFFSET (unsigned)) (DS is selector of current data segment)
	any signed type	Converts value of any signed type to POINTER	BUILD\$PTR (DS, OFFSET (signed)) (DS is selector of current data segment)
	SELECTOR		BUILD\$PTR (SELECTOR, 0)
	OFFSET		BUILD\$PTR (DS, OFFSET) (DS is selector of current data segment)

Notes:

Conversions from REAL to OFFSET, or POINTER, and vice versa, are not allowed. Under WORD32 (the default), LONGINT is equivalent to INTEGER. ADDRESS is equivalent to OFFSET.

The PL/M-386 LOW, HIGH, and DOUBLE Functions

The PL/M-386 `LOW` built-in function converts `DWORD` values to `WORD` values, `WORD` or `OFFSET` values to `WORD` values, and `WORD` values to `BYTE` values. `LOW` is activated using the following form:

```
LOW (expression)
```

Where:

expression has an unsigned binary number type.

If *expression* has a `DWORD` value, `LOW` returns the value of the low-order (least significant) `WORD` of the *expression* value. If *expression* has a `WORD` or `OFFSET` value, `LOW` returns the value of the low-order (least significant) `WORD` of the *expression* value. If *expression* has an `WORD` value, `LOW` returns the value of the low-order (least significant) `BYTE` of the *expression* value. If *expression* has a `BYTE` value, `LOW` returns this value unchanged.

The PL/M-386 `HIGH` built-in function converts `DWORD` values to `WORD` values, `WORD` or `OFFSET` values to `WORD` values, and `WORD` values to a `BYTE` values. `HIGH` is activated using the following form:

```
HIGH (expression)
```

Where:

expression has an unsigned binary number type.

If *expression* has a `DWORD` value, `HIGH` returns the value of the high-order (most significant) `WORD` of the *expression* value. If *expression* has a `WORD` or `OFFSET` value, `HIGH` returns the value of the high-order (most significant) `WORD` of the *expression* value. If *expression* has an `WORD` value, `HIGH` returns the value of the high-order (most significant) `BYTE` of the *expression* value. If *expression* has a `BYTE` value, then `HIGH` will return a zero.

The PL/M-386 `DOUBLE` built-in function converts `BYTE` values to `WORD` values, `WORD` values to `DWORD` values, and `WORD` or `OFFSET` values to `DWORD` values. `DOUBLE` is activated using the following form:

```
DOUBLE (expression)
```

Where:

expression has an unsigned binary number type.

If *expression* has a `BYTE` value, the `DOUBLE` function appends 8 high-order zero bits to convert the *expression* to an `WORD` value and returns this `WORD` value. If *expression* has an `WORD` value, the `DOUBLE` function appends 16 high-order zero bits to convert the *expression* to a `DWORD` value and returns this `DWORD` value. If *expression* has a `WORD` or `OFFSET` value, the `DOUBLE` function appends 32 high-order bits to convert it to a `DWORD` value and returns this `DWORD` value. If *expression* has a `DWORD` value, the `DOUBLE` function returns this `DWORD` value unchanged.

The `FLOAT` Function

`FLOAT` is a built-in `REAL` function that converts a signed integer to the real number data type. It is activated by a function reference with the following form:

```
FLOAT (expression)
```

Where:

expression is a signed integer.

`FLOAT` converts the signed integer to the corresponding real number data type and returns the real number. `FLOAT` can be replaced with `REAL (expression)`.

The `FIX` Function

`FIX` is a built-in `INTEGER` function that converts a `REAL` value to an `INTEGER` value. It is activated by a function reference with the following form:

```
FIX (expression)
```

Where:

expression has a `REAL` value.

`FIX` rounds the `REAL` value to the nearest `INTEGER`. If both `INTEGER` values are equally near, `FIX` rounds to the even value. The resulting `INTEGER` value is then returned.

For example:

```
FIX(1.4)      /* would result in the INTEGER value 1, */
FIX(-1.8)    /* in -2, */
FIX(3.5)     /* in 4, and */
FIX(6.5)     /* in 6. */
```

If the result calculated by `FIX` is not within the implemented range of `INTEGER` values, the result is undefined.

⇒ **Note**

`FIX` is affected by the rounding mode; see Chapter 10. The default mode (round to the nearest or even value) is used in the previous examples.

`FIX` can be replaced with `INTEGER (expression)`.

The INT Function

`INT` is a built-in `INTEGER` function that converts an unsigned binary value, excluding `DWORD` values, to the signed integer data type. It is activated by a function reference with the following form:

```
INT (expression)
```

Where:

expression has an unsigned binary data type, excluding `DWORD`.

`INT` interprets the *expression* value as a positive number and returns the corresponding `INTEGER` value.

If the result calculated by `INT` is not within the implemented range of `INTEGER` values, the result is undefined (see Chapter 5 for ranges for `INTEGER` values).

The SIGNED Function

For `PL/M-386`, `SIGNED` is a built-in `INTEGER` function that converts a `BYTE`, `WORD`, or `DWORD` value to an `INTEGER` value. `SIGNED` is activated by a function reference with the following form:

```
SIGNED (expression)
```

Where:

expression has an unsigned binary number data type, excluding `DWORD`.

If *expression* has a BYTE or HWORD value, it will be extended by 24 or 16 high-order 0 bits, respectively, to produce a WORD value.

SIGNED interprets the WORD value as a 32-bit two's-complement number and returns the corresponding integer value.

If the highest-order (most significant) bit of the WORD value is a 0, SIGNED interprets the WORD value as a positive number and returns the corresponding INTEGER value. For example:

```
SIGNED (0000$0000$0000$0100B)
```

returns an INTEGER value of 4.

If the highest-order bit of the WORD value is a 1, SIGNED returns a negative INTEGER value whose absolute magnitude is the two's complement of the WORD value. For example:

```
SIGNED(1111$1111$1111$1100B)
```

returns an INTEGER value of -4.

SIGNED can be replaced by INTEGER (*expression*).

The UNSIGN Function

The UNSIGN built-in function converts a signed integer to a WORD value. It is activated by a function reference with the following form:

```
UNSIGN (expression)
```

Where:

expression is a signed integer.

UNSIGN converts the INTEGER value to a WORD value.

If the INTEGER value is positive, the WORD value will be numerically the same as the INTEGER value. However, if the INTEGER value is negative, the WORD value will be the two's complement of the absolute magnitude of the INTEGER value. For example:

```
UNSIGN(-4)
```

returns a WORD value of:

```
1111$1111$1111$1100B
```

UNSIGN can be replaced by WORD (*expression*).

The Unsigned Binary Data Type Built-in Functions

The unsigned binary data type built-in functions convert any expression to the specified unsigned binary data type. For example, the `WORD` and `DWORD` built-in functions convert any expression to a `WORD` or `DWORD` value, respectively.

The built-in functions are activated with the form:

```
built-in (expression)
```

Where:

built-in is the name of the data type to which the given expression is converted (e.g., `BYTE` or `WORD`).

expression has any value.

For example, `WORD (INT1)` converts the value of `INT1` to a `WORD` value.

If *expression* is an unsigned binary number, it is converted by truncation or zero extension, if necessary. If *expression* is a signed integer, it is converted by truncation or sign extension, if necessary. If *expression* is a selector, it is converted by truncation or zero extension. If *expression* is a pointer, the offset portion of the pointer is converted by truncation or zero extension; the selector portion of the pointer is discarded. If *expression* is a real number, it is first converted to a signed integer using the numeric coprocessor's real to integer conversion, then the resulting value is converted to the unsigned binary number data type by truncation, if necessary.

Signed Integer Data Type Built-in Function

The signed integer data type built-in function converts any expression to a signed integer data type. It has the form:

```
INTEGER (expression)
```

For example:

```
INTEGER (D)
```

converts the value of `D` to an `INTEGER` value within the `INTEGER` range.

If *expression* is an unsigned binary number or selector, it is converted by truncation or zero extension. If *expression* is a pointer, the offset portion of the pointer is converted by truncation or zero extension; the selector portion of the pointer is discarded. If *expression* is a real number, it is converted using the numeric coprocessor's real to integer conversion.

Specific to PL/M-386, if *expression* is a signed type, it is converted by sign extension. Shorter data types are converted into longer data types by sign

extending the shorter data type value. Longer data types are converted into shorter data types by sign extension of the bits equivalent to the shorter data type. For example, if a CHARINT built-in is used to convert an INTEGER value, the least significant 8 bits are sign extended and the value returned is guaranteed to be in the CHARINT range.

REAL Built-in Functions

The REAL built-in function converts an expression to a REAL value. Expressions of type SELECTOR, OFFSET, and POINTER cannot be converted. The conversion is done using the numeric coprocessor's INTEGER to REAL conversion. If the expression is an unsigned binary number it is zero extended, if necessary, and interpreted as a signed value.

The SELECTOR Built-in Function

The SELECTOR built-in function converts any expression (except the real number data type) to a SELECTOR value. If the expression is any unsigned binary number, except OFFSET, it is truncated or zero extended to 16 bits. If the expression is a signed integer, it is truncated or sign extended to 16 bits. If the expression is of type POINTER, the selector portion of the pointer is returned. If the expression is of type OFFSET, the current data segment selector is returned. Expressions of type REAL cannot be converted.

The POINTER Built-in Function

The POINTER built-in function converts any expression (except real numbers) to a POINTER value. If the expression is any unsigned binary number or signed integer, it is converted to type OFFSET by truncation, zero, or sign extension, if necessary. This OFFSET value is combined with the SELECTOR value of the current data segment to create the POINTER value. If the expression is of type SELECTOR, it is combined with an OFFSET value of zero to create the POINTER value. Expressions of type REAL cannot be converted.

The OFFSET Built-in Function

The `OFFSET` built-in function converts any expression (except real numbers) to an `OFFSET` value. If the expression is any unsigned binary number or signed integer data type, it is converted to type `OFFSET` by truncation, or by zero or sign extension. If the expression is of type `SELECTOR`, an `OFFSET` value of zero is returned. If the expression is of type `POINTER`, the offset portion of the pointer is returned. `ADDRESS` values are equivalent to `OFFSET`. Expressions of type `REAL` cannot be converted.

The ABS and IABS Functions

The `ABS` built-in function returns the absolute value of a real number. It is activated by a function reference with the following form:

```
ABS (expression)
```

Where:

expression is a real number.

If the value of *expression* is positive, `ABS` returns it unchanged. If the value of *expression* is negative, `ABS` returns $-(expression)$.

The `IABS` built-in function returns the absolute value of a signed integer. It is activated by a function reference with the following form:

```
IABS (expression)
```

Where:

expression is a signed integer.

If the value of *expression* is positive, `IABS` returns it unchanged. If the value of *expression* is negative, `IABS` returns $-(expression)$.

Shift and Rotate Functions

With the shift and rotate functions, bit patterns can be moved to the right and to the left. In a shift, bits moved off one end of the pattern are lost, and zero bits move into the pattern from the other end (except in the case of the algebraic shift right function, *SAR*). In a rotate, bits moved off one end of the pattern are moved onto the other end of the pattern. It is not possible to perform a rotate on a signed integer algebraic pattern.

In PL/M-386, a value is handled as a pattern of 8 bits for a *BYTE* or *CHARINT* value, 16 bits for a *HWORD* or *SHORTINT* value, 32 bits for *WORD*, *OFFSET*, or *INTEGER* values, or 64 bits for a *DWORD* value. The pattern is moved to the right or left by a specified number of bits called the bit count.

Rotation Functions

The type of the rotate left (*ROL*) and rotate right (*ROR*) built-in functions depends on the type of expression given as an actual parameter. These built-ins are activated by function references with the following forms:

```
ROL (pattern, count)
ROR (pattern, count)
```

Where:

pattern and *count* are unsigned binary numbers.

If *count* is any unsigned binary number data type except *BYTE*, all but the low-order bits will be dropped to produce a *BYTE* value. If the value of *count* is 0, no rotation occurs.

The value of *pattern* is handled as an 8-bit, 16-bit, 32-bit, or 64-bit quantity. The type of *pattern* determines which of the unsigned binary number data types is used. This, in turn, determines the value of *pattern*. The number of bit positions by which *pattern* is rotated is specified by *count*.

The following are examples of the action of these procedures:

```
ROR (10011101B, 1)      returns a value of 11001110B
ROL (10011101B, 2)      returns a value of 01110110B
ROR (1101011010011010B, 9) returns a value of 0100110101101011B
```

Logical-shift Functions

The type of the logical-shift left (SHL) and logical-shift right (SHR) built-in functions depends on the type of the expression given as an actual parameter. SHL and SHR are activated by function references with the forms:

```
SHL (pattern, count)  
SHR (pattern, count)
```

Where:

pattern and *count* are expressions using an unsigned binary number data type.

If *count* is any unsigned binary number data type except BYTE, all but the 8 low-order bits will be dropped to produce a BYTE value. If the value of *count* is 0, no shift occurs.

The value of *pattern* can be a BYTE, HWORD, WORD, or DWORD value and the value will not be converted. If *pattern* is a BYTE value, the function will return a BYTE value. If *pattern* is an HWORD value, the function will return an HWORD value. If *pattern* is a WORD value, the function will return a WORD value; if *pattern* is a DWORD value, the function will return a DWORD value.

The value of *pattern* is shifted left (by SHL) or right (by SHR), with the bit count given by *count*.

A shift operation can force one bit out of the pattern. For example:

```
SHL(1000$0001B,1)
```

returns 0000\$0010B, losing the former high-order bit, and:

```
SHR(1000$0001B,1)
```

becomes 0100\$0000B, losing the former low-order bit.

If the specified *pattern* and *count* do not lose information, a shift of one bit position has the effect of multiplication by two for a left shift, or division by two for a right shift. For example, suppose that VAR is a BYTE variable with a value of eight. This is represented as 0000\$1000B. SHL(VAR,1) would return 0001\$0000B, which represents 16, and SHR(VAR,1) would return 0000\$0100B, which represents four.

Casting can be used to ensure that no information is lost in a shift, as in the following example:

```
SHL(WORD(LIT$MASK),3)
```

Algebraic-shift Functions

The type of the algebraic-shift left (SAL) and algebraic-shift right (SAR) built-in functions depends on the type of the expression given as an actual parameter. SAL and SAR are activated by function references with the following forms:

```
SAL (pattern, count)  
SAR (pattern, count)
```

Where:

pattern is an expression using a signed integer data type.

count is an expression using an unsigned binary data type.

If *count* is any unsigned binary data type except BYTE, all but the 8 low-order bits will be dropped to produce a BYTE value. If the value of *count* is zero, no shift occurs.

For PL/M-386, the type of *pattern* can be a CHARINT, SHORTINT, or INTEGER value. All values are converted to INTEGER before the shift operations, and an INTEGER value is returned.

In a left shift (SAL), zero-bits move into the pattern from the right (as in SHL and SHR).

In a right shift (SAR), either zero-bits or one-bits move into the pattern from the left. If the original value of *pattern* is positive, the sign bit (leftmost bit) is a 0, and zero-bits move in from the left. If the original value is negative, the sign bit is a 1, and one-bits move in from the left.

In some instances (as in logical shifts), an algebraic shift of one bit position can have the effect of multiplication by two for a left shift or division by two for a right shift. For example, suppose that VAL is an INTEGER variable with a value of -8. This value is 1111\$1111\$1111\$1000B. SAL(VAL, 1) would return 1111\$1111\$1111\$0000B, which is -16, and SAR(VAL, 1) would return 1111\$1111\$1111\$1100B, which is -4.

Concatenate Functions

The concatenate functions (SHLD and SHRD) are built-in WORD double-shift functions that concatenate two WORD values to form a 64-bit string, shift the concatenated pattern left (SHLD) or right (SHRD) by *count* bits, and return the destination WORD. These built-ins are activated by function references with the following form:

keyword (high pattern, low pattern, count)

Where:

keyword is SHLD or SHRD.

high pattern
is a WORD value.

low pattern
is a WORD value.

count is a BYTE, HWORD, or WORD value that determines how many bits to shift the concatenated pattern.

SHLD concatenates the bit pattern of the WORD value *high pattern* with the bit pattern of the WORD value *low pattern* to form a 64-bit string. *high pattern* is placed in the high 32 bits and *low pattern* is placed in the low 32 bits. The concatenated pattern is shifted left by the number of bits given by *count* MODULO 32. These operands are taken MODULO 32 to provide a number between 0 and 31 by which to shift. This has the effect of shifting the high order bits of *low pattern* into the low order bits of *high pattern*. SHLD returns the high 32 bits of the shifted pattern.

SHRD concatenates the bit pattern of the WORD value *high pattern* with the bit pattern of the WORD value *low pattern* to form a 64-bit string. *high pattern* is placed in the high 32 bits and *low pattern* is placed in the low 32 bits. The concatenated pattern is shifted right by the number of bits given by *count* MODULO 32. These operands are taken MODULO 32 to provide a number between 0 and 31 by which to shift. This has the effect of shifting the low order bits of *high pattern* into the high order bits of *low pattern*. SHRD returns the low 32 bits of the shifted pattern.

String Manipulation Procedures and Functions

The term string is used here in a broader sense than previously, in which string was used to refer to a `BYTE` string. In this section, a string is any contiguously stored set of unsigned binary number data type values (excluding `DWORD` and `OFFSET`). A string can be regarded as if it were an unsigned binary number type (excluding `DWORD` and `OFFSET`) array, and the array items can be referred to as elements.

The word `index` refers to the position of a given element within a string. The index is similar to the subscript of an array reference. Thus, the index of the first element of a string is 0, the index of the second element is 1, and so on.

In the following descriptions, the location of a string always means the location of its first element. In each string manipulation procedure, the location of a string is specified by a parameter called *source* or *destination*, which is an expression with a `POINTER` value. The *source* points to the lowest element. For example, with `MOVB` and `MOVW`, the lowest element (element 0) is the first element to be processed. With `MOVRB` and `MOVRW`, the lowest element is the last element to be processed, as discussed in the following sections.

The length of a string is the number of elements it contains. In each string manipulation procedure, the number of elements to be processed is specified by a parameter called *count*.



Note

If the *source* or *destination* string address is in `SELECTOR` or `WORD` form, use the `@` operator of a variable based on the address. Otherwise, the built-in function `BUILD$PTR` can be used to construct the pointer-parameter for the string built-in.

In PL/M-386, each of the string-manipulation procedures described in the following sections (except `XLAT`) is available for `BYTE`, `HWORD`, and `WORD` strings.

The Copy String in Ascending Order Procedure

`MOVxx` is an untyped procedure that copies a string of length *count* from one location to another. It is activated by a `CALL` statement with the following form:

```
CALL keyword (source, destination, count);
```

Where:

keyword `MOVB`, `MOVHW`, `MOVW`

source and *destination*
expressions with `POINTER` values

count expression with `BYTE`, `HWORD`, `OFFSET`, or `WORD` value

`MOVB` copies a `BYTE` string, `MOVHW` copies an `HWORD` string, and `MOVW` copies a `WORD` string.

The string elements are copied in ascending order (i.e., element 0 is copied first, then element 1, etc.). This order is significant if the *source* string and the *destination* string overlap. If the value of *destination* is higher than the value of *source*, and the two strings overlap, elements in the overlap area will be overwritten before they are copied. To avoid the overwriting, use `MOVRxx` instead of `MOVxx`.

The Copy String in Descending Order Procedure

`MOVRxx` is an untyped procedure that copies a string of length *count* from one location to another. It is activated by a call statement with the following form:

```
CALL keyword (source, destination, count);
```

Where:

keyword `MOVRB`, `MOVRHW`, `MOVRW`

source and *destination*
expressions with `POINTER` values

count expression with `BYTE`, `HWORD`, `OFFSET`, or `WORD` value

The `MOVVB` built-in procedure is similar to the `MOVVB` procedure except that the elements in the `MOVVB source` string are copied to the `destination` string in descending order (i.e., element (count-1) is copied first, then element (count-2), and so on, with element 0 copied last). This order is significant when the two strings overlap. If the value of `destination` is higher than the value of `source`, and an overlap exists, elements in the overlap area will not be overwritten until they have been copied. However, if the value of `source` is higher than the value of `destination`, elements in the overlap area will be overwritten before they are copied.

`MOVHW` performs the same function as `MOVVB` except that `MOVHW` copies an `HWORD` string.

`MOVWR` performs the same function as `MOVVB`, except `MOVWR` copies a `WORD` string instead of a `BYTE` string.

⇒ **Note**

If two strings overlap, use a procedure such as the following to make the correct choice between `MOVVB` and `MOVVB`. This ensures that elements in the overlap area will not be overwritten until after they have been copied.

```
MOVBYTES: PROCEDURE (SRC, DST, CNT);
    DECLARE (SRC, DST) POINTER, CNT HWORD;
    IF (OFFSET(SRC)) > (OFFSET(DST)) THEN
        CALL MOVVB (SRC, DST, CNT);
    ELSE CALL MOVVB (SRC, DST, CNT);
END MOVBYTES;
```

This procedure can be activated without the need to consider whether overlap may occur or whether `source` or `destination` is higher.

The Compare String Function

`CMPxx` is a built-in `WORD` function that compares two strings of length `count`. It is activated by a function reference with the following form:

```
keyword (source1, source2, count)
```

Where:

`keyword` `CMPB`, `CMPHW`, `CMPW`

`source1 source2`

expressions with `POINTER` values

`count` expression with `BYTE`, `HWORD`, `OFFSET`, or `WORD` value

CMPB compares two BYTE strings of length *count* whose locations are *source1* and *source2*. It remains a 32-bit instruction even if the WORD16 control is in effect.

If every element in the string at *source1* is equal to the corresponding element in the string at *source2*, CMPxx returns a WORD value, 0FFFFFFFH, for PL/M-386. Otherwise, CMPxx returns the index (position within the strings) of the first pair of elements found to be unequal.

CMPHW performs the same function as CMPB, except that CMPHW compares two HWORD strings. CMPW performs the same as function as CMPB except that CMPW compares two WORD strings instead of two BYTE strings.

The Find Element Functions

FIND is a built-in WORD function that searches a string to find an element that has a specified value. It is activated by a function reference of the following form:

keyword (*source*, *target*, *count*)

Where:

keyword FINDB, FINDHW, FINDW, FINDRB, FINDRHW, FINDRW

source expression with POINTER value

target expression with BYTE, HWORD, or WORD value

count expression with BYTE, HWORD, OFFSET, or WORD value

FINDB examines each element of the source string (in ascending order) until it finds an element whose value matches the BYTE value of *target*, or until *count* elements have been searched, with none of them having matched the *target*. If the search is successful, FINDB returns the index of the first element of the string that matches *target*. If the search is unsuccessful, FINDB returns a WORD value.

FINDHW performs the same function as FINDB, except that FINDHW searches an HWORD string. If *target* has a BYTE value, it is extended by 8 high-order, 0-bits to produce an HWORD value. If *target* has a WORD value, it is truncated by 16 high-order bits to produce an HWORD value.

FINDW performs the same function as FINDB, except that FINDW searches a WORD string. If *target* has a BYTE or HWORD value, *target* is extended appropriately to produce a WORD value.

FINDRB performs the same function as FINDB, except that FINDRB searches the *source* string in descending order. Thus, if each search is successful, FINDRB returns the index of the last (highest subscript) element that matches the BYTE value of *target*. FINDRHW performs the same function as FINDRB, except that FINDRHW searches an HWORD string (in descending order). FINDRW searches a WORD string (in descending order).

The Find String Mismatch Function

SKIPB is a built-in WORD function that searches the BYTE string of length *count* at a specified location (given by *source*) for the first BYTE value that does not match the target BYTE. This search begins with the first BYTE value of the string. The result is a WORD value, either 0FFFFFFFH if the string contains only BYTE values equal to the target BYTE, equal to the index of the first BYTE value not equal to the target BYTE.

The function is activated by a function reference of the following form:

keyword (*source*, *target*, *count*)

Where:

keyword SKIPB, SKIPHW, SKIPW, SKIPRB, SKIPRHW, SKIPRW

source expression with POINTER value

target expression with BYTE, WORD, or HWORD value

count expression with BYTE, HWORD, OFFSET, or WORD value

SKIPW performs the same function as SKIPB, except that SKIPW searches a WORD source string to find the first element that does not match the WORD value of *target*. SKIPHW performs the same function as SKIPB, except that SKIPHW searches an HWORD source string to find the first element that does not match the HWORD value of *target*.

SKIPRB searches a BYTE string of the length specified by *count*, at the location given by *source*, for the last BYTE value that does not match the target BYTE. This search begins with the last BYTE value in the string. The result is a WORD value (0FFFFFFFH) if the string contains only BYTE values equal to the target BYTE, or the index of the last BYTE value, if the last BYTE value in the string is not equal to the target BYTE.

SKIPRW performs the same function as SKIPRB, except that SKIPRW searches for the last element in the WORD source string that does not match the WORD value of the target. SKIPRHW searches for the last element in the HWORD source string that does not match the HWORD value of the target.

The Translate String Procedure

XLAT is an untyped procedure that uses a translation table to translate a BYTE string to produce another BYTE string. It is activated by a CALL statement of the form:

```
CALL XLAT (source, destination, count, table)
```

Where:

source, *destination*, *table*
expressions with POINTER values

count expression with BYTE, HWORD, OFFSET, or WORD value

XLAT translates the *count* BYTE elements in the *source* string, placing the translated elements in the *destination* string. The value of *table* is assumed to be the location of a BYTE string of up to 256 elements. This string is used as a translation table.

The value of an element in the *source* string is used as an index into the translation table. The index selects one element from the translation table; this element is then copied into the *destination* string.

For example, if the fifth element in the *source* string is 202, then 202 is used as an index for the translation table. The 203rd element of the table is copied into the fifth position in the *destination* string.

The elements of the *source* string are translated into the *destination* string in ascending order.

The Set String to Value Procedure

The SET built-in is an untyped procedure that sets each element of a BYTE string, the length of which is specified by *count*, to a single specified value. SET is activated by a CALL statement with the following form:

```
CALL keyword (newvalue, destination, count)
```

Where:

keyword SETB, SETHW, SETW

newvalue expression with BYTE, HWORD, OFFSET, or WORD value -- the high-order bits are dropped to produce a BYTE, WORD, or HWORD value

destination
expression with POINTER value

count expression with BYTE, HWORD, OFFSET, or WORD value

SETB assigns the BYTE value of *newvalue* to each element of a BYTE string.

SETW performs the same function as SETB except that SETW assigns a single WORD value to each element of a WORD string. If *newvalue* has a BYTE or an HWORD value, it will be extended by 24 or 8 high-order 0 bits, respectively, to produce a WORD value.

For information on WORD32 | WORD16 mapping, see Tables 9-3, 10-1, and 11-3.

PL/M-386 Bit Manipulation Built-ins

The Copy Bit String Procedure

MOVBIT is an untyped built-in procedure that copies a bit string of length *count* from one location to another. MOVBIT is activated by a CALL statement with the following form:

```
CALL MOVBIT (sbase, sbitoffset, dbase, dbitoffset, count);
```

Where:

sbase and *dbase*

are expressions with POINTER values.

sbitoffset

are expressions with BYTE, HWORD, OFFSET,

dbitoffset and *count*

are WORD values.

The MOVBIT built-in procedure moves the number of bits specified by *count* from the bit location given by the base address *sbase*, and the bit offset *sbitoffset*, to the location given by the base address *dbase* and the bit offset *dbitoffset*. These bits are moved beginning with the low-order bit (least significant bit).

The MOVRBIT built-in procedure performs the same function as MOVBIT, except that MOVRBIT moves bits in descending order, beginning with the high-order bit (most significant bit).

The Find Set Bit Function

SCANBIT is a built-in WORD function that searches a bit string to find a set bit (i.e., a bit with the value of 1). SCANBIT is activated by a function reference with the following form:

```
SCANBIT (sbase, sbitoffset, count)
```

Where:

sbase is an expression with a POINTER value.

sbitoffset and *count*

are expressions with BYTE, HWORD, OFFSET, or WORD values.

The `SCANBIT` built-in function searches the bit string of length *count* at the bit location given by the base address *sbase* and the bit offset *sbitoffset* for the first set bit, beginning with the low-order bit (least significant bit) in the string. The result of `SCANBIT` is either a `WORD` value of `0FFFFFFFFH` if the string contains all 0 bits, or the index of the first set bit.

`SCANRBIT` performs the same function as `SCANBIT`, except that `SCANRBIT` starts at the high-order bit (most significant bit) in the string and searches for a set bit, in descending order, and returns the location of the first set bit it encounters. The result of `SCANRBIT` is either a `WORD` value of `0FFFFFFFFH` if the string contains all 0-bits, or the index of the first set bit encountered.

Miscellaneous Built-ins

The Move Bytes Procedure

MOVE is an untyped procedure that moves the number of bytes specified by *count* to the location given by the value of *destination*, starting at the location given by the value of *source*. If the *source* and *destination* fields overlap, the result is undefined. MOVE is provided for compatibility with PL/M-80 programs. MOVE is activated by a CALL statement with the following form:

```
CALL MOVE (count, source, destination)
```

Where:

count expression with BYTE, HWORD, OFFSET, or WORD value

source and *destination*
 expressions with OFFSET values

If either *source* or *destination* is a value other than the value OFFSET, the value will be extended by high-order 0 bits to produce the OFFSET value. The values of *source* and *destination* are assumed to be the addresses of the *source* string and the *destination* string.

The operation of the MOVE procedure differs from the MOVEB procedure, as follows:

- The *source* and *destination* parameters must be OFFSET values or they will be converted. POINTER values cannot be used, nor can values be supplied with the @ operator. Thus, MOVE can handle only strings whose locations can be expressed as OFFSET addresses.
- The parameter order is different from the one used by the other built-in string functions.
- The results are always undefined if the *source* and *destination* strings overlap.

The Time Delay Procedure

TIME is an untyped built-in procedure that causes a time delay specified by its actual parameter. TIME is activated by a CALL statement with the following form:

```
CALL TIME (expression);
```

where the *expression* is converted, if necessary, to an HWORD quantity. The length of time measured by the procedure is a multiple of 100 microseconds. If the actual parameter evaluates to *n*, then the delay caused by the procedure is 100*n* microseconds. For example, the statement:

```
CALL TIME (45);
```

causes a delay of 4.5 milliseconds. For PL/M-386, the maximum delay is 12 hours. If required, longer delays can be obtained by repeated activations. The following block takes one second to execute:

```
DO I = 1 TO 40;  
    CALL TIME (250);  
END;
```

The TIME procedure is based on the microprocessor's CPU cycle times. The TIME procedure assumes 16 MHz for Intel386 and Intel486 microprocessors.

Note that in generating code for a call to TIME, the computer generates a loop rather than using interrupt processing. If a task containing a time delay is swapped out in a multi-tasking environment, the time delay of that task stops executing.

The Lock Set Function

LOCKSET is a built-in BYTE function that enables implementation of a simple software synchronization lock. It is called by a function reference with the following form:

```
LOCKSET (lockptr, newvalue)
```

Where:

lockptr expression with POINTER value

newvalue expression with BYTE, HWORD, or WORD value -- the high-order bits are dropped to produce a BYTE value

The action of LOCKSET is as follows: the *lockptr* parameter is used as a pointer to a BYTE variable; the value of *newvalue* is assigned to this variable, and LOCKSET returns the original value of the variable. During this transaction, the CPU prevents any other process from accessing the same memory location.

To see how this facility can be used, assume a system has more than one microprocessor using the same memory, and has a program in one of these microprocessors. This program uses memory locations that are also used by other microprocessors in the system.

Within certain critical regions of the program, it is critical that no other microprocessor can access the shared memory locations. To achieve this, declare a global `BYTE` variable called `LOCK`, and establish a convention that if `LOCK=0`, any microprocessor in the system can access the shared memory locations. However, if `LOCK=1`, no microprocessor can access the shared memory locations except for the microprocessor that set `LOCK` to 1.

Write the function reference `LOCKSET(@LOCK, 1)`. The value 1 will be assigned to `LOCK`. If the value returned by `LOCKSET` is 0, then `LOCK` has not been set, and this microprocessor is the one that set it. At the end of the critical region, the lock must be released by writing `LOCK=0`.

If `LOCKSET` returns a value of 1, then `LOCK` has been set and this microprocessor was not the one that set `LOCK`. Wait until a `LOCKSET(@LOCK, 1)` function reference returns a value of 0 before accessing the shared memory locations.

Thus, the program could contain the following construction:

```

                                                    /*Begin critical region*/
DO WHILE LOCKSET(@LOCK,1);
    /*Do nothing but repeat until LOCKSET returns 0*/
END;
    /*Now LOCK has been set to 1 by this microprocessor*/
. . .
    /*Critical region of program, where shared
    memory locations are accessed*/
. . .
LOCK=0;
                                                    /*End critical region*/
```

In the simple case just described, only one software lock is used. It is represented by the variable `LOCK`. If more than one set of memory locations need protection at different times, it is possible to establish as many different software locks as necessary, with each lock using a different `BYTE` variable.

Also, note that a software lock can be used for purposes other than protecting memory locations. `LOCKSET` provides a mechanism that can be used to implement various types of synchronization in a multiprocessor system.

The Lock Bit Functions

The `BITLOCK` functions are built-in `BYTE` functions similar to the `LOCKSET` built-in described in the previous section. They are called by a function reference with the form:

keyword (bbase, boffset)

Where:

keyword is `BITLOCKSET`, `BITLOCKRESET`, or `BITLOCKCOMPLEMENT`.

bbase is an expression with a `POINTER` value.

boffset is an expression with a `BYTE`, `HWOR`D, or `WORD` value.

The action of `BITLOCKSET` is as follows: the *bbase* and *boffset* parameters are used as the base address and bit offset to point to a certain bit in memory. The value 1 is assigned to this variable, and `BITLOCKSET` returns a `BYTE`. The returned value is `TRUE` (`OFFH`) if the original content of the bit was 1, otherwise it is `FALSE`. During this transaction, the CPU prevents any other process from accessing the same memory location. `BITLOCKRESET` performs the same function as `BITLOCKSET`, except that `BITLOCKRESET` assigns the value 0 to the bit variable. `BITLOCKCOMPLEMENT` performs the same function as `BITLOCKSET`, except that `BITLOCKCOMPLEMENT` complements the `BYTE` variable; that is, if the value was initially 0, it is set to 1 and vice versa.

POINTER and SELECTOR-related Functions

With the following built-in functions, programs can manipulate `POINTER` and `SELECTOR` values that serve as location addresses in the microprocessor's memory.

The Return `POINTER` Value Function

`BUILD$PTR` is a built-in `POINTER` function that takes the specified segment and offset value and returns a `POINTER` value. It is activated by a function reference with the following form:

```
BUILD$PTR (segment, offset)
```

Where:

segment expression with `SELECTOR` value

offset expression with `OFFSET` value

The Return Segment Portion of `POINTER` Function

`SELECTOR$OF` is a built-in `SELECTOR` function that returns the segment portion of a `POINTER`. It is activated by a function reference with the following form:

```
SELECTOR$OF (pointer)
```

Where:

pointer is an expression with a `POINTER` value.

The Return Offset Portion of `POINTER` Function

`OFFSET$OF` returns the offset portion of a `POINTER`. For PL/M-386, `OFFSET$OF` is a built-in `OFFSET` function. It is activated by a function reference with the following form:

```
OFFSET$OF (pointer)
```

Where:

pointer is an expression with a `POINTER` value.

The Set POINTER Bytes to Zero Variable

`NIL` is a built-in `POINTER` pseudo-variable that represents a pointer with all bytes set to zero. `NIL` is activated by a function reference with the following form:

```
NIL
```

The pointer value `NIL` points to no object. The value `NIL` can be assigned to a pointer to indicate, for instance, the end of a linked list.

Note that pointer values equal to `NIL` cannot be used to de-reference data values. For example, if a program contains the following statements:

```
DECLARE P POINTER;  
DECLARE B BASED P BYTE;  
P = NIL;
```

any subsequent references to `B` are invalid and will cause a trap.

The `NIL` `POINTER` variable also has the property that `@NIL` is equal to `NIL`.

`POINTER` variables can be initialized to `NIL` by using `@NIL` with `INITIAL`. For example:

```
DECLARE ENDOFLIST POINTER  
INITIAL (@NIL);
```

initializes `ENDOFLIST` with the value of `NIL` (i.e., all zeros). `OFFSET$OF(NIL)` and `.NIL` are also equal to zero.

WORD16 Built-in Mapping

The native machine word for Intel386 and Intel486 microprocessors is `WORD32` (a 32-bit `WORD`). The `WORD16` control affects the semantics of some data types and built-ins as listed in Table 3-3. In `PL/M-386`, `WORD16` keywords are mapped to the equivalent `WORD32` keyword. `SELECTOR`, `POINTER`, `OFFSET (ADDRESS)` are the same under both `WORD32` and `WORD16`. Table 11-5 in the discussion of the `WORD32 | WORD16` controls shows the correspondence between default (`WORD32`) built-ins and those available when `WORD16` is in effect. For example, Table 11-5 shows that `HWORD` under `WORD32` corresponds to `WORD` under `WORD16`.



Features Involving the Target CPU and Numeric Coprocessor **10**

The PL/M features described in this chapter make direct or indirect use of the target microprocessor and numeric coprocessor hardware.

Microprocessor Hardware-dependent Statements

The **ENABLE** and **DISABLE** Statements

These statements enable and disable the microprocessor interrupt mechanism.

The **ENABLE** statement has the following form:

```
ENABLE ;
```

ENABLE generates an **STI** instruction, causing the microprocessor to enable interrupts after the next machine instruction is executed.

The **DISABLE** statement has the following form:

```
DISABLE ;
```

DISABLE generates a **CLI** instruction, causing the microprocessor to disable interrupts.

The CAUSE\$INTERRUPT Statement

The CAUSE\$INTERRUPT statement causes a software interrupt to be generated. It has the form:

```
CAUSE$INTERRUPT (constant);
```

Where:

constant is a whole-number constant in the range 0 to 255.

CAUSE\$INTERRUPT generates an INT instruction with the constant as the interrupt type, causing the microprocessor to transfer control to the appropriate interrupt vector. Appendix G contains more information on run-time interrupt processing.

The HALT Statement

The HALT statement causes a microprocessor halt. It has the form:

```
HALT;
```

HALT generates an STI instruction followed by an HLT instruction, causing the microprocessor to halt with interrupts enabled.

Microprocessor Hardware Flags

Optimization and the Hardware Flags

To produce an efficient machine-code program from a PL/M source program, PL/M compilers perform extensive optimizations of the machine code. This means that the exact sequence of machine code produced to implement a given sequence of PL/M source statements cannot be predicted.

Consequently, the state of the microprocessor hardware flags cannot be predicted for any given point in the program. For example, suppose that a source program contains the following fragment:

```
...  
SUM = SUM + 250;  
...
```

Where:

SUM is a BYTE variable.

Now, if the value of SUM before this assignment statement is greater than five, the addition will cause an overflow and the hardware CARRY flag will be set.

If there were no optimization of the machine code, this assignment statement could be followed with one of the PL/M features described in the following sections. This would ensure that the feature would operate in a certain fashion depending on whether or not the addition caused the CARRY flag to be set. However, because of the optimization, some machine code instructions could occur immediately after the addition and change the CARRY flag. It cannot be safely predicted whether this will happen or not.



Note

Accordingly, any PL/M feature that is dependent on the CARRY flag (or any of the other hardware flags) can cause the program to run incorrectly. These features must therefore be used with caution, and any program that uses them must be checked carefully to make sure that it operates correctly.

The CARRY, SIGN, ZERO, and PARITY Functions

These built-in BYTE functions return the logical values of the microprocessor hardware flags. These functions take no parameters, and are activated by function references with the following forms:

```
CARRY
ZERO
SIGN
PARITY
```

An occurrence of one of these activations (in an expression) generates a test of the corresponding condition flag. If the flag is set (=1), a value of 0FFH is returned. If the flag is clear (=0), a value of 0 is returned.

The PLUS and MINUS Operators

In addition to the arithmetic operators described in Chapter 5, PL/M has two more: PLUS and MINUS.

PLUS and MINUS perform similarly to + and -, and have the same precedence. However, PLUS sums two numbers and adds the CARRY bit to the result and MINUS subtracts two numbers and subtracts the CARRY bit from the result.

Carry-rotation Functions

SCL and SCR are built-in rotation functions whose types depend on the type of the expression given as an actual parameter. They are activated by function references with the following forms:

```
keyword (pattern, count);
```

Where:

keyword SCL, SCR

pattern and *count*

expressions with BYTE, HWORD, WORD, OFFSET, or DWORD value
-- for count the high-order bits are dropped to produce BYTE values

If the value of *count* is 0, no shift occurs.

For PL/M-386, the value of *pattern* is handled as an 8-bit, 16-bit, 32-bit, or 64-bit binary quantity. This quantity is rotated to the left (by SCL) or to the right (by SCR). This is similar to the ROL and ROR functions described in Chapter 9. The type of *pattern* determines the type of rotate that is performed. The number of bit positions by which the value of *pattern* is rotated is specified by *count*.

The bit rotated off one end of pattern is rotated into the CARRY flag, and the old value of CARRY is rotated to the other end of pattern. In effect, SCL and SCR perform 9-bit rotations on 8-bit values, 17-bit rotations on 16-bit values, and so on.

For example, if the value of CARRY is 0, then:

SCL(11001010B, 2) returns a value of 00101001B and CARRY is set to 1

SCR(11001010B, 1) returns a value of 01100101B and CARRY remains 0

The Decimal Adjust Function

DEC is a built-in BYTE function that performs a decimal adjust operation on the actual parameter value and returns the result of this operation. For PL/M-386, DEC uses the value of the hardware AUXILIARY CARRY flag internally. It is activated by a function reference with the following form:

DEC (*expression*);

Where:

expression is converted, if necessary, to a BYTE value.

