


Systems Reference Library

IBM Operating System/360

PL/I: Language Specifications

This manual is a description of the full facilities of PL/I to be implemented under Operating System/360. However, the reader should not assume that all facilities will be available at initial release. Manuals for specific System/360 implementations will be released later.

Another publication will be issued specifying a subset of the facilities of the language described in this manual. This subset is planned for implementation under the Basic Operating System/360 and Basic Program Support for System/360.



MAJOR REVISION (JULY, 1965)

This publication, Form C28-6571-1, is a major revision of the previous edition, Form C28-6571-0. This new edition reflects a number of significant changes in the language. The sections containing these changes are indicated by a dot (•) to the left of the subject heading in the Table of Contents.

In addition to the technical changes, certain organizational changes have been made. Additional chapters (Chapters 5 and 6) have been formed from material previously contained in Chapter 1. The information on attributes, declarations, and scope of declarations, also previously contained in Chapter 1, has been incorporated into Chapter 4.

This publication was prepared for production using an IBM computer to update the text and to control the page and line format. Page impressions for photo-offset printing were obtained from an IBM 1403 Printer using a special print chain.

Copies of this and other IBM publications can be obtained through IBM Branch Offices.

A form for readers' comments appears at the back of this publication. It may be mailed directly to IBM. Address any additional comments concerning this publication to the IBM Corporation, Programming Systems Publications, Department D39, 1271 Avenue of the Americas, New York, N. Y., 10020.

INTRODUCTION	9	Base	22
Goals of the Language.	9	Scale.	22
Basic Characteristics of PL/I.	10	Mode	22
New Features.	10	Precision.	22
Block Structure.	10	•String Data	22
Description of Data.	10	Character-String Data.	22
Storage Allocation	10	Bit-String Data.	23
Data Conversion.	10	Statement-Label Data.	23
Data Organization.	10	Scalar Quantities.	23
Input/Output	11	Constants	23
Multi-Task Operations.	11	Real Arithmetic Constants.	23
Compile-Time Facilities.	11	Decimal Fixed-Point Constants.	23
Syntax Notation in This Manual	11	Binary Fixed-Point Constants	23
•CHAPTER 1. PROGRAM ELEMENTS	14	Sterling Fixed-Point Constants	23
Basic Language Structure	14	Decimal Floating-Point Constants	23
Language Character Sets	14	Binary Floating-Point Constants.	23
60-Character Set	14	Precision of Real Arithmetic	
48-Character Set	15	Constants	24
Delimiters.	15	Imaginary Arithmetic Constants	24
Operators.	15	•String Constants	24
Arithmetic Operators	15	Character-String Constants	24
Comparison Operators	15	Bit-String Constants	24
Bit-String Operators	15	Statement-Label Constants.	24
String Operator.	15	Variables	25
Parentheses.	15	Data Aggregates.	25
Separators and Other Delimiters.	15	Arrays.	25
Data Character Set.	16	•Structures.	25
Collating Sequence.	16	Arrays of Structures.	26
Identifiers	16	Naming	26
Length of Identifiers.	16	Simple Names.	26
•Keywords.	16	Subscripted Names	26
Statement Identifiers.	16	Cross Sections of Arrays	27
Attributes	17	Qualified Names	27
Separating Keywords.	17	Subscripted Qualified Names	28
Built-in Function Names.	17	Statement Labels	28
Options.	17	Constant.	28
Conditions	17	Variable.	28
•The Use Of Blanks	17	Array	29
Comments.	17	Initial Values for Label Arrays.	29
•Basic Program Structure.	17	•Task Names	29
Simple Statements	18	•Event Names.	30
Compound Statements	18	•CHAPTER 3: DATA MANIPULATION	31
•Prefixes.	18	Expressions.	31
•Label Prefixes	18	•Scalar Expressions.	31
•Condition Prefixes	18	Arithmetic Operations	31
Groups.	19	•Mixed Characteristics.	31
Blocks.	19	•Results of Arithmetic Operation.	31
Use of the END Statement.	21	•Arithmetic Conversions	32
Programs.	21	Bit-String Operations.	33
•CHAPTER 2: DATA ELEMENTS	22	Comparison Operations.	34
Data Types	22	Concatenation Operations	34
Arithmetic Data	22	•Type Conversion.	34

•Bit String to Character String	34	Standard Attributes	58
•Character String to Bit String	34	•The Storage Equivalence	
•Character String to Arithmetic	34	Attribute	59
•Bit String to Arithmetic	35	•The Function Attributes	59
Arithmetic to Character String	35	•File Organization Attributes	59
•Arithmetic to Bit String	35	•Access Attributes	59
•Array Expressions	35	•The KEYLENGTH Attribute	61
Prefix Operators and Arrays	35	The ZERO Attribute	61
Infix Operators and Arrays	35	The ENVIRONMENT Attribute	62
Scalar - Array Operations	35	•Assignment Of Attributes To	
Array - Array Operations	36	Identifiers	62
Array Expressions Involving		•Application of Default	
Structures	36	Attributes	62
•Structure Expressions	36	•Structure Declarations and Attributes	63
•EVALUATION OF EXPRESSIONS	37	Level Number	63
•Order of the Evaluation of		•Structures and the Dimension	
Expressions	37	Attribute	64
•CHAPTER 4: DATA DESCRIPTION	38	•Structures and Data Attributes	64
•Attributes	38	Structures and Scope Attributes	64
•Declarations	38	Structures and Storage Class	
•Explicit Declarations	38	Attributes	64
•The DECLARE Statement	38	•CHAPTER 5: PROCEDURES, FUNCTIONS,	
•Factoring of Attributes	39	AND SUBROUTINES	65
•Multiple Declarations and		Formal Parameters	65
Ambiguous References	39	•Procedure References	65
•Contextual Declarations	40	•Function References and Function	
Implicit Declarations	41	Procedures	66
Scope of Declarations	41	•Generic Functions	66
Scope of External Names	41	Built-in Functions	66
Basic Rule on Use of Names	43	Subroutine References and Subroutine	
The Attributes	43	Procedures	67
Data Attributes	43	•The Arguments in a Procedure Reference	68
Arithmetic Data	43	•The Use of the ENTRY Attribute	69
Base	43	•Passing Arguments to the Entry	
Scale	44	Point	69
Mode	44	The Special Procedure Attribute	
Precision	44	RECURSIVE	70
•Default Conditions for		•CHAPTER 6: DYNAMIC PROGRAM	
Arithmetic Data	45	STRUCTURE	71
•The PICTURE Attribute	45	Program Control	71
String Attributes	47	Activation and Termination of Blocks	71
•The LABEL Attribute	47	Dynamic Descendance	71
•The TASK Attribute	48	Dynamic Encompassing	72
•The EVENT Attribute	48	Allocation of Data and Storage Classes	72
The dimension Attribute	48	Definitions and Rules	72
The SECONDARY Attribute	49	•Storage Classes	72
•The ABNORMAL and NORMAL Attributes	49	The Static Storage Class	72
•The USES and SETS Attributes	50	•The Automatic Storage Class	72
•Entry Name Attributes	50	The Controlled Storage Class	73
•The ENTRY Attribute	51	•Asynchronous Operations And Tasks	74
•The GENERIC Attribute	51	•Synchronous and Asynchronous	
The BUILTIN Attribute	52	Operations	74
Scope Attributes	52		
•Storage Class Attributes	53		
•The ALIGNED and PACKED Attributes	53		
•The DEFINED Attribute	54		
•The INITIAL Attribute	56		
•Symbol Table Attributes	57		
The LIKE Attribute	58		
•File Description Attributes	58		
The FILE Attribute	58		

•Synchronizing Two Asynchronous Operations	74	•CHAPTER 8: STATEMENTS	95
•Task and Events	74	Relationship Of Statements	95
•The Creation of Tasks	75	Classification.	95
•Termination of Tasks.	75	Assignment Statement	95
•Allocation of Data in Tasks	75	Control Statements	95
•Interrupt Operations	75	•Data Declaration Statement	95
•Purpose of the Condition Prefix	76	Error Control and Debug Statements.	95
•Scope of the Condition Prefix	76	Input/Output Statements.	95
•Use of the ON Statement	76	Program Structure Statements	95
•System Interrupt Action	77	Sorting Statement.	95
Use of the REVERT Statement	78	Storage Allocation Statements.	95
Programmer-Defined ON-Conditions.	79	•Sequence of Control	95
Facilities for Program Checkout	79	•Pseudo-Variables	96
•CHAPTER 7: INPUT/OUTPUT.	80	•Alphabetic List of Statements	97
•Data Transmission.	80	•The ALLOCATE Statement	97
List-Directed Transmission.	80	•The Assignment Statement	98
Data-Directed Transmission.	80	The BEGIN Statement.	102
Format-Directed Transmission.	80	•The CALL Statement	102
•Procedure-Directed Transmission	81	•The CLOSE Statement.	103
Data Specifications.	82	The DECLARE Statement.	104
Data Lists.	82	The DELAY Statement.	104
Repetitive Specification	82	The DELETE Statement	104
Transmission of Data-List Elements.	83	•The DISPLAY Statement.	105
•List-Directed Data Specification.	83	•The DO Statement	105
•List-Directed Input.	83	•The END Statement.	107
•List-Directed Input Format	83	The ENTRY Statement.	107
•List-Directed Output	84	•The EXIT Statement	108
•List-Directed Output Format.	84	•The FETCH Statement.	108
•Length of List-Directed Output Fields.	84	•The FORMAT Statement	108
•Data-Directed Data Specification.	85	The FREE Statement	108
•Data-Directed Data On External Medium.	86	•The GET Statement.	109
Data-Directed Output Format.	87	•The GO TO Statement	109
•Length of Data-Directed Data Fields.	87	•The GROUP Statement.	110
•Format-Directed Data Specification.	88	The IF Statement	111
•Format Lists.	88	•The LAYOUT Statement	111
Data Format Items.	89	The Null Statement	112
•External Mode Format Items	89	•The ON Statement	112
•Internal Mode Format Items	91	•The OPEN Statement	114
Control Format Items	92	•The PAGE Statement	116
Spacing Format Item.	92	•The POSITION Statement	117
Positioning Format Items	92	•The PROCEDURE Statement.	118
Remote Format Item	93	•The PUT Statement.	119
•Procedure-Directed Data Specification.	93	•The READ Statement	119
•Input/Output Statements.	93	•The REPOSITION Statement	121
•File Preparation Statements	93	•The RESTORE Statement.	121
•Data Specification Statement.	93	•The RETURN Statement	122
•Data Transmission Statements.	93	The REVERT Statement	122
•Positioning Statements.	93	•The SAVE Statement	123
•Report Generation Statements.	93	•The SEGMENT Statement.	124
•Record Identification Options.	94	•The SIGNAL Statement	125
The Key Option.	94	•The SKIP Statement	125
The Newkey Option	94	•The SORT Statement	126
The Region Option	94	•The SPACE Statement.	127
		•The STOP Statement	127
		•The TAB Statement.	127
		•The WAIT Statement	128
		•The WRITE Statement.	128
		•CHAPTER 9: PROGRAM MODIFICATION	131
		Macro Variables.	131
		•The Macro DECLARE Statement	131

Macro Expressions.132	•Built-In Functions for Manipulation of Arrays.144
•Executable Macro Statements.132	•Array And Structure Built-In Functions	.145
The Macro Assignment Statement.132	•Condition Built-In Functions145
The Macro Null Statement.132	•Other Built-In Functions145
The Macro GO TO Statement132		
The Macro IF Statement.133	APPENDIX 2: PICTURE SPECIFICATION TABLES.147
•Action of the Macro Processor.133	Digit Point and Subfield Delimiting Characters.147
•CHAPTER 10: SPECIAL TOPICS.134	Zero Suppression Characters.147
•Relationship of Arguments and Parameters.134	Drifting Editing Symbols147
Evaluation of Argument Subscripts134	Drifting Characters148
•Use of Dummy Arguments.134	Editing Character148
•Use of the Entry Attribute.134	Conditional Editing Characters.148
Correspondence Of Parameters And Arguments.135	Sign Characters149
Allocation of Parameters.136	Scaling Factor Specification.149
Parameters, Bounds and Length.136	•Sterling Pictures.149
Asterisk Notation for Bounds or Length.136	•Pictures for Character Strings.149
Expressions as Bounds or Length.136	•APPENDIX 3: ON-CONDITIONS.150
•Prologues.137	•Classification of Conditions150
•Data Allocation Across Tasks137	•Computational Conditions.150
•Allocation of Task and Event Names138	•Input/Output Conditions151
•Abnormality.138	•Program Checkout Conditions152
•Programs138	•Programmer-Named Conditions153
•APPENDIX 1: BUILT-IN FUNCTIONS140	•System Action Conditions.153
•Arithmetic Generic Functions140	APPENDIX 4: PERMISSIBLE KEYWORD ABBREVIATIONS154
•Float Arithmetic Generic Functions142	APPENDIX 5: THE 48-CHARACTER SET155
•String Generic Functions143	INDEX.156

ILLUSTRATIONS

Table 1. Arithmetic Base and Scale Conversion.	33	Figure 3. General Format for the CLOSE Statement	104
Table 2. Scope and Use of Names in Example 1, for "Scope of External Names".	42	Figure 4. General Format for the DO Statement	105
Table 3. Allowable statements for CONSECUTIVE, REGIONAL, and INDEXED organizations of a SEQUENTIAL access file.	59	Figure 5. General Format for the OPEN Statement.	114
Table 4. Allowable Statements for the REGIONAL and INDEXED Organizations of a DIRECT Access File.	61	Figure 6. Format of "option" Allowed in the OPEN Statement	114
Figure 1. General Format for Repetitive Specification.	82	Figure 7. General Format for the PAGE Statement.	116
Figure 2. Example of Data-Directed Transmission, Both Input and Output	88	Figure 8. General Format for the READ Statement.	119
		Figure 9. General Format for the SORT Statement.	126
		Figure 10. General Format for the WRITE Statement	129

GOALS OF THE LANGUAGE

Throughout the relatively brief history of electronic data processing, certain computers have been identified with a particular field of activity, either commercial or scientific.

Programmers generally have specialized in one field or the other. High-level languages, of course, have emphasized this divergence, going in one direction for commercial programming and in another direction for scientific programming.

Until recently, this difference presented few problems. Each language was adequate for its use; the commercial programmer dealt with relatively few computations performed upon great amounts of data; the scientific programmer performed complex calculations using small amounts of data.

Now, however, the situation is changing. Business and industry have discovered new uses for the computer, and the commercial programmer finds himself concerned with more involved computations in statistical forecasting and in linear programming for operations research.

In science and engineering, the programmer needs a language to simplify the preparation of reports, to sort and edit technical data; he finds more need for input and output operations. The engineer specifically wants the ability to handle data at the bit level for applications such as circuit analysis.

Today's new computing systems have been designed to cope with all of these computing problems. They handle commercial and scientific programs with equal ease, with new power and new speed; they provide facilities for such new techniques as shared data processing, asynchronous program execution, and real-time processing.

None of the traditional high-level languages, however, can be used with efficiency across the entire range of ability of these new computers.

That is the reason for PL/I, a multipurpose programming language for use not only by commercial and scientific programmers but by the real-time programmer and the systems programmer as well. It is a language designed for efficiency, a language that enables the programmer to use virtually all the power of his computer.

PL/I is organized so that any programmer, no matter how extensive his experience, can use it easily at his own level.

This manual, because it is a reference manual of the entire language, shows the range and power of PL/I, its ability to handle the most complex computing problems.

Actually, however, PL/I need be no more complex than the program for which it is used.

One of the primary aims in the design of the language was modularity, that is, providing different subsets of the language for different applications and different levels of complexity. A programmer using one subset need not even know about the unused facilities.

Although PL/I is relatively machine independent, this modularity might be compared to a completely equipped data processing center. A novice programmer would use only a small part of the system; he can ignore the rest of the equipment. More complex programs, of course, would require more equipment. Some programs would use certain modules of equipment; other programs, other modules. Rarely, if ever, would a program require use of all the facilities.

In PL/I, every attribute -- or description -- of a variable, every option, and every specification has been given a "default" interpretation. Wherever the language provides for one or more alternatives, a "default" interpretation -- or assumption -- is made by the compiler if no choice is stated by the programmer. And in each case, the assumption that was chosen in the design of the language is the one most likely to be required by the programmer who need not know that alternatives exist.

The "modularity" and the "default" aspects are the bases upon which the simplicity of PL/I has been built. They are also part of its power.

BASIC CHARACTERISTICS OF PL/I

The overall aim in the design of the language was to give the programmer freedom in handling his computing system.

Freedom of expression: If a particular combination of symbols has a useful meaning, that meaning is allowed. Although actual statements in the language must be written using a specified character set, data may be composed of any character allowed by the configuration of the individual computer. PL/I is written in a free-field format; the programmer is free to design his own format for listings.

Full access to machine and operating system facilities: the PL/I programmer rarely, if ever, will need to resort to assembly language coding.

NEW FEATURES

Part of the language is, of course, based on earlier programming languages. Certain aspects are expansions of ideas used previously. Other portions are exclusively a part of PL/I. The following paragraphs briefly introduce some of these new features. All of them are discussed more fully within the text.

Block Structure

The statements of a PL/I program are organized into program sections called "blocks." A program may be made up of one block or many blocks. Blocks may be separate from one another, with no common statements, or they may be nested, one within another.

Blocks provide two important logical functions: (1) they define the scope of applicability of data variables and other kinds of names, so that the same name may be used for different purposes in different blocks without ambiguity, and (2) they allow storage for data variables to be assigned only during execution of the block and freed for other uses at the termination of the block.

Certain blocks, called "procedure" blocks, may be invoked (i.e., called into execution) remotely from different places in the program, and they provide means to handle arguments and to return values.

Description of Data

In the language, data is described as having certain characteristics called attributes. For example, numeric data would have a BINARY attribute or a DECIMAL attribute; string data would be either CHARACTER string or BIT string.

Storage Allocation

The computer storage for any data variable in a PL/I program may be assigned statically, for the entire execution of the program, or dynamically, during execution.

Two classes of dynamic storage are available to the PL/I programmer, automatic and controlled. When a variable has the controlled storage attribute, the programmer may allocate or free storage for that variable at any time he wishes. Storage for a variable having the automatic storage attribute is allocated upon entry to the block and freed upon exit.

Data Conversion

In keeping with the freedom of PL/I, mixed expressions are allowed. In the following example, F is declared to be a fixed-point number, G a floating-point number, and H a character string that is ten characters in length.

```
DECLARE F FIXED, G FLOAT, H CHARACTER
(10);
      H = F + G;
```

In the evaluation of the second statement of the above example, F will be converted to a floating-point value, floating-point addition will be performed, and the result will be converted to a character string of ten characters and assigned as a value to H.

Data Organization

Data variables can be grouped into either arrays or structures. An array is composed of elements of the same characteristics. A structure is a collection of variables and arrays, not necessarily alike in characteristics. Structures may also contain other structures. Individual items of an array are referred to by subscripted names; individual items of a structure are referred to by names that may sometimes

have to be qualified to avoid ambiguity.

In PL/I, arrays and structures are treated as variables in their own right. Either of them may be used as the operand of an expression. The expression is then an array expression or a structure expression, and it returns an array or structure result.

Input/Output

The modularity of PL/I is particularly apparent in the input/output facilities. With PL/I, a programmer may control input/output activity to whatever degree he requires. He may handle normal transmission and conversion simply, or he may use the full capability of the language for control of more complex problems of input and output.

Multi-Task Operations

In PL/I, a collection of procedures is called a program; the execution of a program (or many programs or a part of a program) to perform a particular job is called a task.

PL/I provides facilities for handling two or more tasks concurrently. This facility, of course, is extremely important in the use of any computer system with multiprocessing capabilities. It also is valuable for a single processor system with facilities for real-time operations.

During execution of a procedure, the executing task might specify that a subordinate task begin execution upon certain data (i.e., the executing task invokes another task); the new task, called an attached task, might also invoke another task. All tasks then proceed concurrently and, in effect, simultaneously.

The multi-task facilities of PL/I allow a subordinate task to communicate with its originating, or attaching, task through arguments, and through data allocated in the attaching task. The originating task also may, at any time, test to see if a subordinate task is completed and may, if necessary, delay its own execution to wait for the completion.

Compile-Time Facilities

Most programming languages are written on one level only, as statements to the computer to perform certain operations using the supplied data. PL/I not only directs the computer to operate upon the data, but with a macro facility, it directs the compiler to operate upon the program itself.

The programmer can include in his program information that will aid the compiler to produce more efficient code, documentation, and diagnostics.

SYNTAX NOTATION IN THIS MANUAL

Throughout this manual, wherever a PL/I statement -- or some other combination of elements -- is discussed, the manner of writing that statement or phrase is illustrated with a uniform system of notation.

This notation is not a part of PL/I; it is a standardized notation that may be used to describe the syntax -- or construction -- of any programming language. It provides a brief but precise explanation of the general patterns that the language permits. It does not describe the meaning of the language elements, merely their structure; that is, it indicates the order in which the elements may (or must) appear, punctuation that is required, and options that are allowed.

The following rules explain the use of this notation for any programming language; only the examples apply specifically to PL/I:

1. A notation variable is the name of a general class of elements in the programming language. A notation variable must consist of:
 - a. Lower-case letters, decimal digits, and hyphens and must begin with a letter.
 - b. A combination of lower-case and upper-case letters. There must be one portion in all lower-case letters and one portion in all upper-case letters, and the two portions must be separated by a hyphen.

All such variables used are defined in the manual either formally, using this notation, or are defined in prose.

Examples:

- a. digit. This denotes the occur-

rence of a digit, which may be 0 through 9 inclusive.

- b. file-name. This denotes the occurrence of the notation variable named file-name. An explanation of file-name is given elsewhere in the manual.
 - c. DO-statement. This denotes the occurrence of a DO statement. The upper-case letters are used for emphasis.
2. A notation constant denotes the literal occurrence of the characters represented. A notation constant consists either of all capital letters or of a special character.

Example:

```
DECLARE identifier FIXED;
```

This denotes the literal occurrence of the word DECLARE followed by the variable "identifier," which is defined elsewhere, followed by the literal occurrence of the word FIXED followed by the literal occurrence of the semicolon (;).

3. The term "syntactical unit," which is used in subsequent rules, is defined as one of the following:
- a. a single variable or constant, or
 - b. any collection of variables, constants, syntax-language symbols, and reserved words surrounded by braces or brackets.
4. Braces { } are used to denote grouping.

Example:

```
identifier { FIXED }  
           { FLOAT }
```

The vertical stacking of syntactical units indicates that a choice is to be made. The above example indicates that the variable "identifier" must be followed by the literal occurrence of either the word FIXED or the word FLOAT.

5. The vertical stroke | indicates that a choice is to be made.

Example:

```
identifier {FIXED|FLOAT}
```

This has exactly the same meaning as the above example. Both methods are used in this manual to display alternatives.

6. Square brackets [] denote options. Anything enclosed in brackets may

appear one time or may not appear at all.

Example:

```
CHARACTER (length) [VARYING]
```

This denotes the literal occurrence of the word CHARACTER followed by the variable "length" enclosed in parentheses and optionally followed by the literal occurrence of the word VARYING. If, in rule 4, the two alternatives also were optional, the vertical stacking would be within brackets, and there would be no need for braces.

7. Three dots ... denote the occurrence of the immediately preceding syntactical unit one or more times in succession.

Example:

```
[digit] ...
```

The variable, "digit," may or may not occur since it is surrounded by brackets. If it does occur, it may be repeated one or more times.

8. Underlining is used to denote an element in the language being described when there is conflict between this element and one in the syntax language.

Example:

```
operand {&|_|_} operand
```

This denotes that the variables "operand" are separated by either an "and" (&), an "or" (|), or a "not" (¬) symbol. The constant | is underlined in order to distinguish the "or" symbol in the PL/I language from the "or" symbols in the syntax language.

9. The notation ::= should be read "is defined as."

Example:

```
word ::= letter | word | | letter
```

This denotes that a "word" is defined as a letter or a word concatenated with a letter.

10. min max. The combination of these two words with associated numeric values specifies the minimum and maximum number of times a syntactical unit may occur. When min is used without max, the implied max is infinity. When max

is used without min, the implied min is zero.

Examples:

a. min 2 max 6 {digit|letter}

This denotes that either "digit" or "letter" intermixed in any combination must occur at least two times, but no more than six.

b. min 5 {digit|letter}

The variables "digit" or "letter" intermixed in any combination must occur at least five times, but there is no limit on the number of times over five that they may occur.

c. max 3 label

The variable "label" may not occur more than three times in succession. It may not be present at all, or it may occur one, two, or three times.

CHAPTER 1. PROGRAM ELEMENTS

BASIC LANGUAGE STRUCTURE

PL/I allows the programmer to write the statements of his program in a free-field format. A statement, which is a string of characters, is always terminated by the special character, semicolon. A program which is, in turn, a sequence of statements, can thus be regarded simply as a single string of characters, with no special internal grouping. Hence, a PL/I program can be physically represented and transmitted to a computer in a natural way by means of almost any input medium, including a typewriter at a remote terminal.

Input conventions, depending upon the machine configuration or the compiler, can, of course, be set up so that the program string may be presented to the computer through the familiar medium of fixed-length records, e.g., punched cards. This can be accomplished by using certain predetermined fields of the records for the program string, and other fields for arbitrary purposes.

LANGUAGE CHARACTER SETS

One of two character sets may be used to write a source program: either a 60-character set or a 48-character set. No assumptions are made in the language about external or internal codes for the characters. For a given program, the choice between the two sets is optional. (In practice, this choice will depend upon the available equipment.)

60-Character Set

The 60-character set is composed of digits, special characters, and English language alphabetic characters.

There are 29 alphabetic characters, letters A through Z and three additional characters that are defined as and treated as alphabetic characters. These characters and the graphics by which they are represented are:

Currency symbol	\$
Commercial At-sign	@
Number sign	#

There are ten digits. Decimal digits are the digits 0 through 9. A binary digit (bit) is either a 0 or a 1.

An alphanumeric character is either an alphabetic character or a digit.

There are 21 special characters. The names and graphics by which they are represented are:

<u>Name</u>	<u>Graphic</u>
Blank	
Equal or Assignment symbol	=
Plus	+
Minus	-
Asterisk or Multiply symbol	*
Slash or Divide symbol	/
Left Parenthesis	(
Right Parenthesis)
Comma	,
Decimal Point or Period	.
Quotation mark	'
Percent symbol	%
Semicolon	;
Colon	:
Not symbol	~
And symbol	&
Or symbol	
Greater Than symbol	>
Less Than symbol	<
Break character	-
Question mark	?

Special characters may be combined to create operators, e.g., >=, denoting

"greater than or equal to"; ||, denoting concatenation.

48-Character Set

The characters making up the 48-character set are identical to those of the 60-character set, with restrictions and changes as described in Appendix 5.

DELIMITERS

Certain characters are used as delimiters and fall into three classes:

- operators
- parentheses
- separators and other delimiters

Operators

Operators used by the language are divided into four types:

- arithmetic operators
- comparison operators
- bit-string operators
- string operators

Arithmetic Operators

The arithmetic operators are:

- + denoting addition or prefix plus
- denoting subtraction or prefix minus
- * denoting multiplication
- / denoting division
- ** denoting exponentiation

Comparison Operators

The comparison operators are:

- > denoting greater than
- >= denoting greater than or equal to
- = denoting equal to
- != denoting not equal to

<= denoting less than or equal to

< denoting less than
Bit-String Operators

The bit-string operators are:

- ! denoting not
- & denoting and
- | denoting or

String Operator

The string operator is:

- || denoting concatenation

Parentheses

Parentheses are used in expressions and for enclosing lists.

- (left parenthesis
-) right parenthesis

Separators and Other Delimiters

<u>Name</u>	<u>Graphic</u>	<u>Use</u>
comma	,	separates elements of a list
semicolon	;	terminates statements
assignment symbol	=	used in assignment statement and DO statement
colon	:	follows labels and condition prefixes; also used with dimension specifications
blank		used as a separator
quotation mark	'	encloses string constants
period	.	separates items in qualified names; used as a decimal or binary point in constants
percent	%	precedes macro statement

DATA CHARACTER SET

Although the language character set is a fixed set defined for the language, the data character set has not been limited. Data may be represented by characters from the language set plus any other characters permitted by the particular machine configuration.

Any character that will result in a unique bit pattern is a valid character in the data character set, and may be used in source programs to construct character-string constants and comments.

COLLATING SEQUENCE

In the execution of PL/I programs, comparisons of character data will observe the collating sequence resulting from the representations of involved characters in bytes of System/360 storage, in extended binary coded decimal interchange code (EBCDIC).

IDENTIFIERS

An identifier is a string of alphameric and break characters; the initial character must always be alphabetic.

Identifiers in the language are used for the following:

- scalar variable names
- array names
- structure names
- statement labels
- entry names
- file names
- keywords
- condition names
- headings for external names

Examples:

VARA
BCD320
FILE42

XR20A

STARTA

RATE_OF_PAY

#32_45

\$L32

X@_52

@531

AB12#

Length of Identifiers

Identifiers that a programmer constructs in writing a PL/I program must be composed of not more than 31 characters.

KEYWORDS

A keyword is an identifier which is a part of the language. Keywords are not reserved words. They may be classified as follows:

- statement identifiers
- attributes
- separating keywords
- built-in function names
- options
- conditions

Statement Identifiers

A statement identifier is a sequence of one or more keywords used to define the function of a statement (see "Simple Statements").

Examples:

GO TO
DECLARE
READ

Attributes

Attributes are keywords that specify the characteristics of data, procedures, and other elements of the language.

Example:

```
FLOAT
RECURSIVE
SEQUENTIAL
```

Separating Keywords

The five separating keywords are used to separate parts of the IF and DO statements. They are THEN, ELSE, BY, TO, WHILE.

Built-in Function Names

A built-in function name is a keyword that is the name of an algorithm provided by the language and accessible to the programmer (see "Function References and Function Procedures" in Chapter 5).

Examples:

```
DATE
EXP
```

Options

An option is a specification that may be used by the programmer to influence the execution of a statement.

Examples:

```
TASK
CROSS
```

Conditions

A condition is a keyword used in the ON, SIGNAL, and REVERT statements, and as a prefix to other statements (see "Prefixes"). The programmer may specify special action on occurrence of the condition (see "Interrupt Operations").

Examples:

```
OVERFLOW
ZERCDIVIDE
```

THE USE OF BLANKS

Identifiers, constants, picture specifications, composite operators (e.g., \neq), and the class of dummy variables ISUB (see "The DEFINED Attribute" in Chapter 4) may not contain blanks. Blanks are permitted within a character-string constant.

Identifiers, constants, or picture specifications may not be immediately adjacent. They must be separated by an operator, assignment symbol (i.e., =), parenthesis, colon, semicolon, comma, period, blank, or comment. Moreover, additional intervening blanks or comments are always permitted. Blanks are optional between keywords of a statement identifier (e.g., GO TO), and between a level number and its following identifier (see "Structures" in Chapter 2).

Examples:

```
CALLA      is not equivalent to CALL A
A TO B BY C is not equivalent to ATOBBYC
AB+BC      is equivalent to AB + BC
```

COMMENTS

General form:

```
/* character-string */
```

A comment may be used wherever a blank is permitted (except in a character-string constant). The character string in a comment must not contain the character combination `*/` in that sequence.

Example:

```
LABEL: /* THE BLOCK OF CODING BETWEEN
BEGIN-END IS USED FOR PAYROLL CALCULA-
TIONS */
      BEGIN;
      .
      .
      .
      END;
```

BASIC PROGRAM STRUCTURE

A PL/I program is constructed from basic program elements called statements.

Statements are grouped into larger program-elements, the group and the block. There are two types of statements: simple

and compound.

Label Prefixes

Statements may be labeled to permit reference to them. A labeled statement has the form:

identifier: [identifier:]...statement

The one or more "identifiers" are called labels and may be used interchangeably to refer to that statement.

Labels appearing before PROCEDURE and ENTRY statements are special cases and are known as entry names (see "Procedure References"). All other labels are called statement labels.

A label appearing before a statement is said to be declared, by virtue of its appearance as a label.

Statement labels appearing before DECLARE are ignored.

Condition Prefixes

A condition prefix is a keyword that determines whether or not a program interrupt will result upon the occurrence of the specified condition. (For information regarding the use of the condition prefix see the section "Interrupt Operations" in Chapter 6.)

One or more condition prefixes may be attached to a statement.

Each condition prefix is followed by a colon to separate it from the rest of the statement or from other prefixes; condition prefixes precede the entire statement, including any possible label prefixes for the statement.

A condition prefix is a list of condition names, separated by commas and enclosed in parentheses. Thus, a statement with a set of prefixes has the following general form:

{(condition-name [,condition-name] ...) :}... [label:]... statement

The condition names are chosen from the following fixed set:

- UNDERFLOW
- OVERFLOW
- ZERODIVIDE
- FIXEDOVERFLOW
- CONVERSION
- SIZE

SIMPLE STATEMENTS

A simple statement is defined as:

[[statement-identifier]
statement-body] ;

The "statement identifier," if it appears, is a keyword, characterizing the kind of statement. If it does not appear, and the statement body does appear, then the statement is an assignment statement. If only the semicolon appears, the statement is called a null statement.

Examples:

DO I = J TO (DO is the keyword)
10;

A = B + C; (assignment statement)

; (null statement)

COMPOUND STATEMENTS

A compound statement is a statement that contains other program-elements. There are only two of them. They are:

The IF compound statement

The CN compound statement

The final contained statement of a compound statement is a simple statement and thus has a terminal semicolon. Hence, the compound statement will automatically be terminated by this semicolon.

Examples:

IF A=B THEN GO TO S1; ELSE A=C;

CN CVERFLOW GO TO OVFIX;

Each PL/I statement is described in the alphabetic list of statements in Chapter 8.

PREFIXES

There are two types of prefixes: label prefixes and condition prefixes.

SUBSCRIPTRANGE
CHECK (identifier list)

NOTE: CHECK (identifier list) may be used as a prefix only with the PROCEDURE and BEGIN statements.

The meanings of these conditions are explained in "The ON Statement," in Chapter 8.

Any of these condition names may be preceded by the word NO. If NO is used, there can be no intervening blank between NO and the condition. For example, NOCONVERSICN can be specified in the prefix list.

GROUPS

A group is a collection of one or more statements and is used for control purposes.

A group has one of two forms. The first form, called a DO-group, is:

```
[label:] . . . DO-statement
           program-element-1
           program-element-2
           .
           .
           .
           END [label];
```

The label following END must be one of the labels of the DO statement.

The DO statement is called the heading statement of the DO-group, and may specify iteration.

The second form of a group is simply a single statement, as follows:

```
[label:] . . . statement
```

The "statement" is any statement except DO, END, PROCEDURE, BEGIN, DECLARE, FORMAT, ENTRY, or any compile-time statement.

Example of the first form:

```
ALPHA: DO;
      A=B*C;
      IF A < 0 THEN DO; B=1; C=0; END;
      END ALPHA;
```

In the example above, any of the single statements -- except the DO and END statements -- is an example of the second form of a group.

BLOCKS

A block is a collection of statements that defines the program region -- or scope -- throughout which an identifier is established as a name. It also is used for control purposes.

There are two kinds of blocks, begin blocks and procedure blocks.

A begin block has the general form:

```
[label:] . . . BEGIN-statement
           program-element-1
           program-element-2
           .
           .
           .
           END [label];
```

The label following END must be one of the labels of the BEGIN statement.

A procedure block, or procedure, has the general form:

```
label: [label:] . . . PROCEDURE-statement
           program-element-1
           program-element-2
           .
           .
           .
           END [label];
```

The label following END must be one of the labels preceding the PROCEDURE statement.

The BEGIN statement and the PROCEDURE statement in the above forms are called heading statements.

While the labels of the BEGIN statement are optional, the PROCEDURE statement must have at least one label.

Although the begin block and the procedure have a physical resemblance and play the same role in delimiting scope of names (see "Scope of Declarations," in Chapter 4) and defining allocation and freeing of storage (see "Allocation of Data and Storage Classes," in Chapter 6), they differ in an important functional sense. A begin block, like a single statement, is activated by normal sequential flow, and it can appear wherever a single statement can appear. A procedure can only be activated remotely by CALL statements, by statements in which a CALL option appears, or by function references. When a program containing a procedure is executed, control passes around the procedure, from the statement before the procedure statement to the statement after the END statement of the procedure.

Since a procedure can be activated only by a reference to it, every procedure must have a name. The label required for the heading statement of a procedure serves as the procedure name. More than one label provides more than one name.

The procedure name gives a means of activating the procedure at its primary entry point. Secondary entry points can also be defined for a procedure by use of the ENTRY statement. The labels preceding all ENTRY statements in a given procedure and the heading statement of the procedure are collectively called entry names for the procedure.

As the above definition of block implies, any block A can include another block B, but partial overlap is not possible; block B must be completely included in block A. Such nesting may be specified to any depth.

A procedure that is not included in any other block is called an external procedure. A procedure included in some other block is called an internal procedure.

Every begin block must be included in some other block. Hence, the only external blocks are external procedures.

All of the text of a begin block except the labels preceding the heading statement of the block is said to be contained in the block.

All of the text of a procedure except the entry names of the procedure is said to be contained in the procedure.

That part of the text of a block B that is contained in block B, but not contained in any other block contained in B, is said to be internal to block B.

The entry names of an external procedure are not internal to any procedure and are called external names.

The notion of internal to is vital in the definition of scope (see "Scope of Declarations" in Chapter 4).

Example:

```

A: PROCEDURE;
   statement 1
   B: BEGIN;
      statement 2
      statement 3
      END B;
   statement 4
   C: PROCEDURE;
      statement 5
   X: ENTRY;
      D: BEGIN;
         statement 6
         statement 7
         END D;
      statement 8
      END C;
   statement 9
   END A;

```

In this example, statements 1 through 9 are labeled or unlabeled simple statements.

As the brackets on the right indicate, block A contains block B and block C, and block C contains block D.

Block A is an external procedure. The procedure name is A, which is an external name, and the only entry name for the procedure.

X is an entry name corresponding to a secondary entry point for procedure C.

Blocks B and D are begin blocks.

Block C is an internal procedure.

The text internal to block A consists of:

```

PROCEDURE;
statement 1
B:
statement 4
C:
X:
statement 9
END A;

```

The text internal to block B consists of:

```

BEGIN;
statement 2
statement 3
END B;

```

The text internal to block C consists of:

```

PROCEDURE;
statement 5
ENTRY;

```

```
D:
statement 8
END C;
```

The text internal to block D consists of:

```
BEGIN;
statement 6
statement 7
END D;
```

USE OF THE END STATEMENT

As the examples above imply, the END statement has the form:

```
END [label];
```

and is used to terminate a group or a block.

If the optional label following END is not used, the END statement terminates that unterminated group or block headed by the DO, BEGIN, or PROCEDURE statement that physically precedes, and appears closest to, the END statement.

If, however, a label (e.g., L) is used following END, the statement terminates that unclosed group or block headed by the DO, BEGIN, or PROCEDURE statement with the label L that physically precedes, and appears closest to, the END statement. Any groups or blocks headed by DO, BEGIN, or PROCEDURE statements contained in the terminated block L are also automatically terminated by the END statement END L. This feature eliminates the necessity of writing the intermediate END statements to terminate the contained blocks and groups.

The statement labeled L, which heads the group or block terminated by the END statement END L, is internal to a certain block in the program (see "Blocks," for a definition of internal to). The terminating statement END L, together with its own possible statement-labels, is also considered to be internal to the same block. (If the statement labeled L is a BEGIN or

PROCEDURE statement, this block, is of course, the block L.)