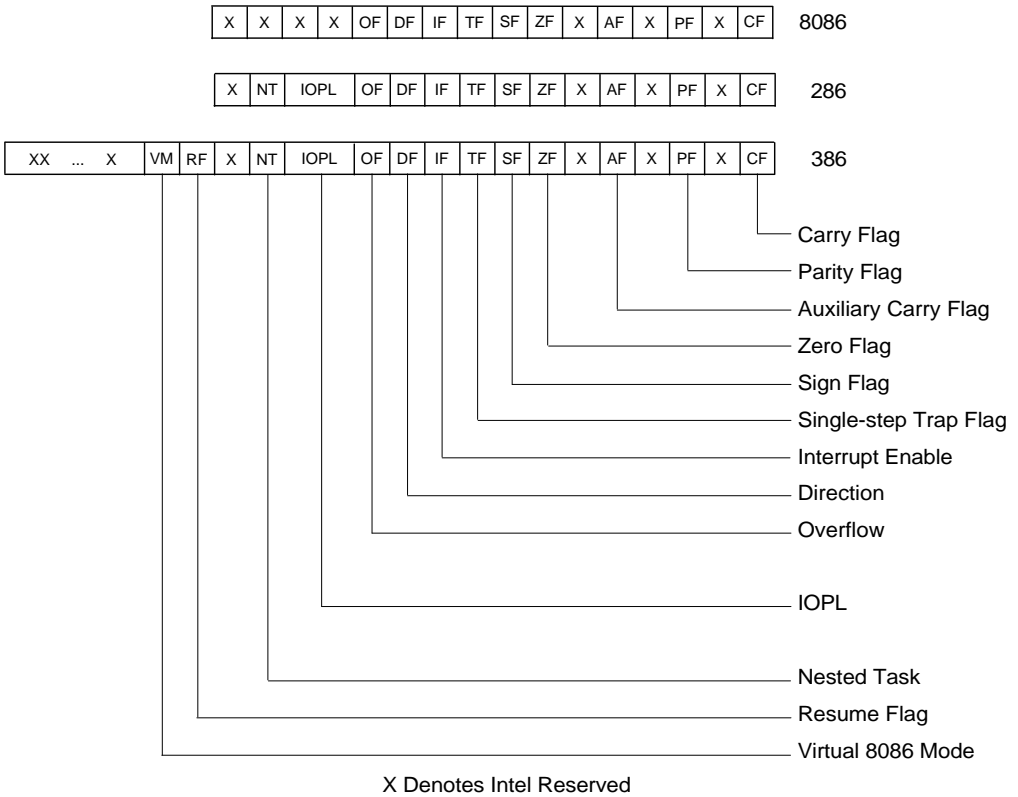
Microprocessor Hardware Registers

The Flags Register Access Variable

FLAGS is a built-in WORD variable that provides access to the microprocessor's hardware flags register (see Figure 10-1, which also has flags registers for the 8086 and 286 registers for comparison). The hardware flags register contains the hardware flags that are altered by the execution of various instructions. The hardware flags register for the Intel386 and Intel486 microprocessors are 32 bits long.

The FLAGS register is assigned to change the setting of the various flags. It can also be read to determine the current flag settings.

For more information on setting the hardware register flags, see the appropriate microprocessor programmer's reference manual.



OSD535

Figure 10-1. The Hardware Flags Register

The STACKPTR and STACKBASE Variables

For PL/M-386, `STACKPTR` is an `OFFSET` variable and `STACKBASE` is a `SELECTOR` variable. They provide access to the microprocessor's hardware stack pointer and stack base registers.

When setting these registers (that is, using `STACKPTR` or `STACKBASE` on the left side of an assignment), care must be exercised because this takes control of the stack away from the compiler. Thus, the compile-time checks on stack overflow and assumptions by the compiler about the run-time status of the stack may be invalid.

Microprocessor Hardware I/O

Input from an I/O port of a single `BYTE`, `HWORD`, `OFFSET`, or `WORD` is performed by the input built-ins as a function invocation in an expression on the right-hand side of an assignment statement. Single `BYTE`, `HWORD`, `OFFSET`, or `WORD` output is achieved by filling the appropriate element of the output array corresponding to the desired output port of the target microprocessor.

Multiple `BYTE`, `HWORD`, `OFFSET`, or `WORD` input is performed as a procedure invocation, reading in a string from the microprocessor's CPU port and storing it in a user-specified memory location. Multiple `BYTE`, `HWORD`, `OFFSET`, or `WORD` output is also performed as a procedure invocation, using a `CALL` statement to send a string from memory into the target microprocessor port.

The Find Value in Input Port Function

The following built-in functions return the values in the specified input port. They are activated by function references with the form:

```
keyword (expression);
```

Where:

keyword `INPUT`, `INHWORD` `INWORD`

expression expression with `BYTE`, `HWORD` or `WORD` value

The value of *expression* specifies one of the input ports of the target microprocessor.

The value returned by *keyword* is the *expression* quantity found in the specified input port.

PL/M-386 also has an `INDWORD` function when the `WORD16` control is used.

The Access Output Port Array

For PL/M-386, `OUTPUT`, `OUTHWORD`, and `OUTWORD` are built-in `BYTE`, `HWORD`, and `WORD` arrays, respectively. They are activated by a function reference with the following form:

```
keyword (expression);
```

Where:

keyword `OUTPUT`, `OUTHWORD`, `OUTWORD`

expression expression with `BYTE`, `HWORD`, or `WORD` value

These functions can access any port from 0 to 65,535, corresponding to the number of output ports on the target CPU. References to these arrays cause the specified *expression* quantity to be latched to the specified hardware output port.

A reference to *keyword* is legal only as the left part of an assignment statement or embedded assignment. For PL/M-386, the right-hand side of the assignment must have a BYTE, HWORD, or WORD value.

Specifying OUTPUT in the assignment statement places the BYTE value of the expression on the right side of the assignment into the specified output port. (Since OUTPUT is a BYTE built-in, the value of the expression is converted automatically to a BYTE type if necessary.)

Specifying OUTWORD in the assignment statement places the WORD (or OFFSET) value of the expression on the right side of the assignment into the specified output port.

Similarly, of OUTHWORD places the HWORD value of the expression on the right side of the assignment into the corresponding output port. PL/M-386 also has an OUTDWORD built-in when the WORD16 control is used.

The Read and Store String Procedure

The read and store string procedures are built in. For PL/M-386, these built-ins read the BYTE, HWORD, OFFSET, or WORD string values latched to the specified hardware input port. The read values, of the length specified by *count*, are then stored at the location specified by *destination*. These procedures are activated by a CALL statement with the following form:

```
CALL keyword (port, destination, count);
```

Where:

keyword BLOCKINPUT, BLOCKINHWORD, BLOCKINWORD

port expression with BYTE or HWORD value

destination
 expression with POINTER value

count expression with BYTE, HWORD, OFFSET or WORD value

The *keyword* specifies the type of string found in the specified input port. The value of *port* specifies one of the input ports of the CPU. The *destination* specifies the location (in memory) at which to store the string. The value of *count* specifies the length of the string.

PL/M-386 also has a BLOCKINDWORD procedure when the WORD16 control is used.

The Write String Procedure

The write string procedures are built-in procedures. For PL/M-386, these built-ins write a BYTE, HWORD, OFFSET, or WORD string to the specified output hardware port. These built-ins are activated by a CALL statement with the following form:

```
CALL keyword (port, source, count);
```

Where:

keyword BLOCKOUTPUT, BLOCKOUTHWORD, BLOCKOUTWORD

port expression with BYTE or HWORD value

source expression with POINTER value

count expression with BYTE, HWORD, OFFSET, or WORD value

The *keyword* specifies the type of string. The value of *port* specifies one of the output ports of the microprocessor CPU. The *source* value specifies the location (in memory) where the string is currently stored. The value of *count* specifies the string length.

PL/M-386 also has a BLOCKOUTDWORD procedure when the WORD16 control is used (see Chapter 10).

The Hardware Protection Model

The Intel386 and Intel486 microprocessors' protection mechanism provides up to four privilege levels within each task. The highest privilege level (level 0) is reserved for the operating system kernel. Below the kernel level, systems can be configured to include a system service level (level 1), an applications service level (level 2), and an application program level (level 3).


The following hardware protection built-in procedures and variables allow access to the protection architecture of these microprocessors.

The Task Register

The TASK\$REGISTER Variable

TASK\$REGISTER is a built-in SELECTOR variable that provides access to the task state register. This register points to a task state segment for the currently executing task.

The format of the task register for the Intel386 microprocessor is:



OSD579


Values are assigned to `TASK$REGISTER` to reset the task state segment for the current task or to enter the protected mode of the microprocessor. However, the selector stored in `TASK$REGISTER` must point to a valid task state segment. Note that values can be assigned to `TASK$REGISTER` only if the program is executed in protection mode at level 0.

`TASK$REGISTER` can also be read to determine the task state segment of the currently executing task.

The Global Descriptor Table Register

The global descriptor table register (GDTR) is a system-wide register used for protected virtual address mode. The GDTR describes a memory area that contains an array of descriptors for the global address space. The register occupies 6 bytes.

Its format for Intel386 and Intel486 microprocessors follows:



OSD580

- LIMIT size of the GDT segment (up to 64K bytes)
- BASE physical memory base address of the GDT segment
- ACCESS access control byte

The SAVE\$GLOBAL\$TABLE Procedure

SAVE\$GLOBAL\$TABLE is a built-in procedure. It is activated by a CALL statement with the form:

```
CALL SAVE$GLOBAL$TABLE (location);
```

Where:

location is an expression with a POINTER value.

SAVE\$GLOBAL\$TABLE saves the contents of the hardware global descriptor table register in the 6-byte save area pointed to by *location*.

The RESTORE\$GLOBAL\$TABLE Procedure

RESTORE\$GLOBAL\$TABLE is a built-in procedure. It is activated by a CALL statement with the form:

```
CALL RESTORE$GLOBAL$TABLE (location);
```

Where:

location is an expression with a POINTER value.


RESTORE\$GLOBAL\$TABLE restores the contents of the hardware global descriptor table register from the save area pointed to by *location*. This save area can be the same area used in a preceding call to SAVE\$GLOBAL\$STATUS.

SAVE\$GLOBAL\$TABLE saves the value of the GDTR in a 6-byte memory area. RESTORE\$GLOBAL\$TABLE restores the value of the GDTR.

The Interrupt Descriptor Table Register

The interrupt descriptor table register (IDTR) is a system-wide register that is used for interrupt processor management. The IDTR describes a segment that contains the linear base address and the size of the interrupt descriptor table (IDT), and a segment containing an array of gate descriptors for the interrupt handlers. The register occupies 6 bytes.

Its format for Intel386 and Intel486 microprocessors follows:



OSD580

- LIMIT size of the segment (up to 64K bytes)
- BASE physical memory base address of the IDT segment
- ACCESS access control byte

The SAVE\$INTERRUPT\$TABLE Procedure

SAVE\$INTERRUPT\$TABLE is a built-in procedure that is activated by a CALL statement with the following form:

```
CALL SAVE$INTERRUPT$TABLE (location);
```

Where:

location is an expression with a POINTER value.

SAVE\$INTERRUPT\$TABLE saves the contents of the hardware interrupt descriptor table register in the 6-byte save area pointed to by *location*.

The `RESTORE$INTERRUPT$TABLE` Procedure

`RESTORE$INTERRUPT$TABLE` is a built-in procedure that is activated by a `CALL` statement with the following form:

```
CALL RESTORE$INTERRUPT$TABLE (location);
```

Where:

location is an expression with a `POINTER` value.

`RESTORE$INTERRUPT$TABLE` restores the contents of the hardware interrupt descriptor table register from the save area pointed to by *location*. This save area can be the same area used in a preceding call to `SAVE$INTERRUPT$TABLE`.

A descriptor can be built that will initialize the interrupt processor operations. `RESTORE$GLOBAL$STATUS` can then be called with a pointer to this descriptor.

The user must ensure that the save area contains a valid descriptor. Note that values can be assigned to the `IDTR` only if the program is executed in protection mode at level 0.

The Local Descriptor Table Register

The `LOCAL$TABLE` Variable

`LOCAL$TABLE` is a built-in `SELECTOR` variable that provides access to the local descriptor table register (`LDTR`). The format of the register is a selector pointing to an `LDT` in the `GDT`. The use of the local descriptor table is like the use of the `GDTR`, except that it defines the local address space.

By assigning a value to `LOCAL$TABLE`, the local address space of the current task is altered. If a task switch occurs, the new contents are not saved in the task state segment. (To ensure proper operation, interrupts must be disabled.)

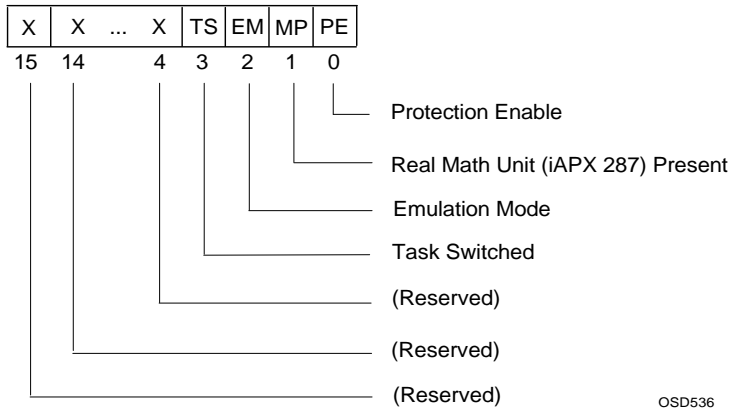
`LOCAL$TABLE` can be read to determine the current active descriptor array segment for the current task.

The user must ensure that the selector in `LOCAL$TABLE` points to a valid descriptor segment. Note that values can be assigned to the `LDTR` only when the program is executed in protection mode at level 0.

The Machine Status Register

The MACHINE\$STATUS Variable

For PL/M-386, MACHINE\$STATUS is a built-in HWORD variable. MACHINE\$STATUS provides access to the machine status word (MSW). The MSW register defines the current status of the processor protection model and the real math unit support. The format of MACHINE\$STATUS is:




MACHINE\$STATUS enables access to the protected mode of the microprocessor. When a value is assigned to this register, the compiler generates a short jump to the next instruction to clear the instruction queue. (Note, however, that values can be assigned to MACHINE\$STATUS only if the program is executed in protection mode at level 0.)

The contents of MACHINE\$STATUS can also be read to determine the current status of various system components.

The CONTROL\$REGISTER, DEBUG\$REGISTER, and TEST\$REGISTER Built-in Arrays

The CONTROL\$REGISTER is a built-in WORD array that provides access to the Intel386 and Intel486 microprocessors' 32-bit control registers that define the current status of the processor and contain page table and page fault information.

The format of CONTROL\$REGISTER (0) is:



OSD575

where:

- PG = paging enabled
- ET = extension type
- TS = task switched
- EM = emulate coprocessor
- MP = numeric coprocessor present
- PE = protection enable

MSW is contained in the low-order 16 bits of CONTROL\$REGISTER (0). However, assigning a value to the MACHINE\$STATUS built-in does not change the ET (extension type) bit.

CONTROL\$REGISTER (2) contains the 32-bit linear address that caused the last detected page fault.

CONTROL\$REGISTER (3) contains the physical page base address for the first level of the page table structure. This address is in the high 20 bits (bits 12 to 31) of CONTROL\$REGISTER (3). The lower 12 bits are ignored when assigning to CONTROL\$REGISTER (3) and are undefined when reading CONTROL\$REGISTER (3). Note that the control registers are accessible only during execution at protected mode level 0. Also note that CONTROL\$REGISTER (1) is not accessible.

The DEBUG\$REGISTER built-in WORD array provides access to six of the eight 32-bit debug registers; DEBUG\$REGISTER (4) and DEBUG\$REGISTER (5) are not accessible. The debug registers are accessible only during execution at protected mode level 0.

TEST\$REGISTER is a built-in WORD array that provides access to the 32-bit test registers of the microprocessor. Of these test registers, only TEST\$REGISTER (6) and TEST\$REGISTER (7) are accessible; these registers are accessible only when executing in protection mode at level 0.

The CLEAR\$TASK\$SWITCHED\$FLAG Procedure

CLEAR\$TASK\$SWITCHED\$FLAG is a built-in procedure that is activated by a CALL statement with the form:

```
CALL CLEAR$TASK$SWITCHED$FLAG;
```

This procedure is used to clear the task switched flag in the machine status word. The processor sets the task switched flag every time a task switch occurs. It can be used to manage the sharing of the math coprocessor.

CLEAR\$TASK\$SWITCHED\$FLAG can be called only when the program is executed in protection mode at level 0.

Segment Information

The GET\$ACCESS\$RIGHTS Function

GET\$ACCESS\$RIGHTS is a built-in WORD function; it is activated by a function reference with the form:

```
GET$ACCESS$RIGHTS (selector);
```

Where:

selector is an expression with a SELECTOR value.


If the segment pointed to by *selector* is visible at the current privilege level, then the hardware ZERO flag is set and a WORD value is returned. If the segment is not visible, or if it is of the wrong type, the hardware ZERO flag is reset, and the returned value is undefined.



Note

The setting of the ZERO flag is guaranteed only if it is tested immediately, before being altered by another operation. (For example, if the value of the function is assigned to an array element indexed by an expression, the value of the ZERO flag may be incorrect.)

Specific to Intel386 and Intel486 microprocessors, the format of the return value is:



OSD576

- X reserved
- G granularity bit
- AVL available for software use
- ACCESS access rights byte

The following example illustrates how the GET\$ACCESS\$RIGHTS function can be used:

```

DECLARE RIGHTS WORD;
DECLARE SEGMENT SELECTOR;
  RIGHTS = GET$ACCESS$RIGHTS (SEGMENT);
  IF ZERO THEN
    /* The segment pointed to by SEGMENT is visible */
    /* and RIGHTS contains the proper access
    /* rights to it. */
  ELSE
    /* SEGMENT is not visible and the contents of */
    /* RIGHTS is undefined. */

```

The GET\$SEGMENT\$LIMIT Function

For PL/M-386, GET\$SEGMENT\$LIMIT is a built-in OFFSET function. GET\$SEGMENT\$LIMIT is activated by a function call of the form:

```
GET$SEGMENT$LIMIT (selector);
```

Where:

selector is an expression with a SELECTOR value.

If the segment pointed to by *selector* is visible at the current protection level, then the hardware ZERO flag is set and the value returned by GET\$SEGMENT\$LIMIT is the size of the segment. If the segment is not visible, the ZERO flag is reset and the returned value is undefined.

Set the ZERO flag with caution. See the note in section for the GET\$ACCESS\$RIGHTS function.

The following example illustrates how the GET\$SEGMENT\$LIMIT function can be used:

```
DECLARE LIMITS OFFSET;  
DECLARE SEGMENT SELECTOR;  
LIMITS = GET$SEGMENT$LIMIT (SEGMENT);  
IF ZERO THEN  
    /* The segment pointed to by SEGMENT is visible */  
    /* and LIMITS contains its proper size.*/  
ELSE  
    /* SEGMENT is not visible and the contents of*/  
    /* LIMITS is undefined.*/
```

Segment Accessibility

It is sometimes helpful to know if the segment pointed to by a selector is readable or writable from the current address space. This becomes particularly important when the selector is a parameter that is passed to the current task.

If an attempt is made to access a segment that is inaccessible, an interrupt will occur. To avoid this interrupt, segment readability and writability can be tested before the segment is accessed.

The SEGMENT\$READABLE Function

SEGMENT\$READABLE is a built-in BYTE function. It is activated by a function reference with the form:

```
SEGMENT$READABLE (selector);
```

Where:

selector is an expression with a SELECTOR value.

SEGMENT\$READABLE returns a value of TRUE (OFFH) if the segment pointed to by selector is reachable and readable from the current privilege level; FALSE (0), if it is not.

The SEGMENT\$WRITABLE Function

SEGMENT\$WRITABLE is a built-in BYTE function. It is activated by a function reference with the form:

```
SEGMENT$WRITABLE (selector);
```

Where:

selector is an expression with a SELECTOR value.

SEGMENT\$WRITABLE returns a value of TRUE (OFFH) if the segment pointed to by *selector* is reachable and writable from the current privilege level; FALSE (0), if it is not.

Adjusting the Requested Privilege Level

The ADJUST\$RPL Function

ADJUST\$RPL is a built-in SELECTOR function that returns the argument of the adjusted requested privilege level (RPL). It is activated by a function reference with the form:

```
ADJUST$RPL (selector);
```

Where:

selector is an expression with a SELECTOR value.

If the requested privilege level (RPL) field of the argument selector is less than the RPL field of the code segment selector for the routine calling the procedure that invoked ADJUST\$RPL, then the hardware ZERO flag is set and the value returned is the argument of an adjusted RPL field. Otherwise, the ZERO flag is reset, and the value returned is the original value of the argument.

Setting the ZERO flag should be done cautiously; see the note in section on the GET\$ACCESS\$RIGHTS function.

The following example illustrates how the ADJUST\$RPL function can be used:

```
P: PROCEDURE (SEGMENT);
    DECLARE SEGMENT SELECTOR;

    SEGMENT = ADJUST$RPL (SEGMENT);
    IF ZERO THEN
        /* The RPL of SEGMENT was less than the RPL of*/
        /* the routine that called P; SEGMENT now has*/
        /* the RPL of the caller.*/
    ELSE
        /* The RPL of SEGMENT was not less than the RPL */
        /* of the routine that called P; SEGMENT is unchanged */
    END P;
```

The REAL Math Facility

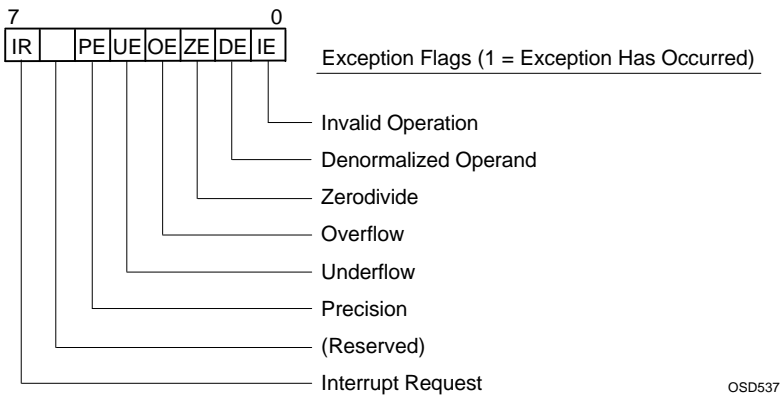
REAL math support for PL/M is provided by the numeric coprocessor. In relation to the program, the REAL math facility consists of the following:

- The REAL stack, used to hold operands and results during REAL operations.
- The REAL error byte (see Figure 10-2), consisting of seven exception flags initialized to all 0s. (The reserved bit is set to 1 by the numeric coprocessor.)

The first six bits in this byte correspond to the possible errors that can arise during REAL operations (see Appendix G). When an error occurs, the facility sets the corresponding bit to 1. A program can invoke a built-in procedure (described in the next section) that reads and clears the REAL error byte.

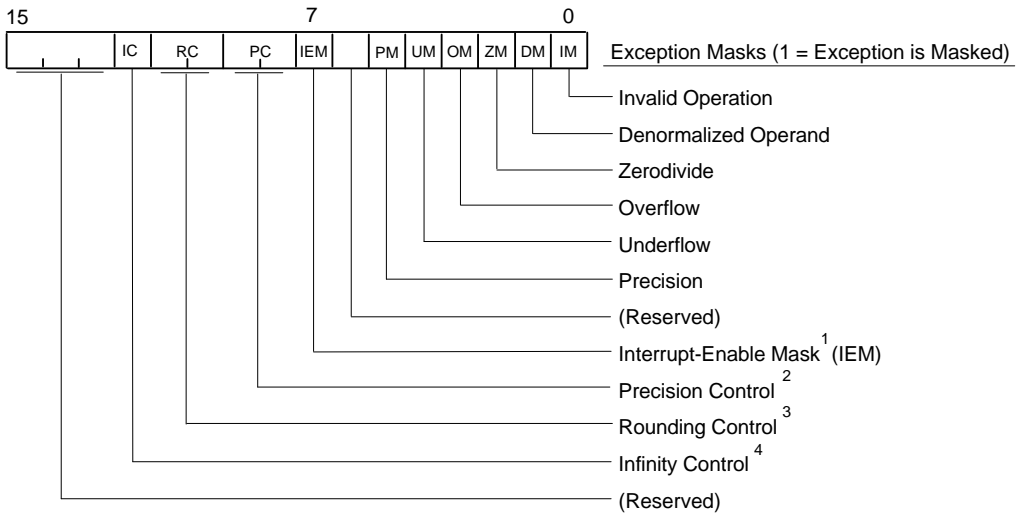
The exception/error categories are discussed in Appendix G.

- The REAL mode word (see Figure 10-3), consisting of 16 bits initialized to 03FFH (or 7FFH for the Intel387™ numeric coprocessor).
 1. Bits 0-7 determine whether the corresponding error condition is to be handled with the default recovery (described next) or with the programmer-supplied exception procedure (see Appendix G for details on writing these). When the bit is 1, the default is used; when it is 0, the user routine is used. In either case, the facility records the error by setting the corresponding bit of the REAL error byte. For most uses, the default recovery is appropriate and less work.



OSD537

Figure 10-2. The REAL Error Byte



- (1) Interrupt - Enable Mask:
 0 Interrupts Enabled
 1 Interrupts Disabled (Masked)
- (2) Precision Control:
 00 24 Bits
 01 (Reserved)
 10 53 Bits
 11 64 Bits
- (3) Rounding Control:
 00 Round To Nearest Or Even
 01 Round Down (Toward $\sim \infty$)
 10 Round Up (Toward $\sim \infty$)
 11 Chop (Truncate Toward Zero)
- (4) Infinity Control:
 0 Projective
 1 Affine

OSD539


Figure 10-3. The REAL Mode Word

This mode word is often called a mask; that is, it lets some signals through (to interrupt processing), but not others. If one of the bits 0-5 is a 0, the corresponding error is said to be unmasked (see the next section for setting the mode word).

If the interrupt is enabled ($IEM = 0$), one of the masked bits is 0, and the corresponding error occurs during floating point processing, then the `REAL` math facility interrupts the host CPU. The numeric coprocessor's interrupt number is dependent on the internal configuration. The exception condition is thus reported and control is passed to the user-written error handling routine. This situation is called an unmasked error. Chapter 8 and Appendix G discuss aspects of interrupt procedures.

Conversely, a masked error means the mode bit corresponding to that error is 1. Masked errors do not cause an interrupt, but are handled as described in Appendix G.

2. Bits 8 and 9 control precision. All intermediate results are held in an internal format of 64-bit precision. The most-significant 24 bits of the final result are returned (plus sign and 7-bit exponent) as the PL/M answer, and rounded, if needed, according to the user-specified control. The default precision setting preserves extended precision and operates slightly faster than the other settings.
3. Bits 10 and 11 control rounding. Rounding introduces an error of less than one unit in the last place to which the result was rounded. Statistically, the default provides the most accurate and unbiased estimate of the true result (i.e., the 64-bit result). In all rounding modes except round down, subtracting a number from itself yields +0; round down yields -0.
4. Bit 12 controls how infinity is handled, as shown below.



Projective Closure

OSD538

Bits 13, 14, and 15 are reserved and are not for PL/M use.

Built-ins Supporting the REAL Math Unit

The INIT\$REAL\$MATH\$UNIT Procedure

INIT\$REAL\$MATH\$UNIT is a built-in untyped procedure activated by a CALL statement, as follows:

```
CALL INIT$REAL$MATH$UNIT;
```

This call is required as the first access to the math coprocessor.

This call initializes the REAL math unit for subsequent operations. This includes setting a default value into the control (REAL mode) word, namely 03FFH or 0000001111111111B. This setting masks all exceptions and interrupts, sets precision to 64 bits, and sets the rounding mode to nearest, with even preferred. This means no interrupts will occur from the REAL math facility regardless of what errors are detected.

Procedures activated after this call has taken effect do not need to do such initialization.

The SET\$REAL\$MODE Procedure

This procedure should only be invoked to change the default mode word (for example, to unmask the invalid exception).

SET\$REAL\$MODE is a built-in untyped procedure, activated by a CALL statement with the following form:

```
CALL SET$REAL$MODE (modeword);
```

Where:

modeword expression with HWORD value

The value of *modeword* becomes the new contents of the REAL mode word (see Figure 10-3). The suggested value for *modeword* is 033EH, (0000001100111110B). This value provides maximum precision, default rounding, and masked handling of all exception conditions except an invalid operation, which can alert the user to errors of initialization or stack usage (see Appendix G for facts and references on writing an interrupt handling procedure).

The GET\$REAL\$ERROR Function

GET\$REAL\$ERROR is a built-in BYTE function activated by a function reference with the following form:

```
GET$REAL$ERROR
```

The BYTE value returned is the current contents of the REAL error byte (see Figure 10-2). This function also clears the error byte in the REAL math facility.

Saving and Restoring REAL Status

If an interrupt procedure performs any floating-point operation, it will change the REAL status. If such an interrupt procedure is activated during a floating-point operation, the program will be unable to continue the interrupted operation correctly after returning from the interrupted procedure. Therefore, it is first necessary for any interrupt procedure that performs a floating-point operation to save the REAL status and subsequently restore it before returning. The built-in procedures SAVE\$REAL\$STATUS and RESTORE\$REAL\$STATUS make this possible. SAVE\$REAL\$STATUS also initializes the numeric coprocessor.

Additionally, these procedures can be used in a multi-tasking environment where a running task using the numeric coprocessor can be preempted by another task that also uses the numeric coprocessor. The preempting task must call SAVE\$REAL\$STATUS before it executes any statements that affect the numeric coprocessor, that is, before calling SET\$REAL\$MODE and before any arithmetic or assignment of REALS (other than GET\$REAL\$ERROR, if needed).

New vectors will be required for the interrupt handlers appropriate to each new task (e.g., to handle unmasked exception conditions). These vectors must be initialized by the operating system.

After its processing is complete and it is ready to terminate, the preempting task must call RESTORE\$REAL\$STATUS to reload the state information that applied at the time the former running task was preempted. This enables that task to resume execution from the point where it relinquished control.



Note

REAL functions without REAL parameters should not call GET\$REAL\$ERROR or SAVE\$REAL\$STATUS before executing at least one floating-point instruction. To do so may result in loss of processor synchronization.

The SAVE\$REAL\$STATUS Procedure

SAVE\$REAL\$STATUS is a built-in untyped procedure activated by a CALL statement with the form:

```
CALL SAVE$REAL$STATUS (location);
```

Where:

location is a pointer to a memory area 108 bytes long where the REAL status information will be saved.

The REAL status is saved at the specified location, and the REAL stack and error bytes are reinitialized.

If the state of the REAL math unit is unknown to this procedure when it is called, as in the case previously mentioned for preempting tasks, then an initialization will destroy existing error flags, masks, and control settings. To avoid this, the appropriate action (except for error-recovery routines, discussed in Appendix G) is to issue:

```
CALL SAVE$REAL$STATUS (@location_1);
```

before any REAL math usage, and

```
CALL RESTORE$REAL$STATUS (@location_1);
```

prior to the procedure's return. The save automatically reinitializes the math unit and the error byte.

This protects the status of preempted tasks or prior procedures and establishes a known initialization state for the current procedure's actions. The microprocessor interrupts are disabled during the save.

⇒ **Note**

The microprocessor must be able to acknowledge numeric coprocessor interrupts or loss of synchronization occurs.

The RESTORE\$REAL\$STATUS Procedure

RESTORE\$REAL\$STATUS is a built-in untyped procedure activated by a CALL statement with the form:

```
CALL RESTORE$REAL$STATUS (location);
```

Where:

location is a pointer to a memory area where the REAL status information was previously saved by a call to the SAVE\$REAL\$STATUS procedure.

This procedure should be called prior to returning from an interrupt procedure where the real math unit's status was saved using SAVE\$REAL\$STATUS.

Interrupt Processing

The WAIT\$FOR\$INTERRUPT Procedure

WAIT\$FOR\$INTERRUPT is a built-in procedure that is activated by a CALL statement with the form:

```
CALL WAIT$FOR$INTERRUPT;
```

This procedure is used to generate an IRET instruction in a nested interrupt task; if it is used elsewhere, the results are undefined. The IRET instruction causes the microprocessor to perform a task switch, saving the status of the outgoing task in its TSS. The next time the interrupt task is activated, execution will begin at the instruction immediately following the IRET, with all the registers unchanged.

The following example illustrates how the WAIT\$FOR\$INTERRUPT procedure can be used:

```
NEW$INTERRUPT:
  CALL INITIALIZE$INTERRUPT$LIST;
                                /* Start of a list of interrupts */
  DO WHILE 1;
    CALL WAIT$FOR$INTERRUPT;
                                /* Wait for next interrupt within list */
    CALL PROCESS$INTERRUPT;
    IF END$OF$INTERRUPT$LIST THEN DO;
      CALL WAIT$FOR$INTERRUPT;
    /* Wait for start of next interrupt sequence */
      GOTO NEW$INTERRUPT;
    END;
  END;
```

WORD16 Mapping for Built-ins

Table 11-5, in the discussion of the WORD32 | WORD16 control, shows the correspondence between default (WORD32) machine built-ins and those available when WORD16 is in effect. For example, Table 11-5 shows that HWORD (WORD32) corresponds to WORD under WORD16.

Intel486 Processor Built-ins

The following are built-ins specific to the Intel486 processor. Specify the MOD486 control for the PL/M-386 compiler to use these functions:

- **BYTE\$SWAP**: This function generates an Intel486 processor instruction that swaps bytes in a 32-bit expression to convert between big and little endian. The **BYTE\$SWAP** function takes a 32-bit expression and returns a value of the same data type as the argument. An argument of less than 32 bits produces a semantic error. To pass a pointer value, use a data type of **WORD** or **OFFSET** instead of **POINTER**.

Invoke **BYTE\$SWAP** as in the following example:

```
DECLARE (a, b) WORD;  
a = BYTE$SWAP( b );  
b = BYTE$SWAP( b + 10 );
```

TEST\$REGISTER: This variable is an array of 8 elements. Each element is a 32-bit unsigned scalar data type. The available registers of **TEST\$REGISTER** include elements (6) and (7). When the compiler control **MOD486** is specified, elements (3) through (5) are also available.

Use **TEST\$REGISTER** as in the following example:

```
DECLARE a WORD;  
a = TEST$REGISTER ( 4 );
```

See also: *Test Registers, i486 Microprocessor Programmer's Reference Manual*

- **INVALIDATE\$DATA\$CACHE**: This function generates the Intel486 processor instruction to clear the entire data cache.

Invoke **INVALIDATE\$DATA\$CACHE** as in the following example:

```
CALL INVALIDATE$DATA$CACHE;
```

- `WB$INVALIDATE$DATA$CACHE`: This function generates the Intel486 processor instruction to first write out all changed lines to memory and then clear the entire data cache.

Invoke `WB$INVALIDATE$DATA$CACHE` as in the following example:

```
CALL WB$INVALIDATE$DATA$CACHE ;
```

- `INVALIDATETLBENTRY`: This function generates the Intel486 processor instruction to clear a specified entry in the paging cache (TLB). Specify the entry to be cleared as an argument, preceded by an @ sign.

Invoke `INVALIDATETLBENTRY` as in the following example:

```
DECLARE a(10) BYTE ;
DECLARE b WORD ;
CALL INVALIDATE$TLB$ENTRY ( @a(5) );
CALL INVALIDATE$TLB$ENTRY ( @b );
```



Compiler Invocation and Controls 11

This chapter describes compiler controls, optimization, and invocation. There are differences in invocation, depending on whether you are running on iRMX or DOS. This chapter covers both operating systems.

Invocation Syntax on iRMX Systems

The general form of the invocation command is:

```
[ :logical_name:]PLM386 filename [control]...
```

Where:

:logical_name:

is the optional logical name for the directory or device containing the PL/M-386 compiler.

PLM386 is the name of the compiler.

filename is the full filename (with directory path) of the file containing the source code. The compiler accepts only one source file per invocation.

control is zero or more of the compiler controls described later in this chapter. Separate multiple controls with spaces to extend the invocation command over multiple lines, use the ampersand (&) as a continuation character.

The INCLUDE control must be the last control.

Errors detected in the invocation command cause the compiler to abort without processing the source file.

The portion of the path set off with colons (:) is an iRMX logical name. A logical name identifies the directory or device that contains the compiler files. In the examples used here, the compiler resides in the *:lang:* directory. The subdirectory *mydir* resides in the directory *source*. *source* resides in *:home:.* When you are logged on as the user *world*, *:home:* is the logical name for the directory */user/world*.

If the logical name is omitted from the invocation command, the operating system automatically searches several directories for the invocation command. The directories searched and the order of the search are defined in the operating system configuration.

Slashes (/) and carets (^), which are also called circumflexes, are used to move up or down the directory tree. To identify a file, start with a logical name (or assume the default). Continue through the directory tree using the slash to search down one level or the caret to search up one level.

For example, if the source file *textfile.plm* is in directory *source*, and *source* is in the directory identified by logical name *:home:*, use the following pathname:

```
- :HOME:SOURCE/TEXTFILE.PLM
```

If the default directory is *:home:source/mydir*, then the same source file can be identified by specifying the path name as follows:

```
- ^TEXTFILE.PLM
```

The caret instructs the operating system to go up one level to find the file.

When you continue an invocation command over multiple lines by entering an ampersand (&) before the line-feed character, the next line automatically appears with the continuation prompt (**). The ampersand can also be used to insert comments. The PL/M-386 compiler ignores characters that appear after an ampersand. For example:

```
- PLM386 :HOME:SOURCE/TEXTFILE.PLM & Run compiler  
** TITLE ("PROJECT SUPERVISOR") & for this file.  
** OPTIMIZE(2) CODE XREF
```

Invocation Examples and Sign-on/Sign-off Messages under the iRMX OS

The following example specifies compilation of a PL/M-386 source file named *myprog.src*. The list file is sent to *myprog.lst*, with the heading TEST 24 on each new page of output. Both the list and object files are written to the directory */user/world/source*.

```
- :LANG:PLM386 /USER/WORLD/SOURCE/MYPROG.SRC &  
** TITLE("TEST 24")
```

The logical name *:home:* can be used in place of the directory pathname */user/world* if you are currently logged on as the user *WORLD*. One of these two specifications must be used if your current directory is not */user/world*. If your default, or current, directory is */user/world/source* only the actual file name, not including the directory pathname, must be specified in the invocation command. To change the default directory, use the *ATTACHFILE* command. Refer to the *iRMX System Call Reference* for additional information on the *ATTACHFILE* command.

The *:lang:* logical name can be omitted if the default iRMX search path is used (which automatically searches *:lang:* for commands).

When invoked, the compiler signs on with the following message:

```
host PL/M-386 COMPILER Vx.y  
Copyright Intel Corporation, years
```

Where:

host identifies the host system.
x.y identifies the compiler version.
years are the copyright years.

When compilation is complete, the compiler signs off with the following message:

```
PL/M-386 COMPLETE. n WARNINGS, m ERRORS.
```

where *n* and *m* are the numbers of warning and error messages generated during compilation.

Invocation Syntax on DOS Systems

The general form of the invocation command is:

```
PLM386 filename[control]...
```

Where:

PLM386 is the name of the compiler. The directory containing the compiler should be in your DOS `PATH`.

filename is the full filename (with directory path) of the file containing the source code. The compiler accepts only one source file per invocation, unless you use the `INCLUDE` control.

control is zero or more of the compiler controls described later in this chapter. Separate multiple controls with spaces.

Note that DOS limits the command line to 128 characters. You can extend the command over multiple lines with the ampersand (&) continuation character.

The `INCLUDE` control (if used) must be the last control. Subsystem controls and certain other controls, identified in this chapter, cannot be part of the invocation command.

Invocation Examples and Sign-on/Sign-off Messages under DOS

The first example specifies compilation of a PL/M-386 source module called *prog1*. The `XREF` control is used.

```
PLM386 PROG1.SRC XREF
```

The second example specifies compilation of a PL/M-386 source module called *myprog*. The list file is sent to *othrfile.lst*, in which the heading `TEST 24` appears on each new page of output.

```
PLM386 MYPROG TITLE('TEST 24') PRINT(OTHRFILE.LST)
```

When invoked, the compiler signs on with the following message:

```
host PL/M-386 COMPILER Vx.y
Copyright Intel Corporation, years
```

Where:

host identifies the host system.

x.y identifies the compiler version.

years are the copyright years.

When compilation is complete, the compiler signs off with the following message:

```
PL/M-386 COMPLETE.    n WARNINGS,    m ERRORS.
```

where *n* and *m* are the numbers of warning and error messages generated during compilation.

File Usage under DOS and the iRMX OS

The PL/M-386 compiler accepts a single source file as input. The compiler creates and deletes work files, as further described below. By default, the compiler creates two files: a print (or list) file and an object file.

Input Files

The pathname used in the invocation command identifies the source file to be compiled. Other files containing source code can be included with the `INCLUDE` control. The source file name and format must follow the file conventions of the OS.

Work Files

The PL/M-386 compiler uses temporary work files that are deleted after compilation. All of these files are located on the device `:WORK:` under the iRMX OS. Under DOS, use the `set` command for selecting an alternate drive for the work files. The following example specifies that the work files be sent to directory `d:`.

```
SET :WORK: = d:\
```



Note

Using the `set` command to relocate work files is useful when the DOS device driver has created a virtual disk. To change the default location of work files, place the work files drive specification command in the `autoexec.bat` file.

All work files have a `.tmp` extension. Avoid using `.tmp` as the extension on any device used by the compiler. It is possible that an existing file with a `.tmp` extension could be deleted or overwritten by the compiler.

The space required for work files is approximately equal to the space required for the source file plus any included files. Be sure that the selected device provides adequate disk space for the compiler work files.

Print Files

The list file (also called the print file) contains a listing of the source program, the messages collected during compilation, and other printed output specified by the listing selection controls. By default, the list file has the same base name as the source file and an *.lst* extension. Unless otherwise specified, the list file is located on the same drive and in the same directory as the source file.

When the `PRINT` control is used, the compiler creates a list file with the same base name as the source file. If another file exists with the same name, the existing file is overwritten. To save the existing file use the `PRINT` control with a parameter; this saves the new list file under another file name.

Object Files

The object file (also called the object code file or object module) contains the object module format translation of the source code. By default, the object file has the same base name as the source file and an *.obj* extension. Unless otherwise specified, the object file is located on the same drive and in the directory as the source file.

The output of a PL/M-386 compiler is an object file containing a compiled module. This object module may be linked with other object modules using the appropriate linker or binder. A knowledge of the makeup of an object module is not necessary for PL/M programming, but can aid in understanding the controls for program size and linkage.

Object modules output by the PL/M-386 compilers contain three sections:

- Code Section
- Data Section
- Stack Section

These sections can be combined in various ways into memory segments for execution, depending on the size of the program.

Code Section

This section contains the instruction code generated for the source program. If either the `LARGE` control or the `ROM` control is used, this section also contains all variables initialized with the `DATA` attribute, all `REAL` constants, and all constant lists.

In addition, the code section for the main program module contains a main program prologue generated by the compiler. This code precedes the code compiled from the source program, and sets the microprocessor for program execution by initializing various registers.

Data Section

All variables are allocated space in this section with the exception of parameters, based variables, and variables located with an `AT` attribute or local to a `REENTRANT` procedure. If the `RAM` control is used, this section also contains all variables initialized with the `DATA` attribute, as well as all `REAL` constants and all constant lists.

If a nested procedure refers to any parameter of its calling procedure, then all parameters of that calling procedure will be placed in the data section during execution. The compiler reserves enough space during compilation to prepare for this.

Stack Section

The stack section is used in executing procedures, as explained in Appendices F and G. It is also used for any temporary storage used by the program but not explicitly declared in the source module (such as temporary values generated by the compiler).