The END statement may itself be labeled, and a reference to this label can be made from any part of the program where the label is known. (For a definition of known, see "Basic Rule on Use of Names" in Chapter 4).

Example:

A: PROCEDURE;	A: PROCEDURE;
.	.
.	.
B: BEGIN;	B: BEGIN;
.	.
.	.
A: PROCEDURE;	A: PROCEDURE;
.	.
.	.
C: DO;	C: DO;
.	.
.	.
X: END B;	END;
END A;	END;
	X: END B;
	END A;

In the example on the left above, the statement X:END B terminates the DO groups, the internal procedure A, and the block B. The statement END A terminates the external procedure A.

The example on the right is equivalent to the example on the left.

The statement X:END B is internal to block B.

PROGRAMS

A program is a set of external procedures. Thus, by definition, a program is a set of procedure blocks, each of which is completely nested, and separate from the others.

CHAPTER 2: DATA ELEMENTS

DATA TYPES

Information that is operated on in a PL/I object program during execution is called data. Each data item has a definite type and representation.

The permitted data types are arithmetic, string, label, task, and event. The details for the specification of data type attributes are contained in Chapter 4.

ARITHMETIC DATA

An arithmetic data item is one that has a numeric value with characteristics of base, scale, mode, and precision. The data item may be represented either as a numeric field or in a coded form, that is, in an internal representation that is implementation dependent. A numeric field is a string of characters that is given a numeric interpretation by means of the PICTURE attribute (see Chapter 4). The base, scale, and precision are all specified in the picture of the numeric field. A data item in coded form does not have a PICTURE attribute, but has its characteristics given by the attributes specifying base, scale, mode, and precision.

Base (decimal or binary), scale (fixed-point or floating-point), and precision have reference to internal representation of the data described and to the internal arithmetic that is to be used.

Base

Arithmetic data may be specified as having either decimal or binary base.

Scale

Arithmetic data may be specified as having either fixed-point or floating-point scale. Fixed-point data items are rational numbers for which the number of decimal or binary digits is specified; the position of the decimal or binary point may also be specified by a scale factor. Floating-point data items are rational numbers in

the form of a fractional part and an exponent part.

Mode

Arithmetic data may be operated on in either the real or complex mode. In the complex mode, a data item is considered to consist of a number pair, the first member of the pair representing the real part of the complex number and the second, the imaginary part.

Precision

The precision of fixed-point data (w,d) is specified by giving the total number of binary or decimal digits, w , to be maintained and a scale factor, d . The precision of floating-point data is specified by giving only the total number of binary or decimal digits to be maintained (i.e., w).

STRING DATA

String data can be classified as character-string or bit-string. The length of a string data item is equivalent to the number of characters (for a character-string) or the number of binary digits (for a bit-string) in the item. A string data item of length zero is known as the null string.

Character-String Data

Character-string data consists of a string of zero or more characters in the data character set (see "Data Character Set," in Chapter 1). The string may be fixed or varying in length. The actual number of characters must be specified if it is of fixed length, and the maximum length must be specified if it is of varying length.

Bit-String Data

Bit-string data consists of a string of zero or more binary digits (0 and 1). The string may be fixed or varying in length. The actual length of the field must be specified if it is of fixed length, and the maximum length must be specified if it is of varying length.

STATEMENT-LABEL DATA

Statement-label data consists of labels of statements (see "Statement Labels" in this chapter).

SCALAR QUANTITIES

A data item may be either a constant or the value of a scalar variable. Constants and scalar variables are called scalar quantities.

CONSTANTS

A constant is a data item that denotes itself, i.e., its representation is both its name and its value; thus, it cannot change during the execution of a program. Each constant has a type, as described below. A signed constant is a constant preceded by one of the prefix operators + or -. Wherever the word "constant" appears alone, and refers to an arithmetic constant, it is to be assumed to refer to an unsigned constant.

Real Arithmetic Constants

A real arithmetic constant is either binary or decimal.

Decimal Fixed-Point Constants

A decimal fixed-point constant is represented by one or more decimal digits with an optional decimal point.

Examples:

72.192
.308
255.
158

Binary Fixed-Point Constants

A binary-fixed point constant is represented by one or more binary digits with an optional binary point followed by the letter B.

Examples:

11011B
11.1101B
.001B

Sterling Fixed-Point Constants

Sterling quantities may be specified and will be interpreted as decimal fixed-point pence. A sterling fixed-point constant consists of the following concatenated fields:

a pounds field that is a decimal integer
a decimal point
a shillings field that is a decimal integer less than 20
a decimal point
a pence field that is one or more decimal digits with an optional decimal point (the integral part must be less than 12.)
an L

Examples:

101.13.8L
1.10.0L
0.0.2.5L

Decimal Floating-Point Constants

A decimal floating-point constant is represented by one or more decimal digits with an optional decimal point, followed by the letter E, followed by an optionally signed decimal exponent.

Examples:

12.E23
317.5E-16
0.1E+3
.42E+73
32E-5

Binary Floating-Point Constants

A binary floating-point constant is represented by one or more binary digits with an optional binary point, followed by the letter E, followed by an optionally signed binary exponent, followed by the letter B. The exponent is a string of decimal digits specifying a power of two.

Examples:

```
1.1011E3B
.11011E-27B
```

Precision of Real Arithmetic Constants

For purposes of expression evaluation, an apparent precision is defined for real arithmetic constants.

Real fixed-point constants have an apparent precision (p,q) where p is the total number of digits in the constant and q is the number of digits specified to the right of the decimal point.

The precision of a sterling constant is equivalent to the precision of its corresponding value in fixed-point pence. This value is determined as follows: multiply the value of the pounds field by 240; add the product of 12 and the value of the shillings field; add the value of the pence field. The precision of the result (with leading zeros removed) is the precision of the corresponding sterling constant.

The precision of a floating-point constant is (p) where p is the number of digits of the constant left of the E. If only the digit zero is left of the E, the precision is 1.

Examples:

```
3.14 has precision (3,2)
0.012E5 has precision (4)
0.9.0.5L has precision (4,1)
0000001 has precision (7,0)
```

Imaginary Arithmetic Constants

An imaginary constant represents a complex value of which the real part is zero and the imaginary part is the value specified.

It is represented by a real arithmetic constant, other than a sterling constant, followed by the letter I. PL/I does not define complex constants with non-zero real parts, but provides the facility to specify such data through expression, e.g., 10.1+9.2I.

Examples:

```
27I
3.968E10I
```

String Constants

String constants can be classified as character-string constants or bit-string constants.

Character-String Constants

A character-string constant is zero or more characters in the data character set enclosed in quotation marks. If it is desired to represent a quotation mark, it must appear as two immediately adjacent quotation marks. The constant may optionally be preceded by a decimal-integer constant in parentheses to specify repetition. If the constant specifying repetition is zero, the result is the null string.

Examples:

```
'$ 123.45'
'JOHN JONES'
'IT'S'
(3)'TOM'
```

The latter is exactly equivalent to

```
'TOMTOMTOM'
```

Bit-String Constants

A bit-string constant is zero or more binary digits enclosed in quotation marks, followed by the letter B. The constant may optionally be preceded by a decimal-integer constant in parentheses, to specify repetition. If the constant specifying repetition is zero, the result is the null string.

Examples:

```
'0100'B
(10)'1'B
```

The latter is exactly equivalent to

```
'1111111111'B
```

Statement-Label Constants

A statement-label constant is an identifier which appears in the program as a statement label.

The value of a label constant becomes undefined when the block to which it is internal becomes inactive.

VARIABLES

A scalar variable, like a constant, denotes a data item. This data item is called the value of the scalar variable. Unlike a constant, however, a variable may take on more than one value during the execution of a program. The set of values that a variable may take on is the range of the variable. The range of a variable is always restricted to one data type and, if the type is arithmetic, to one base, scale, mode, and precision. If there are no further restrictions declared for the range, the variable may assume values over the entire set of data of that type.

Reference is made to a scalar variable by a name, which may be a simple name, a subscripted name, a qualified name, or a subscripted qualified name (see "Naming" in this chapter).

DATA AGGREGATES

In PL/I, variable data items are grouped into arrays or structures. Rules for this grouping are given below. (For the method of referring to an array or structure or a particular item of an array or structure, see "Naming," in this chapter.)

ARRAYS

An array is an n-dimensional, ordered collection of elements, all of which have identical data declaration. If arithmetic, all of the elements of the array must have the same base, scale, mode, and precision or the same picture. If character-string or bit-string, all of the elements must have the same actual length, if fixed length, or the same maximum length, if varying length. The number of dimensions of an array, and the upper and lower bounds of each dimension, are specified by the use of the dimension attribute.

Example:
DECLARE A (3,4);

This statement defines A as an array with 2 dimensions: 3 rows and 4 columns. The matrix given below illustrates the array A.

A (1,1)	A (1,2)	A (1,3)	A (1,4)
A (2,1)	A (2,2)	A (2,3)	A (2,4)
A (3,1)	A (3,2)	A (3,3)	A (3,4)

The elements of an array may be structures (see "Arrays of Structures").

STRUCTURES

A structure is a hierarchical collection of scalar variables, arrays, and structures. These need not be of the same data type nor have the same attributes.

Structures may contain structures. The outermost structure is the major structure, and contained structures are minor structures. A major structure must be at level one. Contained structures must always have a level number numerically greater than the structure in which they are contained. Identifiers preceded by level numbers but having no components are not considered to be structures. The level number may be followed by an optional blank. (Additional information on structures can be found in the section "Structure Declarations and Attributes" in Chapter 4.)

Examples:

1. DECLARE 1 PAYROLL, 2 NAME, 2 HOURS, 3 REGULAR, 3 OVERTIME, 2 RATE;

takes the form:

```
1 PAYROLL
  2NAME
  2HOURS
    3REGULAR
    3OVERTIME
  2RATE
```

In the above example PAYROLL is defined as the major structure containing the scalar variables NAME and RATE and the structure HOURS. The structure HOURS contains the scalar variables REGULAR and OVERTIME.

2. DECLARE 1 A, 2 B, 2 C, 3 D (2), 3 E, 2 F;

This takes the form:

```
A
  B
  C
    D (1)
    D (2)
  E
  F
```

The decimal integers before the identifiers specify the level; the decimal integer in parentheses specifies the bounds of the one-dimensional array. A is defined as the major structure and contains the minor structure C and the scalar variables B and F. C contains D, a one-dimensional

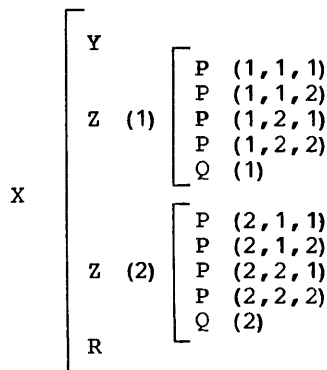
array with two scalar variables, and the scalar variable E.

3. DECLARE 1 A, 3 B, 2 C;

This takes the form:

```
A
  B
  C
```

Note that B and C are at the same level although their level numbers differ.



NAMING

ARRAYS OF STRUCTURES

An array of structures is formed by giving the dimension attribute to a structure. This dimension attribute causes all contained items to be arrays.

Examples:

1. DECLARE 1 CARDIN(3), 2 NAME, 2 WAGES, 3 NORMAL, 3 OVERTIME;

The decimal integers before the identifiers specify the level. The name, CARDIN, represents an array of structures. Because CARDIN has a dimension specified, NAME, NORMAL, and OVERTIME are arrays, and their elements are referred to by subscripted names.

The form of the data is:

```

CARDIN (1)  NAME (1)
            WAGES (1)  NORMAL (1)
                        OVERTIME (1)

CARDIN (2)  NAME (2)
            WAGES (2)  NORMAL (2)
                        OVERTIME (2)

CARDIN (3)  NAME (3)
            WAGES (3)  NORMAL (3)
                        OVERTIME (3)

```

2. DECLARE 1 X, 2 Y, 2 Z (2), 3 P (2,2), 3 Q, 2 R;

X is an undimensioned major structure containing scalar variables, arrays, and a structure.

Y is a scalar variable
 Z is an array of structures
 P is a three-dimensional array
 Q is a one-dimensional array
 R is a scalar variable

The form of the data is:

This section describes the rules for referring to a particular data item, groups of items, arrays, and structures. The permitted types of data names are simple, qualified, subscripted, and subscripted qualified.

SIMPLE NAMES

A simple name is an identifier (see "Identifiers," in Chapter 1) that refers to a scalar, an array, or a structure.

SUBSCRIPTED NAMES

A subscripted name is used to refer to an element of an array. It is a simple name that has been declared to be the name of an array followed by a list of subscripts. The subscripts are separated by commas and are enclosed in parentheses. A subscript is an expression that is evaluated and converted to an integer before use (see "Evaluation of Expressions," in Chapter 3). The number of subscripts must be equal to the number of dimensions of the array, and the value of a specified subscript must fall within the bounds declared for that dimension of the array.

A subscripted name takes the form:

```
identifier (subscript [ , subscript] ...)
```

Examples:

```

A (3)
FIELD (B,C)
PRODUCT (SCOPE * UNIT + VALUE, PERIOD)
ALPHA (1,2,3,4)

```

Cross Sections of Arrays

The concept of cross sections is a logical extension of the subscripting notation. A cross section of an array is referred to by the array name, followed by a list of subscripts, at least one of which is an asterisk. The subscripts are separated by commas, and the entire list is enclosed in parentheses. The number of items in the list must be equal to the number of dimensions of the array. If the array is of dimensionality n , then an asterisk may appear in $k \leq n$ positions. If the j th list position is occupied by an asterisk, the cross section of the array includes elements covered by varying the j th subscript between its bounds. The dimensionality of the cross section is equal to the number of asterisks, k , in the subscript list. If all subscript positions are occupied by asterisks, then this reference to the cross section is equivalent to a reference to the entire array.

A cross section may be used anywhere that the name of an array of dimensionality k is required. Subsequent references to the word "array" in this document should therefore be taken to include cross sections of arrays.

Examples:

1. A (3,*) denotes the third row of the array A.
2. B (*, *, 2) is a two-dimensional cross section and denotes the second plane of the array B.
3. If MATRIX is the array:
1 2 3
4 5 6
7 8 9
MATRIX (*, 2) is the vector:
2
5
8

QUALIFIED NAMES

A simple name usually refers uniquely to a scalar variable, an array, or a structure. However, it is possible for a name to refer to more than one variable, array, or structure if the identically named items are themselves parts of different structures. In order to avoid any ambiguity in referring to these similarly named items, it is necessary to create a unique name; this is done by forming a qualified name. This means that the name common to more than one item is preceded by the name of the structure in which it is contained. This, in turn, can be preceded by the name

of its containing structure, and so on, until the qualified name refers uniquely to the required item. The section "Multiple Declarations and Ambiguous References" in Chapter 4, contains further information on this subject.

Thus, the qualified name is a sequence of structure names specified left to right in order of increasing level numbers; the simple names are separated by periods, and blanks may be placed as desired around the periods. The sequence of names need not include all of the containing structures, but it must include sufficient names to resolve any ambiguity.

The qualified name, once composed, is itself a name. Subsequently, in this document, when the terms scalar variable name, array name, or structure name are used they should also be taken to include qualified names.

A qualified name takes the form:

identifier { . identifier } ...

Examples:

1. A program may contain the structures:

```
DECLARE 1 CARDIN, 2 PARTNO, 2 DESCRIPTION, 2 PRICE;  
DECLARE 1 CARDOUT, 2 PARTNO, 2 DESCRIPTION, 2 PRICE;
```

Elements are then referred to as:

```
CARDIN.PARTNO  
CARDOUT.PARTNO  
CARDIN.PRICE
```

2. A program may contain the structure:

```
DECLARE 1 MARRIAGE, 2 MAN, 3 NAME, 3 DATE,  
2 WOMAN, 3 NAME, 3 DATE;
```

Elements are then referred to as:

```
MAN.NAME  
or MARRIAGE.MAN.NAME
```

```
WOMAN.NAME  
or MARRIAGE.WOMAN.NAME
```

3. If the same program also contains the structure:

```
DECLARE 1 BIRTH, 2 WOMAN, 3 NAME,  
3 DATE, 2 COMPLEXION;
```

Elements are then referred to as:

```
MAN.NAME  
or MARRIAGE.MAN.NAME
```

```
MARRIAGE.WOMAN.NAME
```

BIRTH.NAME
or BIRTH.WOMAN.NAME

COMPLEXION

and the minor structures referred to
as:

MARRIAGE . WOMAN

BIRTH . WOMAN

SUBSCRIPTED QUALIFIED NAMES

The elements of an array contained in a structure and requiring name qualification for identification are referred to by subscripted qualified names. A subscripted qualified name is a sequence of names and subscripted names separated by periods. The order of names is as given for any qualified name. The subscript list following each name refers to the dimensions associated with the name if the name is declared to be the name of an array in the structure description.

As long as the order of the subscripts remains unchanged, subscripts may be moved to the right and attached to names at a deeper level. Unless all of the subscripts are moved to the deepest level, the qualified name is said to have interleaved subscripts.

Provided that sufficient structure names are used to make the name unique, as described for qualified names, and that the total number of subscripts is the same as the total dimensionality of the array, unsubscripted structure names may be omitted in references. Ambiguity of names, however, cannot be resolved by subscripting. A subscripted qualified name takes the general form:

```
identifier [ (subscript [, subscript]
...)]
{. identifier [(subscript [, sub-
script]...)] }...
```

If any subscripts are given in a reference to a qualified name, all those subscripts which apply to dimensions of containing structures must be given.

A subscripted qualified name must have at least one subscript.

Examples:

A is an array of structures with the following description:

```
DECLARE 1 A (10,12) 2 B (5), 3 C (7),
3 D;
```

The following subscripted qualified names refer to the same element, which is the seventh element of C contained in the fifth element of B contained in tenth row and twelfth column of A:

- (1) A (10,12) . B (5) . C (7)
- (2) A (10) . B (12,5) . C (7)
- (3) A (10) . B (12) . C (5,7)
- (4) A . B (10,12,5) . C (7)
- (5) A . B (10,12) . C (5,7)
- (6) A . B (10) . C (12,5,7)
- (7) A . B . C (10,12,5,7)
- (8) A (10,12) . B . C (5,7)
- (9) A (10) . B . C (12,5,7)

If structure B, but not structure A, is necessary for unique identification of this use of C, any of forms (4), (5), (6), or (7) may be used without including the A.

If structure A, but not B, is necessary for identification of C, forms (7), (8), or (9) may be used without including the B.

Except for form (7), all of the qualified names in the above example have interleaved subscripts.

STATEMENT LABELS

Statement-label data is used only in connection with statement labels. Statement-label data may be a constant, a scalar variable, or an array.

CONSTANT

A statement-label constant is an identifier that permits references to be made to statements.

Example:

```
.
.
ROUTINE1: IF X > 5 THEN GO TO EXIT;
.
.
GO TO ROUTINE1;
```

ROUTINE1 is a statement-label constant. EXIT is also a statement-label.

VARIABLE

A label variable is a variable that has as values statement-label constants.

EVENT NAMES

Event names are used only in connection with events (see "Asynchronous Operations and Tasks," in Chapter 6 and "The EVENT Attribute," in Chapter 4). Event names may have the dimension attribute or may be elements of structures. A simple event name has an associated completion status. This status is denoted by the value '0'B

for "not completed" and '1'B for "completed". If the event name has been associated with a given task through the use of an EVENT option in a CALL statement (see Chapter 8), the completion status of the event name will reflect the completion status of the task itself. The completion status of an event name may also be set explicitly by the execution of an assignment statement using the EVENT pseudo-variable (see Chapter 8).

EXPRESSIONS

An expression is an algorithm used for computing a value. Expressions are of the three types: scalar, array, and structure, depending upon the types of the operands involved. The type of the result is also the same as that of the operands. An array (or structure) expression is simply an array (or structure) evaluated by expansion of the expression into a collection of scalar expressions (see "Array Expressions" and "Structure Expressions"). Syntactically, a scalar expression consists of a constant, a scalar variable, a function reference, a scalar expression enclosed in parentheses, a scalar expression preceded by a prefix operator, or two scalar expressions connected by an infix operator. Operands in a scalar expression need not have the same data attributes. If they differ, conversion will be performed before the operation.

SCALAR EXPRESSIONS

A scalar expression returns a scalar value. The type of the value is the type of the expression. The type of the expression is dependent upon the class of operators -- arithmetic, comparison, bit string, and concatenation (see "Operators"). Statement label disignators are not allowed in scalar expressions except as function arguments.

If A and B are expressions, then the operators + and - used in expressions of the form +A or -A, are called prefix operators. When these operators are used in expressions of the form A+B or A-B they are called infix operators.

Arithmetic Operations

An arithmetic expression of any complexity is composed of a set of elementary arithmetic operations.

An elementary arithmetic operation has the general form:

```
{[+|-] operand} | {operand
{+| - | * | / | **} operand}
```

The general form specifies the prefix operations of plus and minus and the infix operations of addition, subtraction, multiplication, division, and exponentiation. Operations are performed only with coded arithmetic data. If necessary, the data will be converted to coded arithmetic type before the operation is performed.

Mixed Characteristics

The two operands of an arithmetic operation may differ in base, scale, mode, and precision. When they differ, conversion takes place according to the following rules:

BASE: If bases differ, the decimal operand is converted to binary.

SCALE: If the scales of the operands differ, the fixed-point operand will be converted to floating-point, except in the case of exponentiation in which the first operand is floating-point and the second is fixed-point with precision (p,0). In the latter case, the second operand is not converted.

MODE: If modes differ, the real operand is converted to complex by assuming an imaginary part of zero.

PRECISION: If precisions differ, no conversion is done; the arithmetic operation is carried out on operands of differing precision in a way consistent with normal mathematical practice. However, a particular implementation may restrict the precision of operands used in arithmetic expressions, and if larger precisions are desired, the built-in functions ADD, SUBTRACT, etc. (see Appendix 1), may be used.

Results of Arithmetic Operation

After the conversions specified above have taken place, the arithmetic operation is performed. Any necessary truncations will be made towards zero, regardless of the base or scale of the operands. Algebraic signs of results will be correct, when an error condition causes truncation or a modulo result (see "OVERFLOW" in Appendix 3).

The base, scale, mode, and precision of the result depend on the operands and the operator in the following ways:

1. Prefix operations: The prefix operations of plus and minus yield a result having the base, scale, mode, and precision of the operand.
2. Floating-point: If the first operand of an infix operation is floating-point the result is floating-point, and the base and mode of the result are the common base and mode of the operands. The precision of the result is the greater of the precisions of the two operands.
3. Fixed-point: If the first operand of a binary operation is fixed, and if the operation is not exponentiation, the result is fixed, and the base and mode of the result are the common base and mode of the operands. If the operation is exponentiation, the second operand is converted to floating point if its scale factor is not zero; and the first operand is converted to floating-point unless the second operand is an unsigned integer constant meeting the conditions of item d below; in these cases, the rules for floating-point apply.

The precision of a fixed-point result depends on the operation and the precisions of the operands, according to rules given below. The following symbols are used:

- N the length of the largest number in the implementation
 m the total number of positions in the result
 n the scale factor of the result
 p the total number of positions in operand one
 q the scale factor of operand one
 r the total number of positions in operand two
 s the scale factor of operand two
 y value of operand two, if it is an unsigned integer constant

a. Addition and subtraction:

$$m = \min(N, \max(p-q, r-s) + \max(q, s) + 1)$$

$$n = \max(q, s)$$

b. Multiplication:

$$m = \min(N, p+r+1)$$

$$n = q*s$$

c. Division:

$$m = N$$

$$n = N-p+q-s$$

d. Exponentiation: if second operand is an unsigned non-zero integer constant,

$$m = (p+1) * y - 1$$

$$n = q * y$$

If $m > N$, however, or y is not an unsigned non-zero integer constant, the first operand is converted to floating-point and rules for floating-point exponentiation apply. If $m \leq N$, the result is obtained by repeated multiplication (see note below for a definition of exponentiation).

e. The above rules hold for both real and complex mode.

NOTE: The operation of exponentiation is defined as follows:

1. Real Mode, $x_1 ** x_2$:

- a. If $x_1 = 0$ and $x_2 < 0$, the ZERO-DIVIDE condition is raised.
- b. If $x_1 = 0$ and $x_2 = 0$, the ERROR condition is raised and the result is set to 0.
- c. If $x_1 = 0$ and $x_2 > 0$, the result is 0.
- d. If $x_1 < 0$ and $x_2 < 0$, the ERROR condition is raised for any x_2 of a form other than fixed-point $(p, 0)$; i.e., anything but a fixed-point integer. When x_2 is a fixed-point integer, the result is

$$1 / [(x_1 * x_1 * x_1 * \dots * x_1), \text{ABS}(x_2) \text{ times}]$$

- e. If $x_1 < 0$ and $x_2 = 0$, the result is 1.
- f. If $x_1 < 0$ and $x_2 > 0$, the limitations are identical to those in item d. When x_2 is a fixed-point integer, the result is $(x_1 * x_1 * x_1 * \dots * x_1), x_2$ times.
- g. If $x_1 > 0$ and $x_2 < 0$, the result is

$$\text{EXP}(x_2 * \text{LOG}(x_1))$$

- h. If $x_1 > 0$ and $x_2 = 0$, the result is 1.
- i. If $x_1 > 0$ and $x_2 > 0$, the result is $\text{EXP}(x_2 * \text{LOG}(x_1))$ for any x_2 except a fixed-point integer, in which case the result is the same as in item f. (See Appendix 1 for a definition of EXP and LOG.)

2. Complex Mode, $z_1 ** z_2$:

If $z_1 = 0$, the ERROR condition is raised in all cases except when the real part of z_2 is > 0 and the imaginary part of z_2 equals 0, in which case the result is 0. Otherwise, the result is $\text{EXP}(z_2 * \text{LOG}(z_1))$.

Arithmetic Conversions

1. Arithmetic Mode Conversion

If a complex value is converted to a real value, the result is the real part of the complex value.

If a real value is converted to a complex value, the result is a complex value that has the real value as the real part and zero as the imaginary part.

2. Integer conversion

If conversion to integer is specified, as in the evaluation of subscript expressions, the conversion will be to fixed-point binary (x,0). Here x is the total number of positions in the field and depends upon the implementation. The scale factor is zero. Truncation, if necessary, will be toward zero.

3. Arithmetic Base and Scale Conversion

Table 1 defines the precision resulting from base and scale conversion. CEIL refers to the ceiling of the expression. (The "ceiling" of a number is the smallest integer equal to or greater than the number.)

Conversion from floating-point scale to fixed-point scale will occur only when a destination precision is known, as in an assignment to a fixed-point variable. If the destination precision is incapable of holding the floating point value, truncation on both left and right will occur, and

the SIZE error condition will be raised (unless disabled).

Bit-String Operations

Bit-string operations have the following general forms:

```

operand
operand & operand
operand | operand
    
```

The prefix operation "not" and the infix operations "and" and "or" are specified above. The operands will be converted to bit-string type before the operation is performed. The result will be of bit-string. If the operands are of different lengths after conversion, the shorter is extended on the right with zeros to the length of the longer. The length of the result will be of this extended length. The result is of varying length if either operand is of varying length or is a reference to the SUBSTR built-in function. Otherwise, the result is of fixed length.

The operations are performed on a bit-by-bit basis. As a result of the operations, each bit position has the value defined in the following table:

Table 1. Arithmetic Base and Scale Conversion

After	Before Conversion			
	Binary Fixed (p,q)	Decimal Fixed (p,q)	Binary Float (p)	Decimal Float (p)
Binary Fixed	(p,q)	(MIN (CEIL (p*3.32) +1, N) , CEIL (ABS (q) *3.32) *SIGN (q))		
Decimal Fixed	(CEIL (p/3.32) +1, CEIL (ABS (q) /3.32) *SIGN (q))	(p,q)		
Binary Float	(p)	(MIN (CEIL (p*3.32) , N))	(p)	(MIN (CEIL (p*3.32) , N))
Decimal Float	(CEIL (p/3.32))	(p)	(CEIL (p/3.32))	(p)

A	B	NOT A	NOT B	A AND B	A OR B
1	1	0	0	1	1
1	0	0	1	0	1
0	1	1	0	0	1
0	0	1	1	0	0

Examples:

If field A is '010111'B, field B is '111111'B, and field C is '101'B, then

A yields '101000'B
 C & B yields '101000'B
 A | C yields '010111'B
 (A | C) & B yields '101111'B

For a discussion of how these expressions are evaluated, see "Evaluation of Expressions," in this chapter.

Comparison Operations

Comparison operations have the general form:

operand {<|<=|=|=|>|=|>} operand

There are three types of comparisons:

1. Algebraic, which involves the comparison of signed numeric values in coded arithmetic form. Conversion of numeric fields will be performed.
2. Character, which involves left-to-right, pair-by-pair comparisons of characters according to a collating sequence. If the operands are of different length, the shorter is extended to the right with blanks.
3. Bit, which involves the left-to-right comparison of binary digits. If the strings are of different lengths, the shorter is extended on the right with zeros.

The result of a comparison is a bit string of length one; the value is '1'B if the relationship is true or '0'B if it is false.

If the operands of a comparison are of different types, the operand of the lower type is converted to the operand of the higher type. The priority of types is (1) arithmetic (highest), (2) character string, (3) bit string.

As a result of the conversion, both operands will then be arithmetic or character string, and algebraic or character comparison will be performed.

Only the operations = and & are defined when either operand is complex.

Concatenation Operations

Concatenation operations have the following general form:

operand||operand

If both operands are of bit-string type, no conversion is performed, and the result is of bit type. In all other cases, the operands are converted where necessary to character-string type before the concatenation is performed, and the result is of character type. The result is of varying length if either operand is of varying length or is a reference to the SUBSTR built-in function. Otherwise, the result is of fixed length.

Examples:

If A is '010111'B, B is '101'B, C is 'XY,Z' and D is 'AA/BB', then

A||B yields '010111101'B
 A||A||B yields '010111101111101'B
 C||D yields 'XY,ZAA/BB'
 D||C yields 'AA/BBXY,Z'

Type Conversion

Bit String to Character String

The bit 1 becomes the character 1, and the bit 0, the character 0. The length is unchanged. The null bit string becomes the null character string.

Character String to Bit String

The characters 1 and 0 become the bits 1 and 0. The conversion is illegal if the character string contains characters other than 0 and 1. The null character string becomes the null bit string.

Character String to Arithmetic

The character string is interpreted according to the rules of list directed

input (see "List-Directed Input Format," in Chapter 7). The value is converted directly to an operand with the same base, scale, mode, and precision that a decimal real fixed-point variable of default precision would have been converted to if it had appeared. The null string is converted to the value zero. Sterling constants are not permitted.

Bit String to Arithmetic

The bit string is interpreted as an unsigned binary integer, and converted to fixed-point binary, precision (S,0), where S depends upon the implementation. The null string is converted to the value zero.

Arithmetic to Character String

CODED ARITHMETIC AND BINARY NUMERIC FIELDS: The arithmetic value is converted to a character string according to the rules of list-directed output specified in Chapter 7.

DECIMAL NUMERIC FIELDS: The numeric field is interpreted as a character string (see Appendix 2).

Arithmetic to Bit String

CODED ARITHMETIC AND DECIMAL NUMERIC FIELDS: The arithmetic value is converted to real then to fixed-point binary, precision (p,0), where p is related to the precision before conversion as follows (with ceilings of expressions used):

BINARY FIXED (r,s) p = min(N,max(r-s,0))
 BINARY FLOAT (r) p = r
 DECIMAL FIXED (r,s) p = min(N,max(CEIL
 (r-s)*3.32),0))
 DECIMAL FLOAT (r) p = min(N,CEIL(r*3.32))

The resulting binary fixed-point value is interpreted as a bit string of length p.

BINARY NUMERIC FIELDS: The numeric field is interpreted as a bit string.

ARRAY EXPRESSIONS

If the operands of an expression refer to arrays or to a combination of scalars and arrays, the expression is an array expression.

An array expression returns an array result. That is, all operations performed on arrays are performed on an element-by-element basis. Therefore, all arrays

referred to in an array expression must be of identical bounds.

Note: Array expressions are not always expressions of conventional matrix algebra.

The appearance of a function reference (other than a built-in function) will imply a scalar result. For example, if A is an array, PROCEDURE (A) is a scalar function with an array argument.

The built-in functions listed under "Arithmetic Generic Functions," "Float Arithmetic Generic Functions," and "String Generic Functions," in Appendix 1 may participate in array expressions with array results. An array may be substituted for any of the arguments of these functions except those arguments which are required to be integer constants, or those which must be converted to integers.

Prefix Operators and Arrays

The result of the operation of a prefix operator or a built-in function upon an array is an array of identical bounds, each element of which is the result of the operation having been performed upon each of the corresponding elements of the original array.

Example:

If A is the array	5	3	-9
	1	-2	7
	6	3	-4
then -A is the array	-5	-3	9
	-1	2	-7
	-6	-3	4

Infix Operators and Arrays

Scalar - Array Operations

The result of an operation in which a scalar and an array are connected by an infix operator is an array of bounds identical to the original, each element of which is the result of the operation performed upon the scalar and upon each of the corresponding elements of the original array.

Example:

If A is the array	5	10	8
	12	11	3

then 3*A is the array 15 30 24
 36 33 9

A (1) .P.Q+B (1) .C.D*2,
A (1) .P.R.+B (1) .C.E*2,
.
.
.
A (3) .S.U+B (3) .F.H*2

Array - Array Operations

The result of an operation in which two arrays of identical bounds are connected by an infix operator is an array of bounds identical to the original arrays, each element of which is the result of the operation performed upon the corresponding elements of the two original arrays by the infix operator.

Example:

If A is the array 2 4
 3 6
 1 7
 4 8

and if B is the array 1 5
 7 8
 3 4
 6 3

then A + B is the array 3 9
 10 14
 4 11
 10 11

A*B is the array 2 20
 21 48
 3 28
 24 24

and MAX (A+B,A*B) is the array
 3 20
 21 48
 4 28
 24 24

Array Expressions Involving Structures

An array expression may involve an array of structures.

Example:

Let A and B be arrays of structures:

1 A (3)	1 B (3)
2 P	2 C
3 Q	3 D
3 R	3 E
2 S	2 F
3 T	3 G
3 U	3 H

Then, A+B*2 is a valid expression that will result in each element of the array B being multiplied by the constant 2 and added to the corresponding element of the array A. The above expression, A+B*2, is equivalent to the following:

STRUCTURE EXPRESSIONS

The operands of a structure expression are structures, or a combination of structures and scalars. A structure expression returns a structure result. Array operands are not allowed in structure expressions.

All operations performed on structures are performed on an element-by-element basis. Thus, all structures appearing in a structure expression must have identical structuring. This means that the structure must have the same number of contained scalars and arrays. The positioning of the scalars and arrays within the structure must be the same, and arrays similarly positioned must have identical dimensions and bounds. The data types need not be the same.

When an operation has one structure and one scalar operand, it is interpreted as many operations, one for each scalar element in the structure. Each sub-operation involves a structure element and the scalar operand.

A structure expression is a shorthand method of applying an expression to each item of a structure.

Example:

If there are two structures:

1 A	1 B
2 PART1	2 PART1
3 SUBPART1	3 SUBPART1
3 SUBPART2	3 ALPHA
3 SUBPART3	3 SUBPART2
2 PART2	2 PART2
3 SUBPART4	3 ALPHA
3 BETA	3 SUBPART4
3 SUBPART5 (3)	3 SUBPART5 (3)

Then the expression A-2*B is shorthand for the following expressions:

A . SUBPART1 - 2*B . SUBPART1
A . SUBPART2 - 2*B . PART1 . ALPHA
A . SUBPART3 - 2*B . SUBPART2
A . SUBPART4 - 2*B . PART2 . ALPHA
A . BETA - 2*B . SUBPART4
A . SUBPART5 - 2*B . SUBPART5

Note that the last expression is an array expression.

```

unit ::= negation | unit1

unit1 ::= constant | scalar-
           variable | function-
           reference |
           (expression)

```

EVALUATION OF EXPRESSIONS

In the following syntactical definition of an expression, operator precedence is indicated as extending from prefix +, prefix -, and ** at highest precedence, down through || at lowest precedence. This hierarchy is modified as indicated by parentheses, and by , which as a unary operator has precedence over any operators immediately to its left, but has lower precedence than relational and arithmetic operators to its right.

Recursion on the right of an operator indicates right-to-left evaluation (prefix operators and **), while recursion on the left of an operator indicates left-to-right evaluation. The operators + and * are commutative, but not associative, as low-order rounding errors will depend on the order of evaluation of an expression. Thus A+B+C is not necessarily equal to A+(B+C).

```

expression ::= union | {expression ||
                        union}

union ::= intersection | {union|
                        intersection}

intersection ::= negation | { intersec-
                        tion&negation}

negation ::= {, negation} | relation

relation ::= sum1|{relation=sum}
            |{relation,=sum}
            |{relation>sum}
            |{relation>=sum}
            |{relation<sum}
            |{relation<=sum}

sum ::= negation | sum1

sum1 ::= product | {sum1 +
                  product} | {sum1 -
                  product}

product ::= negation | product1

product1 ::= factor | {product1 *
                    factor} | {product1 /
                    factor}

factor ::= negation | factor1

factor1 ::= unit | {+ factor} | {-
                factor} |
            {unit**factor}

```

The rules relating to abnormal functions and abnormal data should be noted (see "Abnormality," in Chapter 10).

ORDER OF THE EVALUATION OF EXPRESSIONS

Consider any scalar expression in which two operands are separated by an infix operator in the form of A op B, where "op" denotes any operator. Then either A or B, or both, may be a "composite operand", that is, a subscripted name, a function reference, or a subexpression of the form op C, C op D, or (C). In such cases, the subscripts and arguments that must be evaluated and the functions that must be invoked before the operator is applied, are termed the "elements" of the operand. For the purposes of this discussion, an operand that is an unsubscripted name or a constant is termed a "simple operand."

If A is a simple operand and B is not a simple operand, then A will not be accessed until all the elements of B are accessed. Otherwise all elements of A are accessed before B is accessed. Subscript lists are evaluated and accessed, left to right, immediately before the accessing of the array elements. Argument lists are evaluated and accessed, left to right, immediately before the function is invoked.

Array expressions are evaluated by performing, in turn, a complete scalar evaluation of the expression for each position of the array. The evaluations proceed in row-major order. The result of an evaluation for an earlier position can alter the values of scalar elements for the evaluation of a later position (see Example 1, for "The Assignment Statement," in Chapter 8).

Structure expressions are evaluated by performing a complete scalar evaluation of the expression for each eligible field, in the order in which the fields in the structures are stored. The results of an evaluation for an earlier position can alter the result for the evaluation of a later position.

CHAPTER 4: DATA DESCRIPTION

ATTRIBUTES

An identifier appearing in a PL/I program may refer to one of many classes of objects. It may, for example, represent a variable referring to a complex number expressed in fixed-point form with decimal base; it may refer to a file; it may represent a variable referring to a character string; it may represent a statement label or represent a variable referring to a statement label, etc.

Those properties that characterize the object represented by the identifier, and other properties of the identifier itself (such as scope, storage class, etc.), together make up the set of attributes which can be associated with an identifier.

There are a number of classes of attributes. These classes and the attributes in each class are described further on in this chapter.

When an identifier is used in a given context in a program, attributes from certain of these attribute-classes must be known in order to assign a unique meaning to the identifier. For example, if an identifier is used as a data variable, the data type must be known; if the data type is arithmetic, the base, scale, mode, and precision must be known.

Examples of Attributes:

CHARACTER (50) Association of this attribute with an identifier defines the identifier as representing a variable referring to a string 50 characters in length.