The exact size of the stack is automatically determined by the compiler except for possible multiple invocations of reentrant procedures. You can override this computation of stack size and explicitly state the stack requirement during the binding (linking) process.



Note

When using reentrant procedures or interrupt procedures, be sure to allocate a stack section large enough to accommodate all possible storage required by multiple invocations of such procedures. The stack space requirement of each procedure is shown in the listing produced by the `SYMBOLS` or `XREF` control. This information can be used to compute the additional stack space required for reentrant or interrupt procedures.

Executable Programs

After the source file is compiled, related object modules must be combined to form executable modules. The libraries that provide the necessary run-time support for the application must also be combined with the object modules. To do this you use the BND386 utility, described in the *Intel386 Family Utilities User's Guide*.

DOS offers two ways of automatically invoking and executing multiple programs: batch files and command files. For more information, refer to your DOS operating system manuals.

Introduction to Compiler Controls

Use the compiler controls described in this chapter either in the command that invokes the compiler or as control lines in the source input file.

A control line contains a dollar sign (\$) in the left margin. Normally, the left margin is set at column one, but you can change this with the `LEFTMARGIN` control. Control lines allow selective control over sections of the program. For example, it may be desirable to suppress the listing for certain sections of the program, or to cause page ejects at certain places.

A line in a source file is considered to be a control line by the compiler if there is a dollar sign in the left margin, even if the dollar sign appears to be part of a PL/M comment or character string constant. Control lines within the source code must begin with a dollar sign and can contain one or more controls, each separated by at least one blank. Only the left margin column of a control line should contain a dollar sign.

The following are examples of control lines:

```
$NOCODE XREF
$EJECT CODE
```

There are two types of controls: primary and general. Primary controls must occur either in the invocation command or in a control line that precedes the first noncontrol line of the source file. Primary controls cannot be changed within a module. General controls can occur either in the invocation command or in a control line anywhere in the source input, and can be changed freely within a module. Certain controls can be negated by prefacing the control word with a `NO`. The control descriptions in this chapter indicate that option by showing both options in the headings.

Many controls are available, but a set of defaults is built into the compilers. The controls are summarized in alphabetic order in Table 11-1.

A control consists of a control-name and, in some cases, a parameter. Parameters in control lines must be enclosed in parentheses. Enclosing control parameters on the invocation line in parentheses may be illegal, depending on the host operating system.

The rest of this chapter is organized in the following manner:

- Controls, default settings, abbreviations, and effects are listed in Table 11-1.
- Compiler controls are categorized and an overview for each of the categories is provided.
- Compiler control descriptions are provided in alphabetical order, as listed in Table 11-1. For example, the `NOSYMBOLS` description is located with the `SYMBOLS` description.
- A sample program listing is provided with a description of the listing.

Table 11-1. Compiler Controls

Control	Default	Abbreviation	Effect
CODE NOCODE	NOCODE	CO NOCO	Enables or disables listing of pseudo-assembly code.
COND NOCOND	COND	none none	Determines whether text skipped during compilation appears in the listing.
*DEBUG *NODEBUG	NODEBUG N	DB ODB	Generates debug records in the object module.
EJECT	automatic paging	EJ	Forces a new print page.
IF ELSEIF ELSE ENDIF	not applicable	none	Enables the conditional compilation capability by testing for conditions that use the value of compile-time switches.
INCLUDE	not applicable	IC	Includes other source files as input to the compiler.
*INTERFACE	none	ITF	Enables calls to other high-level languages and to source code translators.
LEFTMARGIN	LEFTMARGIN(1)	LM	Specifies that only input beginning at position <i>n</i> should be processed by the compiler.
LIST NOLIST	LIST	LI NOLI	Enables or disables listing of source program.
*MOD486	none	none	Enables use of the Intel486 instruction set.
*OBJECT *NOOBJECT	OBJECT (<i>source.obj</i>)	OJ NOOJ	Specifies a filename for an object module, or prevents creation of an object module.
*OPTIMIZE	OPTIMIZE(1)	OT	Determines the optimization level during code generation.

* Denotes primary control

continued

Table 11-1. Compiler Controls (continued)

Control	Default	Abbreviation	Effect
OVERFLOW NOOVERFLOW	NOOVERFLOW	OV NOOV	Enables or disables overflow detection during signed arithmetic.
*PAGELENGTH	PAGELENGTH(60)	PL	Specifies the maximum number of lines per page.
*PAGEWIDTH	PAGEWIDTH(120)	PW	Specifies the maximum number of characters per line.
PAGING NOPAGING	PAGING	PI NOPI	Specifies whether the program listing should be page formatted with a heading that identifies the compiler and page number. A user-specified title can also be included (see TITLE).
PRINT NOPRINT	PRINT	PR NOPR	Enables or disables printed output, or selects the device or file to receive the printed output.
*RAM *ROM	**RAM	none	Specifying RAM places the CONSTANT section within the DATA segment in all segmentation. Specifying ROM places constants in the CODE segment.
SAVE RESTORE	none	SA RS	Enables the settings of certain controls to be saved on the stack and restores the control settings after the included file.
SET RESETaaa	RESET(0)	none	Controls the value of switches. SET establishes a value. RESET restores the value to 0.

* Denotes primary control

continued

Table 11-1. Compiler Controls (continued)

Control	Default	Abbreviation	Effect
*SMALL *COMPACT *FLAT *MEDIUM *LARGE	SMALL	SM CP MD FL LA	Determines the segmentation model.
*SUBTITLE	no subtitle	ST	Puts a subtitle on each page of printed output and causes a page eject.
*SYMBOLS *NOSYMBOLS	NOSYMBOLS	SB NOSB	Specifies to the compiler whether or not to produce a listing of identifiers and attributes.
*TITLE	module name in the source code	TT	Places a title on each page of the printed output.
*TYPE *NOTYPE	TYPE	TY NOTY	Specifies whether or not to include type records in the object module.
*WORD16 *WORD32	WORD32	W16 W32	Defines the data type terminology.
*XREF *NOXREF	NOXREF	XR NOXR	Enables or disables a cross-reference listing of source program identifiers.

* Denotes primary control

Input Format Control

The `LEFTMARGIN` control specifies the left margin of the source file.

Code Generation and Object File Controls

These controls determine what type of object file is to be produced and in which directory it is to appear. Object file controls include the following:

```
DEBUG | NODEBUG
INTERFACE
MOD486
OBJECT | NOOBJECT
OPTIMIZE
OVERFLOW | NOOVERFLOW
RAM | ROM
SMALL | COMPACT | FLAT | MEDIUM | LARGE
TYPE | NOTYPE
WORD32 | WORD16
```

Segmentation Controls

For PL/M-386, the segmentation controls influence how locations are referenced in the compiled program, which leads to certain programming restrictions for each of the segmentation controls. These are primary controls. They have the following form:

```
SMALL
COMPACT
MEDIUM
LARGE
FLAT
```

The segmentation controls `SMALL` and `COMPACT` determine the maximum allowable size of the segments produced in the object program as well as the grouping of object types (code, data, constants, and stack). These controls affect the operation of the compiler in various ways and impose certain constraints on the source module being compiled.

The `MEDIUM` control is equivalent to the `SMALL` control. The `LARGE` control is equivalent to the `COMPACT` control except when `LARGE` is used to indicate a subsystem whose name is unknown at compile time.

The `FLAT` control generates an object module containing separate code, data, and stack segments, with constants in the code segment. You can use the `BLD386` `FLAT` control to link the segments together in a single segment up to 4 GB.

For maximum efficiency of the object code, the smallest possible size should be used for any given program. Also, all modules of a program should be compiled with the same segmentation control.

The segmentation controls are described later in this chapter; extensions to these controls, i.e., the use of subsystems, are described in Chapter 13.

Listing Selection and Content Controls

These controls determine what types of listings are produced and where they appear. The controls are as follows:

```
CODE | NOCODE  
LIST | NOLIST  
PRINT | NOPRINT  
SYMBOLS | NOSYMBOLS  
XREF | NOXREF
```

Listing Format Controls

Format controls determine the format of the listing output of the compiler. These controls are as follows:

```
EJECT  
PAGELENGTH  
PAGEWIDTH  
PAGING | NOPAGING  
SUBTITLE  
TITLE
```

Source Inclusion Controls

With these controls, the input source can be changed to a different file. The controls are:

```
INCLUDE  
SAVE | RESTORE
```

Conditional Compilation Controls

These controls cause selected portions of the source file to be skipped by the compiler if specified conditions are not met. Figure 11-1 shows an example program using conditional compilation and Figure 11-2 shows the same example program using the NOCOND control.

The conditional compilation controls are:

```
COND | NOCOND  
IF | ELSEIF | ELSE | ENDIF  
SET | RESET
```

```

system-id PL/M-386 Vx.y  COMPILATION OF MODULE EXAMPLE
OBJECT MODULE PLACED IN cex.obj
COMPILER INVOKED BY:  plm386 cex.src PW(78) SET(DEBUG=3)
 1      EXAMPLE: DO
 2  1    DECLARE BOOLEAN LITERALLY 'BYTE',
        TRUE LITERALLY '0FFH',
        FALSE LITERALLY '0';
 3  1    PRINT$DIAGNOSTICS: PROCEDURE (SWITCHES, TABLES) EXTERNAL;
 4  2    DECLARE (SWITCHES, TABLES) BOOLEAN;
 5  2    END PRINT$DIAGNOSTICS;

 6  2    DISPLAY$PROMPT: PROCEDURE EXTERNAL; END DISPLAY$PROMPT;

 8  2    AWAIT$CR: PROCEDURE EXTERNAL; END AWAIT$CR;

        $IF DEBUG = 1
            CALL PRINT$DIAGNOSTICS (TRUE, FALSE);
        $   RESET (TRAP)
        $ELSEIF DEBUG = 2
            CALL PRINT$DIAGNOSTICS (TRUE, TRUE);
        $   RESET (TRAP)
        $ELSEIF DEBUG = 3
10  1    CALL PRINT$DIAGNOSTICS (TRUE, TRUE);
11  1    CALL PRINT$DIAGNOSTICS (TRUE, TRUE);
        $   SET (TRAP)
        $ENDIF

        $IF TRAP
12  1    CALL DISPLAY$PROMPT;
13  1    CALL AWAIT$CR;
        $ENDIF

14  1    END EXAMPLE;

```

Figure 11-1. Sample Program Using Conditional Compilation (SET control)

system-id PL/M-386 Vx.y COMPILATION OF MODULE EXAMPLE
 OBJECT MODULE PLACED IN cex.obj
 COMPILER INVOKED BY: plm386 cex.src PW(78) SET(DEBUG=3) NOCOND

```

1            EXAMPLE: DO;

2 1        DECLARE IS LITERALLY 'LITERALLY',
            BOOLEAN IS 'BYTE',
            TRUE IS '0FFH'
            FALSE IS '0';

3 1        PRINT$DIAGNOSTICS: PROCEDURE (SWITCHES, TABLES) EXTERNAL;
4 2        DECLARE (SWITCHES, TABLES) BOOLEAN;
5 2        END PRINT$DIAGNOSTICS;

6 2        DISPLAY$PROMPT: PROCEDURE _EXTERNAL; END DISPLAY$PROMPT;

8 2        AWAIT$CR: PROCEDURE EXTERNAL; END AWAIT$CR;

            $IF DEBUG = 1
            $ELSEIF DEBUG = 3
10 1            CALL PRINT$DIAGNOSTICS (TRUE, TRUE);
11 1            CALL PRINT$DIAGNOSTICS (TRUE, TRUE);
            $    SET (TRAP)
            $ENDIF

            $IF TRAP
12 1            CALL DISPLAY$PROMPT;
13 1            CALL AWAIT$CR;
            $ENDIF
14 1        END EXAMPLE;
  
```

Figure 11-2. Sample Program Showing the NOCOND Control

Language Compatibility Control

The `INTERFACE` control enables PL/M to call procedures written in other languages and vice versa. For PL/M-386, this control also enables the use of external procedures compiled with PL/M-286 (or another OMF286 compiler).

Predefined Switches

If one of the switch names (in the following list) appears in an `IF` or `ELSEIF` condition and has not been explicitly assigned a value using the `SET` or `RESET` control, its default value is its primary control value.

<code>SMALL</code>	<code>MEDIUM</code>	<code>WORD16</code>
<code>COMPACT</code>	<code>RAM</code>	<code>WORD32</code>
<code>LARGE</code>	<code>ROM</code>	

If a predefined switch is assigned a value using the `SET` or `RESET` control, it functions from that point on like any other switch. A primary control value is not affected by setting or resetting the predefined switch with the same name.

The four model switches are distinct. Even though the primary controls `SMALL` and `MEDIUM` have the same control interpretation, specifying the `MEDIUM` control sets the `MEDIUM` switch only, and specifying the `SMALL` control sets the `SMALL` switch only (similarly for `COMPACT` and `LARGE`).

For example, given the following sequence of PL/M-386 control lines:

```
$RAM WORD16 MEDIUM ; line 1
$IF RAM ; line 2
.
$ELSEIF WORD32
.
$ELSEIF SMALL
.
$ENDIF
.
$SET (SMALL, WORD32) ; line x
```

At line 2, the switches `RAM` and `WORD16` are true and their counterparts `ROM` and `WORD32` are false. The switch `MEDIUM` is true and the switches `SMALL`, `COMPACT`, and `LARGE` are false. Therefore, the `IF` condition is true and the two `ELSEIF` conditions are false. After line `x`, the switches `RAM`, `WORD16`, `MEDIUM`, `SMALL`, and `WORD32` are true; `ROM`, `COMPACT`, and `LARGE` remain false. The setting of `SMALL` and `WORD32` compile time switches (whether set or reset) does not affect the existing segmentation control or any of the other switches.

Compiler Control Encyclopedia

The following sections present each of the PL/M-386 compiler controls. Note that the segmentation controls are grouped under the `SMALL` control.

CODE | NOCODE

Form `CODE | NOCODE`

Default `NOCODE`

Type General

The `CODE` control specifies that listing of the generated object code in pseudo-assembly language format is to begin. This listing is placed at the end of the program listing in the listing file. Note that the `CODE` control cannot override a `NOPRINT` control.

The `NOCODE` control specifies that listing of the generated object code is to be suppressed until the next occurrence, if any, of a `CODE` control.

COND | NOCOND

These controls determine whether text within an `IF` element will appear in the listing if it is skipped during compilation.

Form `COND | NOCOND`

Default `COND`

Type General

The `COND` control specifies that any text that is skipped is to be listed (without statement or level numbers). Note that a `COND` control cannot override a `NOLIST` or `NOPRINT` control, and that a `COND` control will not be processed if it is within text that is skipped.

The `NOCOND` control specifies that text within an `IF` element that is skipped is not to be listed; however, the controls that delimit the skipped text will be listed. This provides an indication that something has been skipped. Note that a `NOCOND` control will not be processed if it is within text that is skipped.

Figure 11-1 shows an example in which the program was compiled using the `COND` (by default) and `SET` controls with the `SET` switch assignment `DEBUG=3`. Figure 11-2 is the same program, but it was compiled using the `NOCOND` control. These figures demonstrate the use of conditional compilation. See also the description of `SET | RESET`.

DEBUG | NODEBUG

Form DEBUG | NODEBUG

Default NODEBUG

Type Primary

The `DEBUG` control specifies that the object module is to contain the statement number and relative address of each source program statement, information about each local symbol (including based symbols and procedure parameters), and block information for each procedure. This information may be used later by a source level debugging tool.

⇒ **Note**

`OPTIMIZE(0)` is the only level of optimization that does not optimize code between program lines. Thus, it is the only one that gives guaranteed results when debugging programs.

EJECT

Form EJECT

Default None

Type General

`EJECT` stops printing on the current page and starts a new page of printed output.

IF | ELSE | ELSEIF | ENDIF

These controls provide conditional compilation capability based on the values of switches.

These controls cannot be used in the invocation of the compiler, and each must be the only control on its control line. There are no default settings or abbreviations for these controls.

An IF control and an ENDIF control delimit an IF element, which can have several different forms. The simplest form of an IF element is:

```
$IF condition
  text
$ENDIF
```

Where:

condition is a limited form of a PL/M expression in which the only valid operators are OR, XOR, NOT, AND, <, <=, =, < >, >=, and >, and the only valid operands are switches and whole-number constants with a range of 0 to 255. If the switch does not appear in a SET control, a value of false (0) is assumed (except for predefined switches). Parenthesized subexpressions cannot be used. Within these restrictions, *condition* is evaluated according to the PL/M rules for expression evaluation. Note that *condition* must be followed by an end-of-line.

text is text that will be processed normally by the compiler if the least significant bit of the value of *condition* is a 1, or skipped if the bit is a 0. Note that text can contain any mixture of PL/M source and compiler controls. If the text is skipped, any controls within it are not processed.

The second form of the IF element contains an ELSE element:

```
$IF condition
  text 1
$ELSE
  text 2
$ENDIF
```

In this construction, *text 1* will be processed if the least significant bit of the value of *condition* is a 1, and *text 2* will be skipped. If the bit is a 0, *text 1* will be skipped and *text 2* will be processed.

Only one ELSE control can be used within an IF element.

With the most general form of the IF element, one or more ELSEIF controls can be introduced before the ELSE (if any):

```
$IF condition 1
    text 1
$ELSEIF condition 2
    text 2
$ELSEIF condition 3
    text 3
.
.
.
$ELSEIF condition n
    text n
$ELSE
    text n+1
$ENDIF
```

where any of the ELSEIF elements can be omitted, as can the ELSE element.

The conditions are tested in sequence. As soon as one of them yields a value with a 1 as its least significant bit, the associated text is processed. All other text in the IF element is skipped. If none of the conditions yields a least significant bit of 1, the text in the ELSE element (if any) is processed and all other text in the IF element is skipped.

Parentheses cannot be used on a conditional control line. For example, the following line is illegal:

```
$IF A+(B+C)
```

INCLUDE

Form INCLUDE(*pathname*)

Default None

Type General

An INCLUDE control must be the right-most control in a control line or in the invocation command.

The INCLUDE control causes the specified file to be included during compilation. Input continues from this file until an end-of-file is detected, and then processing resumes in the file containing the INCLUDE control.

An included file may also contain INCLUDE controls. Note that such nesting of included files cannot exceed the depth given in Appendix B.

INTERFACE

INTERFACE is a primary control that enhances the compatibility of PL/M with other programming languages. The INTERFACE control enables PL/M programs to call procedures written in other languages, such as iC-386, if those procedures use the variable parameter list (VPL) calling convention. Additionally, with the INTERFACE control, procedures written in PL/M can be called by procedures written in other languages. The calling conventions for procedures written in Pascal, FORTRAN, and PL/M are identical.

There are two types of calling conventions in iC-386. One is the FPL and the other is VPL. The fixed parameter list (FPL) type is the default calling convention of the iC-386 compiler. So whenever the C procedures are defined to be FPL, no special designation is needed. But whenever the C procedure is defined to follow the VPL convention, you must use the INTERFACE control. Note that INTERFACE cannot be part of an invocation command.

For PL/M-386, INTERFACE is a primary control that enables PL/M-386 programs to call or be called by procedures compiled with an Intel386 translator, such as iC-386 or ASM386. It can also be used to provide compatibility with procedures compiled by a 286 translator, such as Fortran-286 or Pascal-286.

The INTERFACE control has the following form:

Form `INTERFACE(lang[/machine[/model[/ram|rom]]]
 [=id[id]. . .])`

Default None

Type Primary

Where:

lang is the name of the language, eg. C, that requires a different calling convention for VPL procedures.

machine is Intel386 when calling VPL iC-386 procedures, and 286 when calling VPL procedures compiled using a 286 translator. Only references to 286 ids from Intel386 modules are supported; Intel386 ids cannot be referenced from 286 modules. Therefore, if *machine* is 286 then all the identifiers in the *id* list must be declared EXTERNAL. If *machine* is 286 and an *id* is PUBLIC, it is an error.

<i>model</i>	is SMALL, COMPACT, MEDIUM, or LARGE and defines the model of segmentation for the specified <i>ids</i> . <i>model</i> determines whether Intel386 POINTER variables are offset-only or select-offset, as defined by the PL/M-386 models of segmentation (see Chapter 13, Table 13-1). If <i>machine</i> is 286, <i>model</i> defaults to LARGE. <i>model</i> should be specified as the same model of segmentation used to compile the 286 code being referenced. If <i>machine</i> is Intel386, <i>model</i> is ignored.
<i>ram rom</i>	is RAM or ROM and defines the placement of constant variables in either the code or data segment. When used with the SMALL model, <i>ram/rom</i> also defines whether POINTER variables are offset-only or selector-offset. The default is RAM unless <i>model</i> is LARGE, in which case the default is ROM. <i>ram/rom</i> is ignored if <i>machine</i> is Intel386.
<i>id</i>	specifies the procedures and variables that are implemented using the specified language interface convention.

When the `INTERFACE` control is used to call procedures compiled with a 286 translator, the program switches from using 32-bit stack offsets to 16-bit offsets. Therefore, the stack pointer for the called procedure must point within the lowest 64K of the stack segment, or else a gate must be used to switch to such a stack segment. Parameters must fit within this boundary as well.

Except as noted above, the calling conventions for Intel386-based languages other than VPL iC-386 procedures are identical to PL/M-386 and therefore do not require the use of the `INTERFACE` control. Because the calling conventions differ for iC-386, `INTERFACE` must be used to call or to be called for VPL iC-386 procedures. The C (VPL) interface convention differs from the PL/M calling convention in the following ways:

- Parameters are evaluated and pushed onto the stack in reverse order.
- A parameter whose size is less than two bytes (for Intel386 and Intel486 processors, four bytes) is zero-extended or sign-extended according to its type.
- Real parameters for iC-386 are always 64-bit double floating-point numbers and are passed on its stack.
- The caller clears the parameters off the stack after return and the callee does not pop parameters off the stack.

If you define a function to be a C calling convention procedure, you can call it with more arguments than the number of parameters you specify in the external declaration. Thus, you can make variable parameter list (VPL) procedure calls to functions such as the C function `printf`. This feature is similar to the ANSI C prototyped function declarations using ellipses (`, . . .`). Type checking occurs for arguments passed to the parameters you specify in the external declaration, not for any additional arguments. For example, no type checking is done on a call to a procedure declared with no parameters.

For example, the following is valid:

```
/* Define SAMP as a C Calling Convention procedure. */
$INTERFACE (C=SAMP)
/* Declare SAMP; specify a parameter of type WORD. */
SAMP: PROCEDURE (P) EXTERNAL;
    DECLARE P WORD;
END SAMP;
/* Declare variables to pass to SAMP. */
DECLARE (A, B, C) WORD;
/* Pass arguments to SAMP. Type checking occurs */
/* for the argument A (parameter P) but not for */
/* arguments B and C. */
CALL SAMP (A, B, C);
```

Constant arguments are typed as in PL/M. This typing can affect the value of an argument passed to a C routine, as demonstrated in the following example, where SAMP2 and SAMP3 are C interface functions:

```
SAMP2: PROCEDURE (D) EXTERNAL;
    DECLARE D INTEGER;
END SAMP2;
SAMP3: PROCEDURE EXTERNAL;
END SAMP3;

CALL SAMP2 ( 113 );
    /* passes 113 as an INTEGER since D is
    declared as type INTEGER. */

CALL SAMP3 ( 115 );
    /* passes 115 as a BYTE
    (since 0 < 115 < 255), but high
    byte(s) are undefined since C
    does integer promotion, even if
    the first argument of SAMP3 is
    an unsigned char argument stack. */

CALL SAMP3 ( INTEGER(115) );
    /* uses an explicit cast to
    ensure the constant is passed correctly. */
```

The following example demonstrates the use of the INTERFACE control to call a PL/M-386 procedure:

```
$WORD32
$INTERFACE (PLM/386/FLAT/ROM=EXAMP)

EXAMP: PROCEDURE (A, B) EXTERNAL;
    DECLARE A WORD, B POINTER;
END EXAMP;

DECLARE X WORD, Y POINTER;

CALL EXAMP (X, Y);
```

In the preceding example, the `INTERFACE` control specifies the procedure `EXAMP` to be defined as Intel386-compatible. The actual parameters `X` and `Y` will be automatically converted to a 32-bit `WORD` and an Intel386 (48-bit) `POINTER`, respectively.

Variables and formal parameters of Intel386-based procedures should be declared the same as in the PL/M-386 code. The PL/M-386 compiler is also able to interpret the terms in 286 context and perform the following mapping:

Term Used	Maps to Data Type
<code>BYTE</code>	8-bit, unsigned
<code>HWORD</code>	8-bit, unsigned
<code>WORD</code>	16-bit, unsigned
<code>DWORD</code>	32-bit, unsigned
<code>QWORD</code>	32-bit, unsigned
<code>CHARINT</code>	8-bit (interpretation dependent on 286 code)
<code>SHORTINT</code>	8 bit (interpretation dependent on 286 code)
<code>INTEGER</code>	16-bit, signed integer
<code>LONGINT</code>	32-bit(interpretation dependent on 286 code)
<code>REAL</code>	32-bit, real
<code>SELECTOR</code>	16-bit, selector
<code>POINTER</code>	see the following paragraphs
<code>OFFSET</code>	16-bit, unsigned

The PL/M-386 compiler converts Intel386-style 48-bit long `POINTERS` to 286 `POINTERS` by truncating the offset portion to 16 bits. In `SMALL RAM`, a `POINTER` is the same as an `OFFSET`, and is treated as such by the compiler.

Note that this mapping is independent of `WORD16` | `WORD32` (defined in Tables 9-3 and 11-4). This means that there is a third mapping of scalar terms to scalar data types.

LEFTMARGIN

This is the only control for specifying the format of the source input.

Form LEFTMARGIN(*n*)

Default LEFTMARGIN(1)

Type General

All characters to the left of position *n* on subsequent input lines are not processed by the compiler (but do appear on the listing). The first character on a line is in column 1.

The new setting of the left margin takes effect on the next input line. It remains in effect for all input from this source file and any included files until it is reset by another LEFTMARGIN control.

Note that a control line is one that contains a dollar sign in the column specified by the most recent LEFTMARGIN control.

LIST | NOLIST

Form LIST|NOLIST

Default LIST

Type General

The LIST control specifies that listing of the source program is to resume with the next source line read. The PL/M-386 compiler numbers all source lines, incrementing the line number for each new-line character. Note that the LIST control cannot override a NOPRINT control. If NOPRINT is in effect, no listing is produced.

The NOLIST control specifies that listing of the source program is to be suppressed until the next occurrence, if any, of a LIST control.

When LIST is in effect, all input lines (from the source file or from an included file), including control lines, are listed, provided there is not a NOPRINT control in effect. When NOLIST is in effect, only source lines associated with error messages are listed.

MOD486

The MOD486 control, recognized by only the PL/M-386 compiler, is a switch governing the instruction set available to the compiler. Use this control to compile source text containing the following built-ins specific to the Intel486 processor:

NAME	USAGE
BYTE\$SWAP	Byte swap function to convert between big and little endian. The endian of a stored value indicates whether the most-significant bit is in the highest (big endian) or lowest (little endian) address of the location.
TEST\$REGISTER	Built-in variable extending the number of available TEST\$REGISTER elements
INVALIDATE\$DATA\$CACHE	Function to clear the entire data cache
WB\$INVALIDATE\$DATA\$CACHE	Function to write-back changed lines to memory and to clear the data cache
INVALIDATE\$TLB\$ENTRY	Function to invalidate a single entry in the paging cache

OBJECT | NOOBJECT

Form	OBJECT(<i>pathname</i>) NOOBJECT
Default	OBJECT(<i>sourcefilename</i> .OBJ)
Type	Primary

The OBJECT control specifies that an object module is to be created during compilation. The *pathname* is a standard host operating system *pathname* that specifies the file to receive the object module. If the control is absent or if an OBJECT control appears without a *pathname*, the object module is directed to a file that has the same name as the source input file, but with the extension .OBJ.

The NOOBJECT control specifies that no object module is to be produced.

OPTIMIZE

This control governs the level of optimization to be performed in generating object code. The *n* parameter can be 0-3, representing the lowest to highest levels of optimization. Figures 11-3 to 11-6 illustrate the different levels of optimization. The same program was compiled for each level, but the source file was printed only for OPTIMIZE(0).

Form OPTIMIZE(*n*)

Where: *n* = 0, 1, 2 or 3

Default OPTIMIZE(1)

Type Primary

OPTIMIZE(0) specifies only folding of constant expressions. Folding means recognizing, during compilation, operations that are superfluous or combinable, and removing or combining them so as to save memory space or execution time. Examples include addition with a zero operand, multiplication by one, and logical expressions with true or false constants.

OPTIMIZE(0) is the only level of optimization that is guaranteed to not optimize code between lines. Figure 11-3 illustrates the OPTIMIZE(0) level of optimization.

system-id PL/M-386 Vx.y COMPILATION OF MODULE
EXAMPLES_OF_OPTIMIZATIONS

OBJECT MODULE PLACED IN *example.obj*

COMPILER INVOKED BY: *plm386 example.src PW(78) FLAT CODE*

OPTIMIZE(0)

```

1           EXAMPLES_OF_OPTIMIZATIONS: DO;
2 1         DECLARE (A,B,C) WORD,
3 1         D(100) WORD,
4 1         (PTR_1, PTR_2) POINTER,
5 1         ABASED BASED PTR_1 (10) WORD;
6 1         DO WHILE D(A+B) < D(A+B+1);
7 2         IF (OFFSET(PTR_1) < (OFFSET(PTR_2))) THEN DO;
8 3         A = A * 2;
9 3         ABASED(A) = ABASED(B);
10 3        ABASED(B) = ABASED(C);
11 3        END;
12 2        ELSE A = A + 1;
13 2        END;
14 1        END EXAMPLES_OF_OPTIMIZATIONS;
    
```

ASSEMBLY LISTING OF OBJECT CODE

```

00000000 8BEC          MOV     EBP,ESP
                @1:
00000002 8B0500000000    MOV     EAX,A
00000008 030504000000    ADD     EAX,B
0000000E 8B0D00000000    MOV     ECX,A
00000014 030D04000000    ADD     ECX,B
0000001A 41              INC     ECX
0000001B 8B04850C000000    MOV     EAX,D[EAX*4]
00000022 3B048D0C000000    CMP     EAX,D[ECX*4]
00000029 0F8375000000    JAE     @2
                                ; STATEMENT # 7
0000002F 8B059C010000    MOV     EAX,PTR_1
00000035 8B0DA0010000    MOV     ECX,PTR_2
    
```

Figure 11-3. Sample Program Showing the OPTIMIZE(0) Control

```

0000003B 3BC1          CMP     EAX,ECX
0000003D 0F834F000000 JAE     @3
                                           ; STATEMENT # 8
00000043 8B0500000000 MOV     EAX,A
00000049 D1E0          SHL     EAX,1
0000004B 890500000000 MOV     A,EAX
                                           ; STATEMENT # 9
00000051 8B059C010000 MOV     EAX,PTR_1
00000057 8B0D04000000 MOV     ECX,B
0000005D 8B0488        MOV     EAX,[EAX].ABASED[ECX*4]
00000060 8B0D9C010000 MOV     ECX,PTR_1
00000066 8B1500000000 MOV     EDX,A
0000006C 890491        MOV     [ECX].ABASED[EDX*4],EAX
                                           ; STATEMENT # 10
0000006F 8B059C010000 MOV     EAX,PTR_1
00000075 8B0D08000000 MOV     ECX,C
0000007B 8B0488        MOV     EAX,[EAX].ABASED[ECX*4]
0000007E 8B0D9C010000 MOV     ECX,PTR_1
00000084 8B1504000000 MOV     EDX,B
0000008A 890491        MOV     [ECX].ABASED[EDX*4],EAX
0000008D E90D000000    JMP     @4
                                           ; STATEMENT # 12
                                @3:
00000092 8B0500000000 MOV     EAX,A
00000098 40            INC     EAX
00000099 890500000000 MOV     A,EAX
                                           ; STATEMENT # 13
                                @4:
0000009F E95EFFFFFF    JMP     @1
                                @2:
                                           ; STATEMENT # 15

```

Figure 11-3. Sample Program Showing the OPTIMIZE(0) Control (continued)

OPTIMIZE(1) specifies strength reduction, elimination of common subexpressions and short-circuit evaluation of some Boolean expressions, as well as the optimizations of level (0).

Strength reduction means substituting quick operations (e.g., shifting by 1 instead of multiplying by 2). This instruction requires less space and executes faster. Adding identical subexpressions may also generate left shift instructions.

Elimination of common subexpressions means that if an expression reappears in the same block, its value is re-used rather than recomputed. The compiler also recognizes commutative forms of subexpressions (e.g., $A + B$ and $B + A$ are seen as the same). Intermediate results during expression evaluation are saved in either registers or on the stack for later use. For example:

```
A = B + C * D / 3 ;
C = E + D * C / 3 ;
```

The value of $C*D/3$ will not be recomputed for the second statement.

Optimizing the evaluation of Boolean expressions uses the fact that in certain cases some of the terms are not needed to determine the value of the expression. For example, in the expression:

```
( A > B AND I > J )
```

if the first term ($A > B$) is false, the entire expression is false, and it is not necessary to evaluate the second term. The use of PL/M built-in procedures does not change this optimization. However, if a user-written function or an embedded assignment is part of the expression, this short evaluation is not done. For example:

```
( A > B AND ( UFUN ( A ) > J ) )
```

is evaluated in full.

Figure 11-4 illustrates the OPTIMIZE(1) level of optimization.

PL/M-386 COMPILER EXAMPLES_OF_OPTIMIZATIONS *date time* PAGE 1

system-id PL/M-386 Vx.y COMPILATION OF MODULES

EXAMPLES_OF_OPTIMIZATIONS

OBJECT MODULE PLACED IN *example.obj*

COMPILER INVOKED BY: *plm386 example.src PW(78) FLAT MODE*

OPTIMIZE(1) NOLIST

PL/M-386 COMPILER EXAMPLES_OF_OPTIMIZATIONS *date time* PAGE 2

ASSEMBLY LISTING OF OBJECT CODE

```

; STATEMENT # 6
00000000 8BEC          MOV     EBP,ESP
                @1:
00000002 8B0500000000 MOV     EAX,A
00000008 8B0D04000000 MOV     ECX,B
0000000E 03C1          ADD     EAX,ECX
00000010 50            PUSH   EAX      ; 1
00000011 40            INC     EAX
00000012 5A            POP     EDX      ; 1
00000013 8B14950C000000 MOV    EDX,D[EDX*4]
0000001A 3B14850C000000 CMP    EDX,D[EAX*4]
00000021 0F8356000000 JAE    @2
                ; STATEMENT # 7
00000027 8B059C010000 MOV    EAX,PTR_1
0000002D 8B15A0010000 MOV    EDX,PTR_2
00000033 3BC2          CMP    EAX,EDX
00000035 0F8337000000 JAE    @3
                ; STATEMENT # 8
0000003B 8B0500000000 MOV    EAX,A
00000041 D1E0          SHL    EAX,1
00000043 890500000000 MOV    A,EA
                ; STATEMENT # 9
00000049 8B159C010000 MOV    EDX,PTR_1
0000004F 8B0C8A        MOV    ECX,[EDX].ABASED[ECX*4]
00000052 890C82        MOV    [EDX].ABASED[EAX*4],ECX
                ; STATEMENT # 10
00000055 8B059C010000 MOV    EAX,PTR_1
0000005B 8B0D08000000 MOV    ECX,C
```

Figure 11-4. Sample Program Showing the OPTIMIZE(1) Control

```

00000061 8B0C88      MOV     ECX,[EAX].ABASED[ECX*4]
00000064 8B1504000000 MOV     EDX,B
0000006A 890C90      MOV     [EAX].ABASED[EDX*4],ECX
0000006D E906000000 JMP     @4
                                ; STATEMENT # 12
                                @3:
00000072 FF0500000000 INC     A
                                ; STATEMENT # 13
                                @4:
00000078 E985FFFFFF JMP     @1
                                @2:
                                ; STATEMENT # 15

```

MODULE INFORMATION:

```

CODE AREA SIZE      = 0000007DH      125D
CONSTANT AREA SIZE = 00000000H      0D
VARIABLE AREA SIZE = 000001A4H     420D
MAXIMUM STACK SIZE = 00000004H      4D
15 LINES READ
0 PROGRAM WARNINGS
0 PROGRAM ERRORS

```

DICTIONARY SUMMARY:

```

410KB MEMORY AVAILABLE
8KB MEMORY USED (1%)
0KB DISK SPACE USED

```

END OF PL/M-386 COMPILATION

Figure 11-4. Sample Program Showing the OPTIMIZE(1) Control (continued)

OPTIMIZE(2) includes OPTIMIZE(0) and OPTIMIZE(1), plus the following:

- Machine code optimizations (e.g., short jumps, moves)
- Elimination of superfluous branches
- Reuse of duplicate code
- Removal of unreachable code and reversal of branch conditions

Optimizing machine code means saving space by using shorter forms for identical machine instructions. This is possible because some instructions have multiple forms. For example:

```
MOV RESULT1,AX; /* move accumulator value to location RESULT1*/
```

can be generated by using three or four bytes for PL/M-86 and PL/M-286, and using five or six bytes for PL/M-386. The former choice saves a byte of storage for the program. Similarly, jumps that the compiler can recognize as within the same segment or closer (within 127 bytes) permit the use of fewer byte instructions.

Elimination of superfluous branches means optimizing consecutive or multiple branches into a single branch. For example:

```
    JZ     LAB1;                /* jump on zero to LAB1 */
    JMP    LAB2;                /* unconditional jump to LAB2 */
LAB1:  .....
      ...
      ...
LAB2:  .....
```

will be transformed into:

```
    JNZ    LAB2;                /* jump on non-zero to LAB2 */
LAB1  .....
      ...
      ...
LAB2:  .....
```

Similarly, multiple branches like the following are eliminated:

```
LAB0:JMP  LAB1
      ...
      ...
LAB1:JPM  LAB2
      ...
      ...
LAB2:.....
```

and transformed into:

```
LAB0:JMP    LAB2
        ...
        ...
LAB1:JMP    LAB2
        ...
        ...
LAB2:.....
```

Reuse of duplicate code can refer to identical code at the end of two converging paths. In such a case, the code is inserted in only one path, and a jump to that path is inserted in the other path. For example:

```
DECLARE A BYTE, SPOT POINTER;
DECLARE S BASED SPOT STRUCTURE (B BYTE, C BYTE);
IF A = 1 THEN
S.C = INPUT (0F7H) AND 07FH;
ELSE
S.C = INPUT (0F9H) AND 07FH;
```

	Before	After
	CMP A, 1H	CMP A, 1H
	JZ & + 5H	JMP @1
	JMP @1	
	IN 0F7H	IN 0F7H
	AND AL, 7FH	JMP @2
	MOV BX, SPOT	
	MOV S [BX+1H], AL	
	JMP @2	
@1:	IN 0F9H	@1: IN 0F9H
	AND AL, 7FH	@2: AND AL, 7FH
	MOV BX, SPOT	MOV BX, SPOT
	MOV S [BX+1H], AL	MOV S [BX+1H],AL
	@2:	

Reuse of duplicate code can also refer to machine instructions, immediately preceding a loop, that are identical to those ending the loop. A branch can be generated to reuse the code generated at the beginning of the loop. For example:

Before	After
ADD AX, BX	LAB0: ADD AX, BX
MOV ANS, AX	MOV ANS, AX
LAB0: MOV AL, DUM1	MOV AL, DUM1
CMP AL, DUM2	CMP AL, DUM2
JNZ LAB1	JNZ LAB1
...	...
...	...
ADD AX, BX	JMP LAB0
MOV ANS, AX	LAB1: ...
JMP LAB0	
LAB1: ...	

This is safe so long as LAB0 is not the target of a jump instruction. The compiler normally handles a whole procedure at a time, and is aware of such a condition. This optimization cannot be safely applied to labels in the outer level of the main program module. This optimization will not change the program and will save code space.

Second level optimization removes unreachable code, takes a second look at the generated object code, and finds areas that can never be reached due to the control structures created earlier.

For example, if the following code were generated before optimization:

```

MOV AX, A
RCR AL, 1
JB @1
JMP @2

@1: MOV AX, 0FFFFH
OUTW 0F6H
JMP @2
MOV AX, B
ADD A, AX
JMP @3

@2: ....
....
....
....

@3: .....
.....

```

Then the removal of unreachable code would produce:

```
        MOV     AX, A
        RCR     AL, 1
        JB     @1
        JMP     @2

@1:     MOV     AX, 0FFFFH
        OUTW   0F6H
        JMP     @2

@2:     ...
        ...
@3:     .....
```

This can be further optimized by reversing the branch condition in the third instruction and removing the unnecessary `JMP @2`:

```
        MOV     AX, A
        RCR     AL, 1
        JNB    @2

@1:     MOV     AX, 0FFFFH
        OUTW   0F6H

@2:     ...
        ...
@3:     .....
```

Figure 11-5 illustrates the `OPTIMIZE(2)` level of optimization.

```
system-id PL/M-386 Vx.y COMPILATION OF MODULES
EXAMPLES_OF_OPTIMIZATIONS
OBJECT MODULE PLACED IN example.obj
COMPILER INVOKED BY: plm386 example.src PW(78) FLAT MODE
OPTIMIZE(2) NOLIST
```

```
                                ; STATEMENT # 6
00000000 8BEC                MOV     EBP,ESP
                                @1:
00000002 A100000000            MOV     EAX,A
00000007 8B0D04000000        MOV     ECX,B
0000000D 03C1                ADD     EAX,ECX
0000000F 50                    PUSH   EAX        ; 1
00000010 40                    INC     EAX
00000011 5A                    POP     EDX        ; 1
00000012 8B14950C000000        MOV     EDX,D[EDX*4]
00000019 3B14850C000000        CMP     EDX,D[EAX*4]
00000020 7348                JNB    @2
                                ; STATEMENT # 7
00000022 A19C010000            MOV     EAX,PTR_1
00000027 8B15A0010000        MOV     EDX,PTR_2
0000002D 3BC2                CMP     EAX,EDX
0000002F 7331                JNB    @3
                                ; STATEMENT # 8
00000031 A100000000            MOV     EAX,A
00000036 D1E0                SHL     EAX,1
00000038 A300000000            MOV     A,EAX
                                ; STATEMENT # 9
0000003D 8B159C010000        MOV     EDX,PTR_1
00000043 8B0C8A                MOV     ECX,[EDX].ABASED[ECX*4]
00000046 890C82                MOV     [EDX].ABASED[EAX*4],ECX
                                ; STATEMENT # 10
00000049 A19C010000            MOV     EAX,PTR_1
0000004E 8B0D08000000        MOV     ECX,C
00000054 8B0C88                MOV     ECX,[EAX].ABASED[ECX*4]
```

Figure 11-5. Sample Program Showing the OPTIMIZE(2) Control

```

00000057 8B1504000000    MOV     EDX,B
0000005D 890C90          MOV     [EAX].ABASED[EDX*4],ECX
00000060 EBA0           JMP     @1
                                           ; STATEMENT # 12
                                @3:
00000062 FF0500000000    INC     A
                                           ; STATEMENT # 13
00000068 EB98           JMP     @1
                                @2:
                                           ; STATEMENT # 15

MODULE INFORMATION:

CODE AREA SIZE      = 0000006AH      106D
CONSTANT AREA SIZE  = 00000000H      0D
VARIABLE AREA SIZE  = 000001A4H     420D
MAXIMUM STACK SIZE  = 00000004H      4D
15 LINES READ
0 PROGRAM WARNINGS
0 PROGRAM ERRORS

DICTIONARY SUMMARY:

410KB MEMORY AVAILABLE
8KB MEMORY USED    (1%)
0KB DISK SPACE USED

END OF PL/M-386 COMPILATION

```

Figure 11-5. Sample Program Showing the OPTIMIZE(2) Control (continued)

OPTIMIZE(3) includes all of the preceding optimizations. It also optimizes indeterminate storage operations (e. g., those using based variables or variables declared with the AT attribute).

⇒ **Note**

The assumption validating this new optimization is that based variables (or AT variables) do not overlay other user-declared variables.

On this optimization level, all Boolean expressions are short-circuited except those containing embedded assignments. (For a description of how this optimization occurs, see OPTIMIZE(1).)

The benefits of this optimization level include more efficient use of code space only if needed values are not overlaid.

Caution in variable-declaration and usage is essential. For example, the sequence:

```
DECLARE (I, J) WORD;
DECLARE THETA (19) AT (@I);
DECLARE A BASED J (10);
STRUCTURE (F1 BYTE, F2 WORD);
..
J=.I;
....
..
A(I).F1 = 7;
A(I).F2 = 99;
THETA(I) = 31;
..
..
```

violates this caution because it causes the values being used as pointers and subscripts to be overlaid. The compiler normally takes steps to avoid the difficulties implied here. But, in OPTIMIZE(3), these steps are omitted due to the implicit requirement that such situations must not be present at this level of optimization.

Figure 11-6 illustrates the OPTIMIZE(3) level of optimization.

system-id PL/M-386 Vx.y COMPILATION OF MODULES

EXAMPLES_OF_OPTIMIZATIONS

OBJECT MODULE PLACED IN example.obj

COMPILER INVOKED BY: plm386 example.src PW(78) FLAT MODE

OPTIMIZE(3) NOLIST

ASSEMBLY LISTING OF OBJECT CODE

```

; STATEMENT # 6
00000000 8BEC          MOV     EBP,ESP
                @1:
00000002 A100000000     MOV     EAX,A
00000007 8B0D04000000   MOV     ECX,B
0000000D 03C1          ADD     EAX,ECX
0000000F 50           PUSH   EAX      ;
00000010 40           INC     EAX
00000011 5A           POP     EDX      ; 1
00000012 8B14950C000000 MOV     EDX,D[EDX*4]
00000019 3B14850C000000 CMP     EDX,D[EAX*4]
00000020 733C          JNB    @2
                ; STATEMENT # 7
00000022 A19C010000     MOV     EAX,PTR_1
00000027 8B15A0010000   MOV     EDX,PTR_2
0000002D 3BC2          CMP     EAX,EDX
0000002F 7325          JNB    @3
                ; STATEMENT # 8
00000031 A100000000     MOV     EAX,A
00000036 D1E0          SHL    EAX,1
00000038 A300000000     MOV     A,EAX
                ; STATEMENT # 9
0000003D 8B159C010000   MOV     EDX,PTR_1
00000043 8B1C8A        MOV     EBX,[EDX].ABASED[ECX*4]
00000046 891C82        MOV     [EDX].ABASED[EAX*4],EBX
                ; STATEMENT # 10
00000049 A108000000     MOV     EAX,C
0000004E 8B0482        MOV     EAX,[EDX].ABASED[EAX*4]
00000051 89048A        MOV     [EDX].ABASED[ECX*4],EAX
00000054 EBAC          JMP    @1

```

Figure 11-6. Sample Program Showing the OPTIMIZE(3) Control


```

; STATEMENT # 12
      @3:
00000056  FF0500000000    INC    A
; STATEMENT # 13
0000005C  EBA4                JMP    @1
      @2:
; STATEMENT # 15

```

MODULE INFORMATION:

```

CODE AREA SIZE      = 0000005EH      94D
CONSTANT AREA SIZE = 00000000H      0D
VARIABLE AREA SIZE = 000001A4H     420D
MAXIMUM STACK SIZE = 00000004H      4D
15 LINES READ
0 PROGRAM WARNINGS
0 PROGRAM ERRORS

```

DICTIONARY SUMMARY:

```

410KB MEMORY AVAILABLE
8KB MEMORY USED   (1%)
0KB DISK SPACE USED

```

END OF PL/M-386 COMPILATION

Figure 11-6. Sample Program Showing the OPTIMIZE(3) Control (continued)

OVERFLOW | NOOVERFLOW

Form `OVERFLOW | NOOVERFLOW`

Default `NOOVERFLOW`

Type `General`

These controls specify whether to detect overflow when performing signed arithmetic. If the `NOOVERFLOW` control is specified, no overflow detection is implemented in the compiled module and the results of overflow in signed arithmetic are undefined. If the `OVERFLOW` control is specified, overflow in signed arithmetic results in a nonmaskable interrupt 4, and it is the programmer's responsibility to provide an interrupt procedure to handle the interrupt. Failure to provide such a procedure may result in unpredictable program behavior when overflow occurs.

If this control is nested within a program statement, overflow detection will begin when the next complete statement is evaluated.

Note that the use of the `OVERFLOW` control results in some expansion of the object code.

Specific to the Intel386 and Intel486 microprocessors, in-line checking code is inserted for detecting machine overflow (32-bit arithmetic overflow) on signed expressions, and value overflow on assignments to `SHORTINT` or `CHARINT` variables.

To save code space and execution time, avoid using `SHORTINT` and `CHARINT` when compiling with the `OVERFLOW` control.

PAGELength

Form `PAGELength (n)`

Default `PAGELength (60)`

Type `Primary`

Pagelength is a non-zero, unsigned number specifying the maximum number of lines to be printed per page of listing output. This number includes the page headings printed on the page.

The minimum value for *n* is 5; the maximum value is 255.

PAGEWIDTH

Form	PAGEWIDTH(<i>n</i>)
Default	PAGEWIDTH(120)
Type	Primary

Pagewidth is a non-zero, unsigned number specifying the maximum line width, in characters, to be used for listing output. The minimum value for *n* is 60; the maximum value is 132.

PAGING | NOPAGING

Form	PAGING NOPAGING
Default	PAGING
Type	Primary

The PAGING control specifies that the listed output is to be formatted onto pages. Each page carries a heading identifying the compiler and a page number, and possibly a user-specified title.

The NOPAGING control specifies that page ejecting, page heading, and page numbering are not to be performed. Thus, the listing appears on one long page as would be suitable for a slow serial output device. If NOPAGING is specified, a page eject is not generated if an EJECT control is encountered.

PRINT | NOPRINT

Form	PRINT(<i>pathname</i>) NOPRINT
Default	PRINT(<i>sourcefilename.LST</i>)
Type	Primary

The PRINT control specifies that printed output is to be produced. The parameter is a standard host operating system pathname that specifies the file to receive the printed output. Any output-type device, including a disk file, can also be given. If the control is absent, or if a PRINT control appears without a pathname, printed output is sent to a file that has the same name as the source input file but with the extension .LST.

The NOPRINT control specifies that no printed output is to be produced, even if implied by other listing controls such as LIST and CODE.

RAM | ROM

Form RAM|ROM

Default RAM

Type Primary

For PL/M-386, the RAM setting places the `CONSTANT` section within the `DATA` segment in all segmentation models.

For all targets, the ROM setting places constants in the `CODE` segment. Under this setting, the `INITIAL` attribute on a variable produces a warning message. Do not use the dot operator for variable references under the ROM option because constants and variables will be relative to different segment registers. If `SMALL` is specified with the ROM control, then PL/M-386 pointers will be six bytes instead of four (see also Appendix F).

If the keyword `DATA` is used in a `PUBLIC` declaration when compiling with the ROM control, `DATA` must also be used in the `EXTERNAL` declaration of program modules that reference it. However, no value list is then permitted because the data is defined elsewhere.

SAVE | RESTORE

Form SAVE|RESTORE

Default None

Type General

With these controls the settings of certain general controls can be saved on a stack and then restored. The main usage of these controls is saving the controls before an included file and restoring them after inclusion of that file is complete. The controls whose settings are saved and restored are:

```
CODE|NOCODE
COND|NOCOND
LEFTMARGIN
LIST|NOLIST
OVERFLOW|NOOVERFLOW
```

The `SAVE` control saves all of these settings on a stack. The maximum capacity of this stack corresponds to the maximum nesting depth for the `INCLUDE` control (the maximum nesting depth is given in Appendix B).

The `RESTORE` control restores the most recently saved set of control settings from the stack.

SET | RESET

These are general controls. The SET control has the following general form:

```
SET (switch_assignment_list)
```

Where:

switch_assignment_list

consists of one or more switch assignments separated by commas.

A switch assignment has the form:

```
switch[=value]
```

Where:

switch is a name which is formed according to the PL/M rules for identifiers. Note that a switch name exists only at the compiler control level, and therefore a switch can have the same name as an identifier in the program; no conflict is possible. Note however that no PL/M reserved word other than a predefined switch can be used as a switch name.

value is a whole-number constant in the range 0 to 255. This value is assigned to the switch. If the value and the equal sign (=) are omitted from the switch assignment, the default value OFFH (true) is assigned to a switch.

The following is an example of a SET control line:

```
$SET(TEST, ITERATION = 3)
```

This example sets the switch TEST to true (OFFH) and the switch ITERATION to 3. Switches do not need to be declared.

Figure 11-1 and 11-2 are examples of a program that was compiled using the SET control.

The RESET control has the form:

```
RESET (switch_list)
```

Where:

switch_list

consists of one or more switch names that have already occurred in SET controls.

Each switch in the switch list is set to false (0).

SMALL | COMPACT | MEDIUM | LARGE | FLAT

The following sections describe the `SMALL`, `COMPACT`, `MEDIUM`, `LARGE` and `FLAT` controls (also called the segmentation controls). For application development under the iRMX Operating System, see the note under the `COMPACT` model description.

SMALL

Form	<code>SMALL</code>
Default	<code>SMALL</code>
Type	Primary

Modules compiled with the `SMALL` control have three sections: code, data, and stack (see the `OBJECT` control). When these modules are bound (linked), similar sections from each module are combined to form two segments: code and data. For the Intel386 and Intel486 microprocessors, the maximum size of each segment is 4G bytes.

In the default `SMALL` case (`RAM`), the code sections from all modules are allocated space within the code segment, which is addressed relative to the `CS` register. Constants are combined with all the data and stack sections in the data segment. For the Intel386 and Intel486 microprocessors, this segment is addressed relative to the `DS` register, with an identical copy in the `SS` and `ES` registers. None of the segment registers are changed during the course of program execution except `ES`, which is used to perform string operations, and `FS` and `GS`, which are used to address data exported by another subsystem. Subsystems are described in Chapter 13.

Therefore, the `SMALL` control can be used only if the total size of all code sections does not exceed 4G bytes. The total size of the constants plus all data and stack sections also cannot exceed 4G bytes.

If the `ROM` control is used, the constants from all the modules are placed with the code in the code segment. The data segment then contains only the data and stack sections from all the modules.

Because only one code segment exists, its segment selector (the `CS` register) is never updated during program execution. (However, an interrupt will update the `CS` register.) Likewise, when `RAM` is used, only one segment exists for all constant, data, and stack sections. The segments' selectors (the `DS` and `SS` registers) are never updated (except when an interrupt occurs, as explained in Appendix G). Therefore, when any location is referenced, a 32-bit offset is calculated and used in conjunction with the appropriate segment selector. `POINTER` values in the `SMALL` (`RAM`) case are 32-bit values for the Intel386 and Intel486 microprocessors.

The following restrictions must be observed:

1. Do not use the @ and dot operations with variables based on SELECTOR. For example:

```
DECLARE SEL SELECTOR;  
DECLARE R BASED SEL BYTE;  
DECLARE PO POINTER;  
PO = @R;                /* invalid under SMALL RAM */
```

2. Do not use the built-in function BUILD\$PTR (see Chapter 9).

COMPACT



Note

The iRMX Operating System supports only the **COMPACT** model.

Form COMPACT

Default SMALL

Type Primary

Modules compiled with the **COMPACT** control have three sections: code, data, and stack (see the **OBJECT** control). When these modules are linked, similar sections from each module are combined to form three segments: code, data, and stack. The maximum size of each segment is 4G bytes for the Intel386 and Intel486 microprocessors.

In the default **COMPACT** case (**RAM**), the code sections from all modules are allocated space within the code segment, which is addressed relative to the CS register. Constants and all data sections are combined in the data segment, which is addressed relative to the DS register and an identical copy is stored in the ES register. The stack is addressed relative to SS. None of the segment registers are changed, except ES, which is used to perform string operations, as well as FS and GS, which are used to address data exported by another subsystem.

If the **ROM** control is used, the constants from all the modules are placed with the code in the code segment. The data segment then contains only the data sections from all the modules.

Since the code, data, and stack segments are fully defined by the time the program is loaded, the segment selectors in the CS and SS registers are never changed.

All six segment registers are initialized by the loader, with ES, FS, and GS initialized to DS. The DS and ES registers are also saved and reinitialized in each interrupt procedure prologue and epilogue to enable distinct interrupt environments. The FS and GS registers are volatile after initialization. References to any location require only a 32-bit offset against these segment selectors.

Observe the following restrictions when using COMPACT.

1. When an exported procedure is indirectly activated, a `POINTER` variable must be used in the `CALL` statement. For example:

```
$COMPACT(SUBSYS HAS MOD1, MOD2, MOD3; EXPORTS PROC)
MOD: DO
    DECLARE P POINTER, W WORD;
    PROC: PROCEDURE PUBLIC;
    .
    .
    .
    END PROC;
    P = @PROC; CALL P;           /* POINTER must be used */
    W = .PROC; CALL W;          /* not allowed */
END MOD1;
```

2. When a procedure that is not exported is indirectly activated, an `OFFSET` variable must be used. Note that `OFFSET` variables do not range over the entire microprocessor address space, but are restricted to offsets within the current code segment. For example:

```
DECLARE P POINTER, O OFFSET;
LPROC: PROCEDURE;                               /* local */
.
.
.
END LPROC;
P = @LPROC; CALL P;                             /* not allowed */
O = .LPROC; CALL O;                             /* OFFSET must be used */
```


MEDIUM

Form	MEDIUM
Default	SMALL
Type	Primary

For PL/M-386, the MEDIUM control is provided for PL/M-86 and PL/M-286 compatibility. The MEDIUM control is interpreted exactly like the SMALL control. For more information, refer to the SMALL control entry in this chapter.

LARGE

Form	LARGE
Default	SMALL
Type	Primary

The LARGE control is provided for PL/M-86 and PL/M-286 compatibility. The LARGE control is interpreted exactly like the COMPACT control in most cases. For more information, refer to the COMPACT control entry in this chapter. When the LARGE control is used in a PL/M-386 subsystem definition, it behaves differently from the COMPACT control. For more information about subsystems, see Chapter 13.

FLAT

The FLAT control is a member of the group of segmentation controls including SMALL and COMPACT. Compiling with the FLAT control generates an object module containing separate code, data, and stack segments, with constants included in the code segment. The FLAT control overrides the RAM or ROM control. Using the `-CONST IN CODE-` or `-CONST IN DATA-` attribute for extended segmentation definition does not result in an error when you specify the FLAT control; however, `-CONST IN CODE-` is redundant and `-CONST IN DATA-` is ignored when FLAT is in effect.

Linking object modules compiled with the FLAT control produces the following linked segments:

- A single code segment (CODE32) containing all the code segments of the object modules
- A single data segment (DATA32) containing all the data segments of the object modules
- A single stack segment (STACK) containing all the stack segments of the object modules

Use the BLD386 `FLAT` control to map the three linked segments together to a single segment of up to 4 Gigabytes.

Since only one segment exists during run-time, all pointers are short (a 32-bit offset with no selector).. Also, compiling the following code with the `FLAT` control does not result in the semantic error generated when compiling this code with the `ROM` control and any other segmentation control:

```
DECLARE B WORD;  
DECLARE A WORD AT (@B) DATA (10);
```

SUBTITLE

Form `SUBTITLE(" subtitle ")`

Default No subtitle

Type General

The subtitle character sequence (truncated on the right to fit, if necessary) is printed on the subtitle line of each page of listed output. Note that a subtitle specified on the invocation line must be enclosed in quotation marks.

The maximum length for subtitle is 60 characters, but a narrow pagewidth may restrict this number.

When a `SUBTITLE` control appears before the first noncontrol line in the source file, it causes the specified subtitle to appear on the first page and all subsequent pages until another `SUBTITLE` control appears.

A subsequent `SUBTITLE` control causes a page eject, and the new subtitle appears on the next page and all subsequent pages until the next `SUBTITLE` control.

SYMBOLS | NOSYMBOLS

Form `SYMBOLS | NOSYMBOLS`

Default `NOSYMBOLS`

Type Primary

The `SYMBOLS` control specifies that a listing of all identifiers in the PL/M source program and their attributes is to be produced in the listing file.

The `NOSYMBOLS` control suppresses such a listing.

Note that the `SYMBOLS` control cannot override a `NO PRINT` control.

TITLE

Form	TITLE(" <i>title</i> ")
Default	TITLE (" <i>modulename</i> ")
Type	Primary

The title character sequence, truncated on the right to fit, if necessary, is placed on the title line of each page of listing output. Note that the character sequence for a title must be enclosed in quotation marks when entered on the invocation line.

The maximum length for the title is 60 characters, but a narrow pagewidth may restrict this number.

TYPE | NOTYPE

Form	TYPE NOTYPE
Default	TYPE
Type	Primary

The TYPE control specifies that the object module is to contain information on the variable types output in symbol records. TYPE records provide a mechanism for promoting type compatibility between subprograms. This information may be used later for type checking when the program modules are combined, or by a debugger.

The NOTYPE control specifies that such type definitions are not to be placed in the object module.

WORD32 | WORD16

Form	WORD32 WORD16
Default	WORD32
Type	Primary

The WORD32 | WORD16 control determines how the compiler interprets the unsigned binary number and signed integer scalar types (as well as the built-ins that specify these data types) in the code being compiled.

When compiling PL/M-286, PL/M-86, or PL/M-80 source code with the PL/M-386 compiler, there are several points to consider before accepting the default (WORD32) or choosing WORD16. See Chapter 3 for a discussion of these points.

Table 11-4 lists the data types as interpreted by the compiler under `WORD32` and `WORD16`. The `WORD16` control does not mean creating PL/M-286 code, but rather that PL/M-386 data types are mapped to the equivalent PL/M-286 data type. It affects only the data types, it does not affect the operation of PL/M-386 functions.

Table 11-4. WORD32 | WORD16 Data Type Mapping

Unsigned Binary Number Data Types	WORD32 (default)	WORD16
BYTE	8-bit	8-bit
HWORD	16-bit	8-bit
WORD	32-bit	16-bit
DWORD	64-bit	32-bit
QWORD	64-bit	64-bit
Signed Integer Data Types	WORD32	WORD16
CHARINT	8-bit	8-bit
SHORTINT	16-bit	8-bit
INTEGER	32-bit	16-bit
LONGINT	32-bit	32-bit

Note that all built-ins that specify data types are different for `WORD16`. Table 11-5 lists the `WORD32 | WORD16` mapping for these built-ins. For example, the `HWORD` built-in is a 16-bit, unsigned binary number under `WORD32`, whereas under `WORD16`, the 16-bit, unsigned binary type is `WORD`.

Table 11-5. WORD32 | WORD16 Built-in Mapping

WORD32	WORD16
(type conversions) BYTE HWORD WORD DWORD, QWORD CHARINT SHORTINT INTEGER	(type conversions) BYTE, HWORD WORD DWORD QWORD SHORTINT, CHARINT INTEGER LONGINT
BLOCKINPUT BLOCKOUTPUT MOVB MOVRB FINDB FINDRB INPUT OUTPUT SKIPB SKIPRB CMPB SETB	BLOCKINPUT BLOCKOUTPUT MOVB, MOVHW MOVRB, MOVRHW FINDB, FINDHW FINDRB, FINDRHW INPUT, INHWOR OUTPUT, OUTHWORD SKIPB, SKIPHW SKIPRB, SKIPRWH CMPB, CMPHW SETB, SETHW
BLOCKINWORD BLOCKOUTWORD MOVHW MOVRHW FINDHW FINDRHW INHWOR OUTWORD SKIPHW SKIPRHW CMPHW SETHW	BLOCKINWORD BLOCKOUTWORD MOVW MOVRW FINDW FINDRW INWORD OUTWORD SKIPW SKIPRW CMPW SETW

continued

Table 11-5. WORD32 | WORD16 Built-in Mapping (continued)

WORD32	WORD16
(type conversions)	(type conversions)
BLOCKINWORD	BLOCKINDWORD
BLOCKOUTWORD	BLOCKOUTDWORD
MOVW	MOVD
MOVRW	MOVRD
FINDW	FINDD
FINDRW	FINDRD
INWORD	INDWORD
OUTWORD	OUTDWORD
SKIPW	SKIPD
SKIPRW	SKIPRD
CMPW	CMPD
SETW	SETD

XREF | NOXREF

Form XREF | NOXREF

Default NOXREF

Type Primary

The XREF control specifies that a cross-reference listing of source program identifiers is to be produced in the listing file.

The NOXREF control suppresses the cross-reference listing.

Note that the XREF control cannot override a NOPRINT control.

Program Listing

Sample Program Listing

During the compilation process, a listing of the source input is produced. Each page of the listing carries a numbered page-header that identifies the compiler, prints a time and date as designated by the host operating system, and optionally gives a title and a subtitle, and/or a date (see Figure 11-7).

The first part of the listing contains a summary of the compilation, beginning with the compiler identification and the name of the source module being compiled. The next line names the file receiving the object code. The next line contains the command used to invoke the compiler. The listing of the program itself is shown in Figure 11-7.

The listing contains a copy of the source input plus additional information. Two columns of numbers appear to the left of the source image. The first column provides a sequential numbering of PL/M statements. (Note that the PL/M-386 compiler treats each new-line character as a line terminator; therefore, blank lines are counted.) Error messages, if any, refer to these statement numbers. The second column gives the block nesting depth of the corresponding statement.

Lines included with the `INCLUDE` control are marked with an equal sign (=) just to the left of the source image. If the included file contains another `INCLUDE` control, lines included by this nested `INCLUDE` are marked with an =1. For yet another level of nesting, =2 is used to mark each line, and so forth up to the compiler's limit of nesting levels (see Appendix B). These markings make it easy to see where included text begins and ends.

system-id PL/M-386 Vx.y COMPILATION OF MODULE STACK
OBJECT MODULE PLACED IN stack.obj
COMPILER INVOKED BY: plm386 stack.src CODE XREF TITLE("Stack
Module")

```
1      STACK: DO;
2      /* This module implements a BYTE stack with
3      push and pop */
3 1    DECLARE S(100) BYTE,
4      /* Stack Storage */
4 1    T BYTE PUBLIC INITIAL(-1);
5      /* Stack Index */
5 1    PPUSH: PROCEDURE (B) PUBLIC;
6      /* Pushes B onto the stack */
6 2    DECLARE B BYTE;
7 2    S(T:=T+1) = B;
8      /* Increment T and store B */
8 2    END PPUSH;
9 1    PPOP: PROCEDURE BYTE PUBLIC;
10     /* Returns value popped from stack */
10 2   RETURN S((T:=T-1)+1);
11     /* Decrement T, return S(T+1) */
11 2   END PPOP;
12 1   END STACK;
12     /* Module ends here */
```

Figure 11-7. Program Listing

Should a source line be too long to fit on the page in one line, it is continued on the following line. Such continuation lines are marked with a hyphen (-) just to the left of the source image.

The `CODE` control can be used to obtain the assembly code produced in the translation of each PL/M statement. Figure 11-8 shows the assembly code listing for the program given in Figure 11-7. This code listing appears in six columns of information in a pseudo-assembly language format:

1. Location counter (hexadecimal notation)
2. Resultant binary code (hexadecimal notation)
3. Label field
4. Opcode mnemonic
5. Symbolic arguments
6. Comment field

```

; STATEMENT # 5
        PPUSH  PROC  NEAR
00000000  55          PUSH  EBP
00000001  8BEC        MOV   EBP,ESP
; STATEMENT # 7
00000003  8A0564000000 MOV  AL,T
00000009  FEC0        INC  AL
0000000B  880564000000 MOV  T,AL
00000011  0FB6C0      MOVZX EAX,AL
00000014  8A4D08      MOV  CL,[EBP].B
00000017  888800000000 MOV  [EAX].S,CL
; STATEMENT # 8
0000001D  5D          POP   EBP
0000001E  C20400     RET   4H
        PPUSH  ENDP
; STATEMENT # 9
        PPOP   PROC  NEAR
00000024  55          PUSH  EBP
00000025  8BEC        MOV   EBP,ESP
; STATEMENT # 10
00000027  8A0564000000 MOV  AL,T
0000002D  FEC8        DEC  AL
0000002F  880564000000 MOV  T,AL
00000035  FEC0        INC  AL
00000037  0FB6C0      MOVZX EAX,AL
0000003A  8A8000000000 MOV  AL,[EAX].S
00000040  5D          POP   EBP
00000041  C3          RET
; STATEMENT # 11
        PPOP   ENDP
; STATEMENT # 12

```

Figure 11-8. Code Listing (continued)

Not all six of the columns will appear on all lines of the code listing. Compiler generated labels (e.g., those that mark the beginning and ending of a DO WHILE loop) are preceded by an AT sign (@). The comments appearing on PUSH and POP instructions indicate the stack depth associated with the stack instruction.

Symbol and Cross-reference Listing

Specifying the XREF or SYMBOLS control adds a summary of all identifier usage in the program listing. Figure 11-9 shows the cross-reference listing of the program given in Figure 11-7. The addresses in ADDR have four leading zeros.

```

PL/M-386 COMPILER Stack Module date time PAGE 3
                                CROSS-REFERENCE LISTING

DEFN  ADDR      SIZE  NAME, ATTRIBUTES, AND REFERENCES
-----
5  0008H        1  B. . . . . BYTE IN PROC(PPUSH) PARAMETER AUTOMATIC      6 7
9  0024H        30  PPOP . . . . . PROCEDURE BYTE PUBLIC STACK=00000004H
5  0000H        33  PPUSH. . . . . PROCEDURE PUBLIC STACK=000000008H
3  0000H       100  S. . . . . BYTE ARRAY(100)      7* 10
1  0000H         1  STACK. . . . . MODULE STACK=00000000H
3  0064H         1  T. . . . . BYTE PUBLIC INITIAL      7 7* 10 10*

```

Figure 11-9. Cross-reference Listing

Depending on whether the SYMBOLS or XREF control was used to request the identifier usage summary, five or seven types of information are provided in the symbol or cross-reference listing. They are as follows:

1. Statement number where the identifier was defined.
2. Relative address associated with the identifier.
3. Size of the object identified (in bytes).
4. The identifier.
5. Attributes of the identifier (including expansion for LITERALLYs and scoping information for local variables and parameters). These attributes reflect the WORD32 | WORD16 terminology of the source file.
6. Statement numbers where the identifier was referenced (XREF control only).
7. Statement numbers where the identifier was assigned a value (XREF control only).

A single identifier can be declared more than once in a source module (i.e., an identifier defined twice in different blocks). Each such unique object, even though named by the same identifier, appears as a separate entry in the listing.

The address given for each object is the location of that object relative to the start of its associated section. The object's attributes determine which section is applicable.

Identifiers in the `SYMBOLS` or `XREF` listing are given in alphabetical order with the following exception: members of structures are listed, in order of declaration, immediately following the entry for the structure itself. Indentation is used to differentiate between these entries.

The `XREF` listing differentiates between items 6 and 7 by adding the asterisk (*) character to statement numbers where a value is assigned. For example, if statement 17 reads as follows:

```
I = I + 1;
```

The list of statement numbers for `I` would include 17 and 17*, indicating a reference and an assignment in statement 17.

The `AUTOMATIC` attribute indicates that the identifier was declared as a parameter or as a local variable in a `REENTRANT` procedure and therefore is allocated dynamically on the stack.

Compilation Summary

Following the listing (or appearing alone if `NOLIST` is in effect) is a compilation summary. Eight pieces of information are provided:

- Code area size gives the size in bytes of the code section of the output module (not including constants, if any).
- Constant area size gives the size in bytes of the constant area. The constant area will be included with either the code or data section in the output module, depending on the specified compiler controls.
- Variable area size gives the size in bytes of the data section of the output module (not including constants, if any).
- Maximum stack size gives the size, in bytes, of the stack section allocated for the output module.
- Lines read gives the number of source lines processed during compilation.
- Program warnings give the number of warning messages issued during compilation.
- Program errors give the number of error messages issued during compilation.
- Dictionary summary gives the actual memory and disk space used by the dictionary during compilation.

Figure 11-10 is an example of the compilation summary.

MODULE INFORMATION:

```
CODE AREA SIZE      = 00000042H      66D
CONSTANT AREA SIZE = 00000000H      0D
VARIABLE AREA SIZE = 00000065H     101D
MAXIMUM STACK SIZE = 00000008H      8D
12 LINES READ
0 PROGRAM WARNINGS
0 PROGRAM ERRORS
```

DICTIONARY SUMMARY:

```
410KB MEMORY AVAILABLE
8KB MEMORY USED   (1%)
0KB DISK SPACE USED
```

END OF PL/M-386 COMPILATION

Figure 11-10. Compilation Summary



Sample Program 12

Introduction

This chapter discusses a sample program consisting of three modules named `FREQ`, `OPEN`, and `PRINT`. The purpose of this program is to illustrate the use of the PL/M language. The program is written in PL/M-386 and compiled with the PL/M-386 compiler.

The program takes an input file, counts the uppercase and lowercase alphabetic characters, and determines the percentage of use for each character. This is printed either to the screen or, if one is specified, to an output file. The program's output lists the number of times each character is used (for uppercase, for lowercase, and for both uppercase and lowercase), and the percentage of use for each character. The source program listings are shown in Figures 12-1 through 12-3.

In addition to the main program modules (`FREQ`, `OPEN`, and `PRINT`), this program also has two include files. The include files, *defns.inc* and *udi.inc* (see Figures 12-4 and 12-5), contain definitions that are used in the program modules. The *defns.inc* include file consists of global variable definitions. The *udi.inc* include file consists of the universal development system interface (UDI) definitions. The UDI definitions are used for operating system interfaces (e.g., file manipulation). Figure 12-6 is an example of the program output.

The following sections describe the source code in each of the program modules. The line numbers in the figures are not part of the source code; they have been added to simplify the discussion of the source code.

FREQ Module

`FREQ` is the main module. The source code is shown in Figure 12-1. As indicated, the line numbers in the figure have been added to simplify the discussion of the source code.

The program lines that begin with a dollar sign (\$) are compiler control lines. Lines that begin with a dollar sign instruct the compiler and are not part of the source program. In any position other than the first character (or the position specified with the `LEFTMARGIN` control), the dollar sign is an insignificant character and can be used as a separator to simplify the reading of variable names.

```

1  $DEBUG PW(75)
2  freq:DO;

3  $INCLUDE (defns.inc)

4  $NOLIST
5          /** LIST of UDI procedures is in OPEN.PLM ***/
6  $INCLUDE (udi.inc)
7  $LIST

8  open$files:PROCEDURE EXTERNAL;
9      END open$files;

10 print$stats:PROCEDURE(arr$ptr, arr$len) EXTERNAL;
11     DECLARE arr$ptr POINTER;
12     DECLARE arr$len WORD;
13     END print$stats;

14     DECLARE buf(80) BYTE;
15     DECLARE console CONNECTION EXTERNAL;
16     DECLARE i BYTE;
17     DECLARE infile CONNECTION EXTERNAL;
18     DECLARE lfreq(26) Freq_Struct;
19     DECLARE num$read BYTE;
20     DECLARE outfile CONNECTION EXTERNAL;
21     DECLARE quit$time BYTE INITIAL(False);
22     DECLARE status WORD;
23     DECLARE total WORD PUBLIC INITIAL (0);

```

Figure 12-1. Source Code for FREQ Module


```

24 $EJECT
25 main:
26 CALL open$files;
27 CALL init$real$math$unit;

28 DO i = 0 to LENGTH(lfreq);
29     lfreq(i).let.low = 0;
30     lfreq(i).let.up = 0;
31     lfreq(i).percent = 0.0;
32 END;
33                                     /*** Now, read the files ***/
34 read$file:DO WHILE (NOT quit$time);
35     num$read = dq$read(infile,@buf,LENGTH(buf),@status);
36     IF num$read <> LENGTH(buf) THEN quit$time = True;

37     DO i = 0 to num$read;
38         total = total + 1;           /*** Total keeps track of ALL
                                     characters ***/
39                                     /*** Read, not just the letters. ***/
40         sh_which_letter:IF (buf(i) >= 'A' AND buf(i) <= 'Z') THEN
41             lfreq(buf(i)-'A').let.up = lfreq(buf(i)-'A').let.up + 1;
42         ELSE IF (buf(i) >= 'a' AND buf(i) <= 'z') THEN
43             lfreq(buf(i)-'a').let.low = lfreq(buf(i)-'a').let.low +
44             1;

44     END;                             /*** Loop i = 0 to num$read ***/

45 read$file:END;

46 stats:
47     CALL print$stats(@lfreq,LENGTH(lfreq) );
48     CALL dq$exit(0);

49 END freq;

```

Figure 12-1. Source Code for FREQ Module (continued)

Line 1 specifies the `DEBUG` control and the pagewidth. The `DEBUG` control instructs the compiler to collect debug information such as the statement number and relative address of each source program module. `PW(75)` specifies an output page 75 characters wide.

Line 2 names the module and establishes the beginning of the module's `DO` block. As stated in Chapter 1, a module must begin with a labeled `DO` statement and end with an `END` statement.

Lines 3 through 6 specify the include files to be used in the program module. Line 4 indicates to the compiler to not list anything until the `LIST` control is encountered, which happens at line 7.

Line 5 is a user comment and will not be interpreted by the compiler. User comment lines begin with a slash/asterisk (`/*`) combination and end with an asterisk/slash (`*/`) combination.

Lines 8 through 23 are the procedure and variable declarations used in the `FREQ` module. Note the `EXTERNAL` declarations in lines 8 through 13. These procedures are declared `EXTERNAL`, which means that the procedure is defined in another module. The calling module must declare the procedure as `EXTERNAL`. The module in which these procedures are defined must declare the procedures as `PUBLIC`.

The variable declarations (see lines 15, 17, and 20) are also `EXTERNAL`. The same rules apply for variables as for procedures. The calling module must declare the variable as `EXTERNAL` and the defining module must declare the variable as `PUBLIC`. If the variable definition is included in the calling module, the definition must be identical to the definition in the declaring module.

Line 18 declares the `lfreq` structure, which is declared in the `defns.inc` file (see Figure 12-4). Line 21 declares `quit$time` as a variable (with the `INITIAL` attribute) of type `BYTE`. In an initialization, the initialization attribute must be placed after the variable attributes. In line 23, `total` is declared as a variable of type `BYTE`. Note also the `PUBLIC` declaration. This indicates that this variable can be used by other modules within the program (if it is declared `EXTERNAL` within the module which uses it).

Line 24 specifies the beginning of a new page (used when the program listing is printed).

The program begins at line 25. Line 26 calls the `open$files` procedure (declared as `EXTERNAL` in line 8). This procedure opens the input file, and if one is specified, the output file. Line 27 calls the compiler built-in procedure, `init$real$math$unit`. This call is required to initialize the `REAL` math facility for subsequent operations.

Lines 28 through 32 consist of more initializations. These lines set (or reset) the values of the structure variable used in the module. `freq_struct` is an array of nested structures (see Chapter 4). `freq_struct` is a 26 element array (one element for each letter in the alphabet). Each element of the `freq_struct` array contains the `let` structure, which consists of a letter and a percent. Nested within the `let` structure is another structure (`low` and `up`). This structure holds the count of uppercase and lowercase characters. To see how `freq_struct` is declared, refer to Figure 12-4.

Lines 34 through 45 show an example of a nested `DO` block. With PL/M, `DO` blocks can be nested up to 18 levels. Line 37 begins a second `DO` block within the `DO` block that begins at line 34. The `DO` block nested within the first `DO` block ends at line 44. The first `DO` block ends at line 45.

Lines 34 through 36 use the UDI function, `dq$read`, to read from a file (`infile`). A specified number of characters are read from the file into an array. The array is `buf` and the number of characters read is `LENGTH(buf)`. The value of `buf` was set in line 14. `LENGTH` is a built-in function (see Chapter 11) that returns the number of elements in an array. The UDI function, `dq$read`, returns the number of characters read (`num$read`) and an error code (`status`).

The nested loop (lines 37 through 44) keeps totals for all the characters read, the uppercase letters read, and the lowercase characters read. This entire loop repeats until the number of characters read in from the input file is less than 80 (this indicates that the input file is empty).

Line 47 calls the external procedure `print$stats`. This procedure is defined in the `PRINT` module. Line 48 calls a UDI procedure, `dq$exit`. Finally, line 49 ends the `FREQ` module.

OPEN Module

The `OPEN` module takes care of the majority of the file-handling procedures for the program. This module makes extensive use of the UDI procedures provided by the run-time support library. The source code is shown in Figure 12-2. Note that the line numbers in the figure are not part of the source code, nor are they the line numbers that the compiler would assign. The line numbers have been added to simplify the discussion of the source code.

```

1  $DEBUG PW(75)

2  open:DO;

3  $NOLIST
4  $INCLUDE (defns.inc)
5  $LIST

6  $EJECT
7  $INCLUDE(udi.inc)

8  $EJECT
9  DECLARE console CONNECTION PUBLIC;
10 DECLARE infile CONNECTION PUBLIC;
11 DECLARE outfile CONNECTION PUBLIC;

12 DECLARE NeedFile(*) BYTE INITIAL('Enter input file name: ');
13 DECLARE OpenError(*) BYTE INITIAL ('Error opening input
    file',CR,LF);

14 open$files:PROCEDURE PUBLIC;
15     DECLARE delim BYTE;
16     DECLARE console$in CONNECTION;
17     DECLARE buffer(80) BYTE;
18     DECLARE status WORD;
19     DECLARE in$buf(81) BYTE;
20     DECLARE i BYTE;
21     DECLARE num$read BYTE;

22     console = dq$create(@4,':CO:'),@status);
23     CALL dq$open(console,WriteOnly,0,@status);

24     /** Process the command line.  It consists of three parts,
25         1) the program name (lf.exe)
26         2) the input file name, if this is not present then
27            ask for it
28         3) the output file name, if this is not present then
29            the output goes to the console ***/

```

Figure 12-2. Source Code for OPEN Module

```

30                                     /*** Read past the program name ***/
31  delim = dq$get$argument(@buffer,@status);
32                                     /*** Find out name of the input file ***/
33  IF delim = CR THEN
34      DO;
35          /*** No input file specified, ask for it ***/
36          CALL dq$write(console,@NeedFile,LENGTH(NeedFile),@status);
37          console$in = dq$attach(@('4,':CI:'),@status);
38          CALL dq$open(console$in,ReadOnly,0,@status);
39          sch001:num$read =
              dq$read(console$in,@in$buf,LENGTH(in$buf),@status);
40          CALL dq$close(console$in,@status);
41
42          /*** Convert the read in buffer to the infile buffer ***/
43          sh_infile:buffer(0) = num$read;
44          DO i = 0 to num$read;
45              IF (in$buf(i) <> CR) AND (in$buf(i) <> LF)
46                  THEN buffer(i+1) = in$buf(i);
47              ELSE
48                  buffer(0) = buffer(0) - 1;    /*** Adjust count for
49                                                  CR/LF ***/
50          END;    /*** End of DO loop to Convert buffer ***/
51      ELSE
52          delim = dq$get$argument(@buffer,@status);
53
54          /*** END; get file name to process ***/
55
56          /*** Open input file ***/
57          infile = dq$attach(@buffer,@status);
58          CALL dq$open(infile,ReadOnly,2,@status);
59          IF status <> E$OK THEN DO;
60              CALL dq$write(console,@OpenError,LENGTH(OpenError),
61                  @status);
62              CALL dq$exit(1);
63          END;    /** Status is not ok **/

```

Figure 12-2. Source Code for OPEN Module (continued)

```
60         /*** Find out if an output file was specified.  If so, ***/
61             /*** open it, if not use the console output ***/
62     IF delim = CR THEN
63         outfile = console;
64     ELSE DO;
65         delim = dq$get$argument(@buffer,@status);
66         outfile = dq$create(@buffer,@status);
67         CALL dq$open(outfile,WriteOnly,2,@status);
68         END;
69     END open$files;
70     END open;
```

Figure 12-2. Source Code for OPEN Module (continued)

Line 1 instructs the compiler to collect debug information and sets the page width for printed output. Line 2 names the module and establishes the beginning of the module's DO block. Lines 3 through 8 specify the inclusion of the program's include files, turn the listing function on and off, and specify a few new pages for printed output (`$EJECT`).

Lines 9 through 11 define and declare some `PUBLIC` variables. Because these variables are declared `PUBLIC`, they can be used in another module. The calling module must declare the variable as `EXTERNAL`. The variable definition is included in the calling module, and it is the same as the definition in the defining module.

Lines 12 and 13 are error messages to be used by the `OPEN` module if the necessary information is not included in the invocation line (which causes an error). Note the use of the asterisk in each of these lines. The asterisk is used as an implicit dimension specifier. The implicit dimension specifier can be used when the size of the array is either unknown or insignificant. In this instance, the size of the array is unknown. The implicit dimension specifier in lines 12 and 13 specifies that the `NeedFile` array and the `OpenError` array will have the same number of elements as the value list (the number of characters in the message).

Line 14 begins the `open$files` procedure. This procedure is declared as `PUBLIC` (it is called by the `FREQ` module) and continues until the end of the module (line 69).

Lines 22 and 23 get and open a connection with the console using predefined UDI procedures. Note the use of the `@` operator in these two lines. The first `@` operator in line 22 allocates storage for the constants 4 and `:CO:`. The other `@` operators are for location references. This means that the value of the reference (e.g., the value of `@status`) is the actual run-time location of the variable.