FLOAT Association of this attribute with an identifier defines the identifier as representing a variable referring to arithmetic data, where the data is represented internally in floating-point form.

EXTERNAL Association of this attribute with an identifier defines the identifier as a name with a certain special scope.

DECLARATIONS

A given identifier is established as a name, which holds throughout a certain scope in the program (see "Scope of Declarations" in this chapter), and a set of attributes may be associated with the identifier by means of a declaration.

If a declaration is internal to a certain block, then the declared identifier is said to be declared in that block.

In a given program, an identifier may represent more than one name. In this case, each different name represented by the identifier is said to be a different use of the identifier. For example, an identifier may represent an arithmetic variable in one part of a program and an entry name in another part. These two parts, of course, cannot overlap.

Each different use of the identifier is established by a different declaration. References to different uses are distinguished by the rules of scope (see "Scope of Declarations").

Declarations may be explicit, contextual, or implicit.

EXPLICIT DECLARATIONS

Explicit declarations are made through use of the DECLARE statement, by which an identifier can be established as a name and given a certain set (possibly empty) of attributes.

Only one DECLARE statement can be used to establish a given use of a given identifier, and all of the explicitly declared attributes for this use must be specified in the DECLARE statement.

The DECLARE Statement

Function:

The DECLARE statement is a non-executable statement used for the specification of attributes of simple names.

General Format:

```
DECLARE [level] name [attribute] ...
[, [level] name [attribute] ...] ...;
```

Syntax rules:

1. Any number of identifiers may be declared as names in one DECLARE statement and must be separated by commas.
2. Attributes must follow the names to which they refer. (Note that the above format does not show factoring of attributes, which is allowable as explained later).
3. "Level" is a non-zero decimal integer constant. If it is not specified, level 1 is assumed. A blank space is not required to separate a level number from the name following it.

General Rules:

1. All of the attributes for a particular name must be declared together in one DECLARE statement.
2. Attributes of EXTERNAL names, declared in separate blocks and compilations, must not conflict or supply explicit information that was not explicit or implicit in other declarations.

Example:

```
DECLARE JOE FLOAT, JIM FIXED (5,3),
        JACK BIT (10);
```

JOE is declared to be a floating-point scalar variable, JIM a five-position, fixed-point scalar variable with three places to the right of the decimal, and JACK a scalar variable of ten bits.

Factoring of Attributes

Attributes common to several name declarations can be factored to eliminate repeated specification of the same attribute for many identifiers. This factoring is achieved by enclosing the name declarations in parentheses, and following this by the set of attributes which are to apply. In the case of a factored level number, the level number precedes the parenthesized list of name declarations.

Example:

```
DECLARE ((A FIXED, B FLOAT) STATIC,
        C CONTROLLED) EXTERNAL SYMBOL;
```

This declaration is equivalent to the following:

```
DECLARE A FIXED STATIC EXTERNAL
        SYMBOL,
        B FLOAT STATIC EXTERNAL SYMBOL,
        C CONTROLLED EXTERNAL SYMBOL;
```

Multiple Declarations and Ambiguous References

Two or more declarations of the same identifier, internal to the same block, constitute a multiple declaration of that identifier only if they have identical qualification (including the case of two or more declarations of an identifier at level 1, i.e., scalars or major structures). In a multiple declaration, only the first declaration (by physical appearance) of the identifier is legal; the others are in error.

Reference to a qualified name is always taken to apply to the identifier (for which the reference is valid) declared in the innermost block containing the reference. Within this block, the reference is unambiguous if either of the following is true:

1. The reference gives a valid qualification for one and only one declaration of the identifier.
2. The reference represents the complete qualification of only one declaration of the identifier. The reference is then taken to apply to this identifier.

Otherwise, the reference is ambiguous and in error.

Examples:

1. DECLARE 1A, 2C, 2D, 3E;
BEGIN;
 DECLARE 1A, 2B, 3C, 3E;
A.C refers to C in the inner block.
D.E refers to E in the outer block.
2. DECLARE 1A, 2B, 2B, 2C, 3D, 2D;
B has been multiply declared.
A.D refers to the second D, since A.D is a complete qualification of only the second D; the first D would have to be referred to as A.C.D.
3. DECLARE 1A, 2B, 3C, 2D, 3C;
A.C is ambiguous because neither C is completely qualified by this reference.
4. DECLARE 1A, 2A, 3A;
 A refers to the first A.
 A.A refers to the second A.
 A.A.A refers to the third A.
5. DECLARE X; DECLARE 1Y, 2X, 3Z, 3A, 2Y, 3Z, 3A;
 X refers to the first DECLARE
 Y.Z is ambiguous
 Y.Y.Z refers to the second Z
 Y.X.Z refers to the first Z

CONTEXTUAL DECLARATIONS

The syntax of PL/I allows identifiers appearing in certain contexts to be recognized without an explicit declaration. The various cases are described below.

1. An identifier may occur in a context where only a file name may appear. In some of these cases, the identifier is said to be declared as a file name, with the default attribute EXTERNAL (see "Application of Default Attributes" in this chapter).

Example:

```
READ FILE (INFILE) DATA;
```

Here, INFILE is declared contextually with the attribute FILE.

2. An identifier may occur in a context where only a task (or event) name (see "The CALL Statement" in Chapter 8 and "Asynchronous Operations and Tasks" in Chapter 6) may appear. In some of these cases, the identifier is said to be declared as a task (or event) name (see "Application of Default Attributes").

Example:

```
WAIT (EVENT2);
```

Here, EVENT2 is declared contextually as an event identifier.

3. An identifier may occur in a context where only a programmer-specified condition name (see Appendix 3) may appear. In this case, the identifier is said to be declared as a condition name, with the attribute EXTERNAL.

Example:

```
ON CONDITION (TEST1) GO TO CHECK;
```

Here, TEST1 is declared contextually as a condition name.

4. An identifier may appear within a statement in a context where only an entry name may appear. That is, an identifier is contextually declared as an entry name if it appears as a label to a PROCEDURE or ENTRY statement or if it appears following the keyword CALL or as the function name in a function reference whose argument list is non-empty. If the occurrence of the identifier does not lie within the scope of the same identifier used to label a PROCEDURE or ENTRY statement, the identifier is given a default attribute of EXTERNAL.

Example:

```
CALL EXPRI;
```

5. An identifier may appear as a label of a statement, i.e., as a statement label or an entry name. (A statement label variable must be explicitly declared with the attribute LABEL.)

In this case, the label or name is said to be declared in the block to which it is internal (for the definition of internal to, see "Blocks"). This implies that every statement label except the label of a BEGIN statement is declared in the block to which its associated statement is internal. It further implies that a label appearing before a BEGIN, ENTRY, or PROCEDURE statement is declared in the immediately containing block.

In the special case where the label is an entry name for an external procedure, the name is said to be declared externally, and has the EXTERNAL attribute (see "Scope of Declarations").

Example:

```
A: PROCEDURE;
    .
    .
    .
P: PROCEDURE;
    .
    .
    .
    LOOP:DOI=1 TO N;
        .
        .
        .
    Q: BEGIN;
        LOOP:DO J=0 TO I;
            .
            .
            .
        END LOOP;
    END Q;
    END LOOP;
END P;
END A;
```

In this example:

A is declared as an external entry name.

P is declared as an entry name in block A.

LOOP in its first use is declared as a statement label in block P. Q is declared as a statement label in block P.

LOOP in its second use is declared as a statement label in block Q.

6. An identifier may appear in a formal parameter list in a PROCEDURE or ENTRY statement. In this case, the identifier is said to be declared in the block to which the list is internal. Attributes may be explicitly declared for the identifier in a DECLARE statement internal to the same block, in which case both the contextual and explicit declarations are regarded as constituting a single declaration.

Example:

```
PAY: PROCEDURE (HOURS, RATE);
      DECLARE HOURS FIXED (6,2);
      .
      .
      .
      END PAY;
```

In this example, HOURS is declared explicitly and RATE contextually in the block PAY.

7. An identifier may appear in list for data-directed transmission (see Chapter 7). In this case it is given the attribute SYMBOL.

Example:

```
READ DATA (A,B);
```

Note: Arithmetic or string attributes of constants are determined contextually.

IMPLICIT DECLARATIONS

An identifier may be used in a block without being explicitly declared or contextually declared. In this case the identifier is said to be implicitly declared in the containing external procedure. As will be seen in the discussion of scope, this implicit declaration will then apply to the entire external procedure block except for any contained blocks where the identifier might be re-declared.

Example:

```
B1: PROCEDURE (Z1,Z2);
      TEMP1=ABS (Z1**2+Z2**2);
      B2: BEGIN;
            TEMP2= 1/(TEMP1+Z2)**2;
            IF TEMP2>TEMP1 THEN RETURN
              (TEMP2);
            END B2;
      RETURN (TEMP1);
      END B1;
```

In this example, TEMP1 and TEMP2 are both implicitly declared in block B1.

SCOPE OF DECLARATIONS

When a declaration of an identifier is made in a program, there is a certain well-defined region of the program over which this declaration is applicable. This region is called the scope of the declaration or the scope of the name established by the declaration.

The scope of a declaration of an identifier is defined as that block B to which the declaration is internal, but excluding from block B all contained blocks to which another declaration of the same identifier is internal.

This definition of scope can be applied to all identifier declarations except the declaration of entry names of external procedures (see "Declarations," in this chapter). The appearance of an identifier as the entry name of an external procedure is regarded as an explicit declaration of the identifier as an entry name with the EXTERNAL attribute. The scope of such a declaration is defined to be the entire external procedure, excluding all contained blocks to which another declaration of the same identifier is internal.

Scope of External Names

In general, distinct declarations of the same identifier imply distinct names with distinct non-overlapping scopes. It is possible, however, to establish the same name for distinct declarations of the same identifier by means of the EXTERNAL attribute. The EXTERNAL attribute is defined as follows:

An explicit or contextual declaration of an identifier that declares the identifier as EXTERNAL is called an external declaration for the identifier. All external declarations for the same identifier in a program will be linked and considered as establishing the same name. The scope of this name will be the union of the scopes of all the external declarations for this identifier.

In all of the external declarations for the same identifier, the attributes declared must be consistent, since the declarations all involve a single name. For example, it would be an error if the identifier ID were used as a file name in some READ statement in a program, and in the same program to declare ID as EXTERNAL ENTRY, since a file name always has the scope attribute EXTERNAL (see "Default Attributes") and the attribute FILE, which conflicts with the attribute ENTRY.

The `EXTERNAL` attribute can be used to communicate between different external procedures or to obtain non-continuous scopes for a name within an external procedure.

An external name is a name that has the scope attribute `EXTERNAL`. If a name is not external, it is said to be an internal name and has the scope attribute `INTERNAL`.

Example 1:

```

1  A:  PROCEDURE;
2      DECLARE (X,Z) FLOAT;
      .
      .
3  B:  PROCEDURE (Y) ;
4      DECLARE Y BIT (6) ;
5      C:  BEGIN;
6          DECLARE (A,X) FIXED;
          .
          .

```

```

7          Y:  RETURN;
          END C;
      END B;
8  D:  PROCEDURE;
9      DECLARE X FILE;
10     Y = Z;
      .
      .
      .
      END D;
END A;

```

The numbers on the left are for reference only, and are not part of the procedure. See Table 2 for an explanation of the scope and use of each name.

Since entry names of external procedures and file names have the attribute `EXTERNAL`, the scope of the entry name A and of the file name X above may include parts of other external procedures of the program.

Table 2. Scope and Use of Names in Example 1, for "Scope of External Names"

<u>Reference Line</u>	<u>Name</u>	<u>Use</u>	<u>Scope (by block names)</u>
1	A	external entry name	all of A except C
2	X	floating-point variable	all of A except C and D
2	Z	floating-point variable	all of A
3	B	internal entry name	all of A
4	Y	bit string	all of B except C
5	C	statement label	all of B
6	A	fixed-point variable	all of C
6	X	fixed-point variable	all of C
7	Y	statement label	all of C
8	D	internal entry name	all of A
9	X	file name	all of D
10	Y	floating-point variable	all of A except B

Example 2:

```

A: PROCEDURE;
1  DECLARE X EXTERNAL;
   .
   .
   .
2  B: PROCEDURE;
   DECLARE X FIXED;
   .
   .
   .
3  C: BEGIN;
   DECLARE X EXTERNAL;
   .
   .
   .
   END C;
   END B;
   END A;
D: PROCEDURE;
4  DECLARE X FIXED;
   .
   .
   .
5  E: PROCEDURE;
   DECLARE X EXTERNAL;
   .
   .
   .
   END E;
   END D;

```

The reference numbers on the left are not part of the procedure.

In example 2, there are five declarations for the identifier X.

Declaration 2 declares X as a fixed-point variable name; its scope is all of block B except block C.

Declaration 4 declares X as another fixed-point variable name, distinct from that of declaration 2; its scope is all of block D except block E.

Declarations 1,3,5 all establish X as a single name; its scope is all of the program except the scopes of declarations 2 and 4.

Basic Rule on Use of Names

A name is said to be known only within its scope. This definition suggests a basic -- and almost self-evident -- rule on the use of names:

All appearances of an identifier which are intended to represent a given name in a program must lie within the scope of that name.

There are many implications to the above rule. One of the most important is the limitation of transfer of control by the statement GO TO A, where A is a statement label.

The statement GO TO A, internal to a block B, can cause a transfer of control to another statement internal to block B or to a statement in a block containing B, and to no other statement. In particular, it cannot transfer control to any point within a block contained in B.

THE ATTRIBUTES

Attributes are used to give characteristics to their associated identifiers. The attributes of the language are divided into the following classes:

- Data attributes
- Dimension attribute
- SECONDARY attribute
- ABNORMAL/NORMAL attributes
- USES and SETS attributes
- Entry name attributes
- Scope attributes
- Storage Class attributes
- ALIGNED and PACKED attributes
- DEFINED attribute
- INITIAL attribute
- Symbol table attributes
- Structure attributes
- LIKE attribute
- File description attributes

DATA ATTRIBUTES

Arithmetic Data

Variables are declared to be of arithmetic type if they are given any of the attributes base, scale, mode, or numeric picture.

Base

The base attribute specifies that the data is in binary or decimal form.

General format:

BINARY|DECIMAL

Rules:

This attribute may not be specified in combination with the PICTURE attribute.

Default:

See "Default Conditions for Arithmetic Data."

Examples:

```
DECLARE A DECIMAL, B BINARY;
```

Scale

Function:

The scale attribute specifies that the data is in fixed-point or floating-point form.

General format:

```
FIXED|FLOAT
```

Rules:

This attribute may not be given in combination with the PICTURE attribute.

Default:

See "Default Conditions for Arithmetic Data."

Examples:

```
DECLARE A FIXED, B FLOAT;
```

Mode

Function:

The mode attribute specifies that the mode of the data is real or complex.

General format:

```
REAL|COMPLEX
```

Rules:

This attribute may be given in combination with the PICTURE attribute, to specify a complex numeric field.

Default:

See "Default Conditions for Arithmetic Data."

Example:

```
DECLARE A COMPLEX, B REAL;
```

Precision

Function:

The precision attribute specifies the number of significant binary or decimal digits to be maintained for both fixed-point and floating-point data, as well as the scale of the data.

General format:

```
(number-of-digits [,scale-factor])
```

Rules:

1. The precision attribute must immediately follow a scale, base, or mode attribute and may never appear alone or separated from one of these attributes.
2. "Number-of-digits" is a decimal integer constant specifying the number of binary or decimal digits to be maintained and is used with both fixed-point and floating-point data.
3. The "scale-factor" is an optionally signed decimal integer constant that defines the position of the point with respect to an integer data item of the specified number of digits. The scale factor is used only with fixed-point data.
4. When the scale is fixed and no scale factor is given, it is assumed to be zero.
5. The scale factor may be negative, and it may be larger than the number of digits.
6. The scale factor effectively multiplies the integer data by the base raised to the power of the scale factor with the sign reversed. For example, decimal data of precision (5,2) represents numbers from .01 to 999.99 or zero in magnitude; decimal data of precision (5,-2) represents numbers from 100 to 9999900 or zero in magnitude.
7. This attribute may not be given in combination with the PICTURE attribute.

Examples:

```
DECLARE A FLOAT (3), B REAL (10)
        FLOAT, X FIXED (5,2);
```

The following table shows the meaning of the scaling for fixed-point variables:

Integer	Scale	Precision	Value
00123	FIXED	(5,2)	1.23
00123	FIXED	(5,-2)	12300
123	FIXED	(3,4)	.0123
123	FIXED	(3,-4)	1230000

Default Conditions for Arithmetic Data

If the base, scale, and mode are not specified, the arithmetic default attributes are dependent upon the first letter of the name. If the first letter of the name is I through N, FIXED REAL BINARY is assumed; otherwise, FLOAT REAL DECIMAL is assumed.

If arithmetic data attributes are partly specified, the remaining attributes are assumed as follows:

```
Base: DECIMAL
Scale: FLOAT
Mode: REAL
```

If precision is not specified, the assumed precision is that which is defined for the particular implementation of the language that is being used, where the definition depends on the scale and base.

The PICTURE Attribute

Function:

The PICTURE attribute is used to define the internal and external formats of numeric and character-string data fields and to specify the editing of data. This discussion is limited to the use of the PICTURE attribute with numeric data. The use of the PICTURE attribute with character-string data is described in "String Attributes." The picture characters are described in Appendix 2.

General format:

```
PICTURE 'numeric-picture-specifications'
```

General rules:

1. PICTURE may not be specified in combination with the base, scale, or precision attributes.

Numeric fields have mode, base, scale, and precision; these are specified by the picture characters used in describing the field, and by the use of the mode attribute if COMPLEX. Note the exception that sterling pictures are treated as a separate category, although they are real fixed-point decimal fields.

2. A "picture specification" is composed of a string of picture characters. It must be enclosed in quotation marks. Individual picture characters may be preceded by an iteration factor, which

is a decimal integer constant, *n*, enclosed in parentheses, to indicate repetition of the character *n* times. If *n* is zero, the character is omitted. This iteration factor specification may not follow the picture character F.

3. Numeric picture specifications must include at least one digit position.
4. The following paragraphs indicate the combination of picture characters that show mode, scale, base, and precision. In this discussion, a fixed-point field has one field, and a floating-point field has two subfields.
 - a. Real binary fixed-point fields take the following general forms:

```
PICTURE '[S|1] ... [V]
           [S|1] ... [F([+|-] integer)]'
PICTURE '[2] ... [V] [2] ... [F([+|-]
           integer)]'
PICTURE '[3] ... [V] [3] ... [F([+|-]
           integer)]'
```

Only one V, representing a point, may be present in a picture specification, but it may be in any position. When a sign character (S) is specified, the field will contain a binary 1, if the value is negative, or a zero, if the value is positive, for each S in the picture.

- b. Real binary floating-point fields take the following general forms:

```
PICTURE '[S|1] ... [V] [S|1] ...
           K[S|1] ...'
PICTURE '[2] ... [V] [2] ...
           K2 ...'
PICTURE '[3] ... [V] [3] ...
           K3 ...'
```

Each sign character allowed to the right of the V in the first form represents the sign of the exponent.

- c. Real decimal fixed-point fields take the following general form:

```
PICTURE '[9] ... [V] [9] ...
           [F([+|-] integer)]'
```

Sign, editing, and zero-suppression picture characters, as explained in Appendix 2, may be included. The V may not appear more than once in a picture specification. If no V is given, the decimal point will be assumed to appear to the right of the last digit. No attempt has been made to show the use of all valid picture characters in the general format above. These are explained in Appendix 2.

- d. Real decimal floating-point fields take the following general form:

```
PICTURE '[9]... [V] [9]...{E|K}
[9]...'
```

Sign, editing, and zero-suppression picture characters may be included. Sign characters refer to the subfield in which they appear, except a CR or a DB, which refers to the first subfield.