Lines 31 through 51 use the UDI procedure, `dqgetargument`, to parse the input line. Line 31 gets the first part of the command line, as well as the delimiter used to separate this part of the command line from the next part (if there is any). Line 33 tests the delimiter. If the delimiter is a carriage return then lines 34 through 49 are processed. Lines 34 through 49 request a file name. If the delimiter is not a carriage return then `dqgetargument` is called again. This routine is also called by lines 62 through 68 to determine whether the program output should go to a file or to the console.

Line 31 passes the invocation line to the following `IF/THEN/ELSE` construct (lines 33 through 51). The `IF/THEN/ELSE` construct checks for an input file name. If no input file is specified, line 36 uses the `NeedFile` string declared in line 12. This prompts the user to enter an input file name. If no input file name is specified in response to the prompt, the program aborts. Otherwise, the string is converted as discussed in the preceding paragraph.

Lines 43 through 48 convert the file name to a UDI call.

Lines 50 and 51 are the `ELSE` clause of `IF delim = CR`.

Lines 53 through 59 open the input file. Lines 62 through 68 open an output file, if one is specified. Otherwise, the program data is sent to the console.

Line 69 is the `END` statement for the `open$files` procedure and line 70 is the `END` statement for the `OPEN` module.

PRINT Module

The `PRINT` module performs the program calculations and prints the information (either to the console or to the specified output file). The source code is shown in Figure 12-3.


```

1  $DEBUG PW(75)
2  print:DO;

3  $NOLIST
4  $INCLUDE (defns.inc)
5  $INCLUDE (udi.inc)
6  $LIST

7  DECLARE BLANK$OUT$LINE LITERALLY
8      'DO j = 0 TO LENGTH(line);line(j) = SPACE;END';
9  DECLARE LETTER LITERALLY '3';
10 DECLARE LOWER LITERALLY '24';
11 DECLARE PCT LITERALLY '33';
12 DECLARE SUM LITERALLY '8';
13 DECLARE UPPER LITERALLY '16';

14 DECLARE outfile CONNECTION EXTERNAL;
15 DECLARE topline(*) BYTE INITIAL
16     ('LETTER TOTAL UPPER LOWER % ',CR,LF);
17     /** ( A 00000 00000 00000 000.0 ***/
18     /** ( 123456789 123456789 123456789 123456789 123456789 ***/
19 DECLARE total WORD EXTERNAL;
20 DECLARE total$str (5) BYTE INITIAL ('TOTAL');

21 int2asc:PROCEDURE(number,stg$ptr,count) BYTE;
22     DECLARE number WORD;
23     DECLARE stg$ptr POINTER;
24     DECLARE count BYTE;

25     DECLARE i BYTE, j BYTE;
26     DECLARE max DWORD;
27     DECLARE string BASED stg$ptr(1) BYTE;
28     DECLARE tmpstg(10) BYTE;

29     max = 1;
30     DO i = 1 TO count;
31         max = 10 * max;
32     END;
33     max = max - 1;
34     DO i = 0 TO LAST(tmpstg);
35         tmpstg(i) = SPACE;
36     END;

```

Figure 12-3. Source Code for PRINT Module

```

37     IF number <= max THEN DO;
38         i = 0;
39         loop:
40             tmpstg(i) = (number MOD 10) " '0' ;
41             i = i " 1;
42             number = number/10;
43             IF number 0 THEN GOTO loop;

44         DO j = 0 TO count;
45             string(count-j) = tmpstg(j);
46             END;
47         END;
48     ELSE DO;
49         DO i = 0 to count;
50             string(i) = '*';
51             END;
52     END;

53     RETURN(i);
54     END int2asc;

55 real2asc:PROCEDURE(number,stg$ptr,count);
56     DECLARE number REAL;
57     DECLARE stg$ptr POINTER;
58     DECLARE count WORD;

59     DECLARE i BYTE, j BYTE;
60     DECLARE int$len BYTE;
61     DECLARE string BASED stg$ptr(1) BYTE;
62     DECLARE tmpnum DWORD;
63     DECLARE tmpstg(10) BYTE;

64         /*** Convert the number to an INTEGER to convert
65             it, assume one ***/
66             decimal place ***/

66     tmpnum = DWORD(number*10.0);

```

Figure 12-3. Source Code for PRINT Module (continued)

```

67  int$len = int2asc(tmpnum,@tmpstg, LAST(tmpstg) );
68  IF int$len = 1 THEN DO;          /*** Handle the case where
                                   the number ***/
69                                   /*** is less than 1.0 ***/
70      int$len = 2;
71      tmpstg(LAST(tmpstg)-1) = '0';
72      END;
73  DO i = 0 TO int$len-2;
74      string(count-i) = tmpstg(LAST(tmpstg)-i);
75      END;
76  string(count-int$len) = '.';
77  string(count-int$len-1) = tmpstg(LAST(tmpstg)-int$len+1);
78  END real2asc;
79  $EJECT
80  print$stats:PROCEDURE (arr$ptr, arr$len) PUBLIC;
81      DECLARE arr$ptr POINTER;
82      DECLARE arr$len WORD;
83      DECLARE array BASED arr$ptr(1) Freq_Struc;
84      DECLARE i BYTE, j BYTE;
85      DECLARE line(50) BYTE;
86      DECLARE status WORD;
87      DECLARE tmp BYTE;
88      DECLARE ii BYTE;
89
89      call dq$write(outfile,@topline,LENGTH(topline),@status);
90
90      printlines:DO ii = 0 TO arr$len-1;
91
91          BLANK$OUT$LINE;
92          line(LETTER) = ii + 'A';
93          /*** Get the total and convert number to ascii ***/
94          tmp = int2asc (array(ii).let.low + array(ii).let.up),
94          @line(SUM),5);
95          tmp = int2asc (array(ii).let.low, @line(LOWER),5);
96          tmp = int2asc (array(ii).let.up, @line(UPPER),5);
97
97          array(ii).percent = REAL((array(ii).let.low +
97          (array(ii).let.up)) /
98              REAL(total) * 100.0;
99  CALL real2asc (array(ii).percent, @line(PCT),5);

```

Figure 12-3. Source Code for PRINT Module (continued)

```

100     line(LAST(line)-1) = CR;
101     line(LAST(line)) = LF;

102     CALL dq$write(outfile,@line,LENGTH(line),@status);

103     END printlines;                               /*** print loop ***/

104     BLANK$OUT$LINE;
105     DO i = 0 TO LAST(total$str);
106         line(LETTER-2"i) = total$str(i);
107     END;

108     tmp = int2asc( total, @line(SUM),5);

109     call dq$write(outfile,@line,LENGTH(line),@status);

110 END print$stats;

111 END print;

```

Figure 12-3. Source Code for PRINT Module (continued)

Line 1 instructs the compiler to collect debug information and sets the page width for printed output. Line 2 names the module and establishes the beginning of the module's DO block. Lines 3 through 6 specify the inclusion of the program's include files and turn the listing function on and off.

Lines 7 through 13 are a group of `literally` definitions; each one creates an alternate name for a sequence of characters. Lines 7 and 8 declare `BLANKOUTLINE` as the alternate name for the DO loop used to blank out the output line buffer. Additionally, after line 13, the number 16 will reference `UPPER` (for uppercase character). This is a useful function to eliminate keystrokes, to make the program more readable, and to declare quantities that may be fixed in one module, but subject to change in another module.

Lines 14 through 20 contain more declarations, as well as the header string for the output (line 16).

Lines 21 through 54 perform an integer-to-ASCII translation. Lines 55 through 78 convert real numbers to ASCII characters.

Line 80 is the beginning of the `print$stats` procedure. The `print$stats` procedure is called by the `FREQ` module, therefore it is declared `PUBLIC` in this module. Note the based variable in line 83. In this instance, the location of array is based on the address of `arr$ptr`, which is passed into the `print$stats` procedure. The size of the array is unknown (except through the parameter). The 1 enclosed in parentheses enables the use of `arr$ptr` as an array (any number can be used).

Line 89 calls a UDI procedure that writes to an external connection declared in the `OPEN` module. Note the use of `BLANKOUTLINE` in line 91.

Lines 90 through 103 are a `DO` loop that is repeated for each letter in the alphabet. For each character, the `line(LETTER)` array is filled with the letter, the total, the total uppercase, the total lowercase, and the percent. This information is then sent to the specified output device (the console or a file).

Lines 93 through 96 call the procedure to convert the total into ASCII characters. Lines 97 and 98 figure the percentage of use for each character. Line 99 calls the procedure to convert the percentage to ASCII characters. Lines 100 and 101 insert a carriage return and a line feed in the console display or in the output file.

Line 110 ends the `print$stats` procedure and line 111 ends the `PRINT` module.

Include Files

As stated earlier, there are two include files with this program (see Figures 12-4 and 12-5).

```
DECLARE DCL LITERALLY 'DECLARE';
DCL LIT          LITERALLY 'LITERALLY';

DCL CR          LITERALLY '0DH';
DCL LF          LITERALLY '0AH';

DCL True        LITERALLY '0FFH';
DCL False       LITERALLY '000H';

DCL Freq_Struc LITERALLY 'STRUCTURE (let STRUCTURE
                               (low WORD, up WORD),
                               percent REAL)';
DCL SPACE          LITERALLY '020H';
```

Figure 12-4. Include File -- defns.inc

Figure 12-4 is the *defns.inc* file. It contains definitions for terms used in common by all of the modules in the program (excluding the UDI definitions). Note the declaration of a structure in this include file (Freq_Struc). This structure is used in the PRINT module and the FREQ module. This structure declaration illustrates several levels of nesting. Structures can be nested up to 32 levels.

Figure 12-5 is the *udi.inc* file. It contains UDI definitions that are used throughout the modules. The UDI is a predefined set of procedure calls that enables use of operating system functions.

```

DECLARE CONNECTION literally 'WORD';
DECLARE ReadOnly LITERALLY '1';
DECLARE WriteOnly LITERALLY '2';
DECLARE E$OK LITERALLY '0H';

dq$attach:procedure (path$p,except$p) CONNECTION external;
    declare path$p pointer; declare except$p pointer;
    end dq$attach;

dq$close:procedure (aftn,exception$ptr) external;
    declare aftn CONNECTION, exception$ptr pointer;
    end dq$close;

dq$create:procedure (path$p,exception$ptr) CONNECTION external;
    declare (path$p,exception$ptr) pointer;
    end dq$create;

dq$exit:procedure (completion$code) external;
    declare completion$code word;
    end dq$exit;

dq$get$argument:PROCEDURE (arg$ptr, ex$ptr) BYTE EXTERNAL;
    declare arg$ptr POINTER, ex$ptr POINTER;
    END dq$get$argument;

dq$open:procedure (aftn,mode,num$buf,exception$ptr) external;
    declare aftn CONNECTION, exception$ptr pointer;
    declare (mode,num$buf) byte;
    end dq$open;

dq$read:PROCEDURE(aftn,buf$ptr,count,ex$ptr) WORD EXTERNAL;
    declare aftn CONNECTION;
    declare buf$ptr POINTER;
    declare count WORD;
    declare ex$ptr POINTER;
    END dq$read;

dq$write:procedure (aftn,buffer,count,exception$ptr) external;
    declare aftn CONNECTION;
    declare count word;
    declare (buffer,exception$ptr) pointer;
    end dq$write;

```

Figure 12-5. Include File -- udi.inc



Overview

Program segmentation is the division of a program into memory segments. It is a technique used to optimize the code produced by the compiler. The segmentation controls (`COMPACT`, `LARGE`, `MEDIUM`, `SMALL`, and `FLAT`) manage program segmentation by defining the physical relationship in memory of a program's code, data, constants, and stack. They determine which (if any) segments get combined. For example, specifying the `SMALL` segmentation control for a program module locates all of the module's code, data, constants, and stack in two segments, `CODE` and `DATA`. When the program's modules are combined, sections from the separately compiled modules are combined into segments according to the specified segmentation controls. This optimizes code because references to locations in the same memory segment are more efficient.

Extended segmentation models are a super-set of the segmentation controls. The extended segmentation models (which consist of the `SMALL`, `COMPACT`, and `LARGE` subsystems) provide enhanced program speed and aid in the construction of large programs. An extended segmentation model consists of a number of subsystems. A subsystem is a collection of program modules that use the same segmentation controls. A program is made up of one or more subsystems. With subsystems, program modules that are compiled with different segmentation controls can be combined.

This chapter defines the use of extended segmentation models, and contains the following sections:

- Introduction
- Segmentation controls architecture overview
- Using subsystems
- Syntax
- Exporting procedures
- Large matrix example

Introduction

Extended segmentation models provide the following programming advantages:

- Efficient use of memory.
- Access to the microprocessor's segmented architecture.
- Storage reduction for external references to pointers and code.
- Increased program execution speed for intersegment calls and data access.

Additionally, to simplify the development of large programs, the segmentation controls can be used to partition the program into a collection of related subsystems.

Partitioning a large program into a series of subsystems isolates code references within the same segment. The compiler processes each program module individually, assigning code, data and stack segments for each module (according to the specified segmentation control). As a source file is translated, the compiler generates a `STACK` segment for the program stack, as well as a `DATA` segment for the program data and a `CODE` segment for the program's executable code. When the program modules are combined, the `CODE`, `DATA` and `STACK` segments from all of the individual program modules are combined. Use of the segmentation controls ensures that the segment names generated by the compiler are combined according to the overall structure of the program.

A subsystem is either open or closed. An extended segmentation model can have only one open subsystem, but any number of closed subsystems.

An open subsystem does not have a name and claims the program modules that are not claimed by another subsystem. Effectively, a program that uses only the segmentation controls is an open subsystem. Modules can be added to the open system without having to change the subsystem definition.

A closed subsystem has a name and, optionally, a list of program modules used in the subsystem. To add a module to a closed subsystem, the subsystem definition must be changed.

Segmentation Controls Architecture Overview

The segmentation controls described in Chapter 11 define the physical relationship in memory of program code, data, constants, and stack during program execution. When a PL/M source file is compiled, an object module conforms to a particular extended segmentation model.

There are three extended segmentation models: `SMALL`, `COMPACT`, and `LARGE`. For Intel386 and Intel486 microprocessors, each segment can be as large as 4G bytes.

There are two submodels within each model: `RAM` and `ROM`. Specifying `RAM` places the program constants in the `DATA` segment. Specifying `ROM` places the program constants in the `CODE` segment.

Tables 13-1 and 13-2 define the memory partitions and the placement of pointers in the various architectural models available with the segmentation controls. Table 13-1 shows how memory is partitioned. Table 13-2 defines the register addresses and the pointer values. Table 13-3 defines the register addresses and the pointer values for the Intel386 and Intel486 microprocessor-specific `ES` register. Note that the `POINTER` variable value for these microprocessors, when using the `SMALL ROM` extended segmentation controls, is 6 bytes.

Table 13-1. Segmentation Controls and Memory Partitions

Control	CODE	Segment Name	
		DATA	STACK
SMALL RAM	code	data constants stack	
SMALL ROM	constants	data code	stack
COMPACT RAM	code constants	data	stack
COMPACT ROM	constants code	data	stack
MEDIUM RAM*	separate CODE segment for each module's code	data constants stack	
MEDIUM ROM*	separate CODE segment for each module's code and constants	data stack	
LARGE RAM*	separate CODE segment for each module's code	separate DATA segment for each module's data and constants	stack
LARGE ROM*	separate CODE segment for each module's code and constants	separate DATA segment for each module's data	stack

* The Intel386 and Intel486 microprocessors use only the SMALL and COMPACT segmentation controls. For the segmentation controls (not subsystems), MEDIUM is equivalent to SMALL and LARGE is equivalent to COMPACT.

Table 13-2. Segmentation Controls, Register Addresses and Pointer Values

Control	Register Address			Pointer Variable Value
	CS	DS	SS	
SMALL RAM	CODE seg. Offset-reference relative to DS	DATA seg. Offset-reference relative to DS	DATA seg. Has same value as DS Offset-reference	4-byte offset only
SMALL ROM	CODE seg. Constant reference requires selector-offset containing CS value and offset within CODE segment Code reference requires offset-reference relative to DS	DATA seg. Offset reference	DATA seg. Has same value as DS Offset-reference	6-byte selector-offset
COMPACT RAM	CODE seg. Selector-offset reference	DATA seg. Selector-offset reference	STACK seg. Selector-offset reference	6-byte selector-offset
COMPACT ROM	CODE seg. Selector-offset reference	CODE seg. Selector-offset reference	STACK seg. Selector-offset reference	6-byte Selector-offset
MEDIUM RAM	Current CODE seg. Selector-offset reference Updated when PUBLIC or EXTERNAL procedure is activated	DATA seg. Selector-offset reference	DATA seg. Selector-offset reference	6-byte Selector-offset

continued

**Table 13-2. Segmentation Controls, Register Addresses and Pointer Values
(continued)**

Control	Register Address			Pointer Variable Value
	CS	DS	SS	
MEDIUM ROM	Current CODE seg. Selector-offset reference Updated when PUBLIC or EXTERNAL procedure is activated	DATA seg. Selector-offset reference	DATA seg. Selector-offset reference	6-byte Selector-offset
LARGE RAM	Current CODE seg. Selector-offset reference Updated when PUBLIC or EXTERNAL procedure is activated	Current DATA seg. Selector-offset reference Updated when PUBLIC or EXTERNAL procedure is activated	STACK seg. Selector-offset reference	6-byte Selector-offset
LARGE ROM	Current CODE seg. Selector-offset reference Updated when PUBLIC or EXTERNAL procedure is activated	Current CODE seg. Selector-offset reference Updated when PUBLIC or EXTERNAL procedure is activated	STACK seg. Selector-offset reference	6-byte Selector-offset

The values given in Tables 13-1 and 13-2 are identical for Intel386 and Intel486 microprocessors. Additionally, these microprocessors have the ES register address. Table 13-3 states the values for the ES register.

Table 13-3. Intel386 and Intel486 Microprocessor-specific ES Register Segmentation Controls, Register Addresses and Pointer Values

Control	ES Register Address	POINTER Variable Value
SMALL RAM	DATA seg. Offset reference	4-byte offset only
SMALL ROM	DATA seg. Offset reference	6-byte selector offset
COMPACT RAM	DATA seg. Selector-offset reference	6-byte selector-offset
COMPACT ROM	DATA seg. Selector-offset reference	6-byte selector-offset

The `SMALL RAM` segmentation control is the most efficient. Because all of the code resides in one segment, jumps and calls are always within the same segment (intra-segment). However, the `SMALL RAM` segmentation control provides less protection and cannot be used to pass pointers to library procedures unless the library procedure is also a `SMALL RAM` model.

Use the `COMPACT` segmentation controls (`COMPACT RAM` and `COMPACT ROM`) for separate management of the code, data, and stack, or to improve segment-limit protection. To reference stack-based variables, the `COMPACT` segmentation controls use selector-offset references. This is less efficient than using offset-only references. However, data and constant references within a `COMPACT` segmentation control module are within the same segment (intra-segment). Note that if a `COMPACT` program must pass a data address to a procedure in a different subsystem, it must use a selector-offset reference.

Using Subsystems

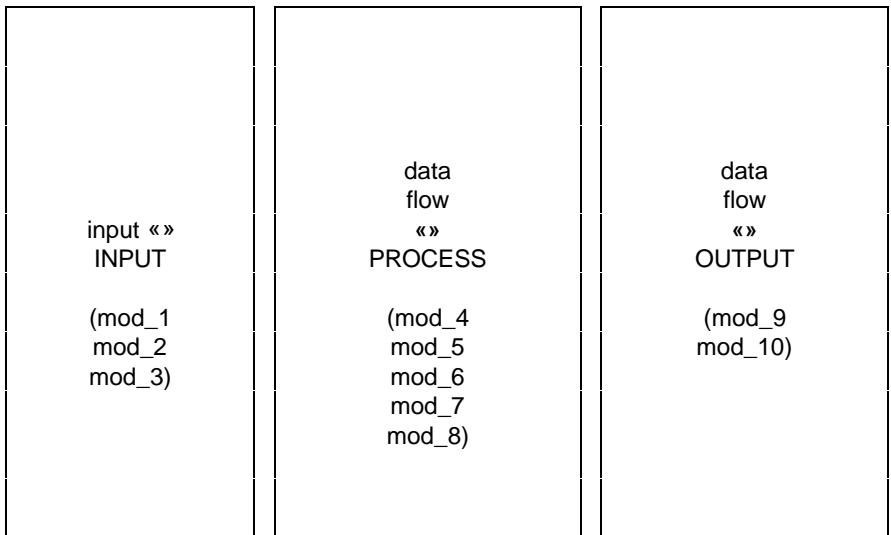
Subsystems offer an efficient way to manage programs with large amounts of data, to share data between program modules, and to communicate with other programs.

For example, subsystems are useful when several programmers are each writing a separate module for a highly structured program in which sharing data between modules is accomplished with parameter passing, by value only. To maintain the integrity of each section's data requires that each section have its own `DATA` segment. In this way, code in one module of the program cannot mistakenly destroy data belonging to another section of the program. In this instance, each module could be a `COMPACT` subsystem, with its own `CODE` and `DATA` segments.

As another example, a program performing I/O usually requires operating system support routines. In many cases, the operating system will operate at a higher protection level than the application program. Thus, operating system procedure calls are intersegment calls. The application program views the operating system as a separate subsystem. Usually, operating system interface libraries are supplied to application programmers; these libraries perform the inter-subsystem communication details. If a program needs to make a direct operating system call without using a presupplied library, the program itself must define the necessary subsystem environments at compile time.

It is usually more efficient to structure a large program with subsystems. With subsystems the code and data can be partitioned into manageable pieces bigger than one module. Within each subsystem, calls and jumps are near (4 byte offset), references can be offset only, and the data of each subsystem is protected from being overwritten by other subsystems. Calls and jumps between subsystems are still far, and references between subsystems need to be selector-offset. In general, a program's structure is such that it is possible to break the program into pieces with a minimum number of intersegment calls, jumps, and references.

For example, consider a program consisting of 10 modules, mod_1 through mod_10. Modules 1 through 3 deal with input and initial processing. Modules 4 through 8 do the main data processing. Modules 9 and 10 output the data. The following figure illustrates the structure of the program:



The total code space required by this program exceeds 64K bytes, and the total data space also exceeds 64K bytes. The `LARGE` segmentation control can be used. This control provides each module with its own `CODE` and `DATA` segment. For this example, this results in a total of 21 segments (10 `CODE`, 10 `DATA`, and 1 `STACK`). For the `LARGE` segmentation control, all calls and jumps are far, and all intermodule references must be through selector-offset `POINTERS`.

If, for example, `COMPACT` subsystems are used instead of the `LARGE` segmentation control, modules 1 through 3 can form one subsystem, which you could call `SUB_INPUT`. Modules 4 through 8 can form subsystem `SUB_PROCESS`. Finally, modules 9 and 10 can form subsystem `SUB_OUTPUT`. The number of segments has been reduced to seven: 3 `CODE`, 3 `DATA`, and 1 `STACK`. Since most of the calls, jumps, and references now take place within only one of the subsystems, the program is much more efficient. The only far calls and jumps, and the only selector-offset references needed are those in the interfaces between the subsystems.

A typical program does not require subsystems. The code space of 4 Gigabytes and the data space of 4 Gigabytes is quite sufficient for most programs. However, consider a program that processes a large amount of data such as a $10 \times 1,000,000,000$ `REAL` matrix. A `REAL` scalar consists of 4 bytes, so the total memory needed is 40 billion bytes. Rows could be used to partition the matrix. Each row would be 4 billion bytes, which would fit into a single `DATA` segment.

Ten `COMPACT` subsystems (named `ROW1`, `ROW2`, etc.) could be created, each containing a 1-billion element `REAL` array. Procedures to store and retrieve particular matrix elements can be written and called from the normal matrix processing code. An example of such a program is shown later in this chapter.

It is not just dividing a program into subsystems that increases its efficiency. If all the even numbered modules had been placed in one subsystem, for instance, and all the odd numbered ones into another, the efficiency of the program would not have improved as it did when the modules were grouped into subsystems according to the logical structure of the program.

Note also the following points:

1. Not all subsystems must use the same segmentation control. For instance, if `SUB_PROCESS` in the preceding example is small enough, it could be a `SMALL` subsystem.
2. If a `SMALL` subsystem is mixed with subsystems using other segmentation controls, the main program must be in `SMALL`. This is because anything compiled in `SMALL` assumes that `DS` and `SS` are identical. This will be so only if the main program is `SMALL`. Notice that in this case, the `STACK` segment resulting from the `COMPACT` and `LARGE` subsystems will not be used, since the stack of the main program is in the combined `DATA-STACK` segment of the `SMALL` model.
3. `SMALL` RAM subsystems have the limitation that the `SMALL` segmentation control uses short (offset only) pointers. A `SMALL` RAM subsystem cannot receive a pointer from another subsystem, because it cannot save the selector portion. A `SMALL` RAM subsystem can, however, pass a pointer to a subsystem that is not `SMALL` RAM, because its own `DS` is known to it. However, a `SMALL` RAM subsystem cannot pass a pointer (which points to a procedure), since `DS` is assumed as the selector to all pointers.
4. `MEDIUM` is a segmentation control only, not an extended segmentation model.

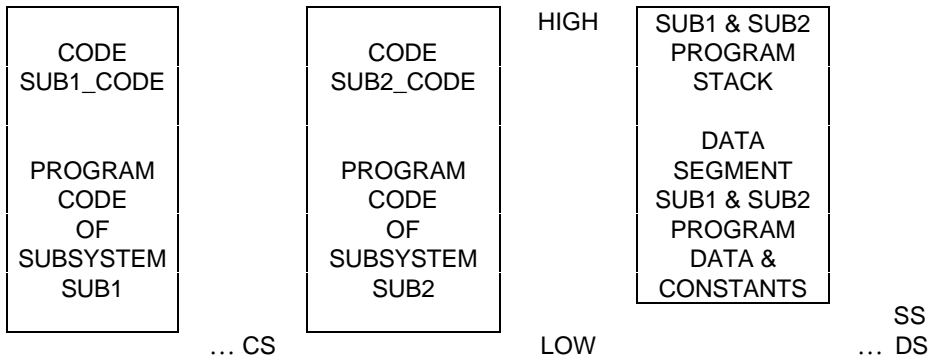
Later sections describe the memory layouts of programs using the standard segmentation controls: `FLAT` | `COMPACT` | `LARGE` | `MEDIUM` | `SMALL`. To understand the memory layouts of programs structured with subsystems, it is necessary to make the distinction between compiling modules and combining modules into a program.

The compiler compiles only one module at a time. When modules are combined into a program, many `CODE`, `DATA` and `STACK` segments, which were generated during separate compilations, are combined. When combining program modules, all segments with the same name are combined. The segmentation controls work by controlling the names of the segments generated by the compiler. This ensures that the segment names will be combined as desired when the modules are combined into a program.

The standard `SMALL` segmentation control causes the compiler to name the `CODE` segment `CODE`, and the `DATA-STACK` segment `DATA`. Since under the standard `SMALL` model all `CODE` segments have the same name, and all `DATA-STACK` segments have the same name, they are combined when the modules are combined.

A module belonging to a `SMALL` subsystem, on the other hand, takes the name of its `CODE` segment from the name of the subsystem. The name of its `DATA-STACK` segment is still `DATA`. Thus, a `SMALL` subsystem named `SUB1` contains one `CODE` segment named `SUB1_CODE`, and one `DATA-STACK` segment named `DATA`. A `SMALL` subsystem named `SUB2` contains one `CODE` segment named `SUB2_CODE`, and one `DATA-STACK` segment named `DATA`. When the program modules are combined, all segments with the same name are combined.

The memory layout of the loaded program containing the two subsystems `SUB1` and `SUB2` is as follows (it is assumed that both subsystems are `SMALL` `RAM`):



Note that a program using the `MEDIUM` segmentation control is equivalent to a program in which each module is declared to be in a unique `SMALL` subsystem.

A module belonging to a `COMPACT` subsystem takes the name of its `CODE` segment and the name of its `DATA` segment from the subsystem name. So a `COMPACT` subsystem named `SUB1` contains one `CODE` segment named `SUB1_CODE`, one `DATA` segment named `SUB1_DATA`, and one `STACK` segment named `STACK`. A `COMPACT` subsystem named `SUB2` contains one `CODE` segment named `SUB2_CODE`, one `DATA` segment named `SUB2_DATA`, and one `STACK` segment named `STACK`. The loaded program will contain five segments, two `CODE` segments, two `DATA` segments, and one `STACK` segment. Note that a program using the `LARGE` segmentation control is equivalent to a program in which each module is declared to be in a unique `COMPACT` subsystem.

A `LARGE` subsystem can be simulated by a `COMPACT` subsystem containing only one module. However, `LARGE` subsystems are useful for the following reason. A `LARGE` subsystem named `SUB1`, which contains the modules `MOD1`, `MOD2`, and `MOD3`, has three `CODE` segments named `MOD1_CODE`, `MOD2_CODE`, and `MOD3_CODE`, and three `DATA` segments named `MOD1_DATA`, `MOD2_DATA`, and `MOD3_DATA`. As usual, it contains one `STACK` segment named `STACK`. It is possible to use a `LARGE` subsystem instead of inventing names for many `COMPACT` subsystems, each containing only one module. Note that the segment name in the `LARGE` subsystem is derived from the module names and not from the subsystem name.

The `LARGE` segmentation control is identical to the `COMPACT` segmentation control. However, there is a difference between `LARGE` and `COMPACT` subsystems. In a `LARGE` subsystem, the external definition of all symbols in the `EXPORTS` list have their segment field set to an unknown value. This enables the creation of external far objects with public locations that are unknown at compile time. In all other respects, a `LARGE` subsystem is identical to a `COMPACT` subsystem.

Open Subsystems

Compiling files using only the segmentation controls and using no other subsystem controls produces open subsystems. When object modules are combined, all modules created from compilations specifying a particular segmentation control are automatically combined. Segments are created according to the rules for the segmentation control. A list of modules belonging to an open subsystem is therefore not needed at compile time. Modules can be freely added to or deleted from an open subsystem at any time during program development.

Note that both `RAM` and `ROM` modules are combined into the single open subsystem. For a `SMALL` subsystem, be careful when combining `RAM` and `ROM` modules, particularly concerning the passing of pointer parameters and the accessing of constants not in the current module.

It is not possible to pass pointer parameters between `SMALL RAM` and `SMALL ROM` modules, because pointers are defined differently in each submodel. Also, it is not possible to directly reference constants defined in a `ROM` module from a `RAM` module, and vice versa, because `RAM` modules define constants to be in the data segment, and `ROM` modules define constants to be in the code segment.

In the `COMPACT` model, passing pointer parameters between `RAM` and `ROM` modules is not a problem, because pointers are always long. As in `SMALL`, the restriction on direct reference to constants applies.

The names of the segments in both `SMALL` and `COMPACT` models are identical: `CODE32` for the code segment, `DATA` for the data segment. This means that if `SMALL` and `COMPACT` modules are combined, they will also be combined to form a single open subsystem consisting of the `CODE32`, `DATA`, and `STACK` segments. Care must be taken regarding stack references, because `COMPACT` defines a separate stack segment and `SMALL` does not. For more information on Intel386 microprocessor segment combining, see the binder chapter in the *Intel386 Family Utilities User's Guide*.

Closed Subsystems

A closed subsystem differs from an open subsystem in two ways: it has a name and it consists of a specific list of modules. The compiler must know the name of the subsystem and the modules belonging to the subsystem in order to create a closed subsystem.

The need for a closed subsystem name is simply to differentiate a particular closed subsystem from another closed subsystem or from the open subsystem. This is done as follows: the name of the subsystem is added to the beginning of the segment names to create unique code and data segments.

For example, if a subsystem is named `PHASE1`, then the code sections from all modules belonging to the `PHASE1` subsystem are combined into a single code segment called `PHASE1_CODE32`; similarly for `COMPACT` subsystems the data sections are combined into a single data segment called `PHASE1_DATA`. When using `COMPACT`, however, the stack sections are still combined into a segment called `STACK` because only one execution-time stack is usually necessary. Using `SMALL` all data and stack segments are combined in one segment called `DATA`, as usual.

A closed subsystem module list is needed for differentiation. For instance, if the compiler is not informed that module `SCANNER` belongs to subsystem `PHASE1`, then the compiler has no choice but to assume that module `SCANNER` belongs to the open subsystem.

Thus, every module in a program either is specified as part of a closed subsystem or, by default, becomes part of the open subsystem. A program can consist of only closed subsystems, or of both closed subsystems and the open subsystem, or of only the open subsystem (by default). There is only one open subsystem per program; all open subsystems are treated as one subsystem by the utility used to combine the program modules.

Communication Between Subsystems

Within a subsystem there can be code and/or data items (procedures and variables) that must be known by other subsystems; that is, they are meant to be referenced from other subsystems. Such items are said to be exported. The export of a symbol is not directed at any one particular subsystem; it is directed at all subsystems in the program, including its own subsystem.

It is important to realize that the subsystem definitions are additions to normal intermodule `PUBLIC/EXTERNAL` definitions, not replacements.

For instance, module `MOD1` belongs to subsystem `SUB1` and makes a reference to symbol `SYM2`; `SYM2` belongs to subsystem `SUB2`. `SYM2` must be declared as `EXTERNAL` in `MOD1`, as usual, and must also be declared as `PUBLIC` and exported from `SUB2`. Using this information, the compiler generates an intersegment reference to `SYM2`.

Syntax

Defining subsystems means telling the compiler what extended segmentation model each subsystem uses, and which modules belong to each subsystem. In addition, it means telling the compiler which procedures and data are accessible from outside the subsystem.

Making everything available to all subsystems defeats the purpose of subsystems. For example, if a procedure is declared to be accessible from outside the subsystem, it is a far procedure. This means that all calls are far calls, even if the procedure is never actually accessed from outside its subsystem.

Each subsystem in a PL/M program has one extended segmentation model definition, which takes one of the following forms:

1. `$ model (subsystem-id [submodel] [x])`
2. `$ model ([submodel] [x])`
3. `$ model (submodel [x])`

where `[x]` is of the form;

```
[HAS module-list]
```

or

```
[HAS module-list; EXPORTS public-list]
```

or

```
[EXPORTS public-list]
```

Where:

model is SMALL, COMPACT, or LARGE and specifies the extended segmentation model for the subsystem. All modules in the subsystem must be compiled with the same extended segmentation model.

submodel is -CONST IN CODE- or -CONST IN DATA- and specifies the placement of constants. -CONST IN CODE- corresponds to the ROM submodel; -CONST IN DATA- corresponds to the RAM submodel.

The default depends on the segmentation control and corresponds to the defaults of RAM|ROM for each model. The use of the RAM and ROM controls (see Chapter 11) can create conflicts when subsystems are defined. RAM is specified by -CONST IN DATA-; ROM is specified by -CONST IN CODE-.

subsystem-id is any PL/M identifier that can be used as a module name, and specifies the name of the subsystem. This ID does not conflict with any IDs used within the program. A subsystem control without *subsystem-id* defines the open subsystem.

HAS *module-list* is a list of module names, separated by commas, specifying the modules belonging to the subsystem. These module names must exactly match the module names from each source file comprising the subsystem. (A module name is the name of the outermost DO block of a source file.) A particular module name can appear in only one *module-list*. There are no default modules in the *module-list*. Any module for which a name does not appear in a *module-list* becomes part of the open subsystem.

EXPORTS *public-list* is a list of procedure, variable, and constant IDs, specifying the code and data objects exported by the subsystem (i.e., accessible outside of the subsystem). Using a dollar sign (\$) in a procedure name within a subsystem definition will cause an error. Any symbol in the exports list may be declared PUBLIC in at most one of the modules belonging to the subsystem, and should be declared EXTERNAL in all modules in and out of the subsystem that access the symbol.

A particular exported symbol can appear in only one *public-list*.

The *public-list* is exhaustive. Only the symbols in the *public-list* can be referenced from other subsystems. Symbols in the subsystem declared PUBLIC but not appearing in the *public-list* are accessible only from within the subsystem itself. Conversely, PUBLIC symbols that are not intended to be referenced from outside the subsystem should not appear in the *public-list*. These symbols are called domestic symbols.

In most applications of the subsystem controls, the HAS and EXPORTS lists will have several dozen entries apiece. To accommodate lists of this length, a subsystem control may be continued over more than one control line. (The continuation lines must be contiguous, and each must begin with a dollar sign (\$) in the first column.) Keep in mind that using a dollar sign in a procedure name within a subsystem definition will cause an error. Also, note that any number of HAS and EXPORTS lists can appear in a control, in any order. This enables formatting of the subsystem specification so it can be easily read and maintained.

Consider the following subsystem definition:

```
$COMPACT(SUB_INPUT -CONST IN CODE- HAS mod_1, mod_2, mod_3; EXPORTS input)
$SMALL(SUB_PROCESS HAS mod_4, mod_5, mod_6, mod_7, mod_8)
$COMPACT(SUB_OUTPUT HAS mod_9, mod_10; EXPORTS format, output)
```

This sample program contains three subsystems: SUB_INPUT, SUB_PROCESS, and SUB_OUTPUT. SUB_INPUT and SUB_OUTPUT use the COMPACT extended segmentation model. SUB_PROCESS uses the SMALL extended segmentation model. Constants are stored with the code in SUB_INPUT. The SUB_INPUT subsystem contains the modules mod_1, mod_2, and mod_3, and exports one symbol, input. SUB_PROCESS contains modules 4 through 8. SUB_PROCESS contains the main program, as it must, since it is the only SMALL subsystem in the program. (Recall that when mixing SMALL with other models, the main program must be SMALL.) For this reason it does not need to export any symbols. A subsystem containing the main program can export symbols (for instance, global variables). But other subsystems must export at least one symbol, otherwise they are totally inaccessible to the main program, and therefore useless to the program of which they are a part.) SUB_OUTPUT supplies two symbols called format and output.

The preceding subsystem definition should appear in all 10 modules (mod_1 through mod_10), even though not all the exported symbols are used by all subsystems. It is recommended that the subsystem definition be kept in an include file, then included in each module compiled. This avoids any problems in maintaining consistency between the subsystem definitions of all source modules.

Consider another example, this time containing an open subsystem. Start from an existing COMPACT program that does not use extended segmentation models, but whose code has grown too large. Assume that the following modules from the original program (ATTACH, OPEN, CLOSE, ERRORS, ALLOCATE, FREE) were compiled with the following segmentation control:

```
$COMPACT
```


If the modules `ALLOCATE` and `FREE` are factored out from the original program, creating `SUBSYS1`, the subsystem definition is as follows:

```
$COMPACT(SUBSYS1 HAS ALLOCATE, FREE)
```

Now, suppose that the modules remaining in the open subsystem reference entry points `AllocBuff` and `FreeBuff` in `SUBSYS1`. These must be exported from `SUBSYS1` as follows:

```
$COMPACT(SUBSYS1 HAS ALLOCATE, FREE;
$          EXPORTS AllocBuff, FreeBuff)
```

or

```
$COMPACT(SUBSYS1 HAS ALLOCATE; EXPORTS AllocBuff:
$          HAS FREE; EXPORTS FreeBuff)
```

The second form illustrates how multiple `HAS` and `EXPORTS` lists can be used to document the items exported from each module.

If a routine in `SUBSYS1` references the procedure `FatalError` in the module `ERRORS`, the definition of the open subsystem is as follows:

```
$COMPACT (EXPORTS FatalError)
```

No data structures need to be changed, because data reference values can be two bytes. All procedures except `AllocBuff` and `FreeBuff` use the short call and return mechanism.

Placement of Segmentation Controls

The segmentation controls have special restrictions associated with their placement. These rules are as follows:

- The segmentation controls are primary controls. They must appear before the `DO` statement of the module name.
- Only the definition of the open subsystem (with no submodel and no `EXPORTS` list) can be placed on the invocation line; definitions of all other subsystems must occur inside the source program.

The subsystem definitions for the entire program can be included in the compilation of each module using the `INCLUDE` control. The compiler extracts the information needed to correctly and efficiently compile each module's intrasubsystem and inter-subsystem references.

Exporting Procedures

A symbol included in a subsystem's `EXPORTS` list must be declared `PUBLIC` in one of the modules in that subsystem. The symbol, called an exported symbol, can be referenced by modules in other subsystems. A `PUBLIC` symbol defined within a subsystem but not listed in its `EXPORTS` list is called a domestic symbol. It should be referenced only by modules within the same subsystem.

A procedure should be exported only if it must be referenced outside the defining subsystem, because accessing exported procedures will, in general, require more code and time than is required for domestic procedures.

Exported procedures have the following characteristics:

- The long form of call and return is used.
- The caller's `DS` and `ES` registers are saved and restored upon entry and exit.
- The `DS` and `ES` registers are loaded with the associated data segment upon entry.

Note that if a `SMALL` or `MEDIUM` module calls a procedure that is exported from a `COMPACT` or `LARGE` subsystem, the stack sections of the two will not be combined when the modules are combined because the segments containing them have different names (see Chapter 11). To get the proper stack size, the `SEGSIZE` control on the utility used to combine the program modules must be used to increase the size of the `DATA` segment. This segment must be increased by the sum of the stack requirements for both the `SMALL` or `MEDIUM` module and the subsystem.