- e. Complex fields may contain those picture characters that are valid for real fields as described above. They take the general form:

real-picture

The "real-picture" represents both portions of the complex number. The attribute COMPLEX must also be specified. The real-picture may not specify a sterling field.

- f. Sterling fields are considered to be real fixed-point decimal fields. When involved in arithmetic operations, they will be converted to a value representing fixed-point pence. Sterling pictures have the general form:

```
PICTURE 'G editing-character-1
pounds-field separator-1
shillings-field separator-2
pence-field'
```

"Editing character 1" may be one or more of the following picture characters:

\$ + - S

The "pounds field" may contain the following picture characters:

Z Y * 9 T I R \$ + - S

"Separator 1" may be one or more of the following picture characters:

/ . B V

The "shillings field" may be:

```
{99|ZZ|Y9|Z9|ZY|8}
```

The 9s may be replaced by T, I, or R.

The picture character Z may occur only if the whole of the field to the left of this character (including the pounds field) is

also suppressed using the character Z.

"Separator 2" may be one or more of the picture characters:

/ . B V H

The "pence field" takes the form:

```
{99|ZZ|Y9|7|Z9|ZY|6} [V|V.|.V]
[9|Z|Y]... [B]... [P] [B]... [CR|DB
|S|-|+] [B]...
```

Any of the nines may be replaced by one of the following:

T I R

In a sterling picture, there can not be more than one of the following characters:

T I R CR DB S + -

Zero suppression characters can only appear after the decimal point in the pence field if all digits are suppressed in the field.

5. The precision of picture specifications is described below. In this discussion, the following picture characters, actual and conditional, are defined as digit positions:

```
1 2 3 9 Z * Y T I R
and the drifting
$ S + -
```

The precision of a fixed-point numeric field is (m,n), where m is the total number of digit positions in the field and n is the number of digit positions following the V. If a drifting string contains n drifting characters, this specifies n-1 digit positions. For sterling pictures, m is 3 + the number of digits in the pounds field + the number of fractional digits in the pence field.

The precision of a floating-point field is (p), where p is the total number of digit positions before the E or K.

Decimal or binary fixed-point pictures may have a scaling factor. This may be achieved by placing the following at the extreme right of the picture subfield:

F ([+|-] integer)

with the "integer" value represented by q, this specifies that the decimal or binary point should be assumed to

be g places to the right (or left, if negative) of the position assumed in the absence of the scaling factor. The precision of the numeric field is then (m,n-g).

These precisions may not exceed the limits for decimal fixed-point values, as defined for the particular implementation of PL/I.

String Attributes

Function:

The string attributes specify string data to be either in bit-string form or in character-string form with a specified length. The form of character-string data may also be specified.

General format:

$$\left\{ \begin{array}{l} \left\{ \begin{array}{l} \text{BIT} \\ \text{CHARACTER} \end{array} \right\} (\text{length}) \text{ [VARYING]} \\ \text{PICTURE 'character-picture-} \\ \text{specifications' } \end{array} \right\}$$

Rules:

1. BIT specifies bit-string data, CHARACTER specifies character-string data, and PICTURE specifies character-string data in picture form.
2. The "length" specifies the actual length of fixed-length strings and the maximum length of varying-length strings, in which case the word VARYING is used. If VARYING is specified, then either BIT or CHARACTER must also be specified.
3. The length specification may be an expression or an asterisk.
4. If the length specification is an expression, it will be converted to an integer at the point of allocation or upon entry to the declaring block for parameters.
5. An asterisk may be used when the length is to be taken from a previous allocation for parameters or controlled variables or if it is to be specified in a subsequent ALLOCATE statement for CONTROLLED variables.
6. The length of strings declared STATIC must be a decimal integer constant.
7. Since PICTURE is an attribute that also may apply to arithmetic data, a separate explanation is in the section entitled "The PICTURE Attribute." Additional picture characters are provided when the PICTURE attribute is used to declare character-string data.

- These may be found in Appendix 2.
8. BIT, CHARACTER, or VARYING may not be specified if PICTURE is specified.

Example:

```
DECLARE A BIT (10), B CHARACTER (5), C
        PICTURE 'XAA9AA', D BIT(*) VARYING;
```

A is a field of ten bits; B is a field of five characters; C is a field of characters, letters, and a decimal digit; and D is a field of bits with a maximum length to be taken from a previous allocation or to be specified in a subsequent ALLOCATE statement.

The LABEL Attribute

Function:

The LABEL attribute specifies that the associated variable will have statement labels as values. To aid optimization of the object program, it may also specify the values a label variable may have during execution of the program.

General format:

```
LABEL [(statement-label-constant
        [, statement-label-constant]...)]
```

Rules:

1. If no statement-label constants are specified following the LABEL attribute, the value of the variable may be any of the statement labels known in the scope of the variable.
2. If the variable is a parameter, the value can also be any statement label that could be passed as an argument, or any value permitted for any label variable that may be specified as an argument.
3. If a list of statement-label constants is specified, the variable may have as values only members of the list. The label constants in the list must be known in the block containing the declaration.
4. An entry name cannot be a value of a label variable.

Example:

```
DECLARE START LABEL (LABEL1, LABEL2,
                    LABEL3);
```

The TASK Attribute

Function:

The TASK attribute specifies that the associated identifier is used as a task name (see "Asynchronous Operations and Tasks," in Chapter 6, the general rules under "The CALL Statement," in Chapter 8, and "Task Names," in Chapter 2).

General format:

TASK

Rules:

1. An identifier may be explicitly declared with the TASK attribute in a DECLARE statement. It may be contextually declared by its appearance in a TASK option appended to a CALL statement (see Chapter 8).
2. Task names may also have the following attributes:

Dimension attribute
Scope attribute (the default is INTERNAL)
Storage class attribute (the default is AUTOMATIC)
DEFINED attribute (task names may only be defined on other task names)
ABNORMAL attribute (all task names are ABNORMAL)
SECONDARY attribute

3. A task name can appear only in a TASK option (see "The CALL Statement," in Chapter 8) or as the argument in the PRIORITY built-in function or in the PRIORITY pseudo-variable.

The EVENT Attribute

Function:

The EVENT attribute specifies that the associated identifier is used as an event name (see "Asynchronous Operations and Tasks," in Chapter 6, the general rules under "The CALL Statement," in Chapter 8, and "Event Names," in Chapter 2).

General format:

EVENT

Rules:

1. An identifier may be explicitly declared with the EVENT attribute in a DECLARE statement. It may be contextually

declared by its appearance in an EVENT option appended to a CALL statement (see Chapter 8), or by its appearance in a WAIT statement (see Chapter 8).

2. Event names may also have the following attributes:

Dimension attribute
Scope attribute (the default is INTERNAL)
Storage class attribute (the default is AUTOMATIC)
DEFINED attribute (event names may only be defined on other event names)
ABNORMAL attribute (all event names are ABNORMAL)
SECONDARY attribute

3. An event name can appear only in an EVENT option (see "The CALL Statement," in Chapter 8), a WAIT statement (see Chapter 8), or as the argument in the EVENT built-in function or in the EVENT pseudo-variable.

THE DIMENSION ATTRIBUTE

Function:

The dimension attribute defines the bounds of an array.

General format:

(bound [, bound] ...)

where "bound" is
{[lower-bound :] upper-bound} | *

Rules:

1. The number of "bounds" specifies the number of dimensions in an array.
2. Bounds that are expressions are evaluated and converted to integer data when storage is allocated for the array or when linkage is established for parameters.
3. The bounds are indicated as follows:
 - a. If only the upper bound is given, the lower bound is assumed to be one.
 - b. When the actual bounds for each dimension are to be taken from a previous allocation for that identifier or are to be specified in a subsequent ALLOCATE statement, an asterisk must be used to represent each of the dimension bounds. Thus, asterisks may be used only for parameters and CONTROLLED variables.
 - c. The lower bound must be less than

- or equal to the upper bound.
- The bounds of arrays declared static must be optionally signed decimal integer constants.

Examples:

- DECLARE TABLEA (5,8) , TABLEB (-5:5,10) ;

TABLEA is a two-dimensional array with 5 rows and 8 columns (subscripts 1 to 5 and 1 to 8). TABLEB is a two-dimensional array with 11 rows and 10 columns (subscripts -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5 for the rows and 1 through 10 for the columns).
- DECLARE MATRIX (*,*) ;

MATRIX is a two-dimensional array. The bounds are to be taken from a previous allocation for MATRIX or are to be subsequently specified in an ALLOCATE statement.

THE SECONDARY ATTRIBUTE

Function:

The SECONDARY attribute is used to specify that certain data normally does not require efficient storage.

General format:

SECONDARY

Rules:

- This attribute may be declared only for major structures, arrays, and variables not contained in structures or arrays, i.e., for variables at level 1.
- The attribute specifies that where possible and necessary, less than normally efficient storage may be allocated to the variable.

THE ABNORMAL AND NORMAL ATTRIBUTES

Function:

The ABNORMAL and NORMAL attributes are used to specify procedures and/or data as being either normal or abnormal.

General format:

ABNORMAL|NORMAL

Rules for abnormality of procedures:

- Abnormality is a property of both external and internal procedures. Blocks invoking procedures that are abnormal must be within the scope of an ABNORMAL, USES, or SETS declaration for the invoked entry name. However, the invocation of an abnormal procedure does not make the invoking procedure itself abnormal. These attributes enable program optimization to be performed.
- An external procedure is abnormal if it or any procedures invoked by it:
 - Access, modify, allocate or free external data.
 - Modify, allocate, or free their arguments.
 - Return inconsistent function values for the same argument values.
 - Maintain any kind of history.
 - Perform input/output operations.
 - Return control from the procedure by means of a GO TO statement.
- An internal procedure is abnormal:
 - Under any of the conditions listed above for external procedures.
 - If it, or any procedures called by it, access, modify, allocate, or free variables declared in an outer block.
- Abnormal external procedures invoked as functions must be declared with at least one of the attributes, ABNORMAL, USES, or SETS. The scope of this declaration must include the invoking block.
- ABNORMAL used alone specifies that all possible types of abnormality should be assumed. It is unnecessary to specify ABNORMAL for the built-in functions, TIME and DATE.
- The NORMAL attribute specifies that the entry name is for a procedure that is not abnormal.

Rules for abnormal data:

- The ABNORMAL attribute may be declared for any variable.
- The ABNORMAL attribute specifies that a variable may be altered or otherwise accessed at an unpredictable time during the execution of a program. This situation might occur, for example, during the execution of an ON-unit as described in "The ON Statement," in Chapter 8.
- Every time ABNORMAL data is referred to, its associated storage contains its current value.

Default for abnormality of procedures:

If an external entry name appears only as a function reference, the entry name is

assumed to have the NORMAL attribute; otherwise, the entry name is assumed to be ABNORMAL. Entry names of all internal procedures and entry names of external procedures invoked in CALL statements are assumed to have the ABNORMAL attribute.

Default for abnormality of data:

Variables are assumed to be NORMAL, except structures containing ABNORMAL elements; such structures may not be declared NORMAL.

THE USES AND SETS ATTRIBUTES

Function:

The USES and SETS attributes are used to specify, for an entry name, the nature of an abnormality due to data manipulation.

General format:

```
USES (item[,item] ...)  
SETS (item[,item] ...)
```

Rules:

1. The items of the list following a USES or SETS attribute may be as follows:
 - a. A decimal integer *n*, specifying the *n*th argument of any invocation of the procedure at the declared entry name.
 - b. An unsubscripted data name known to both the block containing the declaration and the invoked procedure.
 - c. An asterisk indicating all identifiers described in b.
2. An item in the USES list specifies the following:
 - a. That the invoked procedure or procedures invoked by it access that item.
 - b. That neither the invoked procedure nor procedures invoked by it reassign that item unless it is also specified in a SETS attribute.
 - c. That neither the invoked procedure nor procedures invoked by it access any other data known to the block, except data designated by explicit arguments in either a CALL statement, a statement with a CALL option, or a function reference.
3. An item in the SETS list specifies the following:
 - a. That the invoked procedure or procedures invoked by it reassign, allocate, or free that item.
 - b. That neither the invoked procedure nor procedures invoked by it

access that item other than to reassign, allocate, or free it, unless it is also specified in a USES attribute.

- c. That neither the invoked procedure nor procedures invoked by it reassign, allocate, or free any other data known for the block, except data designated by explicit arguments in the case of a CALL statement.
4. The USES and SETS attributes may be declared for any entry name used to invoke a procedure. The scope of this declaration must include the invoking block. If the ENTRY attribute is not declared, ENTRY is implied. If either USES or SETS is declared in the invoking procedure, complete information must be given about the data that is used and/or set by the invoked procedure.
5. If an item in a USES or SETS list, as described in 1b above is defined on a base (see "The DEFINED Attribute") and if the base and any other items defined on it are known both to the invoking and invoked blocks, the base and the other items must also be specified in the list.
6. A structure or array name appearing in a USES or SETS list implies that the names of all items contained in the structure or array also are on the list. It does not imply that items defined on elements of the structure are in the list; these must be declared as in rule 5, above.
7. If the USES or SETS attribute is specified and the invoked procedure is abnormal in any other way, the ABNORMAL attribute must still be specified (unless it is given by default). If the USES or SETS attribute is specified and the invoked procedure is not otherwise abnormal, the ABNORMAL attribute should not be specified.

ENTRY NAME ATTRIBUTES

An identifier may be declared to be an entry name by giving it the ENTRY attribute. It may be declared to have any of the attributes SETS, USES, and BUILTIN. These attributes all imply ENTRY and thus ENTRY need not be specified. The entry name also may have the attributes ABNORMAL or NORMAL and, with the exception of LABEL, SECONDARY, TASK, EVENT, or a dimension attribute, any of the data attributes listed in the beginning of this section.

The data attributes specify the characteristics of the value returned when the entry name is invoked as a function. If

data attributes are not specified, default or implicit characteristics will be assumed (see "Assignment of Attributes to Identifiers" in this chapter).

An explicit declaration of an internal entry name and the procedure block having the entry name must both be internal to the same block.

The ENTRY Attribute

Function:

The ENTRY attribute is used to declare, within a procedure, entry names that are referred to in that procedure.

General format:

```
ENTRY [ (parameter-attribute-list
        [,parameter-attribute-list] ...)]
```

Rules:

1. When ENTRY is used, it specifies that the identifier being declared is an entry name. An entry name must be declared with the ENTRY attribute unless the entry label is known in the same block, or unless a reference is made to the entry name in a CALL statement or in a function reference with arguments, or if it is declared to have any of the attributes SETS, USES, GENERIC, and BUILTIN. INTERNAL entries may only be declared in the block to which the procedure is internal. ENTRY without a parameter attribute list specifies nothing about the number or nature of the parameters.
2. When ENTRY is used with parameter attribute lists, each parameter attribute list is a succession of attributes describing the parameters of the entry point. Permitted attributes are those allowed for parameters.
3. The number of parameter attribute lists must be the same as the number of parameters required by the entry point. If a parameter attribute list is null, its place must be kept by a comma.
4. Parameter attribute lists are not necessary if the parameters of the entry name are not to be described.
5. The dimension attribute may be specified for array parameters. It must be the first attribute specified for the parameter.
6. The structuring for a structure parameter is specified by a structure description using level numbers without identifiers, the level number being immediately followed by the list of

attributes for that level of the structure. The first item in the description of the structure parameter must be at level one.

7. Expressions occurring in ENTRY attributes for length or dimension bounds are evaluated upon entering the block to which the declaration of the ENTRY attribute is internal. If an argument position specifies an entry with no data attributes, no default data attributes are provided.

Default:

If no attributes or level numbers are given for a parameter, no assumptions are made about it. When any attributes are specified, the remaining required attributes are deduced according to the default rules given in "Assignment of Attributes to Identifiers."

The GENERIC Attribute

Function:

The GENERIC attribute is used to define a name as a family of entry names, each of which is referred to by the name being declared. When the generic name is referred to, the proper entry name is selected, based upon the arguments specified for the generic name in the procedure reference.

General format:

```
GENERIC (entry-name-declaration
        [,entry-name-declaration] ...)
```

Rules:

1. No other attributes may be specified for the name being given the GENERIC attribute.
2. Each "entry name declaration" following the GENERIC attribute corresponds to one member of the family.
3. Each entry name declaration must have the ENTRY attribute. It may optionally have ABNORMAL, NORMAL, USES, SETS, BUILTIN, and data attributes. No entry name declaration may have the GENERIC attribute.
4. Each entry name declaration must specify attributes or level numbers for every parameter of the associated entry name. Attributes unspecified but required for full definition will be deduced from default rules.
5. When a generic name is referred to, the attributes of the arguments must match exactly the list following the entry name declaration of one and only

one member of the family. The reference is then interpreted as a reference to that member. Thus, the selection of a particular entry name is based upon the arguments of the reference to the generic name.

6. The selection of a particular entry name is first based on the number of arguments in the reference to the name. The following attributes are then considered in choice of generic members:

Base
 Scale
 Mode
 Precision
 PICTURE
 LABEL (but not range list)
 Dimensionality (but not bounds)
 CHARACTER (but not length)
 BIT (but not length)
 VARYING
 TASK
 EVENT
 ENTRY (but not parameter description or other attributes of entry names other than data attributes of the value returned by a function)
 FILE (but no other FILE attributes) structuring, including only the attributes listed above for the structure members.

7. Generic entry names (as opposed to references) may be specified as arguments to non-generic procedures if the invoked entry name is declared with the ENTRY attribute (explicit or implicit for internal procedures). This ENTRY attribute must specify that the appropriate parameter is an entry name and specify by means of a further ENTRY attribute the attributes of all its parameters. This enables a choice to be made of which family member is to be passed.

Example:

```
DECLARE
  CALCULATE GENERIC (FIXCALC ENTRY (FIXED),
    FLTCALC ENTRY (FLOAT)), Y FLOAT
  INITIAL (50);
X=Y + CALCULATE (Y);
```

The assignment statement results in the invocation of the procedure FLTCALC, since the argument Y matches the entry attribute of the FLTCALC member of the family.

The BUILTIN Attribute

Function:

The BUILTIN attribute specifies that the reference to the associated identifier within the scope of the declaration is interpreted as a reference to the built-in function or pseudo-variable of the same name.

General format:

BUILTIN

Rules:

1. BUILTIN is used to refer to a built-in function or pseudo-variable in a block that is contained in another block in which this name has been declared to have another use.
2. If the BUILTIN attribute is declared for an entry name, it may have no other attributes.
3. For a list of built-in functions see Appendix 1.

SCOPE ATTRIBUTES

Function:

The scope attributes are used to specify the scopes in which declared identifiers are known.

General format:

```
{ INTERNAL
  { EXTERNAL [(identifier)] }
```

Rules:

1. For a full discussion of the INTERNAL and EXTERNAL attributes, see "Scope of Declarations".
2. In the form EXTERNAL (identifier), the identifier specifies a heading for the declared name. The scope of the name is then the union of the scopes of all EXTERNAL declarations of the same name with the same heading.

Default:

If the scope is unspecified for variable names, INTERNAL is assumed.

Example:

```
DECLARE SUM EXTERNAL (X);
```

The variable SUM is declared external with the heading X. In other declarations,

the heading distinguishes this variable from other variables named SUM with no heading or other headings.

STORAGE CLASS ATTRIBUTES

Function:

Storage class attributes are used to allocate a particular class of storage to variables.

General format:

STATIC|AUTOMATIC|CONTROLLED

Rules:

1. STATIC specifies that storage is allocated at the start of execution of the program and is not released until program execution has been completed.
2. AUTOMATIC specifies that storage is allocated on each entry to the block to which the storage declaration is internal. The storage is released on leaving the block. If the block is a procedure that is invoked recursively, the previously allocated storage is "pushed down" on entry, and the latest allocation of storage is "popped up" on termination. (For a discussion of "pushed down" and "popped up" storage, see "Allocation of Data and Storage Classes" in Chapter 6.)
3. CONTROLLED specifies that full control will be maintained over the allocation and freeing of storage by means of the statements ALLOCATE and FREE.
4. AUTOMATIC variables may have INTERNAL scope only. STATIC and CONTROLLED variables may have INTERNAL or EXTERNAL scope.
5. Storage class attributes may not be specified for entry names, file names, members of structures, or DEFINED data.
6. STATIC and AUTOMATIC attributes may not be specified for parameters.
7. Variables declared with adjustable lengths and dimensions may not have the STATIC attribute.
8. If a procedure involving static storage is invoked from within or as a separate task, the static storage is common to all invocations.
9. If, during execution of a statement, controlled data is allocated or freed (by an abnormal function, for example), any reference in the statement to that data produces an undefined result.
10. Storage class attributes may only be given for variables at level 1. The storage class applies to all elements

of a structure or array of structures. If a structure is CONTROLLED, only the major structure, and not the elements, may be allocated and freed.

Default:

1. If storage class is unspecified and the scope is EXTERNAL, STATIC is assumed.
2. If storage class is unspecified and the scope is INTERNAL, AUTOMATIC is assumed.
3. If neither storage class nor scope is specified, AUTOMATIC is assumed.

Example:

```
EXAMPLE: PROCEDURE;  
        DECLARE A STATIC INITIAL  
            (0), B CONTROLLED, C (10);  
        ALLOCATE B;  
        A = A + 1;  
        .  
        .  
        .  
        FREE B;  
        WRITE LIST (A);  
        END EXAMPLE;
```

The variable A is of the static storage class and is used to count the number of times the procedure is invoked. The variable B is of the controlled storage class, and storage is allocated and freed by use of the ALLOCATE and FREE statements. The variable C is of the automatic storage class by default.

THE ALIGNED AND PACKED ATTRIBUTES

Function:

The ALIGNED and PACKED attributes are used to specify in storage the arrangement of string or numeric field data elements within data aggregates.

General format:

ALIGNED|PACKED

Rules:

1. These attributes may be specified for the following:
 - a. Names of major structures.
 - b. Names of arrays that are not themselves part of a structure.
2. PACKED specifies that each string or numeric field element is packed in storage contiguous with the string or numeric field elements that surround it. There should be no unused storage between two adjacent elements, provid-

ed all data elements of the aggregates are string or numeric field variables of the same class. In other cases, some unused space may appear but storage is to be conserved when possible.

3. ALIGNED specifies that each string data element within the aggregate may start at a storage boundary to be defined individually for each implementation of PL/I. This implies that two adjacent string or numerical field elements of a homogeneous aggregate may not necessarily occupy contiguous storage, if a more efficient program is possible.
4. Arguments to the STRING generic function must be PACKED structures.

Default:

1. The default for major structures is PACKED.
2. The default for arrays that are not part of structures is ALIGNED.

Examples:

```
DECLARE
  1 A (10) PACKED, 2 B BIT
    (200), 2 C BIT (500), 2 D BIT
    (300), E (10,15) ALIGNED BIT (15);
```

All elements of A, an array of structures, will occupy a continuous area of storage. Each element of the array E will start at a storage boundary defined for that implementation of PL/I. There may be unused storage between the elements of the latter array.

THE DEFINED ATTRIBUTE

Function:

The DEFINED attribute specifies that scalar, array, or structure data is to occupy the same storage as that already assigned to other data.

General format:

```
DEFINED base-identifier
      [subscript list]
```

Rules for defining:

1. In general, the defined item must have the same characteristics as the "base identifier." However, mixed defining is permitted in the following two classes:
 - a. The bit class is composed of the following:
 - (1) numeric fields of binary base
 - (2) fixed-length bit strings

- (3) packed arrays or structures of either or both (1) and (2)
 - b. The character class is composed of the following:

- (1) numeric fields of decimal base
 - (2) fixed-length character strings
 - (3) packed arrays or structures of either or both (1) and (2)
2. The INITIAL, SYMBOL, storage class, and scope attributes must not be specified for the defined item. The VARYING attribute must not be specified for either the defined item or the base. It should be noted that although the base may have the EXTERNAL attribute, the defined item is always INTERNAL. The name of the base, if declared external, will be known in all blocks in which it is declared external, but the name of the defined item will not. However, the value of the defined item will be changed if the value of the base item is changed in an external block.
3. The LIKE attribute must be specified for a defined item if it is to apply.
4. The defined item must always be specified as a subset (including the full set) of the base identifier. Thus, the dimensions and string length of a defined item must be a subset of those of the base.
5. No other attribute conflicts, except those mentioned in the above rules, are allowed between the defined item and the base.
6. Expressions specified in base-identifier subscript lists are evaluated when the defined item is referred to and not when it is declared. Use of a defined item in an argument list is interpreted as a reference to the defined item.
7. The base identifier must always be known within the block where the defined identifier is declared and cannot have been declared with the DEFINED attribute.
8. Expressions in attributes of the defined data other than in the DEFINED attribute are evaluated on entry to the declaring block. The current generation of the base at each point of reference to the defined item is normally taken as the defining base. If, however, a defined item is passed as a parameter and the base is reallocated, the parameter will be based on the generation current at the time of invocation.
9. Data defined on a CONTROLLED base refers to the most recent generation of base data.

Rules for scalar defining:

1. Both the defined item and base identifier must be scalars.

2. The base identifier may be subscripted, in order to specify a scalar element of an array, but the defined term may not be an element of a structure or an array.
3. Permitted forms are as follows:

<u>Defined Item</u>	<u>Base Identifier</u>
coded arithmetic	coded arithmetic of the same base, scale, mode, and precision
label	label
binary numeric field or bit string	binary numeric field or bit string
decimal numeric field or character string	decimal numeric field or character string
task	task
event	event

4. The POSITION attribute may be specified when the base is a fixed-length string.

General Format:

POSITION (decimal-integer-constant)

This specifies the position (n) relative to the start of the base where the defined item begins. If omitted, POSITION (1) is assumed. The position (n) is restricted as follows: $n + \text{length}(\text{defined-item}) - 1$ must be less than or equal to the length (base identifier). If POSITION is given, then the DEFINED attribute must also be given.

Rules for array defining:

1. Both the defined item and the base identifier must be arrays.
2. The defined item must have a dimension attribute, and may not be an element of a structure.
3. The permitted forms are the same as those for scalar defining.
4. In array defining there is a relationship between each element of the defined array and a corresponding element of the base identifier.
5. The elements of the defined array must have lengths less than, or equal to, the lengths of the base array elements.
6. The POSITION option may be given when the base identifier refers to an array of strings. It specifies that each element of the defined array begins at the nth bit or character position of the corresponding element of the base array.
7. In array defining, both the defined array and the base array may be arrays of structures. In this case, one of the following conditions must be

satisfied:

- a. Both strings must be PACKED and composed of string or numeric field elements of the same class.
- b. Both arrays must have identical structure descriptions.

In this form of defining, the base may not be a parameter.

8. Two classes of array defining are permitted, simple and subscripted.

Rules for simple array defining:

1. The base identifier must be an unsubscripted array name having the same number of dimensions as the defined array.
2. The dimension bounds of the defined array must be a subset of the bounds of the base array.
3. A subsequent subscripted reference to the defined array is interpreted as a reference to the base array with identical subscripts.
4. A subsequent unsubscripted reference to the defined array is interpreted as a reference to the declared subset of the base array specified by the dimension bounds.

Rules for subscripted array defining:

1. The base identifier must be an array name followed by a defining subscript list. The base array need not have the same number of dimensions as the defined array.
2. The defining subscript list defines the relationship between the elements of the defined array and the base array and must have as many subscripts as the base array has dimensions.
3. The defining subscripts may be an expression, including the dummy variable iSUB, and take the following form:

$$a_1 \text{ *1SUB} \pm a_2 \text{ *2SUB} \dots \pm a_n \text{ *nSUB}$$

In the expression iSUB, i is a decimal integer constant in the range 1 to n, with n being the dimensionality of the defined array. The symbol a is any scalar expression involving variables known within the block containing the DEFINED declaration. The integer value of the expression will be used. If any a is zero, that iSUB may be omitted.

4. The subscripted reference to the defined array is then interpreted as follows:
 - a. Each iSUB in the defining subscript is replaced by the integer value of the ith subscript given for the defined array. Before replacement, the subscript conceptually is enclosed in parentheses.
 - b. The reference to the defined array

elements is a reference to the base array element specified by the generated subscript.

5. For example, in the following statement, B is a vector having as elements the diagonal of matrix A:

```
DECLARE A(10,10), B(10) DEFINED
A(1SUB, 1SUB);
```

6. A subsequent unsubscripted reference to the defined array is interpreted as a reference to the entire array as defined by the mapping.
7. If a defined array name is specified as an argument to an invoked procedure, the expressions in the defining subscript list are evaluated before the invocation. The invoked procedure can still reassign values to elements of the defined array by a parameter, but the relationship between the defined array elements and the base elements is frozen on entry.
8. The POSITION option cannot be used with subscripted array defining.

Rules for mixed defining (structures and arrays):

1. Major structures and arrays not contained in structures, having elements all of the same class as described earlier in this section, may be defined on scalar strings of the same class and on structures or arrays having elements all of the same class and which are not part of a major structure having the ALIGNED attribute.
2. When scalar strings are defined items, base items may be any of the following:
 - a. major structures that are PACKED
 - b. minor structures contained in a PACKED major structure
 - c. structure elements with the base being specified as a subscripted structure name
 - d. unsubscripted PACKED arrays not contained in structures
 - e. unsubscripted arrays not contained in undimensional structuresAll of the elements of the base item must be of the same class as the defined string.
3. When the base item is a scalar string, the POSITION option may be specified to indicate that the defined array or structure is offset from the start of the string. It may not be specified with mixed defining when the base is an array or structure.
4. Defining subscript lists may not be used with mixed defining.
5. The base in mixed defining may not be a parameter.

Examples:

```
DECLARE 1 P, 2 Q CHARACTER (10), 2 R
CHARACTER (100),
PSTRING1 CHARACTER (110) DEFINED P;
```

```
DECLARE LIST CHARACTER (40), ALIST
CHARACTER (10)
DEFINED LIST, BLIST CHARACTER (20)
DEFINED LIST
POSITION (21), CLIST CHARACTER (10)
DEFINED LIST
POSITION (11);
```

```
DECLARE ALL (16), EVEN (8) DEFINED ALL
(2*1SUB);
```

THE INITIAL ATTRIBUTE

Function:

The INITIAL attribute either specifies constant values to be assigned to data when storage is allocated to it, or it specifies a procedure to be invoked to perform initialization at allocation.

General format:

1. INITIAL (item [, item]...)
2. INITIAL CALL entry-name [argument-list]

Rules for form 1:

1. In this discussion, the term constant denotes either an optionally signed constant or a complex expression of the following form:
real-constant {+|-} imaginary-constant
2. One constant value is required for a scalar; more may be given for an array.
3. Constant values specified for an array are assigned to successive elements of the array in row-major order (final subscript varying most rapidly).
4. If too many constant values are specified, excess ones are ignored; if not enough are specified, the remainder of the array is not initialized.
5. The items in the list may be an optionally signed constant, an asterisk denoting no initialization for a particular element, or an iteration specification.
6. The iteration specification has one of the following general forms:
(iteration-factor) constant
(iteration-factor) (item [, item] ...)

(iteration-factor) *

7. The "iteration factor" may be any expression that satisfies the rules stated in the section on "Prologues" in Chapter 10. When storage is allocated, the expression is evaluated to give an integer that specifies the number of repetitions.
8. Only constants are permissible as iteration factors for STATIC data.
9. A negative or zero iteration factor yields no initialization.
10. Iterations may be nested.
11. See "Statement Label Data," in Chapter 2, for an alternative method of specifying initial values for label arrays.
12. The INITIAL attribute may not be given for the following:

entry names
file names
DEFINED data
structures
parameters
TASK data
EVENT data

13. If only one parenthesized scalar expression precedes a string initial value, it is interpreted as a replication factor for the string. If two appear, the first is taken to be an initialization iteration factor, the second, a string replication factor. For example:

(2) 'A' is equivalent to ('AA')
(2) (1) 'A' is equivalent to ('A', 'A')

Rules for form 2:

1. The entry name and arguments passed must satisfy the conditions stated in "Prologues."
2. This form may not be used to initialize STATIC data.

Examples:

1. DECLARE SWITCH INITIAL ('1'B);
2. DECLARE MAXVALUE INITIAL (99),
MINVALUE INITIAL (-99);
3. DECLARE A (100,10) INITIAL ((920) 0,
(20) ((3) 5, 9);
4. DECLARE TABLE (20,20) INITIAL CALL
INITIALIZE (X,Y);

The third example results in the following: each of the first 920 elements of A is set to 0, the next 80 elements consist of 20 repetitions of the sequence 5,5,5,9.

In the last example, INITIALIZE is the name of a procedure that sets the initial values of elements in TABLE.

X and Y are arguments passed to INITIALIZE.

SYMBOL TABLE ATTRIBUTES

Function:

The symbol table attributes are used in conjunction with data-directed input/output operations. They specify whether or not the names of data-directed input/output elements are to be placed in a symbol table.

General format:

{ SYMBOL [(identifier)] }
{ NOSYMBOL }

Rules:

1. SYMBOL specifies that the declared name is to be placed in the symbol table.
2. SYMBOL (identifier), used when the declared name must be qualified to make it unique, specifies that the identifier in parentheses is to appear in the symbol table as a synonym for the name to which it refers.
3. A variable whose name or synonym appears in the symbol table may have its values transmitted under data-directed directed input without a data list.
4. NOSYMBOL specifies that the declared name is not to appear in the symbol table.
5. The appearance of an identifier as an element to be transmitted in a data-directed input or output list is a contextual declaration of the attribute SYMBOL.
6. It is illegal to specify SYMBOL without a synonym for two structure elements declared internal to the same block and having the same identifier.
7. Symbol table attributes may not be declared for entry names, files, labels, task and event data, DEFINED data, or parameters.

Reference:

See Chapter 7 for a complete discussion of data-directed input and output.

Default:

The default attribute is NOSYMBOL unless the name appears in a list for data-directed input or output.

THE LIKE ATTRIBUTE

Function:

The LIKE attribute specifies that the name being declared is given the same structuring as the name following the attribute LIKE.

General format:

LIKE structure-name

Rules:

1. The "structure name" may be unqualified or qualified, but it may not be subscripted.
2. The structure must be known to the block containing the LIKE attribute.
3. Neither the structure name nor any of its substructures can be declared with the LIKE attribute.
4. The LIKE attribute specifies that the name being declared is a structure with a substructure having elements with attributes and names identical to the names and attributes of the elements of the named structure. Contained dimension and length attributes are recomputed. Attributes of the structure name itself do not carry over, only its elements enter into this process.
5. If the structure description of the named structure has been declared, and if a direct application of the description to the structure being declared LIKE would cause an incorrect discontinuity in level numbers, then the level numbers will be modified by a constant before application.
6. The number that immediately follows the member that has the LIKE attribute must have a level-number that is equal to or less than that of the member that has the LIKE attribute.

Examples:

1. DECLARE 1 A(10),
2 FIELD1,
3 DTL1 PIC '\$ZZ.99',
3 DTL2 CHAR (10),
2 FIELD2 BIT (50),
1 X,
2 FIELD1,
3 SUBFLD1 (20) LIKE A . FIELD1
3 TABLES (3),
2 FIELD2 LIKE A . FIELD1;

The above is equivalent to:

```
DECLARE 1 A(10),  
2 FIELD1,  
3 DTL1 PIC '$ZZ.99',
```

```
3 DTL2 CHAR (10),  
2 FIELD2 BIT (50),  
1 X,  
2 FIELD1,  
3 SUBFLD1 (20),  
4 DTL1 PIC '$ZZ.99',  
4 DTL2 CHAR (10),  
3 TABLE (3),  
2 FIELD2,  
3 DTL1 PIC '$ZZ.99',  
3 DTL2 CHAR (10);
```

2. DECLARE 1 A EXTERNAL, 2 (B,C,D), 1 E
LIKE A;

The above is equivalent to :

```
DECLARE 1 A EXTERNAL, 2 (B,C,D), 1 E,  
2 (B,C,D);
```

FILE DESCRIPTION ATTRIBUTES

File description attributes are used to describe data files. Declarations of the same file in more than one external procedure must not conflict. For a complete discussion of data files see Chapter 7.

The FILE Attribute

Function:

The FILE attribute specifies that the associated identifier is a file name.

General form:

FILE

Standard Attributes

Function:

The standard attributes are used to assign a file name as a synonym for the standard input/output file.

General format:

STANDIN|STANDOUT

Rules:

No other attributes except the ZERO attribute are used to describe a file if it has been given a STANDIN or STANDOUT attribute.

The Storage Equivalence Attribute

Function:

The storage equivalence attribute is used to specify the sharing of storage used for transmission to external media by two files.

General format:

POOL (file-name)

Rules:

1. The "file name" specifies a file that may share the storage area necessary for transmission to an external medium with the file currently being described.
2. The technique of pooling is implementation defined.

The Function Attributes

Function:

The function attributes specify the function of a file.

General format:

INPUT | OUTPUT | INOUT

Rules:

The particular attribute used specifies that the file is to be used for input, output, or for both input and output. A declaration of INOUT for a file with SEQUENTIAL access denotes the update-in-place mode. Such files must be accessed in the sequence READ, then WRITE.

File Organization Attributes

Function:

The file organization attributes, which are used in conjunction with the access attributes, specify the manner in which the records comprising a file are located.

General format:

{ CONSECUTIVE
REGIONAL (maximum-number-of-records)
INDEXED }

Rules:

1. CONSECUTIVE specifies that the location of records within the file is dependent upon the current physical position of the device containing the file.
2. REGIONAL specifies that the location of records within the file is determined by relative physical regions of the device containing the file. The "maximum number of records" indicates the maximum number of records within a region and must be a decimal integer constant. If this is omitted there is one record per region. Such a file may be accessed in the DIRECT mode with (optionally) only a REGION option, and no KEY option.
3. INDEXED specifies that the location of records within a file is determined by means of an ordered index that addresses areas of the device containing the file.

Default:

If DIRECT access is specified, then INDEXED is assumed. Otherwise the default is CONSECUTIVE.

Access Attributes

Function:

The access attributes specify the manner in which the records within a file are accessed.

General format:

SEQUENTIAL|DIRECT

Rules:

1. SEQUENTIAL normally specifies that the next record to be accessed is determined by the order implicitly given the file by virtue of its organization. If, however, a KEY is provided, direct access is performed.
2. DIRECT specifies that the next record to be accessed is determined by an explicitly stated identification (see the KEY, NEWKEY, and REGION options with the READ and WRITE statements, in Chapter 8). Direct access is permitted only on files organized in the REGIONAL or INDEXED fashion.

Default:

If REGIONAL or INDEXED organization is specified, then DIRECT is assumed. Otherwise the default is SEQUENTIAL.

Table 3. Allowable statements for CONSECUTIVE, REGIONAL, and INDEXED organizations of a SEQUENTIAL access file.