The SMALL RAM segmentation control uses short pointers. Therefore, care must be taken when calling procedures that have pointer parameters and are exported from a SMALL subsystem. In these cases, the compiler always uses the value of the current DS register as the selector portion of the long pointer. This means that passing a pointer to any data items declared in the SMALL module will produce the proper result, but the following restrictions must be observed for the special cases:

1. If the actual parameter is the NIL pointer, DS:0 will be passed to the exported procedure. Consequently, the procedure executes differently if it is called from a SMALL module than if it had been called from a COMPACT, MEDIUM, or LARGE module. For example:]

```

$COMPACT (FOO HAS N; EXPORTS FOO)
$SMALL
M: DO;
    DECLARE PTR POINTER;
    FOO: PROCEDURE (P) EXTERNAL;
        DECLARE P POINTER;
    END FOO;
    CALL FOO (NIL);                /* Wrong, will pass DS:0 */
    PTR = NIL;
    CALL FOO (PTR);                /* Wrong, will pass DS:0 */
END M;
$COMPACT (FOO HAS N; EXPORTS FOO)
N: DO;
    FOO: PROCEDURE (P) PUBLIC;
        DECLARE P POINTER;
        DECLARE B BYTE;
        B = (P=NIL);             /* Will assign FALSE (000H) to B */
                                /* if FOO is called from SMALL; Will */
                                /* assign TRUE (0FFH) to B otherwise */
    END FOO;
    CALL FOO (NIL);              /* Right, will pass 0:0 */
END N;

```

2. If the actual parameter is a pointer to a procedure, the compiler extends the short pointer with DS and then passes the value of DS:(offset of procedure) to the exported procedure. This situation should be avoided because the result of any reference through such a pointer is undefined. For example:

```

$COMPACT (FOO HAS N; EXPORTS FOO)
$SMALL
M: DO;
    DECLARE PTR POINTER;
    DECLARE TABLE(10) BYTE;
    FOO: PROCEDURE (P) EXTERNAL;
        DECLARE P POINTER;
    END FOO;
    BAZ: PROCEDURE;
        ...
    END BAZ;
    CALL FOO (@BAZ);                                /* Wrong, will pass */
                                                    /* DS:offset-of-BAZ */

    PTR = @BAZ;
    CALL FOO (PTR);                                  /* Wrong, will pass DS:PTR */
    CALL FOO (@TABLE);                               /* Right, will pass pointer */
END M;                                              /* to TABLE */

```

Large Matrix Example

The large REAL matrix example can now be fully developed (see Using Subsystems). Recall that one module for each row is needed, with each module containing a 1-billion element REAL array. Running such an application is possible only on systems having virtual memory management for supporting such large data. The first module could be:

```

ROW0_MOD: DO;                                /* ROW0_MOD is the module name */
    DECLARE ROW0 (1000000000) REAL PUBLIC;
    END ROW0_MOD;

```

The modules for ROW1 through ROW9 are similar. The subsystem definition at this point is:

```
$COMPACT ( ROW0_SYS HAS ROW0_MOD; EXPORTS ROW0 )
$COMPACT ( ROW1_SYS HAS ROW1_MOD; EXPORTS ROW1 )
$COMPACT ( ROW2_SYS HAS ROW2_MOD; EXPORTS ROW2 )
$COMPACT ( ROW3_SYS HAS ROW3_MOD; EXPORTS ROW3 )
$COMPACT ( ROW4_SYS HAS ROW4_MOD; EXPORTS ROW4 )
$COMPACT ( ROW5_SYS HAS ROW5_MOD; EXPORTS ROW5 )
$COMPACT ( ROW6_SYS HAS ROW6_MOD; EXPORTS ROW6 )
$COMPACT ( ROW7_SYS HAS ROW7_MOD; EXPORTS ROW7 )
$COMPACT ( ROW8_SYS HAS ROW8_MOD; EXPORTS ROW8 )
$COMPACT ( ROW9_SYS HAS ROW9_MOD; EXPORTS ROW9 )
```

Now define the program:

```
MATRIX_MOD: DO;
DECLARE ROW0 (1000000000) REAL EXTERNAL;
DECLARE ROW1 (1000000000) REAL EXTERNAL;
DECLARE ROW2 (1000000000) REAL EXTERNAL;
DECLARE ROW3 (1000000000) REAL EXTERNAL;
DECLARE ROW4 (1000000000) REAL EXTERNAL;
DECLARE ROW5 (1000000000) REAL EXTERNAL;
DECLARE ROW6 (1000000000) REAL EXTERNAL;
DECLARE ROW7 (1000000000) REAL EXTERNAL;
DECLARE ROW8 (1000000000) REAL EXTERNAL;
DECLARE ROW9 (1000000000) REAL EXTERNAL;
```

```

DECLARE ROW_POINTERS (10) POINTER INITIAL (
@ROW0, @ROW1, @ROW2, @ROW3, @ROW4,
@ROW5, @ROW6, @ROW7, @ROW8, @ROW9 );
RETRIEVE_ELEMENT: PROCEDURE (ROW,COL) REAL PUBLIC;
    DECLARE (ROW,COL) WORD;
    DECLARE ROW_PTR POINTER,
    ROW_ARRAY BASED ROW_PTR (1) REAL;
    ROW_PTR = ROW_POINTERS (ROW);
    RETURN ROW_ARRAY (COL);
END RETRIEVE_ELEMENT;
STORE_ELEMENT: PROCEDURE (ROW,COL,VAL) PUBLIC;
    DECLARE (ROW,COL) WORD;
    DECLARE VAL REAL;
    DECLARE ROW_PTR POINTER,
    ROW_ARRAY BASED ROW_PTR (1) REAL;
    ROW_PTR = ROW_POINTERS (ROW);
    ROW_ARRAY (COL) = VAL;
END STORE_ELEMENT;
/* the matrix processing code inserted here */
END MATRIX_MOD;

```

Now assume that other modules will be added to this program later. In this case, it is better to put `MATRIX_MOD` and these other modules in the `COMPACT OPEN` subsystem. This way modules can freely be added or deleted without having to redefine the overall subsystem structure. Also assume the need to calculate sines and cosines of various matrix elements. The functions `SINE` and `COSINE` are supplied in an external math package. The only thing known about this package is that all its routines require long calls.

The final subsystem definition is now:

```

$LARGE ( EXPORTS SINE, COSINE )
$COMPACT ( ROW0_SYS HAS ROW0_MOD; EXPORTS ROW0 )
$COMPACT ( ROW1_SYS HAS ROW1_MOD; EXPORTS ROW1 )
$COMPACT ( ROW2_SYS HAS ROW2_MOD; EXPORTS ROW2 )
$COMPACT ( ROW3_SYS HAS ROW3_MOD; EXPORTS ROW3 )
$COMPACT ( ROW4_SYS HAS ROW4_MOD; EXPORTS ROW4 )
$COMPACT ( ROW5_SYS HAS ROW5_MOD; EXPORTS ROW5 )
$COMPACT ( ROW6_SYS HAS ROW6_MOD; EXPORTS ROW6 )
$COMPACT ( ROW7_SYS HAS ROW7_MOD; EXPORTS ROW7 )
$COMPACT ( ROW8_SYS HAS ROW8_MOD; EXPORTS ROW8 )
$COMPACT ( ROW9_SYS HAS ROW9_MOD; EXPORTS ROW9 )

```

The `COMPACT` control should appear in the invocation line. The first control line indicates that the symbols `SINE` and `COSINE` require long references and belong to some unknown subsystem. The next ten lines define the ten closed subsystems, each containing a row of the matrix. The `COMPACT` control is specified on the invocation line when compiling `MATRIX_MOD` (and when compiling any other module in the program except the `ROW` modules).

Every subsystem definition should be consistent. For example, `ROW0_MOD` must reside in the same subsystem in each definition. It is convenient to put control lines, such as those shown above, in an include file. If any changes to the subsystem definitions are made later, only one file needs to be updated.



Error and Warning Messages 14

The compiler may issue these kinds of error and warning messages:

- PL/M program error messages
- Fatal command tail and control error messages
- Fatal input/output error messages
- Fatal insufficient memory error messages
- Fatal compiler failure error messages
- Insufficient memory warning messages

The source errors are reported in the program listing; the fatal errors are reported on the console device.

PL/M Program Error and Warning Messages

Nearly all of the source PL/M program error messages are interspersed in the listing at the point of error and follow the general format:

```
*** ERROR mmm IN ppp (LINE ppp), NEAR 'aaa', message
```

or:

```
*** WARNING mmm IN ppp (LINE ppp), NEAR 'aaa', message
```

Where:

mmm is an error number from the following list.

ppp is the actual source line number where the error occurs.

aaa is the source text near where the error is detected.

message is a message from the following list.

The following source error messages may be encountered.

```
*** ERROR 1 INVALID CONTROL
```

An unrecognized control in the control line; for example:

```
$NXCODE; /* probably intended NOCODE */
```

- *** ERROR 2 PRIMARY CONTROL FOLLOWS NON-CONTROL LINE
Primary controls can be control lines in the source program, but they must come first. No other statements can precede them.
- *** ERROR 3 MISSING CONTROL PARAMETER
Certain controls (e.g., INCLUDE), require a parameter.
- *** ERROR 4 INVALID CONTROL PARAMETER
Examples are an illegal pathname for a control such as OBJECT or a string where a number is required.
- *** ERROR 5 INVALID CONTROL FORMAT
See Chapter 11 for correct formatting of control lines.
- *** ERROR 7 INVALID PATHNAME
The pathname for a file is incorrectly specified; see the host-system operating instructions.
- *** WARNING 8 ILLEGAL PAGELENGTH, IGNORED
The pagelength specified is less than 5 or greater than 255; the default is 60.
- *** ERROR 9 ILLEGAL PAGEWIDTH, IGNORED
The pagewidth specified is less than 60 or more than 132; the default is 120.
- *** WARNING 10 RESPECIFIED PRIMARY CONTROL, IGNORED
Primary controls can be specified only once and cannot alter a previous setting.
- *** ERROR 11 MISPLACED ELSE OR ELSEIF CONTROL
ELSE or ELSEIF control occurred without a corresponding IF control.
- *** ERROR 12 MISPLACED ENDIF CONTROL
ENDIF control occurred without a corresponding IF control.
- *** ERROR 13 MISSING ENDIF CONTROL
End of source file without an ENDIF control to match a previous IF.
- *** ERROR 14 NAME TOO LONG(31), TRUNCATED
Switch variable name in IF, ELSE, SET, or RESET statement is too long.
- *** ERROR 15 MISSING OPERATOR
Two operands in an expression must be separated by an arithmetic, logical, or relational operator.
- *** WARNING 16 INVALID CONSTANT, ZERO ASSUMED
The constant specified by SET, IF, or ELSEIF is invalid.
- *** ERROR 17 INVALID OPERAND
SET, RESET, IF, or ELSEIF is used in an invalid position.
- *** WARNING 18 PARENTHESES IGNORED WITHIN CONDITIONAL COMPILATION CONDITION
Parentheses within conditional compilation conditions are ignored and the expression is evaluated according to the regular precedence rules.

- *** ERROR 19 LIMIT EXCEEDED: SAVE NESTING
See Appendix B for the correct limit.
- *** ERROR 20 LIMIT EXCEEDED: INCLUDE NESTING
See Appendix B for the correct limit.
- *** ERROR 21 MISPLACED RESTORE CONTROL
RESTORE works only if there is a preceding SAVE.
- *** ERROR 22 UNEXPECTED END OF CONTROL
A segmentation control requires a continuation line or a right parenthesis.
- *** ERROR 23 SYMBOL EXISTS IN MORE THAN ONE HAS LIST
A module name can occur in only one HAS list.
- *** ERROR 24 SUBSYSTEM ALREADY DEFINED
The subsystem name has already been defined.
- *** ERROR 25 CONFLICTING SEGMENTATION CONTROLS
More than one segmentation control affecting the module being compiled was encountered. One common cause is specifying both `-CONST IN CODE-` and `ROM` in a module with a subsystem definition.
- *** ERROR 26 ILLEGAL PL/M IDENTIFIER
Identifiers can be up to 31 alphanumeric characters or the underscore; the first character must be alphabetic or the underscore.
- *** ERROR 27 PREDEFINED SWITCHES ARE NOT VALID BEFORE MODULE NAME
Predefined switches can be used only after the first `DO` statement.
- *** WARNING 28 INVALID PL/M CHARACTER, IGNORED
Look near the text flagged for an invalid character, or one that is inappropriate in context. Delete it or retype the statement.
- *** WARNING 29 UNPRINTABLE CHARACTER, IGNORED
Retype the line in question using valid characters.
- *** ERROR 30 STRING TOO LONG, TRUNCATED
See Appendix B for the correct limit.
- *** ERROR 31 ILLEGAL CONSTANT TYPE
A constant contains illegal characters. This might reflect missing operators (e.g., `A=4T` instead of `A=4+T`).
- *** ERROR 32 INVALID CHARACTER IN CONSTANT
For example, `107B` and `0ABCD` will cause this error because neither can be valid in any PL/M interpretation; `7` is not a binary numeral, `B` cannot occur in decimal or octal, and neither string ends in `H`.

*** ERROR 33 RECURSIVE MACRO EXPANSION

Following is an example causing this error:

```
DECLARE A LITERALLY 'B';
DECLARE B LITERALLY 'A';
.
.
.
B=4; /* error discovered here */
```

LITERALLYs cannot be declared circularly (i.e., solely in terms of each other).

*** ERROR 34 LIMIT EXCEEDED: MACRO NESTING (5)

This error occurs when too many DECLARE statements refer back through each other to the one that actually supplies a type. See Appendix B for the correct limit. For example:

```
DECLARE A LITERALLY 'B';
DECLARE B LITERALLY 'C';
. . . .
. . . .
DECLARE Y LITERALLY 'Z';
DECLARE Z BYTE INITIAL (77);
.
.
.
A=7; /* error discovered here */
```

*** ERROR 35 LIMIT EXCEEDED: SOURCE LINE LENGTH (128)

See Appendix B for the correct limit.

*** ERROR 37 INVALID REAL CONSTANT

*** WARNING 38 REAL CONSTANT UNDERFLOW

An underflow occurred when conversion into floating-point was attempted.

*** WARNING 39 REAL CONSTANT OVERFLOW

An overflow occurred when conversion into floating-point was attempted.

*** ERROR 40 NULL STRING NOT ALLOWED

Strings of length zero are not supported.

*** ERROR 41 DELETED: "tokens"

The compiler deleted tokens while attempting to recover from a syntax error.

*** ERROR 42 NEAR "syntax" INSERTED: "tokens"

The compiler inserted tokens while attempting to recover from a syntax error.

*** ERROR 43 STATEMENTS FOLLOW MODULE END

Statements follow the logical end-of-module.

*** ERROR 44 CONSTANT TOO LARGE

A constant value (e.g., 999,999,999,999) is too large for the compiler.

- *** WARNING 45 MISMATCHED BLOCK IDENTIFIER
If a label is supplied in an END statement, the label must match the first unmatched DO statement above the END. Sometimes the error involves a module name confused with a procedure name.
- *** ERROR 46 DUPLICATE PROCEDURE NAME
Procedure names must be unique.
- *** ERROR 47 LIMIT EXCEEDED: PROCEDURES
Too many procedures in this module. Break it into smaller modules. See Appendix B for the correct limit.
- *** ERROR 48 DUPLICATE PARAMETER NAME
A parameter must be declared exactly once. This message indicates that the flagged parameter already has a definition at this block level, as in:
YAR: PROCEDURE (YAR77, YAR78);
DECLARE YAR77 BYTE;
DECLARE YAR77 BYTE;
- Perhaps a different spelling was intended.
- *** ERROR 49 NOT AT MODULE LEVEL
The flagged attribute or initialization can be valid only at the module level, not in a procedure.
- *** ERROR 50 DUPLICATE ATTRIBUTE
Attributes should be specified only once. This message means the compiler has found a declaration like:
DECLARE B BYTE EXTERNAL EXTERNAL;
- *** ERROR 51 MISSING OR ILLEGAL INTERRUPT VALUE
Interrupt numbers must be whole-number constants between 0 and 255. Thus -7 or 272 would be invalid.
- *** ERROR 52 INTERRUPT WITH PARAMETERS
No parameters can be used in interrupt procedures.
- *** ERROR 53 INTERRUPT WITH TYPED PROCEDURE
Interrupt procedures must be untyped.
- *** ERROR 54 INVALID DIMENSION
- *** ERROR 55 LIMIT EXCEEDED: NESTED STRUCTURES
See Appendix B for the correct limit.
- *** ERROR 56 STAR DIMENSION WITH STRUCTURE MEMBER
Star dimension (*) must not be used with structures. The dimensions for an array that is a structure member must be specified explicitly.
- *** ERROR 57 CONFLICT WITH PARAMETER
Object cannot be a parameter.
- *** ERROR 58 DUPLICATE DECLARATION
The flagged item already has a definition declared at this block level.

- *** ERROR 59 ILLEGAL PARAMETER TYPE
Parameters cannot be declared of type structure or array.
- *** ERROR 60 DUPLICATE LABEL
Each label must be unique within its block or scope. Otherwise, GOTOS and CALLS would have ambiguous targets.
- *** ERROR 61 DUPLICATE MEMBER NAME
Member has been declared twice in the same structure. For example, in:
DECLARE AIR STRUCTURE (F4 BYTE, F4 BYTE);
subsequent references to AIR.F4 would be ambiguous.
- *** ERROR 62 UNDECLARED PARAMETER
A parameter named in the procedure statement was not defined in the body of the procedure.
- *** ERROR 63 CONFLICTING ATTRIBUTES
A variable has been declared with inconsistent attributes (e.g., PUBLIC or EXTERNAL, DATA or INITIAL, AT or BASED).
- *** ERROR 64 LIMIT EXCEEDED: DO BLOCKS
See Appendix B for the correct limit.
- *** ERROR 65 ILLEGAL PARAMETER ATTRIBUTE
Certain attributes cannot be used to declare a parameter (e.g., PUBLIC, EXTERNAL, DATA, INITIAL, AT, or BASED).
- *** ERROR 66 UNDEFINED BASE
A variable was declared BASED using an undeclared identifier.
- *** ERROR 67 INVALID TYPE OR ATTRIBUTE FOR BASE
A base must be a non-subscripted scalar of type ADDRESS, POINTER, WORD, SELECTOR, or OFFSET.
- *** ERROR 68 MISPLACED DECLARATION
Declarations and procedures can be interspersed, but not declarations and executable statements.
- *** ERROR 69 INVALID BASE WITH LABEL OR MACRO
BASED cannot be used with LABEL or LITERALLY types.
- *** ERROR 70 INVALID DIMENSION WITH LABEL OR MACRO
LABEL or LITERALLY cannot be dimensioned.
- *** ERROR 71 INITIALIZATION LIST REQUIRED
A list of initial values is required if the INITIAL attribute, the non-external * dimension form, or the non-external DATA attribute is used.
- *** ERROR 72 BASED CONFLICTS WITH ATTRIBUTES
Examples of attributes conflicting with base include AT, DATA, INITIAL, PUBLIC, and EXTERNAL.

- *** ERROR 73 DATA OR EXECUTABLE STATEMENTS IN EXTERNAL
An EXTERNAL procedure, being defined elsewhere, cannot contain executable statements or data declarations for variables that are not formal parameters.
- *** ERROR 74 MISSING RETURN FOR TYPED PROCEDURE
A typed procedure must return a value; thus, it must include a RETURN statement.
- *** ERROR 75 INVALID NESTED REENTRANT PROCEDURE
Reentrant procedures cannot contain procedures.
- *** ERROR 76 LIMIT EXCEEDED: FACTORED LIST
Too many variables were named in a factored declaration. Break it into several declarations. See Appendix B for the correct value.
- *** ERROR 77 LIMIT EXCEEDED: STRUCTURE MEMBERS
See Appendix B for the correct value.
- *** ERROR 78 MISSING PROCEDURE NAME
Every procedure must have a name.
- *** ERROR 79 MULTIPLE PROCEDURE LABELS
Procedures must have only one name.
- *** ERROR 80 DECLARATIONS MAY NOT BE LABELED
Labels cannot be used on declaration statements.
- *** ERROR 81 STAR DIM WITH FACTORED LIST NOT ALLOWED
Separate the array declarations giving the data initializations for each array separately, or explicitly state the dimensions of the factored array declarations as in the following examples:

```

DECLARE (A,B) (*) BYTE DATA ('abcd', 'xyzw'); /* illegal */
DECLARE (A) (*) BYTE DATA ('abcd'); /* legal */
DECLARE (B)(*) BYTE DATA ('xyzw'); /* legal */

```

or

```

DECLARE (A,B) (4) BYTE DATA ('abcd', 'xyzw'); /* legal */

```
- *** ERROR 82 SIZE EXCEEDS *nn* BYTES
Storage for the declared item exceeds the maximum storage for the microprocessor. For the Intel386 and Intel486 microprocessor *nn* is 4G bytes.
- *** WARNING 83 PROCEDURE CONTAINS NO EXECUTABLE STATEMENTS
This procedure does nothing, but executes successfully.
- *** ERROR 85 INITIAL USED WITH ROM OPTION
Variables declared with INITIAL are not initialized until load-time. Thus, if the program is in ROM, these initializations will never occur.
- *** ERROR 86 LIMIT EXCEEDED: NUMBER OF PARAMETERS
The procedure declaration includes too many parameters. See Appendix B for the correct limit.

- *** ERROR 88 LIMIT EXCEEDED: PROGRAM TOO COMPLEX
The program has too many complex expressions, cases, or procedures. Break it into smaller procedures.
- *** ERROR 89 COMPILER ERROR: BAD ERROR RECOVERY
An unrecoverable error occurred. Trying a different copy of the compiler on a different drive might reveal that the first copy has been damaged. Contact your Intel representative.
- *** ERROR 90 COMPILER ERROR: MULTIPLE PARSE ARGS
See source error message number 89.
- *** ERROR 91 LIMIT EXCEEDED: PROGRAM TOO COMPLEX
The program has too many complex expressions, cases, or procedures. Break it into smaller procedures.
- *** ERROR 92 COMPILER ERROR: PARSE ARG STACK UNDERFLOW
See source error message number 89.
- *** ERROR 93 LIMIT EXCEEDED: PROGRAM TOO COMPLEX
The program has too many complex expressions, cases, or procedures. Break it into smaller modules.
- *** ERROR 94 COMPILER ERROR: PARSE STACK UNDERFLOW
See source error message number 89.
- *** ERROR 95 COMPILER ERROR: PARSE BUFFER OVERFLOW
See source error message number 89.
- *** ERROR 96 LIMIT EXCEEDED: BLOCK NESTING
The program has too many nested DO blocks. Break it into smaller procedures. See Appendix B for the correct limit.
- *** ERROR 97 COMPILER ERROR: SCOPE STACK UNDERFLOW
See source error message number 89.
- *** ERROR 98 LIMIT EXCEEDED: STATEMENT TOO COMPLEX
The statement is too large for the compiler. Break it into several smaller statements.
- *** ERROR 99 COMPILER ERROR: SEMANTIC UNDERFLOW
See source error message number 89.
- *** ERROR 100 STRING CONSTANT TOO LONG
String constants used as scalars have a maximum of four characters.
- *** ERROR 101 UNSUBSCRIPTED ARRAY
The array reference requires a subscript.
- *** WARNING 102 UNQUALIFIED STRUCTURE
This statement is ambiguous as to which structure or member it references.

- *** ERROR 103 NOT AN ARRAY
Subscripts are permitted only on identifiers declared as arrays. Check spelling consistency.
- *** ERROR 104 MULTIPLE SUBSCRIPTS
PL/M has only single dimension arrays. Therefore, only one subscript is permitted in an array reference. For example, for any array TING references of the form TING(2,4) or TING(3,7,9,6) are invalid.
- *** ERROR 105 NOT A STRUCTURE
For example, a reference of the form GNU.F1, where GNU was not declared a structure.
- *** ERROR 106 UNDEFINED IDENTIFIER
Every identifier must be declared.
- *** ERROR 107 UNDEFINED MEMBER
For example, KAPI.HORN, where KAPI is a valid, declared structure but HORN is an undeclared member of the structure.
- *** ERROR 108 ILLEGAL ITERATIVE DO INDEX TYPE
Only expressions of type BYTE, WORD, and INTEGER can be used.
- *** ERROR 109 UNDEFINED OR NOT A LABEL
The identifier following GOTO must be a label; the flagged item was declared otherwise, or the identifier was declared as a label but was not defined.
- *** ERROR 110 MISSING RETURN VALUE
A typed procedure must return a value that is specified by its RETURN statement.
- *** ERROR 111 INVALID RETURN WITH UNTYPED PROCEDURE
An untyped procedure does not return a value; thus, its RETURN statement cannot specify one.
- *** ERROR 112 INVALID INDIRECT TYPE
Only WORD or POINTER scalars can be used for indirect calls. This excludes WORD or POINTER expressions; BYTE, DWORD, INTEGER, or REAL scalars; all structures; and all arrays.
- *** ERROR 113 INVALID PARAMETER COUNT
The number of actual parameters supplied in a CALL must be equal to the number of formal parameters declared in the procedure.
- *** ERROR 114 QUALIFIED PROCEDURE NAME
Procedure names cannot be qualified.
- *** ERROR 115 INVALID FUNCTION REFERENCE
Typed procedures can be invoked only by use in an expression, not by a CALL.
- *** ERROR 116 INVALID CASE EXPRESSION TYPE
Case expressions must be of type BYTE, WORD, or INTEGER.

- *** ERROR 117 LIMIT EXCEEDED: NUMBER OF ACTIVE CASES
Reduce the number of cases in this case statement; the maximum number has been exceeded.
- *** ERROR 118 TYPE CONFLICT
An example of type conflict is WORD and REAL mixed in a reference.
- *** ERROR 119 INVALID BUILT-IN REFERENCE
Built-in reference was qualified with a member name, or OUTPUT/OUTWORD did not appear on the left side of an assignment.
- *** ERROR 120 INVALID PROCEDURE REFERENCE
Untyped procedures must be invoked by a CALL statement; references to such procedures are not permitted in expressions.
- *** ERROR 121 INVALID LEFT-HAND SIDE OF ASSIGNMENT
The left-hand side of the assignment must be a scalar variable. For example, PROCEDURE=4 or INWORD(7)=9.
- *** ERROR 122 INVALID REFERENCE
Invalid label reference.
- *** ERROR 123 USE OF "." MAY BE UNSAFE
The "dot" operator does not always produce correct results in a PL/M program that contains more than one data segment or more than one code segment.
- *** ERROR 124 PROCEDURE NAME REQUIRED
Procedure name is required for SET\$INTERRUPT and INTERRUPT\$PTR built-ins.
- *** ERROR 125 PROCEDURE NAME ONLY
Parameters are not allowed on the procedure name in SET\$INTERRUPT and INTERRUPT\$PTR.
- *** ERROR 126 BAD INTERRUPT NUMBER
Interrupt numbers in a CAUSE\$INTERRUPT statement must be whole-number constants in the range (0 - 255).
- *** ERROR 127 CONSTANT ONLY
In this instance, a constant is required.
- *** ERROR 128 ARRAY REQUIRED
Some built-ins need an array name as a parameter.
- *** ERROR 129 INTERRUPT PROCEDURE REQUIRED
The name declared in a SET\$INTERRUPT procedure or INTERRUPT\$PTR function must be a previously declared procedure.
- *** ERROR 130 INVALID RESTRICTED OPERAND
Illegal use of a dot operator.
- *** ERROR 131 INVALID RESTRICTED OPERATOR
Only + and - can be used in restricted expressions.

- *** ERROR 133 REFERENCE REQUIRED
A variable reference is required for LENGTH, LAST, and SIZE.
- *** ERROR 134 VARIABLE REQUIRED
The operand to LENGTH, LAST, and SIZE must be a variable.
- *** ERROR 135 VALUE TOO LARGE
A value is too large for its contextually determined type.
- *** ERROR 136 ABSOLUTE POINTER WITH SHORT POINTERS
Two possible causes in the SMALL (RAM) case: pointer variables cannot be initialized with or assigned whole number constants; or the @ operator cannot be used with a variable that was located at an absolute address that was specified by a whole number constant.
- *** ERROR 137 INVALID RESTRICTED EXPRESSION
Only addresses or constant types can be used in restricted expressions.
- *** ERROR 138 PUBLIC AT EXTERNAL
PUBLIC declarations must be fully defined within the procedure. For example:
DECLARE KUN BYTE EXTERNAL;
DECLARE JAN BYTE PUBLIC AT (.KUN);

is illegal.
- *** ERROR 139 PUBLIC AT ABSOLUTE
Absolute locations for PUBLICS are supported only in the LARGE model.
- *** ERROR 140 PUBLIC AT MEMORY
PUBLIC at @MEMORY is not supported by COMPACT.
- *** ERROR 141 AT BASED VARIABLE
Based variables cannot be used in AT clauses.
- *** ERROR 142 ILLEGAL FORWARD REFERENCE
An AT expression cannot have a forward reference. Any location reference in the AT expression must refer to previously declared variables.
- *** ERROR 143 VARIABLE TYPE REQUIRED IN AN AT EXPRESSION
The AT expression must be a variable name. For example:
DECLARE B BYTE AT (.proc_name);

is illegal.
- *** ERROR 144 LIMIT EXCEEDED: DATA OR STACK SEGMENT TOO LARGE
- *** ERROR 145 LIMIT EXCEEDED: CODE OR CONST SEGMENT TOO LARGE
- *** ERROR 146 LIMIT EXCEEDED: NUMBER OF EXTERNALS
See Appendix B for the correct limit.
- *** ERROR 147 LABEL NOT AT LOCAL OR MODULE LEVEL
Label was not used correctly.
- *** ERROR 148 INITIALIZING MORE SPACE THAN DECLARED
The number of initialization values exceeds the number of declared elements.

- *** ERROR 149 ILLEGAL MODULE NAME REFERENCE
Module names cannot be referenced.
- *** WARNING 150 USE OF "." WITH FAR PROCEDURE
A subsequent indirect call made through the respective address/pointer generates the wrong type of call.
- *** WARNING 151 USE OF "@" WITH NEAR PROCEDURE
See source error message number 150.
- *** ERROR 152 INVALID "." OR "@" OPERAND
Must be used with a variable, procedure, or constant list.
- *** ERROR 153 INVALID RETURN IN MAIN PROGRAM
A main program must have no returns.
- *** ERROR 154 STAR DIMENSIONED VARIABLE WITH LENGTH, SIZE OR LAST
The LENGTH, LAST, and SIZE built-in functions cannot be used with variables declared with the implicit dimension specifier (*) and the EXTERNAL attribute.
- *** ERROR 155 SYMBOL EXPORTED FROM ANOTHER SUBSYSTEM
A PUBLIC symbol in this module is also exported by another subsystem.
- *** ERROR 156 LONG POINTER REQUIRED FOR THIS CONSTRUCT
A model with long pointers is required.
- *** ERROR 158 INITIALIZATION CONFLICTS WITH ATTRIBUTES
An external variable cannot be initialized.
- *** ERROR 159 ILLEGAL INTERRUPT PROCEDURE REFERENCE
An interrupt procedure cannot be invoked with the CALL statement.
- *** ERROR 160 INTERRUPT PROCEDURES MUST BE PUBLIC
An interrupt procedure must also be given the PUBLIC attribute.
- *** ERROR 161 ILLEGAL ABSOLUTE POINTER OR SELECTOR
Constants cannot be assigned to POINTERS or SELECTORS, nor used to initialize them. POINTERS and SELECTORS also cannot be passed as actual parameters.
- *** ERROR 162 LIMIT EXCEEDED: STATEMENT TOO COMPLEX
The statement is too large for the compiler. Break it into several smaller statements.
- *** WARNING 162 LIMIT EXCEEDED: PROGRAM COMPLEXITY
Too many complex expressions, cases, etc. Break it into smaller procedures.
- *** ERROR 163 COMPILER ERROR: SEMANTIC UNDERFLOW
See source error message number 89.
- *** ERROR 164 COMPILER ERROR: INVALID NODE
See source error message number 89.
- *** ERROR 165: 286 INTERFACE OBJECT NOT EXTERNAL
If the *machine* parameter is 286, all identifiers in the *id* list must be declared EXTERNAL.

- *** ERROR 166 COMPILER ERROR: INVALID TREE
See source error message number 89.
- *** ERROR 167 COMPILER ERROR: SCOPE STACK UNDERFLOW
See source error message number 89.
- *** ERROR 168 LIMIT EXCEEDED: PROGRAM COMPLEXITY
The program has too many complex expressions, cases, or procedures. Break it into smaller procedures.
- *** ERROR 169 COMPILER ERROR: INVALID RECORD
See source error message number 89.
- *** ERROR 170 INVALID DO CASE BLOCK, AT LEAST ONE CASE REQUIRED
The DO CASE block is described in Chapter 6.
- *** ERROR 171 LIMIT EXCEEDED: NUMBER OF CASES
- *** ERROR 172 LIMIT EXCEEDED: NESTING OF TYPED PROCEDURE CALLS
- *** ERROR 173 LIMIT EXCEEDED: NUMBER OF ACTIVE PROCEDURES AND DO CASE GROUPS
See Appendix B for the correct limit.
- *** ERROR 174 ILLEGAL NESTING OF BLOCKS, ENDS NOT BALANCED
For every DO, an END is needed.
- *** ERROR 175 COMPILER ERROR: INVALID OPERATION
See source error message number 89.
- *** ERROR 176 LIMIT EXCEEDED: REAL EXPRESSION COMPLEXITY
The REAL stack has eight registers. Heavily nested use of REAL functions with REAL expressions as parameters can get excessively complex. See Appendix F.
- *** ERROR 177 COMPILER ERROR: REAL STACK UNDERFLOW
See source error message numbers 89 and 176.
- *** ERROR 178 LIMIT EXCEEDED: BASIC BLOCK COMPLEXITY
There is a very long list of statements without labels, CASES, IFs, GOTOS, and/or RETURNS. Either break the procedure into several smaller procedures, or add labels to some of the statements.
- *** ERROR 179 LIMIT EXCEEDED: STATEMENT SIZE
The statement is too large for the compiler. Break it into several smaller statements.
- *** ERROR 199 LIMIT EXCEEDED: PROCEDURE COMPLEXITY FOR OPTIMIZE (2)
The combined complexity of statements, user labels, and compiler-generated labels is too great. Simplify as much as possible, perhaps breaking the procedure into several smaller procedures.
- *** ERROR 200 ILLEGAL INITIALIZATION OF MORE SPACE THAN DECLARED
The number of initialization values exceeds the number of declared elements.
- *** ERROR 201 INVALID LABEL: UNDEFINED
No definition for this label was found.

- *** ERROR 202 LIMIT EXCEEDED: NUMBER OF EXTERNAL ITEMS
See Appendix B for the correct limit.
- *** ERROR 203 COMPILER ERROR: BAD LABEL ADDRESS
See source error message number 89.
- *** ERROR 204 LIMIT EXCEEDED: CODE SEGMENT SIZE
See Appendix B for the correct limit.
- *** ERROR 205 COMPILER ERROR: BAD CODE GENERATED
See source error message number 89.
- *** ERROR 206 LIMIT EXCEEDED: DATA SEGMENT SIZE
See Appendix B for the correct limit.
- *** ERROR 207 ATTEMPT TO USE 0 AS DIVISOR IN DIVISION/MODULO
Zero cannot be used as a divisor in division/modulo; use 1. This error appears at the end as a semantic error.
- *** ERROR 210 COMPILER ERROR: OBJECT MODULE GENERATION ERROR
- *** ERROR 211 COMPILER ERROR: DEBUG SEGMENT SIZE OVERFLOW
- *** ERROR 212 COMPILER ERROR: ILLEGAL FIXUP
- *** ERROR 230 COMPILER ERROR: INVALID INTERNAL TYPE
See source error message number 89.
- *** ERROR 241 ILLEGAL TYPE CASTING
For example:
pt2=pointer (real_value)

is illegal.
- *** ERROR 242 TRUNCATION OF n BIT OFFSET
OFFSET was assigned to a variable with a size less than 32 bits; the assigned value may not be a valid OFFSET. For the 8086 and 286 microprocessors, n is 16. For the Intel386 and Intel486 microprocessors, n is 32.
- *** ERROR 243 286 INTERFACE OBJECT NOT EXTERNAL
If the *machine* parameter is 286, all identifiers in the *id* list must be declared EXTERNAL.
- *** ERROR 244 SYMBOL REPEATED IN INTERFACE SPECIFICATION
Symbols can be used only once in an INTERFACE control (i.e., a symbol cannot be repeated in the INTERFACE control).
- *** ERROR 245 AT VARIABLE IN DIFFERENT SEGMENT
A variable cannot be declared using both the DATA attribute and the AT attribute when using the ROM option. DATA should be in CODE segments and INITIAL should be in DATA segments.
- *** WARNING 247 INDIRECT CALL THROUGH 16 BIT VARIABLE
An indirect call through a 16-bit variable is not recommended because a 16-bit variable can address only the first 64K of a segment.

- *** WARNING 248 BASE TYPE HAS ONLY 16 BITS OFFSET
Use of a 16-bit base specifier is not recommended because it can address only the first 64K of a segment.
- *** ERROR 251 COMPILER ERROR: INVALID OBJECT
- *** ERROR 252 COMPILER ERROR: SELF NAME LINK
- *** ERROR 253 COMPILER ERROR: SELF ATTR LINK
See source error message number 89.
- *** ERROR 254 LIMIT EXCEEDED: PROGRAM COMPLEXITY
The program has too many complex expressions, cases, or procedures. Break it into smaller modules.
- *** ERROR 255 LIMIT EXCEEDED: SYMBOLS
See Appendix B for the correct limit.



Note

If a terminal error is encountered, program text beyond the point of error is not compiled. A terminal error message will appear at the point of error in the program listing.

Fatal Command Tail and Control Error Messages

Fatal command tail errors are caused by an improperly specified compiler invocation command or an improper control. The errors that can occur are as follows:

```
COMMAND TAIL TOO LONG
COMMAND TAIL BUFFER LIMIT EXCEEDED AT OR NEAR: xxx
ILLEGAL COMMAND TAIL SYNTAX OR VALUE
UNABLE TO PARSE COMMAND TAIL AT OR NEAR: xxx
ILLEGAL COMMAND TAIL SYNTAX OR VALUE
UNRECOGNIZED CONTROL IN COMMAND TAIL
INVOCATION COMMAND DOES NOT END WITH <CR><LF>
ILLEGAL COMMAND TAIL SYNTAX
```

Fatal Input/Output Error Messages

Fatal input/output errors occur when the user specifies an incorrect pathname for compiler input or output. These error messages are of the form:

```
PL/M-386 xxx ERROR --  
FILE:  
NAME:  
ERROR:  
COMPILATION TERMINATED
```

These errors also occur when the device runs out of space (e.g., the list file is larger than the available memory).

Fatal Insufficient Memory Error Messages

The fatal insufficient memory errors are caused by a system configuration with insufficient RAM memory to support the compiler.

The errors that can occur due to insufficient memory are as follows:

```
NOT ENOUGH MEMORY FOR COMPILATION  
DYNAMIC STORAGE OVERFLOW  
NOT ENOUGH MEMORY FOR CODE GENERATION
```


Fatal Compiler Failure Error Messages

The fatal compiler failure errors are internal errors that should never occur. If you encounter such an error, please contact your Intel representative. The errors falling into this class are as follows:

```
*** ERROR 89 COMPILER ERROR: BAD ERROR RECOVERY
*** ERROR 90 COMPILER ERROR: MULTIPLE PARSE ARGS
*** ERROR 92 COMPILER ERROR: PARSE ARG STACK UNDERFLOW
*** ERROR 94 COMPILER ERROR: PARSE STACK UNDERFLOW
*** ERROR 95 COMPILER ERROR: PARSE BUFFER OVERFLOW
*** ERROR 97 COMPILER ERROR: SCOPE STACK UNDERFLOW
*** ERROR 99 COMPILER ERROR: SEMANTIC UNDERFLOW
*** ERROR 163 COMPILER ERROR: SEMANTIC UNDERFLOW
*** ERROR 164 COMPILER ERROR: INVALID NODE
*** ERROR 166 COMPILER ERROR: INVALID TREE
*** ERROR 167 COMPILER ERROR: SCOPE STACK UNDERFLOW
*** ERROR 175 COMPILER ERROR: INVALID OPERATION
*** ERROR 177 COMPILER ERROR: REAL STACK UNDERFLOW
*** ERROR 203 COMPILER ERROR: BAD LABEL ADDRESS
*** ERROR 205 COMPILER ERROR: BAD CODE GENERATED
*** ERROR 210 COMPILER ERROR: OBJECT MODULE GENERATION
*** ERROR 211 COMPILER ERROR: DEBUG SEGMENT SIZE OVERFLOW
*** ERROR 212 COMPILER ERROR: ILLEGAL FIXUP
*** ERROR 230 COMPILER ERROR: INVALID INTERNAL TYPE
*** ERROR 251 COMPILER ERROR: INVALID OBJECT
*** ERROR 252 COMPILER ERROR: SELF NAME LINK
*** ERROR 253 COMPILER ERROR: SELF ATTR LINK
```

It is also possible to receive an UNKNOWN FATAL ERROR message.

Insufficient Memory Warning Messages

The following warnings may occur if there are too many symbols for symbol processing:

```
NOT ENOUGH MEMORY FOR FULL DICTIONARY LISTING
NOT ENOUGH MEMORY FOR ANY XREF PROCESSING
NOT ENOUGH MEMORY FOR FULL XREF PROCESSING
```



PL/M Reserved Words and Predeclared Identifiers **A**

Introduction

These are reserved words in PL/M-386. They cannot be used as identifiers.

ADDRESS	INTEGER
AND	INTERRUPT
AT	LABEL
BASED	LITERALLY
BY	LONGINT
BYTE	MINUS
CALL	MOD
CASE	NOT
CHARINT	OFFSET
DATA	OR
DECLARE	PLUS
DISABLE	POINTER
DO	PROCEDURE
DWORD	PUBLIC
ELSE	REAL
ENABLE	REENTRANT
END	RETURN
EOF	SELECTOR
EXTERNAL	SHORTINT
GO	STRUCTURE
GOTO	THEN
HALT	TO
HWORD	WHILE
IF	WORD
INITIAL	QWORD
	XOR

The following are PL/M-386 identifiers, built-in procedures and predeclared variables. If one of these identifiers is declared in a `DECLARE` statement, the corresponding built-in procedure or predeclared variable becomes unavailable within the scope of the declaration.

ABS	NIL
ADJUSTRPL	OFFSETOF
BLOCKINPUT	OUTHWORD
BLOCKINWORD	OUTPUT
BLOCKOUTPUT	OUTWORD
BLOCKOUTWORD	PARITY
BUILDPTR	RESTOREGLOBALTABLE
CARRY	RESTOREINTERRUPTABLE
CAUSEINTERRUPT	RESTOREREALSTATUS
CLEARTASKSWITCHEDFLAG	ROL
CONTROLREGISTER	ROR
CMPB	SAL
CMPW	SAR
DEBUGREGISTER	SAVEGLOBALTABLE
DEC	SAVEINTERRUPTABLE
DOUBLE	SAVEREALSTATUS
FINDB	SCANBIT
FINDHW	SCANRBIT
FINDRB	SCL
FINDRHW	SCR
FINDRW	SEGMENTREADABLE
FINDW	SEGMENTWRITABLE
FIX	SELECTOROF
FLAGS	SETB
FLOAT	SETHW
GETACCESSRIGHTS	SETREALMODE
GETREALERROR	SETW
GETSEGMENTLIMIT	SHL
HIGH	SHLD
IABS	SHR
INHWORD	SHRD

INITREALMATHUNITSKIPRB	SIGN
INPUT	SIGNED
INT	SIZE
INWORD	SIZE
LAST	SKIPB
LENGTH	SKIPHW
LOCALTABLE	SKIPRHW
LOCKSET	SKIPRW
LOW	SKIPW
MACHINESTATUS	STACKBASE
MOVB	STACKPTR
MOVBIT	TASKREGISTER
MOVE	TESTREGISTER
MOVHW	TIME
MOVRB	UNSIGN
MOVRBIT	WAITFORINTERRUPT
MOVRHW	XLAT
MOVRW	ZERO
MOVW	

Identifiers with WORD16 Control

The following identifiers are specific to PL/M-386 when using the WORD16 control.

BLOCKINDWORD
 BLOCKOUTDWORD
 CMPD
 FINDD
 FINDRD
 INDWORD
 MOVD
 MOVRD
 OUTDWORD
 SETD
 SKIPD
 SKIPRD

Identifiers with MOD486 Control

The following identifiers are specific to PL/M-386 when using the MOD486 control.

BYTESWAP
TESTREGISTER
INVALIDATEDATACACHE
WBINVALIDATEDATACACHE
INVALIDATETLBENTRY



PL/M Program Limits **B**

Feature	PL/M-386
Indirection level (A BASED on B, B BASED on C)	unlimited***
Length of a string constant	255
Nesting of blocks	18
Nesting of INCLUDE controls	5
Nesting of LITERALLY invocations	5
Nesting of structures	32
Number of active cases	255
Number of cases in a DO CASE block	255
Number of DO blocks in a procedure	65536
Number of declared EXTERNAL items	**
Number of elements in a factored list	64
Number of EXTERNAL items used	**
Number of labels on a statement	unlimited*
Number of nested procedures and DO cases	255
Number of nested typed procedures	18
Number of procedures in a module	1016
Numbers of characters in a line	128
Segment size	4G
Size of LITERALLY string	unlimited*
Structure size	4G-1
Symbol capacity	2500
Total number of members in a structure (at all levels)	128

* Limited by the total size of the symbol table.

** Limited by either the number of procedures or the number of symbols, or both.

*** Unlimited means limited only by the amount of free memory allocated by the compiler.



Grammar of the PL/M Language **C**

This appendix lists the entire syntax of the PL/M language in Backus-Naur Form (BNF) notation. Since the semantic rules are not included here, this syntax permits certain constructions that are not actually allowed. The terminology used in the BNF syntax has been designed for convenience in constructing concise and rigorous definitions. Its appearance differs substantially from the main body of the manual.

The notations used here are slightly extended from standard BNF notations. An ellipsis (...) indicates that the syntactic element preceding it can be repeated indefinitely. The vertical bar (|) separates alternatives. Braces ({}) enclose required alternatives and brackets ([]) enclose optional alternatives. The vertical bar within braces and brackets is also a separator of alternatives.

When items are stacked vertically within brackets, only one of the items can be used.

Lexical Elements

Character Sets

<character> ::= <apostrophe> | <non-quote character>

<apostrophe> ::= ' '

<non-quote character> ::= <letter> | <decimal digit> | \$ |
<special character> | blank

<letter> ::= <uppercase letter> | <lowercase letter>

<uppercase letter> ::= 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

<lowercase letter> ::= 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

<decimal digit> ::= 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9

<special character> ::= + | - | * | / | < | > | = | : | ; | . | , |
(|) | @ | _

Tokens

<token> ::= <delimiter> | <identifier> | <reserved word> |
<numeric constant> | <string>

Delimiters

<delimiter> ::= <simple delimiter> | <compound delimiter>

<simple delimiter> ::= + | - | * | / | < | > | = | : | ; | . | , | (|
) | @

<compound delimiter> ::= <> | <= | >= | :=

Identifiers

<first id character> ::= <letter> | _

<identifier> ::= <first id character> [<letter> |
<decimal digit> | \$ | _]...

<reserved word> (For a list of reserved words, see Appendix A.)

Numeric Constants

```
<numeric constant> ::= <binary number> | <octal number> |  
                        <decimal number> | <hexadecimal number> |  
                        <floating point number>  
  
<binary number> ::= <binary digit> [<binary digit> | $]... B | Q  
<octal number> ::= <octal digit> [<octal digit> | $]... { O | Q }  
<decimal number> ::= <decimal digit> [<decimal digit> | $]... [D]  
<hexadecimal number> ::= <decimal digit> [<hexadecimal digit>  
                                         | $]... H  
<floating point number> ::= <digit string> <fractional part>  
                             [<exponent part>]  
<fractional part> ::= [<.digit string>]  
<exponent part> ::= E [+ | -] <digit string>  
<digit string> ::= <decimal digit> [<decimal digit> | $]...  
<binary digit> ::= 0 | 1  
<octal digit> ::= <binary digit> | 2 | 3 | 4 | 5 | 6 | 7  
<decimal digit> ::= <octal digit> | 8 | 9  
<hexadecimal digit> ::= <decimal digit> | A | B | C | D | E | F
```

Strings

```
<string> ::= '<string body element> [<string body element>]... '  
<string body element> ::= <non-quote character> | "
```

PL/M Text Structure: Tokens, Blanks, and Comments

```
<pl/m text> ::= <token> | <separator> [<token> | <separator>]...  
<separator> ::= blank | <comment>  
<comment> ::= /* [<character>]... */
```

Modules and the Main Program

`<compilation> ::= <module> [EOF]`

`<module> ::= <module name> : <simple do block>`

`<module name> ::= <identifier>`

Declarations

<declaration> ::= <declare statement> | <procedure definition>

DECLARE Statement

<declare statement> ::= DECLARE <declare element list>;

<declare element list> ::= <declare element>[, <declare element>]...

<declare element> ::= <factored element> | <unfactored element>

<unfactored element> ::= <variable element> | <literal element> |
 <label element>

<factored element> ::= <factored variable element> |
 <factored label element>

Variable Elements

<variable element> ::= <variable name specifier> [<array specifier>]
 <variable type> | [<variable attributes>]

<variable name specifier> ::= <non-based name> |
 <based name> BASED <base specifier>

<non-based name> ::= <variable name>

<based name> ::= <variable name>

<variable name> ::= <identifier>

<base specifier> ::= <identifier>[.<identifier>]

<variable attributes> ::= [PUBLIC] [<locator>][<initialization>] |
 [EXTERNAL] [<constant attribute>]

<locator> ::= AT(<expression>)

<constant attribute> ::= DATA

<array specifier> ::= <explicit dimension> | <implicit dimension>

<explicit dimension> ::= (<numeric constant>)

<implicit dimension> ::= (*)

```
<variable type>::= <basic type> | <structure type>
<basic type>::= Address | BYTE | HWORD | DWORD | QWORD | CHARINT |
                OFFSET | SHORTINT | INTEGER | REAL | SELECTOR |
                POINTER | OFFSET
```

Label Element

```
<label element>::= <identifier> LABEL [PUBLIC | EXTERNAL]
```

Literal Elements

```
<literal element>::= <identifier> LITERALLY <string>
```

Factored Variable Element

```
<factored variable element>::= ( <variable name specifier>
                                [,<variable name specifier>]...)
                                [<explicit dimension>] <variable type>
                                [<variable attributes>]
```

Factored Label Element

```
<factored label element>::= (<identifier> [,<identifier>]...)
                            LABEL [PUBLIC | EXTERNAL]
```

The Structure Type

```
<structure type>::= STRUCTURE (<member element>
                                [,<member element>]...)
<member element>::= <unfactored member> | <factored member>
<unfactored member>::= <member name> [<explicit dimension>]
                        <variable type>
<member name>::= <identifier>
<factored member>::= (<member name>(,<member name>)...)
                    [<explicit dimension>] <variable type>
```

Procedure Definition

```
<procedure definition> ::= <procedure statement> [<declaration>...]
                          [<unit>...] <ending>

<procedure statement> ::= <procedure name> : PROCEDURE
                          [<formal parameter list>] [<procedure type>]
                          [<procedure attributes>];

<procedure name> ::= <identifier>

<procedure type> ::= <basic type>

<formal parameter list> ::= (<formal parameter>
                             [, <formal parameter>]...)

<formal parameter> ::= <identifier>

<procedure attributes> ::= {EXTERNAL | PUBLIC | <interrupt> |
                           REENTRANT}...
```

Attributes

AT

```
<locator> ::= AT (<expression>)
```

INTERRUPT

```
<interrupt> ::= INTERRUPT
```

Initialization

```
<initialization> ::= {INITIAL | DATA} (<initial value>
                                       [, <initial value>]...)

<initial value> ::= <expression> | <string>
```

Units

```
<unit> ::= <conditional clause> | <do block> | <basic statement> |  
        <label definition><unit>  
  
<basic statement> ::= <assignment statement> | <call statement> |  
                    <goto statement> | <null statement> |  
                    <return statement> |  
                    <microprocessor dependent statement>  
  
<scoping statement> ::= <simple do statement> | <do-case statement> |  
                       <do-while statement> |  
                       <iterative do statement> | <end statement> |  
                       <procedure statement>  
  
<label definition> ::= <identifier>:
```

Basic Statements

Assignment Statement

```
<assignment statement> ::= <left part> =<expression>;  
<left part> ::= <variable reference> [, <variable reference>]...
```

CALL Statement

```
<call statement> ::= CALL <simple variable> [<parameter list>];  
<parameter list> ::= (<expression> [, <expression>]...)  
<simple variable> ::= <identifier> | <identifier>.<identifier>
```

GOTO Statement

```
<goto statement> ::= {GOTO | GO TO} <identifier>
```

Null Statement

```
<null statement> ::= ;
```

RETURN Statement

```
<return statement> ::= <typed return> | <untyped return>  
<typed return> ::= RETURN <expression>;  
<untyped return> ::= RETURN;
```


Microprocessor-dependent Statements

```
<microprocessor dependent statement> ::= <disable statement> |  
                                         <enable statement> |  
                                         <halt statement> |  
                                         <cause interrupt statement>
```

```
<disable statement> ::= DISABLE;
```

```
<enable statement> ::= ENABLE;
```

```
<halt statement> ::= HALT;
```

```
<cause interrupt statement> ::= CAUSE$INTERRUPT (numeric constant);
```

Scoping Statements

Simple DO Statement

<simple do statement> ::= DO;

DO-CASE Statement

<do-case statement> ::= DO CASE <expression>;

DO-WHILE Statement

<do-while statement> ::= DO WHILE <expression>;

Iterative DO Statement

<iterative do statement> ::= DO <index part> <to part> [<by part>];

<index part> ::= <index variable>=<start expression>

<to part> ::= TO <bound expression>

<by part> ::= BY <step expression>

<index variable> ::= <simple variable>

<start expression> ::= <expression>

<bound expression> ::= <expression>

<step expression> ::= <expression>

END Statement

<end statement> ::= END [<identifier>];

Procedure Statement

<procedure statement> ::= <procedure name> : PROCEDURE
 [<formal parameter list>] [<procedure type>]
 [<procedure attributes>;]

Conditional Clause

```
<conditional clause> ::= <if condition><>true unit> |  
                        <if condition><>true element> ELSE  
                        <>false element>  
  
<if condition> ::= IF <expression> THEN <>true unit>  
<>true element> ::= [<label definition>...] <do block> |  
                  [<label definition>...] <basic statement>  
  
<>false element> ::= <unit>  
  
<>true unit> ::= <unit>
```

DO Blocks

```
<do block> ::= <simple do block> | <do-case block> | <do-while block> |  
              <iterative do block>
```

Simple DO Blocks

```
<simple do block> ::= <simple do statement> [<declaration>...]  
                  [<unit>...] <ending>  
  
<ending> ::= [<label definition>...] <end statement>
```

DO-CASE Blocks

```
<do-case block> ::= <do-case statement> [<unit>...] <ending>
```

DO-WHILE Blocks

```
<do-while block> ::= <do-while statement> [<unit>...] <ending>
```

Iterative DO Blocks

```
<iterative do block> ::= <iterative do statement> [<unit>...]  
<ending>
```

Expressions

Primaries

`<primary> ::= <constant> | <variable reference> | <location reference>
| <subexpression>`

`<subexpression> ::= (<expression>)`

Constants

`<constant> ::= <numeric constant> | <string>`

Variable References

`<variable reference> ::= <data reference> | <function reference>`

`<data reference> ::= <name>[<subscript>] [<member specifier>]`

`<subscript> ::= (<expression>)`

`<member specifier> ::= .<member name>[<subscript>]`

`<function reference> ::= <name>[<actual parameters>]`

`<actual parameters> ::= (<expression>[, <expression>]...)`

`<member name> ::= <identifier>`

`<name> ::= <identifier>`

Location References

`<location reference> ::= @<constant list> | @<variable reference>`

`<constant list> ::= (<constant>[, <constant>]...)`

Operators

`<operator> ::= <logical operator> | <relational operator> |
<arithmetic operator>`

`<logical operator> ::= AND | OR | NOT | XOR`

`<relational operator> ::= < | > | <= | >= | <> | =`

`<arithmetic operator> ::= + | - | PLUS | MINUS | * | / | MOD`

Structure of Expressions

```
<expression> ::= <logical expression> | <embedded assignment>
<embedded assignment> ::= <variable reference> := <logical expression>
<logical expression> ::= <logical factor> | <logical expression>
                          <or operator> <logical factor>
<or operator> ::= OR | XOR
<logical factor> ::= <logical secondary> | <logical factor>
                    <and operator> <logical secondary>
<and operator> ::= AND
<logical secondary> ::= [<not operator>] <logical primary>
<not operator> ::= NOT
<logical primary> ::= <arithmetic expression> [<relational operator>
                    <arithmetic expression>]
<relational operator> ::= < | > | <= | >= | <> | =
<arithmetic expression> ::= <term> | <arithmetic expression>
                            <adding operator> <term>
<adding operator> ::= + | - | PLUS | MINUS
<term> ::= <secondary> | <term> <multiplying operator> <secondary>
<multiplying operator> ::= * | / | MOD
<secondary> ::= [<unary minus> | <unary plus>] <primary>
<unary minus> ::= -
<unary plus> ::= +
```

□ □ □

Differences Between PL/M Compilers

D

Differences between PL/M-86 and PL/M-80

PL/M-86 differs from PL/M-80 in the following respects:

- Support for floating-point arithmetic
- Support for signed arithmetic
- Addition of `REAL`, `INTEGER`, `POINTER`, and `SELECTOR` data types
- Addition of the `@` location reference operator
- Support for nested structures
- Expanded set of built-in procedures

In addition, the PL/M-80 reserved word `ADDRESS` is replaced by the PL/M-86 reserved word `WORD`. PL/M-80 has only the `BYTE` and `ADDRESS` data types. However, PL/M-86 has the following data types: `BYTE`, `WORD`, `DWORD`, `INTEGER`, `REAL`, `POINTER`, and `SELECTOR`.

The PL/M-86 rules for expression evaluation are more complete than those of PL/M-80. Other differences stem from the ones noted here. For example, an iterative `DO` block operates differently if its index variable is an `INTEGER` variable.

Compatibility of PL/M-80 Programs and the PL/M-86 Compiler

PL/M-80 programs that operate correctly on an 8080 microprocessor can be recompiled with the PL/M-86 compiler to produce code that will run on an 8086 microprocessor. You may need to edit the PL/M-80 source code to change identifiers that are PL/M-86 reserved words. (It is not necessary to change ADDRESS to WORD; ADDRESS is a PL/M-86 reserved word with the same meaning as WORD.)

Note that where PL/M-86 programs would normally have POINTER variables and location references formed with the @ operator, PL/M-80 programs have ADDRESS (WORD) variables and location references formed with the dot operator. PL/M-80 usage is less restricted than PL/M-86 usage, because arithmetic operations can be used on WORD values. In general, the PL/M-86 compiler supports PL/M-80 usage to provide upward compatibility. Some restrictions affect the types of expressions that can be used in the AT attribute, the INITIAL and DATA initializations, and location references. See also the discussions of size controls and the dot and @ operators in this manual.

Differences between PL/M-286 and PL/M-86

PL/M-286 differs from PL/M-86 in the following respects:

- POINTER and SELECTOR variables cannot be assigned absolute (i.e., constant) values. Only the equals operator (=) can be used with POINTER variables. For SELECTOR variables the logical (AND, OR, NOT, XOR) and relational (<, >, <=, >=, <>, =) operators can be used.
- Access to the hardware flag register is provided with the built-in variable FLAGS.
- Four built-in functions have been added to support multiple byte and word input: BLOCKINPUT, BLOCKINWORD, BLOCKOUTPUT, and BLOCKOUTWORD (available to PL/M-86 via the MOD86 | MOD186 control).
- The type of the STACKBASE variable has been changed from WORD to SELECTOR.
- New built-in procedures and functions have been added to support the 286 hardware protection model.
- Interrupt procedures are no longer assigned numbers in the source program. (This is done by the 286 system builder.) Interrupt procedures also cannot be called directly, and the SET\$INTERRUPT and INTERRUPT\$PTR built-ins have been removed.
- The memory array has been removed.

Compatibility of PL/M-86 Programs and the PL/M-286 Compiler

PL/M-86 programs that operate correctly on an 8086 microprocessor can be recompiled with the PL/M-286 compiler to produce code that will run on an 286 microprocessor. The PL/M-86 source code must be edited as follows:

- Assignments to the `STACKBASE` built-in variable must be changed from `WORD` to `SELECTOR`.
- All absolute pointer and selector assignments must be changed. (Pointers can be assigned a zero value using the new built-in function `NIL`.) Also, relational operations on pointer and selector values for any operation other than equality and inequality must be changed.
- The interrupt numbers on all interrupt procedures must be deleted. Interrupt vectors will be assigned to these procedures by the 286 system builder. Direct calls to interrupt procedures must also be changed.
- References to the `SET$INTERRUPT`, `INTERRUPT$PTR`, and `MEMORY` built-ins must be removed.

Differences between PL/M-386 and PL/M-286

PL/M-386 differs from PL/M-286 in the following respects:

- The string built-ins `FIND`, `CMP`, and `SKIP` return a value of `0FFFFFFFFH` for the not found and string equal results.
- Support for 64-bit unsigned scalars.
- Support for 8-bit and 32-bit signed scalars.
- Addition of `HWORD` and `QWORD` unsigned integers, and the `CHARINT`, `SHORTINT`, and `LONGINT` signed integer data types.
- `ADDRESS` is the same as `OFFSET` (and not as `WORD` as in PL/M-286).
- Support for casting functions.
- Support for `WORD32` and `WORD16` mapping for data type identifiers.
- Addition of the `WORD32|WORD16` primary compiler controls, which ensure PL/M data type and language compatibility.
- `MEDIUM` and `LARGE` segmentation controls no longer indicate unique meaning to the compiler; `MEDIUM` is interpreted as `SMALL` and `LARGE` is interpreted as `COMPACT` except when `LARGE` is used to indicate a subsystem whose name is unknown at compile time.

- Several new built-in procedures and functions have been added to support the new data types (for example, `CMPHW`, `BLOCKINHWORD`; see Chapters 9 and 10), and some bit-string operations (for example, `SCANBIT`, `MOVBIT`).
- The built-ins `CONTROL$REGISTER`, `DEBUG$REGISTER`, and `TEST$REGISTER` have been added to support the Intel386 microprocessor.
- The following built-ins have been added to support the Intel486 microprocessor: `BYTE$SWAP`, `TEST$REGISTER`, `INVALIDATE$DATA$CACHE`, `WB$INVALIDATE$DATA$CACHE`, and `INVALIDATE$TLB$ENTRY`.
- The `FLAT` and `MOD486` compiler controls have been added.

Compatibility of PL/M-286 Programs and the PL/M-386 Compiler

PL/M-286 programs can be compiled with the PL/M-286 compiler to produce code that will run on Intel386 and Intel486 microprocessors in 286 microprocessor mode and interface with PL/M-386 code through `INTERFACE (/286)`. PL/M-286 programs can be recompiled with the PL/M-386 compiler to produce code that will run on an Intel386 and Intel486 microprocessors in their native microprocessor mode.



Character Set E

This appendix lists the ASCII character set and indicates whether the characters are part of the PL/M source character set. Table E-1 is a list of codes.

Table E-1. Character Set

Dec	Hex	PL/M	Character
0	00	NO	NULL
1	01	NO	SOH
2	02	NO	STX
3	03	NO	ETX
4	04	NO	EOT
5	05	NO	ENQ
6	06	NO	ACK
7	07	NO	BEL
8	08	NO	BS
9	09	YES	HT
10	0A	YES	LF
11	0B	NO	VT
12	0C	NO	FF
13	0D	YES	CR
14	0E	NO	SO
15	0F	NO	SI
16	10	NO	DLE
17	11	NO	DC1
18	12	NO	DC2
19	13	NO	DC3
20	14	NO	DC4
21	15	NO	NAK
22	16	NO	SYN

Table E-1. Character Set (continued)

Dec	Hex	PL/M	Character
23	17	NO	ETB
24	18	NO	CAN
25	19	NO	EM
26	1A	NO	SUB
27	1B	NO	ESC
28	1C	NO	FS
29	1D	NO	GS
30	1E	NO	RS
31	1F	NO	US
32	20	YES	SP
33	21	NO	!
34	22	NO	"
35	23	NO	#
36	24	YES	\$
37	25	NO	%
38	26	NO	&
39	27	YES	'
40	28	YES	(
41	29	YES)
42	2A	YES	*
43	2B	YES	+
44	2C	YES	,
45	2D	YES	-
46	2E	YES	.
47	2F	YES	/
48	30	YES	0
49	31	YES	1
50	32	YES	2
51	33	YES	3
52	34	YES	4
53	35	YES	5

Table E-1. Character Set (continued)

Dec	Hex	PL/M	Character
54	36	YES	6
55	37	YES	7
56	38	YES	8
57	39	YES	9
58	3A	YES	:
59	3B	YES	;
60	3C	YES	<
61	3D	YES	=
62	3E	YES	>
63	3F	YES	?
64	40	YES	@
65	41	YES	A
66	42	YES	B
67	43	YES	C
68	44	YES	D
69	45	YES	E
70	46	YES	F
71	47	YES	G
72	48	YES	H
73	49	YES	I
74	4A	YES	J
75	4B	YES	K
76	4C	YES	L
77	4D	YES	M
78	4E	YES	N
79	4F	YES	O
80	50	YES	P
81	51	YES	Q
82	52	YES	R
83	53	YES	S
84	54	YES	T

Table E-1. Character Set (continued)

Dec	Hex	PL/M	Character
85	55	YES	U
86	56	YES	V
87	57	YES	W
88	58	YES	X
89	59	YES	Y
90	5A	YES	Z
91	5B	NO	[
92	5C	NO	\
93	5D	NO]
94	5E	NO	^
95	5F	YES	_
96	60	NO	`
97	61	YES	a
98	62	YES	b
99	63	YES	c
100	64	YES	d
101	65	YES	e
102	66	YES	f
103	67	YES	g
104	68	YES	h
105	69	YES	i
106	6A	YES	j
107	6B	YES	k
108	6C	YES	l
109	6D	YES	m
110	6E	YES	n
111	6F	YES	o
112	70	YES	p
113	71	YES	q
114	72	YES	r
115	73	YES	s

Table E-1. Character Set (continued)

Dec	Hex	PL/M	Character
116	74	YES	t
117	75	YES	u
118	76	YES	v
119	77	YES	w
120	78	YES	x
121	79	YES	y
122	7A	YES	z
123	7B	NO	{
124	7C	NO	
125	7D	NO	}
126	7E	NO	~
127	7F	NO	DEL

□ □ □

Linking to Modules Written in Other Languages

F

Introduction

This appendix describes the calling conventions used by the [x]86 family of languages. These calling conventions are standardized so that a module written in PL/M can freely call procedures, subroutines, and subprograms in other modules written in other [x]86 languages.

The information in this appendix is not necessary to call PL/M procedures and functions from PL/M. See Chapter 8 for information about parameters and arguments.

The calling conventions and stack and register usage described in this appendix are needed to call ASM subroutines. Also, the corresponding data types listed at the end of this appendix are needed to write a subroutine that can pick up the data in the PL/M program. Refer to the ASM macro assembler operating instructions for more information about combining PL/M programs with ASM programs and for examples.

The easiest way to ensure compatibility between assembly-language subroutines that are combined with PL/M programs or procedures is to write a dummy procedure in PL/M. This procedure would have the same argument list and the same attributes as the assembly language subroutine. Then compile the PL/M procedure with the correct segmentation control and the `CODE` control. This will produce a pseudo-assembly listing of the generated microprocessor code, which can then be copied to the prologue and epilogue of the assembly language subroutine.

With PL/M, separate modules can be written and compiled, and combined at a later time. This allows you to create separately tested modules that are combined after they are internally bug-free. Not all modules have to be in PL/M: you can choose the appropriate language for each module. Be sure to combine the modules properly with a binder or a linker in order to satisfy references to externals. Because the [x]86 languages (excluding C) follow the same calling sequence, control will pass to a called module correctly. The standard calling sequence is described in the following section.) However, the called module might not be able to deal intelligently with the data passed to it since languages treat some data structures differently.

By specifying arguments in a reference to an external procedure, data is passed to the external procedure. The number of arguments and the order in which they are specified must match the number and order of the corresponding parameters in the external procedures declaration (see Chapter 8).

All arguments for parameters are passed on the microprocessor's stack, or the numeric coprocessor's register stack, in the order in which they were specified. For Intel386 and Intel486 microprocessors, the space occupied by a parameter pushed on the microprocessor's stack is always a multiple of four bytes. Functions return non-real values in a register, and REAL values on the top of the numeric coprocessor's register stack.

Calling Sequence

The calling sequence for each procedure activation places the procedure's actual parameters (if any) on the stack and then activates the procedure with a `CALL` instruction.

Parameters are placed on the microprocessor's stack or the numeric coprocessor's register stack in left-to-right order. Because the stack grows from higher locations to lower locations, the first parameter occupies the highest position on the stack, and the last parameter occupies the lowest position. Stack representation for the different PL/M parameters is described in Table F-1.

Table F-1. Stack Representation for PL/M Parameters

Parameter	Intel386 and Intel486 CPU Stack Representation
BYTE	Four bytes, with the higher three bytes undefined.
CHARINT	Four bytes, with the higher three bytes undefined.
HWORD	Four bytes, with the high two bytes undefined.
SELECTOR	Four bytes, with the high two bytes undefined.
SHORTINT	Four bytes, with the high two bytes undefined.
WORD	Four bytes, with no undefined bytes.
OFFSET	Four bytes, with no undefined bytes.
INTEGER	Four bytes, with no undefined bytes.
REAL	Four bytes, with no undefined bytes.
DWORD	Eight bytes with the high 32 bits pushed first and the low 32 bits 16 bits pushed second.

For Intel386 and Intel486 microprocessors, a `POINTER` parameter in the `SMALL (ROM)` and `COMPACT` cases consists of a selector and an offset. The 16-bit selector is pushed first, followed by the 32-bit offset.

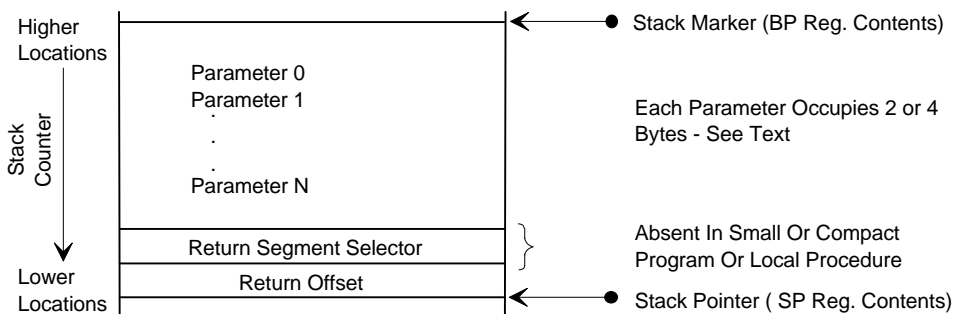
The left-most seven `REAL` parameters are passed on the numeric coprocessor's stack. If more than seven `REAL` parameters are present, the rest (after the left-most seven) are passed on the microprocessor's stack and are intermixed with the other non-real parameters in the order in which all parameters were declared.

After the parameters are passed, the `CALL` instruction places the return address on the stack. In the `SMALL` and `COMPACT` cases with local (or non-exported) procedures, this address is a 32-bit offset (the contents of the EIP register) and occupies four bytes on the stack.

For procedures exported from a subsystem, the return address is a `POINTER` value consisting of a selector and offset; the return address is placed on the stack in the same way a `POINTER` parameter is passed. The 16-bit segment selector (contents of the CS register) is pushed first, then the 32-bit offset (EIP register contents) is pushed.

For all of the microprocessors, control is passed to the code of the procedure by updating the EIP register. For procedures exported from a subsystem, the CS register is also updated.

Figure F-1 shows the stack layout at the point where the procedure gains control.



OSD540

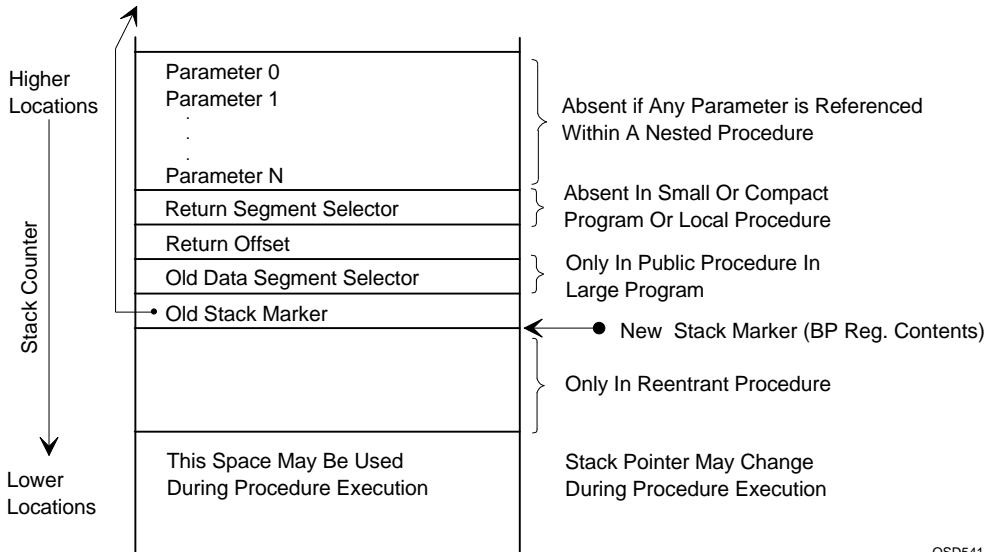
Figure F-1. Stack Layout at Point Where a Non-interrupt Procedure is Activated

Procedure Prologue

In compiling the procedure itself, the compiler inserts a sequence of instructions called the procedure prologue. The procedure prologue varies depending on the type of procedure being compiled as follows:

- If the procedure has the `PUBLIC` attribute and the program size is `LARGE`, or if it is exported from a subsystem, the content of the `DS` register is placed on the stack and is then updated to the data segment of the procedure. `ES` is set to `DS`.
- If any parameter of the procedure is referenced by a nested procedure, all parameters are copied from the stack to space reserved for them in the data segment.
- The stack marker offset (`EBP` register contents) is placed on the stack, and the current stack pointer (`ESP` register contents) is used to update the `BP` or `EBP` register.
- If the procedure has the `REENTRANT` attribute, space is reserved on the stack for any variables declared within the procedure (this does not include based variables, variables with the `DATA` attribute, or variables with the `AT` attribute).

Control then passes to the code compiled from the executable statements in the procedure body. Figure F-2 shows the stack layout at this point.



OSD541

Figure F-2. Stack Layout During Execution of a Non-interrupt Procedure Body

Procedure Epilogue

To return from the procedure, the compiler inserts an instruction sequence called the epilogue. This accomplishes the following:

- If the compiler has used stack locations for temporary storage or local variables during procedure execution, the stack pointer (the ESP register) is loaded with the stack marker (the EBP register), discarding the temporary storage.
- The old stack marker is restored by popping the stored value from the stack into the EBP register.
- If the procedure has the `PUBLIC` attribute and the program size is `LARGE` or it is exported from a subsystem, the old data segment selector is restored by popping the stored value from the stack into the DS register. Additionally, ES is set to DS.
- The stored return address (a 32-bit offset) is popped into the EIP register. If the procedure is exported, the stored return address selector is also popped into the CS register. Any parameters stored on the stack are discarded.

Register Usage

Table F-2 provides a summary register usage.

Table F-2. Summary of the Intel386 Microprocessor Register Usage

Register	Must Preserve	Usage
EAX	No	Return BYTE (AL), HWORD (AX), WORD, DWORD, CHARINT (AL), SHORTINT (AX), INTEGER, SELECTOR (AX), POINTER offset portion, and OFFSET.
EBX	No** (Yes, when using C language interface.)	--
ECX	No	--
EDX	No	Return upper half of DWORD values, POINTER segment selector.
ESP	Yes*	Stack pointer
EBP	Yes	Stack marker
ESI	No (Yes when using C language interface.)	--
EDI	No (Yes when using C language interface.)	--
FLAGS	No	--
CS	Yes	Called procedure's code segment.
DS	Yes	Caller's data segment.
SS	Yes	Caller's stack segment.
ES	Yes	Caller's data segment.
FS, GS	No	--

* ESP must be adjusted so that all arguments are removed from the stack on return (except when using C language interface).

** The C language interface referred to in the table is the variable parameter list format.

The numeric coprocessor's stack contains the first seven `REAL` arguments passed by the calling program. The numeric coprocessor's status word is unknown and does not need to be saved. If the status word is changed, the numeric coprocessor's mode word must be saved on entry and restored before exit.

If an assembly language subroutine alters the `DS` or `SS` registers, and expects to be called by a `PL/M` program, the subroutine must save the contents of these registers upon entry and restore them before returning to the `PL/M` program. Additionally, the `CS` and `ES` registers must be preserved by the called procedure.

`PL/M` uses the `ESP` and `EBP` registers to address the stack. If a called assembly language subroutine uses the stack register, the subroutine must save the contents of the register on entry and restore the register's contents before returning control to the `PL/M` program. Before returning, the called subroutine must also adjust the `ESP` register to remove all parameters from the microprocessor's stack. Additionally, the `CS` and `ES` registers must be preserved by the called procedure.

The `EAX`, `EBX`, `ECX`, `EDX`, `ESI`, `EDI`, `FS`, and `GS` registers do not need to be preserved. A called subroutine can freely use these registers.

An assembly language program calling a `PL/M` procedure cannot expect the contents of the general-purpose registers (`EBP` and `ESP`) to be preserved. If the contents of these registers are needed, they must be saved prior to calling the `PL/M` procedure.

Table F-3 summarizes the microprocessor registers used to hold simple data types that are returned by typed procedures.

Table F-3. Registers Used to Hold Simple Data Types

Intel386 Microprocessor Procedure Type	Register
BYTE	AL
CHARINT	
HWORD	AX
SHORTINT	
DWORD	DX:AX
INTEGER	AX
WORD	EAX
OFFSET	
INTEGER	
DWORD	EDX:EAX
POINTER (SHORT, SMALL RAM)	EAX
POINTER (LONG, COMPACT, SMALL ROM)	EDX:EAX
SELECTOR	AX
REAL	Top of the numeric coprocessor's stack.

Segment Name Conventions

Tables F-4 summarizes the segmentation of a subsystem under the `SMALL` and `COMPACT` program controls. The table shows the name of the segment in which each type of program section is stored for each control and for subsystems.

Table F-4. Summary of PL/M-386 Segment Names

Model	SubModel	Code	Data	Const	Stack
SMALL	IN DATA	CODE32	DATA	DATA	DATA
	IN CODE	CODE32	DATA	CODE32	DATA
SMALL (subsystem)	IN DATA	S_CODE32	DATA	DATA	DATA
	IN CODE	S_CODE32	DATA	S_CODE32	DATA
COMPACT	IN DATA	S_CODE32	S_DATA	S_CODE32	STACK
	IN CODE	S_CODE32	S_DATA	S_CODE32	STACK
COMPACT (subsystem)	IN DATA	S_CODE32	DATA	S_DATA	STACK
	IN CODE	S_CODE32	DATA	S_CODE32	STACK

Notes:

CODE32 denotes a segment name composed of CODE32.

DATA denotes a segment name composed of DATA.

S_CODE32 denotes a segment name composed of the subsystem name and CODE32.

S_DATA denotes a segment name composed of the subsystem name and DATA.

C Language Compatibility

The iC-*n*86 calling conventions, procedure prologue and epilogue, and register usage differ from other Intel *n*86 languages. However, the `INTERFACE` control, described in Chapter 11, allows C procedures to call procedures written in PL/M and vice versa.

The procedure prologue and epilogue and register usage for the VPL (variable parameter list) calling convention for iC-*n*86 differ from other *n*86 languages. These differences are as follows:

- All parameters (real and non-real), are passed on the microprocessor's stack. The last parameter is pushed first and the first parameter is pushed last so that the first parameter is in the lowest memory location.
- An integral parameter that is four bytes must be zero or sign-extended, as required by the C language.
- The space occupied by a parameter pushed on the microprocessor stack is always a multiple of four bytes for Intel386 and Intel486 microprocessors.
- Both short (floating-point) and long (double) real parameters are pushed as long real parameters, as required by the C language. Therefore, all real parameters passed from or to iC-386 procedures must be typed as 64-bit `REAL` in the PL/M-386 code.
- The calling procedure pops the parameters from the microprocessor stack after the called procedure has returned. Except when the called procedure is a function returning real results, the called procedure must not leave any entries in the numeric coprocessor stack.
- The ESP, EBP, CS, DS, ES, and SS registers should be preserved by the called procedure. (They are used for global storage). The EBX, ESI, and EDI registers should also be preserved by the called procedure. These registers can be used by the caller for local data storage.
- The EAX, ECX, EDX, FS and GS registers do not need to be preserved by the called procedure.

Design Guidelines

The following guidelines should be followed when combining C and PL/M modules. These guidelines are demonstrated in the code example which follows afterwards.

1. Identify all C functions which use the VPL calling convention. Library function calling conventions are found by checking the .h include files.
2. Use the PL/M-386 INTERFACE compiler control to allow the PL/M compiler control to generate VPL code.
3. All PL/M functions should be in a "#pragma fixed-params("plmf,...") list. This will guarantee that any call to a PL/M function will use the FPL calling convention.
4. Compile all files and link them in the same way C files are linked.

Code Example

This code example, run under the iRMX Operating System, discusses how a PL/M application makes C function calls. An example of this is when a large PL/M application is being converted to C. Mixing C and PL/M modules allows the converted C modules to be debugged one at a time after conversion. Another example is a PL/M application which needs access to the extensive I/O routines available in C libraries.

The PL/M example, named *ptest.plm*, shows how a PL/M function calls a C function that uses the VPL calling convention. It also includes a C procedure called by the C example named *candplm.c*, which the FPL calling convention.

The code example uses the \$INTERFACE control to signal the compiler that **printf** is a C function that uses the VPL convention. Any function that has the "varparams" attribute should be included in the \$INTERFACE list, such as the **printf** function. C functions compiled under the iC-386 C compiler use the FPL convention by default and should not be included in the "varparams" list.

```

/* PL/M module - ptest.plm */

/* Only printf uses the VPL convention */
  $INTERFACE(C=printf)

    ptest: DO;

/* Uses VPL convention */
    printf:PROCEDURE EXTERNAL;
    END printf;

/* Uses FPL convention */
    c_call_plm_funct:PROCEDURE (i) WORD EXTERNAL;
    DECLARE i WORD;
    END c_call_plm_funct;

/* Uses FPL convention */
    c_fpl_funct:PROCEDURE (i) WORD EXTERNAL;
    DECLARE i WORD;
    END c_fpl_funct;

/* This procedure is called by a C function which uses FPL */
    plmproc:PROCEDURE (i) WORD PUBLIC;
    DECLARE i WORD;
    i = i-1;
/* The string in the PL/M call to printf terminates with 0DH, 0AH,
00H so it conforms to C string conventions. These symbols cause a
<CR>, <LF>, and C end of string. */

/* Call to VPL C function. */
    CALL printf(@(0Dh,'In plmproc, i = %d',0Dh,0aH,00H),i);

/* Call to FPL C function */
    i = c_fpl_funct(i);
    RETURN(i);
    END plmproc;

```

```
/* This main function is written in PL/M and may be used as a
substitute for a C main modules. The function calls printf to
demonstrate how PL/M calls a C function which uses VPL calling
conventions. It also calls a C function which uses FPL calling
conventions. */
```

```
main:PROCEDURE PUBLIC;
DECLARE i WORD, j WORD, k DWORD;
i = 5;
k = 12345678H;
```

```
/* Call to VPL C function */
CALL printf(@('The value of i = %d, and k =
%xH',0Dh,0aH,00H),i,k);
```

```
/* Call to FPL C function */
j = c_call_plm_func(i);
CALL printf(@('The value of j = %d',0aH,00H)j);
END main;
END ptest;
```

The following C example, named *candplm.c*, demonstrates how a PL/M call is made from a C application.

```
/* C module - candplm.c */

#include <stdio.h>
#include <reent.h>
#include <locale.h>

/* Sets FPL for PL/M functions */
#pragma fixedparams("plmproc")
extern unsigned int plmproc(unsigned int);

/* The C function c_fpl_func uses the FPL calling conventions. The
PL/M function "plmproc" calls this function to demonstrate how a
PL/M procedure calls a C function which uses the FPL calling
convention. */

unsigned int
c_fpl_func(unsigned int i)
{
    i -= 1;
    printf("c_fpl_func, i = %d\n",i);
    return (i);
}

/* The C function call_plm_func uses the FPL calling conventions.
This function calls "plmproc" to demonstrate how a C function calls
a PL/M function. */

unsigned int
c_call_plm_func(unsigned int i)
{
    i -= 1;
    printf("c_call_plm_func, i = %d\n",i);
}

/* Call to PL/M function */
i = plmproc(i);
return(i);
}
```


Compiling C and PL/M Modules

The submit file, named *plmsub.csd*, contains the following command syntax to compile and bind *pctest.plm* and *candplm.c*.

```
ic386 candplm.c debug code compact
plm386 pctest.plm debug code compact
bnd386 /intel/obj/cstart32.obj, &
pctest.obj, candplm.obj, &
/intel/lib/clibxf32.lib &
renameseg (code32 to code) &
segsz (stack(2400H)) &
debug object(plmsub) &
rc(dm(4000h,0FFFFFFh))
```

To invoke the submit files, enter the following command at the iRMX "-" prompt:

```
- submit plmsub over plmsub.out echo
```

When the processing stops and the prompt has returned, enter the following to run the example:

```
- plmsub
```

The output of the code example follows:

```
The value of i = 5, and k is 12345678H
c_call_plmfunct, i = 4
In plmproc, i = 3
c_fpl_funct, i = 2
The value of j = 2
```

Note that the variable "i" changes as it is passed as a parameter. The value is initially set to 5. As it passes through each function, it is decremented and its new value is displayed.

□ □ □

Run-time Interrupt Processing **G**

General Information

Interrupts can be initialized when the CPU receives a signal on its maskable interrupt pin from a peripheral device, or when control is transferred to an interrupt vector by the `CAUSE$INTERRUPT` statement. If the program runs under an operating system that traps interrupts, the information in this appendix may not be applicable.

Note that the CPU does not respond to the interrupt signal unless interrupts are enabled. The PL/M-386 compiler do not generate any code to enable or disable interrupts at the start of the main program.

If interrupts are enabled and vectored through an interrupt gate, the following actions take place:

1. The CPU completes any instruction currently in execution.
2. The CPU issues an acknowledge interrupt signal and waits for the interrupting device to send an interrupt number.
3. The CPU flag register is placed on the stack (occupying two bytes of stack storage).
4. Interrupts are disabled by clearing the IF flag.
5. Single stepping is disabled by clearing the TF flag.
6. The CPU activates the interrupt procedure corresponding to the interrupt number sent by the interrupting device.
7. When that procedure terminates, the stack is automatically restored to the state it was in when the interrupt was received, and control returns to the point where it was interrupted.

The mechanism for this activation and restoration are described in the following sections. If interrupts are vectored through a trap gate, the fourth step is not performed; if they are vectored through a task gate, all seven steps are replaced by a task switch.

See also: interrupt processing, *System Concepts*

The Interrupt Descriptor Table

The interrupt descriptor table (IDT) contains descriptors that vector interrupts, traps, and protection exceptions to their respective handling routines.

These descriptors are called gates; they can be either interrupt gates, trap gates, or task gates. Interrupt gates and trap gates point to a particular entry point in the address space of the interrupted user (i.e., to an interrupt procedure). Task gates point to an interrupt processing task state segment (TSS).

BLD386 sets up the IDT and to assign numbers to vector the individual gates to the appropriate interrupt procedure or task. For more information, see the *Intel386 Family Utilities User's Guide*.

The IDT can hold up to 256 gates. Gates 0 through 31 are reserved for internal use.

Procedures and Tasks

For Intel386 and Intel486 microprocessors, when an interrupt is vectored through an interrupt gate, all registers must be pushed onto the stack, and interrupts are automatically disabled. (Interrupts must be explicitly enabled.) The interrupt procedure then begins execution. The interrupt procedure ends with an IRET instruction that acts as a normal return. Hence, execution starts at the beginning of a procedure each time it is entered.

The interrupt process differs for an interrupt vectored through a task gate. The registers for the interrupted task are saved in the TSS, and the microprocessor's registers are loaded from the TSS of the interrupt task. Thus, no explicit register saving is necessary. Interrupts are enabled or disabled depending on the flag settings in the interrupt task's TSS during execution of the interrupt task (unless explicitly changed). This enables interruption of the interrupt task. However, a protection violation occurs if an interrupt task is busy and an attempt is made to vector through the busy interrupt task.

The interrupt task also ends with an IRET instruction, but in this case it acts as a task switch, saving the status of the outgoing interrupt task in memory. When the task is re-entered, execution continues at the first instruction after the IRET instruction.

Interrupt Procedure Prologue and Epilogue

An interrupt procedure begins by declaring its name and its `PUBLIC` or `EXTERNAL` attribute. The following interrupt procedure declaration is the correct form for PL/M-386:

```
HANDLER: PROCEDURE INTERRUPT PUBLIC;
```

This alerts the compiler to create a code prologue appropriate to a routine that will, in general, be invoked by interrupts.

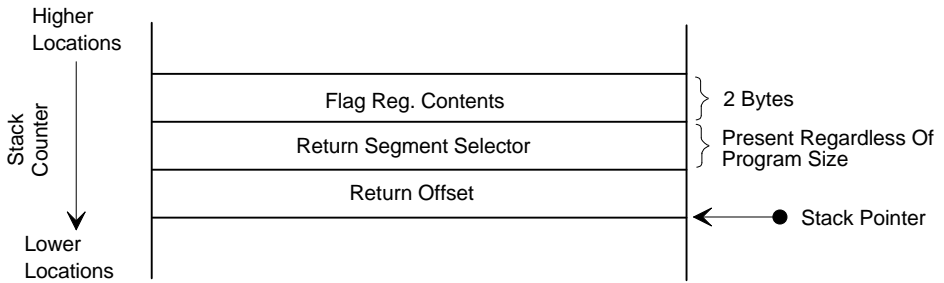
At the beginning of each interrupt procedure, the interrupt procedure prologue inserted by the compiler accomplishes the following tasks:

1. Pushes the CPU registers onto the stack in the following order: EAX, ECX, EDX, EBX, ESP, EBP, ESI, EDI.
2. Pushes the ES, FS, GS, and DS register content on the stack.
3. If the interrupt procedure has the `PUBLIC` attribute, and if it is exported from a subsystem, the contents of the DS register is placed on the stack and is then updated to the data segment of the procedure. In addition, ES is set to DS.
4. The stack marker offset (EBP register contents) is placed on the stack, and the current stack pointer (ESP register contents) is used to update the EBP register.
5. If the procedure has the `REENTRANT` attribute, space is reserved on the stack for any variables declared within the procedure. (This does not include based variables, variables with the `DATA` attribute, or variables with the `AT` attribute.)

⇒ **Note**

The compiler may temporarily use the DS register and the ES register in some cases (e.g., string built-ins), but always restores it. Take care to note this possibility when writing an interrupt procedure in assembly language.

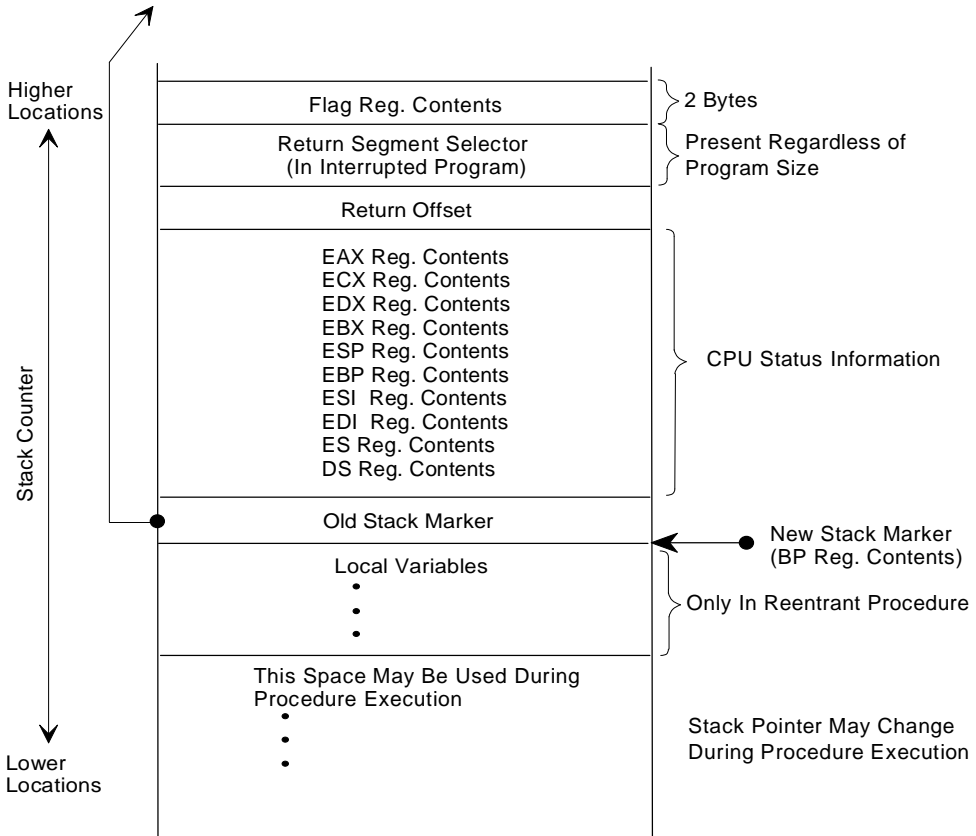
At the point where the interrupt procedure prologue gains control, the stack layout is as shown in Figure G-1.



OSD542

Figure G-1. Stack Layout at Point Where an Interrupt Procedure Gains Control

After the interrupt procedure prologue is executed (at the point where the code compiled from the procedure body gains control), the stack layout is as shown in Figure G-2.



OM02063

Figure G-2. Stack Layout during Execution of Interrupt Procedure Body

The return from the procedure body is called the interrupt procedure epilogue; it restores the stack to the state it was in before the interrupt occurred. The interrupt procedure epilogue contains the following steps:

1. If the compiler has used stack locations for temporary storage or local variables during procedure execution, the stack pointer (the ESP register) is loaded with the current stack marker (the EBP register) discarding the temporary storage.
2. The old stack marker is restored by popping the stored value from the stack into the EBP register.
3. The old data segment is restored by popping the stored value from the stack into the DS register. This step will occur only if the procedure has a `PUBLIC` attribute and it is exported from a subsystem.
4. The stack is popped into the CPU registers in the following order: EDI, ESI, EBP, ESP, EBX, EDX, ECX, EAX. Note that the ESP register value is discarded.
5. An `IRET` instruction is executed to return from the interrupt procedure restoring the IP or EIP, CS, and the flag register contents from the stack.

At this point, the stack has been restored to the state it was in before the interrupt occurred, and processing continues normally.

Interrupt Tasks

A task on the microprocessor is a single thread of execution; that is, a stream of instructions and data with a task state image. The task state image is made up of the contents of the task registers, the task's status word, and the virtual locations of the task's instructions and data segments.

Tasks are initiated with a task switch operation. The CPU stores the task state image of the outgoing task (held in the processor registers) in memory, and loads the task state image of the incoming task into task registers. Because all the registers are reloaded and a new address space is entered, it is impossible to jump directly from one task to another.

Interrupt tasks are frequently written as one loop. At the beginning is the code needed to initialize the task, followed by the steps needed to handle the interrupt. Call the `WAITFORINTERRUPT` built-in procedure (see Chapter 10) to generate an `IRET` instruction. When the task is activated again, execution continues at the instruction following the `IRET`, with all the registers unchanged. At the end of the interrupt task, use a `GOTO` statement to loop back to the top of the interrupt task. Thus, an interrupt task never terminates, unless an operating system function removes the task.

Use of the `WAITFORINTERRUPT` procedure is demonstrated in the following example. This task is designed to handle messages that arrive in pieces, each one being preceded by an interrupt.

```
TASK: DO
    DECLARE local variables;
    local procedures;
NEW$MESSAGE:
    CALL INITIALIZE$MESSAGE$PROCESSING;
    DO FOREVER;
        CALL WAIT$FOR$INTERRUPT;
/* IRET to wait for next interrupt, which continues here */
    CALL PROCESS$PIECE$OF$MESSAGE;
    IF LAST$PIECE$OF$MESSAGE$ THEN DO;
        CALL TERMINATE$MESSAGE$PROCESSING;
        CALL WAIT$FOR$INTERRUPT;
        /* IRET to wait for start of next message */
    GOTO NEW$MESSAGE
    END;
END;
END TASK;
```

Exception Conditions in REAL Arithmetic

Six exception conditions can occur during normal numeric operations:

- Invalid operation
- Denormalized operand
- Zero divide
- Overflow
- Underflow
- Precision

These exceptions are discussed in the following sections. In each case, only a few of the possible causes are described because most are not likely to occur with PL/M usage. To perform sophisticated numeric processing of extreme precision and flexibility, refer to the microprocessor-specific programmer's reference manual.

The six exceptions and their default responses are summarized in Table G-1.

As the following sections indicate, the masked, default response to most exceptions will provide the least abrupt, most appropriate action for PL/M applications. Many real math exceptions that occur in other processors will not occur with the numeric coprocessor because of the extended range of intermediate results it holds. The soft recovery of gradual underflow (described in the denormalized exception section) also extends the range of permissible execution rather than reporting a hard-failure condition.

Table G-1. Exception and Response Summary

Exception	Masked Response	Unmasked Response
Invalid Operation	In one operand is NAN, return it; if both are NAN's return NAN with larger absolute value; if neither is NAN, return indefinite NAN.	Request interrupt. (Numeric coprocessor stack unchanged.)
Zero divide	Return infinity signed with "exclusive or" of operand signs.	Request interrupt. (Numeric coprocessor stack unchanged.)
Denormalized	Memory operand: proceed as usual. Register operand: convert to valid unnormal, then reevaluate for then reevaluate for exceptions.	Request interrupt. (Numeric coprocessor stack unchanged.)
Overflow	Return properly signed infinity.	Register destination: adjust exponent, store result (see note), request interrupt. Memory destination: request interrupt.
Underflow	Denormalize result.	Register destination: adjust exponent, store result (see note), request interrupt. Memory destination: request interrupt.
Precision	Return rounded result.	Return rounded result, request interrupt.

Note:

On overflow, 24,576 decimal is subtracted from the true result's exponent. This forces the exponent back into range and enables a user exception handler to ascertain the true result from the adjusted result that is returned. On underflow, the same constant is added to the true result's exponent.

Programmers who use the recommended setting of the `REAL` mode word (see Chapter 10) need to handle only the invalid exception. Study of the other exception conditions is advised, however, to gain a general understanding of their use.

Invalid Operation Exception

This exception generally indicates a program error. It could be caused by referencing an uninitialized REAL variable or by referencing a location that does not contain a REAL value (as might occur with an out-of-range subscript for a REAL array). Attempting to take the square root of a negative number or to store a number too large for integer format would also generate this exception.

Another interpretation of this exception is stack error. This may be caused by failing to restore the math unit status before returning from an interrupt routine that had saved the status. Another cause is the generation of more than eight intermediate results during REAL arithmetic, which can arise if REAL procedure function calls are nested too deeply. The compiler ensures that no single procedure does this, but cannot check what may occur only at run time. This exception can also occur when REAL functions (typed procedures) are used as operands within longer REAL expressions. For example:

```
DELTA$1 = ALPHA * (BETA/GAMMA) + CHI (PSI, RHO, PI)
```

where all these names are typed REAL and CHI is some previously declared REAL function having three parameters.

The following is less likely to cause an exception condition:

```
EPS = CHI (PSI, RHO, PI)
DELTA$1 = ALPHA * (BETA/GAMMA) + EPS
```

because all REAL arithmetic is performed using the numeric coprocessor's stack, which has eight registers. The first seven REAL parameters supplied in procedure calls are placed on this stack. If the procedure is typed (i.e., invoked as a function), it can be embedded as one operand within a longer REAL expression.

Because the evaluation of such an expression also involves the use of this stack for prior and subsequent arithmetic operations, stack overflow may occur. This overflow amounts to unpredictable destruction of original parameters or intermediate results. It becomes more likely as the complexity of REAL expressions containing REAL functions is increased. Thus, it is safer to use an assignment statement first to store the function's value in a real variable; then use that variable in the larger expression.

If stack error might apply, modify the code for the effected procedures to call the built-in procedures SAVE\$REAL\$STATUS and RESTORE\$REAL\$STATUS as their first and last operations, respectively.

The masked (default) response is to set the result to one of the special bit patterns called Not-A-Number (NaNs), usually the indefinite value, the smallest NaN representable in the specified precision. It also sets bit 0 of the REAL error byte.

If bit 0 of the REAL mode word is 0 (invalid exception unmasked), an interrupt occurs, transferring control to the user-written interrupt handler.

Denormal Operand Exception

This condition arises when numeric operations have resulted in a number whose exponent is literally zero and whose significand is non-zero, or have resulted in a number whose significand does not begin with a one. Denormals usually arise in response to masked underflow. Gradual underflow is the masked, default response to underflow. A small denormal added to a large normal REAL number can give an acceptable, in-range answer if the denormal exception is masked. In practice, denormals are very rare since intermediate results are kept in temporary real format (15-bit exponent).

This condition causes bit 1 of the REAL error byte to be set to 1. If bit 1 of the REAL mode word is 1, the response described previously occurs. If bit 1 is 0, an interrupt occurs, transferring control to the user-written interrupt handler.

Zero Divide Exception

This condition arises when in the course of some REAL computation a divisor turns out to be zero. The masked response, when bit 2 of the REAL mode word is 1, is to return infinity appropriately signed. If bit 1 is 0, an interrupt occurs, giving control to the user-written interrupt handler. In either case, bit 2 of the REAL error byte is set to 1.

Overflow Exception

This error occurs when a real result is too large for the format in use. For assigning to REAL scalar types, it occurs if the result is greater than about 3.37×10^{38} . For intermediate REAL computations, it occurs if the result is greater than about 10^{4932} . The overflow exception can arise during assignment, addition, subtraction, multiplication, division, or conversion to integer.

The masked, default response (bit 3 of REAL mode word = 1) is to return infinity (signed if affine mode is set) and set bit 3 of the REAL error byte to 1. Unmasked overflow must go through a user-written interrupt handler.

Underflow Exception

Underflow exception is caused by an exponent too small for the format in use. For REAL assignments, it occurs if the exponent is less than -127; and for intermediate results if the exponent is less than -16383. Underflow can be caused by the same type of REAL operations as overflow.

The masked, default response (bit 4 of REAL mode word = 1) is to use the denormal number created by adjusting the very small result. It adjusts the significand, moving significant digits off to the right and raising the exponent until the latter becomes non-zero. For example, with single precision values, a 24-bit significand of .01 with an exponent of zero implies the number $1 \times 2^{*-129}$ because a zero exponent in this format means -127. If the denormal exception is masked, this number would be adjusted into a significand of .001 with an exponent of 1 (i.e., $0.001 \times 2^{*-126}$), prior to operation. This number would then be available for use in subsequent REAL operations that might yield valid results. Zero is returned if it is the rounded result. Bit 4 of the REAL error byte is set to 1. Unmasked underflow must go through a user-written interrupt handler.

Precision Exception

This error occurs when the result of an operation is inexact (i.e., rounded) or when an overflow exception occurs. No special action is performed by a masked response (bit 5 of REAL mode word = 1) other than setting bit 5 of the REAL error byte. Unmasked response is as chosen by the user.

Writing a Procedure to Handle REAL Interrupts

This section partially summarizes the information pertaining to interrupts, floating-point usage, and procedures. (Additional facilities for handling REAL interrupts may be provided by the operating system, or can be performed with the system builder.)

An interrupt-handling procedure may, for example, begin as follows:

```
HANDLER: PROCEDURE INTERRUPT PUBLIC;
```

If HANDLER will do any REAL arithmetic or assignments, its first executable statements should be of the form:

```
ERR$INFO = GET$REAL$ERROR;  
/* must declare ERR$INFO$ BYTE earlier */
```

or:

```
CALL SAVE$REAL$STATUS (@Local_Save_Area);  
/* also declare earlier */
```

Each procedure clears the error byte. The latter procedure also clears out the REAL stack. Thus, after either procedure is used, the REAL error byte no longer contains the flagged cause of the exception condition that invoked HANDLER.

Using SAVE\$REAL\$STATUS is a way of avoiding possible stack errors from cumulative usage. This enables errors in HANDLER to be detected independently of the originating exception condition. It also enables HANDLER to restore the state of the interrupted procedure despite HANDLER's own use of the REAL facility. SAVE\$REAL\$STATUS also makes available all the information regarding the state of the numeric coprocessor exceptions, stack, and operations, as shown in the following paragraph.

Thus, the beginning of a typical routine to handle REAL exception conditions could look like this:

```
HANDLER: PROCEDURE INTERRUPT PUBLIC;  
    DECLARE ERR$INFO BYTE;  
    DECLARE LOCAL$SAVE$AREA (94) BYTE;  
    ERR$INFO = GET$REAL$error;
```

or, to perform extensive manipulations on the save area, declare a structure permitting access to the save area's component parts by name and/or byte, as follows:

```
HANDLER: PROCEDURE INTERRUPT PUBLIC;  
    DECLARE ERR$INFO BYTE;  
    DECLARE SAVE$AREA STRUCTURE (  
        CONTROL(2)    BYTE,  
        STATUS(2)     BYTE,  
        TAG            HWORD,  
        INSTR_OFF     WORD,  
        INSTR_SEL     SELECTOR,  
        OPERAND_OFF   WORD,  
        OPERAND_SEL   SELECTOR,  
        STACK_TOP(5)  WORD,  
        STACK_ONE(5)  WORD,  
        STACK_TWO(5)  WORD,  
        STACK_3 (5)   WORD,  
        STACK_4 (5)   WORD,  
        STACK_5 (5)   WORD,  
        STACK_6 (5)   WORD,  
        STACK_7 (5)   WORD);  
    CALL SAVE$REAL$STATUS (@SAVE_AREA);  
    ERR$INFO = SAVE_AREA.STATUS(0);
```



Note

To make use of the TAG word, use the masks and shifts to access the individual fields shown in Figure G-3.

Call either the `SAVE$REAL$STATUS` procedure or the `GET$REAL$ERROR` function, but not both. If the extra information gained by the save is not needed (i.e., only the exceptions are needed), use the `GET$REAL$ERROR` function. If both are called, the second call will produce incorrect results.



Tag Values:

00 = Valid (Normal or Unnormal)

01 = Zero (True)

10 = Special (Not-A-Number, ∞ , or Denormal)

OSD544

11 = Empty

Figure G-3. Tag Word Format

The rest of HANDLER can perform any appropriate actions. This is an application dependent decision. Among the possibilities:

- Incrementing an exception counter for later display
- Printing diagnostic data (e.g., the contents of SAVE\$AREA)
- Aborting further execution of the calculation causing exception
- Aborting all further execution

The format of the LOCAL_SAVE_AREA as it is filled by the save procedure is shown in Figure G-4

The final action prior to returning (if desired) to the interrupted procedure is to restore the status of the REAL math unit:

```
CALL RESTORE$REAL$STATUS (@LOCAL_SAVE_AREA);
```

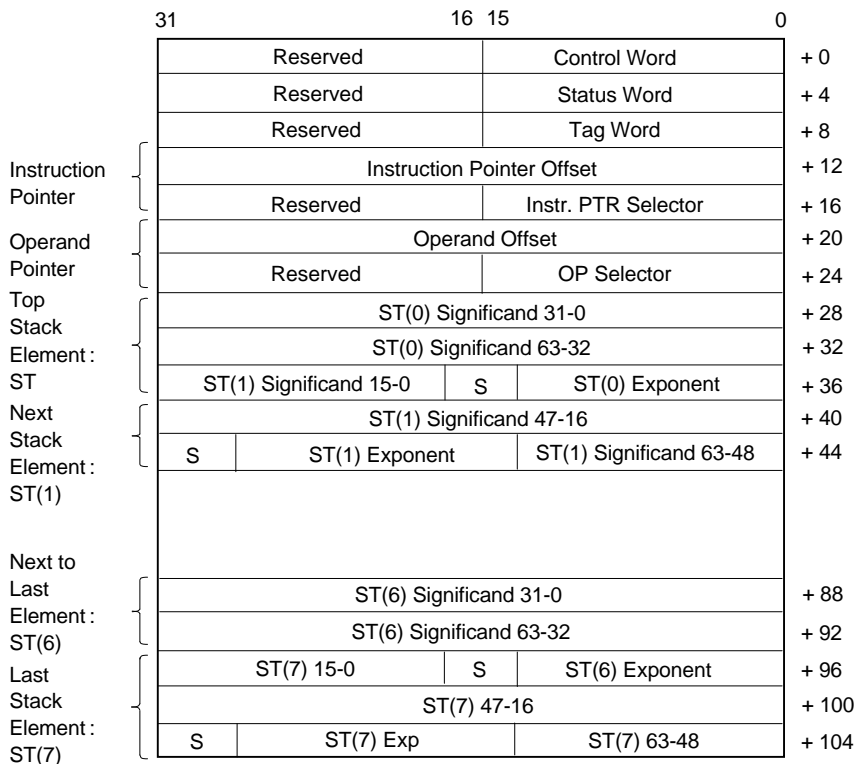
However, if GET\$REAL\$ERROR is not used prior to the SAVE\$REAL\$STATUS call, the local save area will contain the original contents of the error byte. Under these circumstances, first clear the lower byte of the saved status word before the RESTORE statement to avoid retriggering the same exception that invoked HANDLER in the beginning.

To do so, use a command of the form:

```
LOCAL_SAVE_AREA (2) = 0; /* should precede restore */
```

or:

```
SAVE_AREA.STATUS (0) = 0;
```



OSD546

Figure G-4. Memory Layout of the REAL Save Area in Protected Mode for the Intel386 Microprocessor



Run-time Support for PL/M Applications **H**

In addition to tools that support the software development process, Intel provides run-time support for application programs.

Numeric Coprocessor Support Libraries

Three specific libraries contained in the Intel387 numeric coprocessor support directory are of use to PL/M-386 programmers. The directory, */intel/ndp387*, contains these libraries:

dc387n.lib/dc387f.lib (near and far) converts floating-point representations from ASCII decimal format to internal binary format, and vice versa.

cl387n.lib/cl387f.lib (near and far) is a common elementary function library that provides an assortment of common elementary functions, i.e., logarithmic, exponential, trigonometric, and hyperbolic, involving floating-point numbers, such as rounding.

eh387n.lib/eh387f.lib (near and far) includes floating-point exception-handling procedures.

For additional information on the libraries contained with Intel387 numeric coprocessors, see the numeric coprocessor reference manual.

PL/M Support Libraries

The PL/M support libraries contain connection procedures and complex built-ins written in assembly language. The following support library modules are provided:

Interface to 286 CPU code :

INTERFACE286_FAR
INTERFACE286_NEAR

Math function for double words :

LQ_DWORD_DIVIDE
LQ_DWORD_MULTIPLY

Bit manipulation functions:

MOVBIT
MOVRBIT
SCANBIT
SCANRBIT



@ operator, 57

A

ABS function, 165
ADJUST\$RPL function, 203
Algebraic-shift functions, 168
Apostrophe in string, to include, 33
Arithmetic operators, 79, 91
Arrays, 67
ASM interface, 237
Assignment, 75, 96
AT attribute, 43, 61
Attributes
 EXTERNAL, 143
 INTERRUPT, 143
 PUBLIC, 143
 REENTRANT, 143

B

Based variables, 35, 54, 58
Binary number variables, 50
BITLOCK functions, 180
blanks, 29
Block structure, 119
BLOCKINDWORD, BLOCKINHWORd,
 BLOCKINPUT, BLOCKINWORD
 procedures, 192
BLOCKOUTDWORD, BLOCKOUTHWORD,
 BLOCKOUTPUT, BLOCKOUTWORD
 procedures, 193
blocks, 23
BUILD\$PTR function, 183
Built-in arrays, 195, 196
Built-in procedures and variables, 25
Built-ins, 149
BYTE\$SWAP built-in function, 211

C

C language compatibility, 380
cache, clearing, 211
CALL statement, 118, 135
Calling sequence, 371
CARRY flag, 187, 189
CAUSE\$INTERRUPT statement, 186
Character set, 27
character strings, 33
CLEAR\$TASK\$SWITCHED\$FLAG built-in
 procedure, 200
Closed subsystems, 309
CODE control, 232
Comments, 33
Communication between subsystems, 310
COMPACT control, 221, 231, 263
Compilation summary listing example, 271
Compound operands, 78
Concatenate functions, 169
COND control, 232
Constants, 30, 43, 76
CONTROL\$REGISTER built-in array, 198
Cross-reference listing example, 270

D

Data attribute declaration, 35
DATA keyword, 39, 43
Data types, 35, 48
DEBUG control, 233
DEBUG\$REGISTER built-in array, 198
DEC built-in function, 189
Decimal adjust, 189
Declaration statements, 23, 35, 36, 45, 122
Denormal operand exception, 397
Dimension specifier, 42
DISABLE statement, 185

DO block and statement, 23
DO statement, 103
dollar sign, 30

E

EJECT control, 233
ENABLE statement, 185
END statement, 23, 103, 112
Evaluation of expressions, 87
Example of subsystem, 316
Example program, 279
executable statements, 24
Exporting procedures, 314
expressions, 25, 75
Extended segmentation model syntax, 310
EXTERNAL attribute, 123

F

factored declaration, 36
File inclusion with compiler controls, 227
File usage, 217
FIND element functions, 173
Find string mismatch function, 174
Find value in input port function, 191
FIX function, 160
FLAGS function, 189
Flags, hardware, 187, 188
FLAT control, 265
FLOAT function, 160
Floating-point arithmetic, 50
Flow of control, 103
Function references, 78, 135

G

GDT register, 194
GET\$ACCESS\$RIGHTS function, 200
GET\$REAL\$ERROR function, 208
GET\$SEGMENT\$LIMIT function, 201
Global descriptor table register, 194
GOTO restrictions, 126
GOTO statement, 117

H

HALT statement, 186

I

I/O hardware, 192
IABS function, 165
Identifiers, 30
IDTR register, 196, 197
IF control, 234
IF statement, 112
IF|ELSE|ELSEIF|ENDIF controls, 233
Implicit dimension specifier, 42
INCLUDE control, 235
INIT\$REAL\$MATH\$UNIT built-in procedure, 207
INITIAL keyword, 38
Initialization, 39
Input files, 217
INPUT, INHWOR, INWORD functions, 191
Input/Output support, 26
INT function, 161
INTEGER keyword and variables, 51
INTERFACE control, 236
intermodule references, 47
Interrupt
 Mechanism, to enable or disable, 185
 Processing, 210
 Software, to generate, 186
Interrupt descriptor table, 388
Interrupt descriptor table register, 196
Interrupt procedures, 389
Interrupt processing, 387
Invalid operation exception, 396
INVALIDATE\$DATA\$CACHE built-in function, 211
INVALIDATE\$TLB\$ENTRY built-in function, 212
IRET instruction, to generate, 210

L

Label declarations, 35, 46
Languages interface, 231, 236
LAST function, 151
LDT register, 197

LEFTMARGIN control, 221, 226, 241
 LENGTH function, 150
 Line numbers, 241
 Linkage attributes, 119
 Linking to modules in other languages, 369
 LIST control, 241
 Listing example, 271
 LITERALLY declarations, 44, 45
 local descriptor table register, 197
 LOCAL\$TABLE variable, 197
 Location references, 54, 56, 60, 78
 LOCKSET function, 180
 Logical operators, 85
 Logical-shift functions, 167

M

Machine overflow, 258
 Machine status register, 198
 MACHINES\$STATUS built-in variable, 198
 Math facility, 204
 MEDIUM control, 231, 265
 Messages, 321
 MINUS operator, 188
 MOD486 control, 242
 module, 23
 MOV_B, MOV_{HW}, MOV_W procedures, 171
 MOV_{BIT} procedure, 177
 Move bit patterns right or left, 166
 MOVE procedure, 179
 MOV_{RB}, MOV_{RHW}, MOV_{RW}
 procedures, 171
 MSW register, 198

N

NIL function, 184
 null statement, 107
 Number base (binary, decimal, hexadecimal,
 and octal), 31

O

Object files, 218
 OFFSET function, 165
 OFFSET type, 57
 OFFSET\$OF function, 183

Open subsystems, 308
 operands, 75
 Operator precedence, 87
 OPTIMIZE control, 243
 OUTPUT, OUTD_{WORD}, OUTH_{WORD},
 OUT_{WORD} functions, 191
 OVERFLOW control, 258
 Overflow exception, 397

P

PAGEL_{LENGTH} control, 258
 PAGE_{WIDTH} control, 259
 PAGING control, 259
 Parameters, actual and formal, 132
 PARITY flag, 188
 PLUS operator, 188
 POINTER function, 164
 POINTER keyword and type, 48, 54
 Precision exception, 398
 PRINT control, 259
 Print files, 218
 Privilege level, to adjust, 200, 202
 Procedure declarations, 47
 Procedure epilogue, 375
 Procedures, 23, 131
 Activation, 135
 Declaration, 131
 Definition, 132
 Exit from, 139
 Parameters, 132
 Scope, 132
 Typed, 134
 Untyped, 134
 Procedures and tasks, 388
 Protection architecture of the microprocessor,
 to access, 193
 PUBLIC attribute, 123

R

RAM control, 219, 231, 260
 Read string procedure, 192, 194
 REAL functions, 160
 REAL interrupts, 399
 REAL keyword and variables, 51
 REAL math facility, 204

- Recursion, direct and indirect, 146
- Register usage, 376
- Registers, 189
- Relational operators, 83
- Requested privilege level, to adjust, 203
- Reserved words, 339
- RESET control, 228, 231, 232, 261
- RESTORE control, 227, 260
- RESTORE\$GLOBAL\$TABLE built-in procedure, 195
- RESTORE\$INTERRUPT\$TABLE built-in procedure, 197
- RESTORE\$REAL\$STATUS built-in procedure, 208
- RETURN statement, 139
- ROL function, 166
- ROM control, 231
- ROR function, 166
- Rotation functions, 166, 188
- RPL, to adjust, 203
- Run-time support, 405

S

- SAL function, 168
- Sample program, 279
- SAR function, 168
- SAVE control, 260
- SAVE\$GLOBAL\$TABLE built-in procedure, 195
- SAVE\$INTERRUPT\$TABLE built-in procedure, 196
- SAVE\$REAL\$STATUS built-in procedure, 208
- SCANBIT function, 177
- scientific notation, 52
- SCL built-in function, 188
- scope, 119, 123, 126
- SCR built-in function, 188
- Segment information and accessibility functions, 200, 202
- Segment name conventions, 379
- SEGMENT\$READABLE function, 202
- SEGMENT\$WRITABLE function, 202
- Segmentation controls, 299, 313
- SELECTOR function, 164
- SELECTOR keyword and type, 54, 57

- SELECTOR\$OF function, 183
- Separators, 29
- set command, 217
- SET control, 261
- SET procedures, 176
- SET\$REAL\$MODE procedure, 207
- SHL function, 167
- SHLD function, 169
- SHR function, 167
- SHRD function, 169
- SIGN flag, 188
- Signed arithmetic, 51
- SIGNED function, 161
- Signed integer data type built-in function, 160
- SIZE function, 151
- SKIP functions, 174
- SMALL control, 262
- Source code, to insert compiler control line, 221
- special characters, 29
- Stack layout, 392
- Stack representation, 371
- STACKBASE variable, 190
- STACKPTR variable, 190
- Statements
 - CALL, 192, 193, 195, 196, 200, 207, 209, 210
 - CAUSE\$INTERRUPT, 186
 - DISABLE, 185
 - ENABLE, 185
 - HALT, 186
- String manipulation procedures and functions, 170
- Strings, 33
- Structures, 69
- Subexpressions, 78
- Subscripted variables, 68
- Substitution (characters/values/quantities), 44, 45
- Subsystems, 297
- SUBTITLE control, 266
- Support libraries, 405
- SYMBOLS control, 220, 266

T

- TASK\$REGISTER variable, 193
- Tasks, 392
- temporary-real format, 52
- TEST\$REGISTER built-in array, 198, 211
- TIME procedure, 180
- TITLE control, 267
- Tokens, 29
- Translate string procedure, 175
- TYPE control, 267
- Type conversion, 96, 152
- Typed procedures, 134

U

- Underflow exception, 398
- underscore, 30
- UNSIGN function, 162
- Unsigned arithmetic, 50
- Unsigned binary data type built-in functions, 163
- Untyped procedures, 132
- Using subsystems, 303

V

- Value conversion, 152
- Variable declarations, 36
- Variable references, 78

W

- WAIT\$FOR\$INTERRUPT built-in procedure, 210
- WB\$INVALIDATE\$DATA\$CACHE built-in function, 212
- WMOVB function, 171
- WORD16 mapping for built-ins, 211
- WORD32|WORD16 control, 267
- WORD32|WORD16 type mapping, 64
- Work files, 217
- Write string procedure, 193

X

- XLAT procedure, 175
- XREF control, 220, 221, 270

Z

- Zero divide exception, 397
- ZERO flag, 200

WE'D LIKE YOUR OPINION

Please rate the following:	Excellent	Good	Fair	Poor
■ Manual organization	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
■ Technical accuracy	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
■ Completeness	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
■ Clarity of concepts and wording	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
■ Quality of examples and illustrations	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
■ Overall ease of use	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Comments: _____

Please list any errors you found (include page number): _____


Name _____

Company Name _____

Address _____


May we contact you? _____ Phone _____

Thank you for taking the time to fill out this form.



**NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES**

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 79 HILLSBORO, OR



POSTAGE WILL BE PAID BY ADDRESSEE

**OPD Technical Publications, HF2-72
Intel Corporation
5200 NE Elam Young Parkway
Hillsboro, OR 97124-9978**



Please fold here and close the card with tape. Do not staple.

WE'D LIKE YOUR COMMENTS....

This document is one of a series describing Intel products. Your comments on the other side of this form will help us produce better manuals. Each reply will be reviewed. All comments and suggestions become the property of Intel Corporation.

If you are in the United States and are sending only this card, postage is prepaid.

If you are sending additional material or if you are outside the United States, please insert this card and any enclosures in an envelope. Send the envelope to the above address, adding "United States of America" if you are outside the United States.

Thanks for your comments.

International Sales Offices

AUSTRALIA

Intel Australia Pty. Ltd.
Unit 1A
2 Aquatic Drive
Frenchs Forest, NSW, 2086
Sydney

Intel Australia Pty. Ltd.
711 High Street
1st Floor
East Kw. Vic., 3102
Melbourne

BRAZIL

Intel Semicondutores do Brazil LTDA
Avenida Paulista, 1159-CJS 404/405
CEP 01311-Sao Paulo - S.P.

CANADA

Intel Semiconductor of Canada, Ltd.
999 Canada Place
Suite 404, #11
Vancouver V6C 3E2
British Columbia

Intel Semiconductor of Canada, Ltd.
2650 Queensview Drive
Suite 250
Ottawa K2B 8H6
Ontario

Intel Semiconductor of Canada, Ltd.
190 Attwell Drive
Suite 500
Rexdale M9W 6H8
Ontario

Intel Semiconductor of Canada, Ltd.
1 Rue Holiday
Suite 115
Tour East
Pt. Claire H9R 5N3
Quebec

CHINA/HONG KONG

Intel PRC Corporation
China World Tower, Room 517-518
1 Jian Guo Men Wai Avenue
Beijing, 100004
Republic of China

Intel Semiconductor Ltd.
32/F Two Pacific Place
88 Queensway
Central
Hong Kong

FINLAND

Intel Finland OY
Ruosilantie 2
00390 Helsinki

FRANCE

Intel Corporation S.A.R.L.
1, Rue Edison-BP 303
78054 St. Quentin-en-Yvelines
Cedex

GERMANY

Intel GmbH
Dornacher Strasse 1
85622 Feldkirchen bei Muenchen
Germany

INDIA

Intel Asia Electronics, Inc.
4/2, Samrah Plaza
St. Mark's Road
Bangalore 560001

ISRAEL

Intel Semiconductor Ltd.
Atidim Industrial Park-Neve Sharet
P.O. Box 43202
Tel-Aviv 61430

ITALY

Intel Corporation Italia S.p.A.
Milanofiori Palazzo E
20094 Assago
Milano

JAPAN

Intel Japan K.K.
5-6 Tokodai, Tsukuba-shi
Ibaraki, 300-26

Intel Japan K.K.
Hachioji ON Bldg.
4-7-14 Myojin-machi
Hachioji-shi, Tokyo 192

Intel Japan K.K.
Bldg. Kumagaya
2-69 Hon-cho
Kumagaya-shi, Saitama 360

Intel Japan K.K.
Kawa-asa Bldg.
2-11-5 Shin-Yokohama
Kohoku-ku, Yokohama-shi
Kanagawa, 222

Intel Japan K.K.
Ryokuchi-Eki Bldg.
2-4-1 Terauchi
Toyonaka-shi, Osaka 560

Intel Japan K.K.
Shinmaru Bldg.
1-5-1 Marunouchi
Chiyoda-ku, Tokyo 100

Intel Japan K.K.
Green Bldg.
1-16-20 Nishiki
Naka-ku, Nagoya-shi
Aichi 460

KOREA

Intel Korea, Ltd.
16th Floor, Life Bldg.
61 Yoido-dong, Youngdeungpo-
Ku
Seoul 150-010

MEXICO

Intel Technologica de Mexico
S.A. de C.V.
Av. Mexico No. 2798-9B, S.H.
44620 Guadalajara, Jal.,

NETHERLANDS

Intel Semiconductor B.V.
Postbus 84130
3009 CC Rotterdam

RUSSIA

Intel Technologies, Inc.
Kremenchugskaya 6/7
121357 Moscow

SINGAPORE

Intel Singapore Technology, Ltd.
101 Thomson Road #08-03/06
United Square
Singapore 1130

SPAIN

Intel Iberia S.A.
Zurbaran, 28
28010 Madrid

SWEDEN

Intel Sweden A.B.
Dalvagen 24
171 36 Solna

TAIWAN

Intel Technology Far East Ltd.
Taiwan Branch Office
8th Floor, No. 205
Bank Tower Bldg.
Tung Hua N. Road
Taipei

UNITED KINGDOM

Intel Corporation (U.K.) Ltd.
Pipers Way
Swindon, Wiltshire SN3 1RJ

If you need to contact Intel Customer Support

Contacting us is easy. Be sure that you have the following information available:

- Your phone and FAX numbers ready
- Complete description of your hardware or software configuration(s)
- Your product's product code
- Current version of all software you use
- Complete problem description

Type of Service		How to contact us
<p>FaxBACK* fax-on-demand system</p> <p>24 hrs a day, 7 days a week</p>	<p>Using any touch-tone phone, have technical documents sent to your fax machine. Know your fax number before calling.</p>	<p>U.S. and Canada: (800) 628-2283 (916) 356-3105</p> <p>Europe: +44-1793-496646</p>
<p>Intel PC and LAN Enhancement Support BBS</p> <p>24 hrs a day, 7 days a week</p>	<p>Information on products, documentation, software drivers, firmware upgrades, tools, presentations, troubleshooting.</p>	<p>U.S and Canada: (503) 264-7999 Europe: +44-1793-432955</p> <p>Autobaud detect 8 data bits, no parity, 1 stop</p>
<p>CompuServe* Information Service</p> <p>24 hrs a day, 7 days a week</p>	<p>Worldwide customer support: information and technical support for designers, engineers, and users of 32-bit iRMX® OS and Multibus product families.</p>	<p>Worldwide Locations: (check your local listing)</p> <p>Type: GO INTEL C once online.</p>
<p>Customer Support</p>	<p>Intel Multibus Support engineers offering technical advice and troubleshooting information on the latest Multibus products.</p>	<p>U.S. and Canada: (800) 257-5404 (503) 696-5025 FAX: (503) 681-8497 Hrs: M-F; 8-5 PST</p> <p>Europe: +44-1793-641469 FAX: +44-1793-496385 Hrs: M-F; 9-5:30 GMT</p>
<p>Hardware Repair</p>	<p>Multibus board and system repair.</p>	<p>U.S. and Canada: (800) 628-8686 (602) 554-4904 FAX: (602) 554-6653 Hrs: M-F; 7-5 PST</p> <p>Europe: +44-1793-403520 FAX: +44-1793-496156 Hrs: M-F; 9-5:30 GMT</p>
<p>Sales</p>	<p>Intel Sales engineers offering information on the latest iRMX and Multibus products and their availability.</p>	<p>Worldwide: Contact your local Intel office or distributor</p> <p>U.S. and Canada: (800) 438-4769 (503) 696-5025 FAX: (503) 681-8497 Hrs: M-F; 8-5 PST</p>
<p>Correspondence Mail letters to:</p>	<p>Worldwide: Intel Customer Support Mailstop HF3-55 5200 NE Elam Young Parkway Hillsboro, Oregon 97124-6497</p>	<p>Europe: European Application Support Intel Corporation (UK) Ltd. Pipers Way Swindon, Wiltshire England SN3 1RJ</p>

* Third-party trademarks are the property of their respective owners.