	INPUT	OUTPUT	INOUT
CONSECUTIVE	READ	WRITE	READ WRITE
	GET PUT	GET PUT	GET PUT
	SPACE SKIP GROUP SEGMENT	SPACE SKIP GROUP SEGMENT	SPACE (In) SKIP (In) GROUP (In) SEGMENT (In)
	TAB POSITION LAYOUT	TAB POSITION LAYOUT PAGE	TAB POSITION LAYOUT
REGIONAL	READ	WRITE	WRITE
	READ REGION	WRITE REGION	WRITE REGION
	READ REGION and KEY	WRITE REGION and KEY	WRITE REGION and KEY
		WRITE REGION and NEWKEY	WRITE REGION and NEWKEY
	GET PUT	GET PUT	GET READ REGION READ KEY and REGION
	SPACE	SPACE	SPACE (In) GET PUT
TAB POSITION LAYOUT	TAB POSITION LAYOUT	TAB POSITION LAYOUT	
INDEXED	READ	WRITE	READ WRITE
	READ KEY	WRITE KEY WRITE NEWKEY	READ KEY WRITE KEY WRITE NEWKEY
			SPACE (In)
	TAB POSITION GET PUT SPACE LAYOUT	TAB POSITION GET PUT SPACE LAYOUT	TAB POSITION GET PUT LAYOUT

Table 4. Allowable Statements for the REGIONAL and INDEXED Organizations of a DIRECT Access File

	INPUT	OUTPUT	INOUT
REGIONAL	READ REGION	WRITE REGION	READ REGION WRITE REGION
	READ REGION and KEY	WRITE REGION and KEY WRITE REGION and NEWKEY	READ REGION and KEY WRITE REGION and KEY WRITE REGION and NEWKEY
	SPACE GET PUT	SPACE GET PUT	SPACE (In) GET PUT
	TAB POSITION LAYOUT	TAB POSITION LAYOUT	TAB POSITION LAYOUT
INDEXED	READ	WRITE KEY WRITE NEWKEY	READ KEY WRITE KEY
	READ KEY		WRITE NEWKEY
	TAB POSITION GET PUT SPACE LAYOUT	TAB POSITION GET PUT SPACE LAYOUT	TAB POSITION GET PUT SPACE LAYOUT

Note that a declaration of INOUT for a file with SEQUENTIAL access indicates update-in-place. It applies to all file organizations on direct access devices. The order of access to such a file must be first READ, then WRITE.

A CONSECUTIVE file will, on WRITE, replace the last record read. REGIONAL and INDEXED files will, on READ, remember the KEY (and REGION) values and use them for the next WRITE.

The KEYLENGTH Attribute

Function:

The KEYLENGTH attribute specifies the length, in characters, of the keys associated with records within files organized in the REGIONAL or INDEXED mode.

General format:

KEYLENGTH (length)

Rules:

For indexed organization, this attribute must appear.

Default:

For regional organization, absence of this attribute implies a key length of zero (a keylength of zero is only permitted on REGIONAL files which have one record per region).

The ZERO Attribute

Function:

The ZERO attribute specifies that trailing blank characters in data input fields read by format items E, F, or G are to be treated as numeric zeros.

General format:

ZERO

Rules:

The ZERO attribute has no effect on output data.

The ENVIRONMENT Attribute

Function:

The ENVIRONMENT attribute is an implementation defined attribute which specifies various characteristics of a file which are not related to the PL/I language.

General Format:

ENVIRONMENT (option-list)

Rules:

1. The option list will be defined individually for each implementation of PL/I.
2. Information such as the device type, buffering, file organization support, record format, file disposition, etc., may be specified in the option list.
3. Parentheses occurring within the option list must be matched.

ASSIGNMENT OF ATTRIBUTES TO IDENTIFIERS

Identifiers can be given attributes explicitly through DECLARE statements, by occurrences in certain recognizable contexts, and by default rules for identifiers incompletely described by the programmer.

Within an external procedure, statement label constants, internal entry labels, parameters, and identifiers appearing in DECLARE statements are qualified by the respective blocks in which their declarations (contextual or explicit) occur. For an identifier occurring as a parameter, the characteristic, "parameter," is combined with any explicitly declared attributes for the identifier. Default attributes are added as described below. An identifier occurring as an internal entry label is given the attributes INTERNAL ENTRY, which then are also combined with any declared attributes for that identifier, after which defaults are applied.

The following attributes, assigned through context, are recognized in the indicated ways:

1. ENTRY (subroutine): CALL statement or CALL option
2. ENTRY (function): identifier followed by parenthesized list (if the identifier is not initial value for a label array and has not been declared in a containing block as an array).
3. FILE: READ, WRITE, SPACE, GROUP, SKIP, PAGE, LAYOUT, SORT, OPEN, CLOSE, GIVING, POOL, FROM, or ON, REVERT, or SIGNAL (file conditions).
4. TASK: TASK option
5. EVENT: EVENT option or WAIT statement
6. SYMBOL: DATA list
7. (programmer named condition): ON CONDITION, SIGNAL CONDITION, or REVERT CONDITION

If an identifier appearing in one of these contexts has been declared explicitly or contextually in a containing block without the indicated attribute, an error is raised. If it has already been declared with the attribute, then the identifier is taken to be the one in the innermost block in which it has been so declared. In case 1 above, the characteristic, "subroutine," is added to the declaration.

If an identifier found in one of these contexts has not been previously declared as described above, then it is qualified by the containing external procedure and is given the indicated attribute. Defaults are then added.

Remaining undeclared simple identifiers are qualified by the containing external procedure, and default attributes are assigned.

Application of Default Attributes

Default assumptions are as follows, for the identifier classes indicated:

ENTRY type: EXTERNAL is assumed. If the entry is EXTERNAL and is not a subroutine, then NORMAL is assumed. Otherwise, ABNORMAL is

assumed. Scale, base, mode and precision defaults for the value returned are the same as for Arithmetic type given below.

If a procedure has multiple entry names and no data attributes, there is potential ambiguity in the characteristics of the value to be returned. In order to avoid this ambiguity, succeeding labels are interpreted as if they were entry names for successive ENTRY statements. For example, in the following, statement a is interpreted as if both statement b and statement c had been written.

- a. A: B: ENTRY;
- b. A: ENTRY;
- c. B: ENTRY;

FILE type: EXTERNAL scope is assumed. If neither organization nor access method is given, then CONSECUTIVE SEQUENTIAL is assumed. If CONSECUTIVE is given, then SEQUENTIAL is assumed. If INDEXED or REGIONAL is given, then DIRECT is assumed. If DIRECT is given, then INDEXED is assumed. If SEQUENTIAL is given, then CONSECUTIVE is assumed.

TASK type: ABNORMAL is assumed. Scope and storage class defaults are the same as for Arithmetic type given below. ALIGNED is assumed for arrays not in structures.

EVENT type: Defaults are the same as for TASK type.

LABEL type: Range is assumed to be all labels which could be assigned to the variable. NORMAL is assumed. Scope and storage class defaults are the same as for Arithmetic type given below. ALIGNED is assumed for arrays not in structures.

Condition type: EXTERNAL scope is assumed.

String type: NOSYMBOL is assumed. NORMAL is assumed. Scope and storage class defaults are the same as for Arithmetic type given below. ALIGNED is assumed for arrays not in structures.

Major Structure type: PACKED is assumed. NOSYMBOL is assumed. NORMAL is assumed. Scope and storage class defaults are the same as for Arithmetic type given below.

Minor Structure type: NOSYMBOL is assumed. NORMAL is assumed. INTERNAL is assumed.

Elementary Structure Element type: NOSYMBOL is assumed. NORMAL is assumed. INTERNAL is assumed. If Arithmetic type has been indicated, then scale, base, mode, and precision defaults are the same as for Arithmetic type given below.

Arithmetic type: If none of scale, base, and mode has been given, then if the identifier starts with any of the letters I - N, FIXED BINARY REAL is assumed; otherwise FLOAT DECIMAL REAL is assumed. If at least one of these has been given, then the remaining defaults are FLOAT, DECIMAL and REAL. Default precision is implementation defined, dependent on scale and radix. ALIGNED is assumed for arrays not in structures. NOSYMBOL is assumed. NORMAL is assumed. INTERNAL is assumed. If no storage class is given, then AUTOMATIC is associated with INTERNAL and STATIC with EXTERNAL.

STRUCTURE DECLARATIONS AND ATTRIBUTES

This section is a summarization of data declarations and attributes as they apply specifically to structures.

LEVEL NUMBER

The outermost structure is a major structure, and all contained structures are minor structures.

A structure is specified by declaring the major structure name and following it with the names of all contained elements. Each name is preceded by a level number, which is a non-zero decimal integer constant. A major structure is always at level one and all elements contained in a structure (at level n) have a level number that is numerically greater than n , but they need not necessarily be at level $n+1$, nor need they all have the same level number.

A minor structure at level n contains all following items declared with level numbers greater than n up to but not including the next item with a level number

less than or equal to n . A major structure description is terminated by the declaration of another item at level one, by the declaration of an item having no level number, or by the end of a declaration list.

STRUCTURES AND THE DIMENSION ATTRIBUTE

When a structure name is given the dimension attribute, it is an array of structures, and all contained items are arrays (see "Arrays of Structures," in Chapter 2). Contained scalar items, contained structure elements, and cross sections of contained arrays are referred to, respectively, by subscripted names, subscripted qualified names, and the asterisk notation (see "Naming," in Chapter 2).

STRUCTURES AND DATA ATTRIBUTES

Structures and arrays of structures are not given data attributes. These can be

given only to elementary structure elements.

STRUCTURES AND SCOPE ATTRIBUTES

Major structure names may be declared with the EXTERNAL attribute. Items contained in structures may not be declared with the EXTERNAL attribute, and even if INTERNAL is unspecified, they are assumed to be INTERNAL.

STRUCTURES AND STORAGE CLASS ATTRIBUTES

All items in the same structure must be of the same storage class, since only the major structure may be given a storage-class attribute. The storage class of the major structure applies to all elements of the structure. If a structure has the CONTROLLED attribute, only the major structure, not its elements, may be allocated and freed.

FORMAL PARAMETERS

The PROCEDURE statement heading a given procedure and defining the primary entry point to the procedure may specify a list of formal parameters. (For syntax and details of the PROCEDURE statement, see Chapter 8.)

One or more ENTRY statements may also be used in the procedure to define secondary entry points. Like the heading statement of the procedure, each of the ENTRY statements must have at least one label to serve as an entry name for that point, and each may specify a list of formal parameters, unrelated to the parameter lists for the other entry points. (For syntax and details see "The ENTRY Statement.")

The formal parameters are identifiers and may appear in statements of the procedure in the context of scalar variable names, array names, structure names, statement label designators, entry names, file names, task names, or event names.

The appearance of an identifier in a formal parameter list for a procedure constitutes a declaration of the identifier as a parameter. This contextual declaration can be combined with an explicit declaration and other contextual declarations in the procedure that will associate required attributes with the parameter. Required attributes not declared explicitly or contextually will be assigned by default.

No declarations of the parameter can appear outside the procedure. (For further details about the restrictions on attributes of parameters see "Arguments and Parameters," in Chapter 10.)

Example:

```
SBPRIM: PROCEDURE (X, Y, Z);
        DECLARE (X, Y, A, B) FIXED, Z
            FLOAT;
        A = X-1; B = Y+1;
        GO TO COMMON;
SBSEC:  ENTRY (X, Z);
        A = X-2; B = X-3;
        COMMON: Z = A**2+A*B+B**2;
        END SBPRIM;
```

In this example, the procedure may be entered at its primary entry point SBPRIM, where the formal parameter list is (X, Y, Z), or at its secondary entry point SBSEC, where the formal parameter list is (X, Z).

PROCEDURE REFERENCES

At any point in a program where an entry name for a given procedure is known, the procedure may be invoked by a procedure reference, which has the form:

```
entry-name [(argument [,argument]
            ...)]
```

The number of arguments (possibly zero) in the procedure reference must be equal to the number of formal parameters in the list for the entry point denoted by the entry name.

The procedure invoked by the procedure reference may be an external or an internal procedure. If it is an internal procedure, the block to which the entry name is internal must be active at the time of invocation of the procedure (for a definition of "active," see "Activation and Termination of Blocks" in Chapter 6).

When a procedure reference invokes a procedure, each argument specified in the reference is associated with its corresponding formal parameter in the list for the denoted entry point, and control is passed to the procedure at the entry point. The conditions the arguments must satisfy, and the manner of association of each argument with its matching parameter are discussed in "The Arguments in a Procedure Reference."

When a procedure becomes inactive, the association between arguments and parameters is terminated.

There are two distinctly different uses for procedures, determined by one of two contexts in which a procedure reference may appear:

1. A procedure reference may appear as an operand in an expression. (For a complete description of expression, see "Expressions," in Chapter 3). In this case, the reference is said to be a function reference, and the procedure is invoked as a function procedure, or simply a function.
2. A procedure reference may appear following the keyword CALL, either in a CALL statement or in a statement using a CALL option. In this case, the reference is said to be a subroutine reference, and the procedure is

invoked as a subroutine procedure, or simply a subroutine.

(Ordinarily a given procedure will be used exclusively as a function procedure or exclusively as a subroutine procedure.)

FUNCTION REFERENCES AND FUNCTION PROCEDURES

When a function reference appears in an expression, the function procedure is invoked. The procedure is then executed, using the arguments, if any, specified in the function reference. The result of this execution is the required value, which is passed with return of control back to the point of invocation. This returned value is then used, in place of the function reference, to evaluate the expression.

The procedure invoked by a function reference normally will terminate execution with a statement of the form RETURN (expression), where expression is a scalar expression of arithmetic, character-string, or bit-string type (see "The RETURN Statement"). (A GO TO statement may also be used to terminate execution of a procedure invoked by a function reference.) It is the value of this expression that will be returned as the function value. The PROCEDURE or FENTRY statement at the invoked entry point may specify data attributes for the function value (see "The PROCEDURE Statement" and "The ENTRY Statement," in Chapter 8). Just prior to return, the expression is evaluated, and, before being passed back, the value is converted, if necessary, to conform to these attributes, or, if the attributes are not specified, to the default attributes implied by the entry name.

If the invoked function procedure is terminated by a GO TO statement, the evaluation of the expression that invoked the function will not be completed and control will go to the designated statement.

GENERIC FUNCTIONS

A generic function is a family of functions with a single name. A function reference to a generic function causes the selection of a certain member of the family, depending upon the attributes of the arguments. The characteristics of the value returned depend upon the member that is selected.

Generic functions may be built-in (see below) or specified by the programmer, who may, by means of the attribute GENERIC, define a name to be a generic function name. An entry name may be explicitly declared with the GENERIC attribute. The GENERIC attribute requires a list of all of the entry names of the family and the attributes of all of the arguments for each member (different members must have different argument attribute patterns). Then any reference appearing in the scope of this declaration and using the declared generic name as an entry name will result in the use of that member of the declared family that has the same argument attribute pattern as the pattern in the argument list of the reference. For complete details see "Entry Name Attributes" in Chapter 4.

Subroutine procedures may also be generic. The method of selecting a particular subroutine corresponds exactly to that of selecting a particular function.

BUILT-IN FUNCTIONS

Besides function procedures written by the programmer, a function reference may invoke one of a comprehensive set of built-in functions.

The set of built-in functions is an intrinsic part of PL/I. It includes not only the commonly used arithmetic functions but also functions for manipulating strings and arrays, as well as other necessary or useful functions related to special facilities provided in the language. The complete list of these functions and their descriptions can be found in Appendix 1.

A large number of the built-in functions are generic. The built-in generic functions are of considerable convenience to the programmer. He may, for example, always use the same name EXP for the exponential function, regardless of whether the argument is of REAL or COMPLEX mode, regardless of the precision of the argument, etc., and automatically he will obtain that one of the EXP family that fits the requirements.

Each built-in function, whether or not it is generic, has a specified number of arguments given. For some built-in functions only a minimum is specified; additional arguments are optional. For others, a maximum is specified; only one argument is required.

Each of the built-in functions that are not generic has only a single member. When a reference is made to one of these functions, any arguments whose attributes do not match the attributes required by that function are converted to the appropriate form before the function is invoked. The characteristics of the value returned are determined by the function.

Unlike programmer-specified functions, which always return a scalar value, there are many built-in functions that may effectively return an array or structure value when array or structure expressions are used in certain of their argument positions. This facility is useful in combination with the facility of array or structure expressions.

The fixed set of names for the built-in functions is part of the language of PL/I. However, the identifiers corresponding to these names are not reserved; any such identifier can be used by the programmer for other purposes. If the identifier is declared explicitly for some other use, any appearance of the identifier in the scope of this declaration will refer to that other use. The built-in function cannot, of course, be used in this scope. If the identifier appears, but not in the scope of a declaration establishing the identifier for another use, the identifier will be regarded as implicitly declared in the containing external procedure with the attribute BUILTIN, and this appearance will refer to the built-in function.

If an identifier corresponding to a built-in function name is declared to have a use other than as the built-in function in some block, the built-in function can be used in contained blocks by declaring the identifier with the attribute BUILTIN.

SUBROUTINE REFERENCES AND SUBROUTINE PROCEDURES

When a procedure is invoked as a subroutine by the execution of a CALL statement or a statement with a CALL option, the initial action is the same as if the procedure were invoked as a function: the arguments in the procedure reference, if any, are associated with the formal parameters and control is passed to the procedure at the denoted entry point. (If the invocation involves a task option, the procedure will not necessarily be activated immediately; see "Asynchronous Operations and Tasks" in Chapter 6.)

Unlike the function procedure, the subroutine procedure does not return an expli-

citly specified value to the point of invocation, and control need not necessarily be returned to this point. The procedure may terminate in the following ways:

1. Control reaches a RETURN statement for the procedure. When executed, this statement returns control to the first executable statement logically following the invoking statement, unless the invocation specified a task option or the procedure was invoked by a statement with a CALL option. If a task option has been used, control is simply terminated for this task. If the procedure was invoked by a statement having a CALL option, control is returned to that statement at the point immediately following the CALL option.
2. Control reaches an END statement for the procedure, which in this case is treated as a RETURN statement. The effect is as in case 1.
3. Control reaches a GO TO statement in the procedure that transfers control out of the procedure. (This is not permitted if the procedure has been invoked by a statement with a CALL option or in a CALL statement with a task option.) In this case, control will go to the designated statement (see "The GO TO Statement"). The statement label designator of the GO TO statement may be a parameter of type LABEL, which is associated with a label argument passed from the invoking procedure.
4. Control reaches an EXIT or STOP statement.

Example of Function Reference:

```

COMP: PROCEDURE;
      .
      .
      .
      S1: P10=Q5*POLY5(R0, VAL1);
      .
      .
      .
      POLY5: PROCEDURE (C, X);
             RETURN(C+X*(1+X*(2+X*(3+X*(4
             +5*X)))));
             END POLY5;
      .
      .
      .
      END COMP;

```

In this example, the external procedure COMP contains the function procedure POLY5, which is invoked when the expression Q5*POLY5(R0, VAL1) is being evaluated during execution of the assignment statement labeled S1. When POLY5 is invoked, the arguments R0 and VAL1 will be associated with the parameters C and X, respectively.

The returned value for POLY5 (R0, VAL1) will be the value of the expression:

```
R0+VAL1*(1+VAL1*(2+VAL1*(3+VAL1*(4+5*
VAL1))))
```

Examples of Subroutine Reference:

1. COMP: PROCEDURE;

```

.
.
.
.
S1: CALL POLY5 (R0, VAL1);
S2: P10 = Q5*TEMP;
.
.
.
.
POLY5: PROCEDURE (C, X);
TEMP=C+X*(1+X*(2+X*(3+X*
(4+5*X)));
RETURN;
END POLY5;
.
.
.
END COMP;
```

In the above example, the effect is the same as in the previous example using the function reference. The subroutine procedure POLY5 is invoked by the CALL statement labeled S1. The arguments and parameters are associated as in the previous example, but here, the value of the expression (the same as in the previous example) is assigned within the subroutine to the variable TEMP, which is used by the statement labeled S2, after the RETURN statement passes control back to that statement. Thus, communication of the value is by means of the shared variable TEMP, which, of course, remains available for use following the execution of S2.

In some cases the invoked and the invoking procedure may be separated in such a way that sharing a name in the above simple manner is not possible (see "Scope of Declarations"). Another more general method of communicating values from the invoked procedure, which may be applied in these cases, is illustrated in the following alternative example:

2. COMP: PROCEDURE;

```

.
.
.
.
S1: CALL POLY5 (R0, VAL1, TEMP);
S2: P10=Q5*TEMP;
.
.
.
.
POLY5: PROCEDURE (C,X,Z);
```

```
Z=C+X*(1+X*(2+X*(3+X*
(4+5*X)));
```

```
RETURN;
END POLY5;
```

END COMP;

Here, the invocation of POLY5 by the CALL statement will associate the variable TEMP with the parameter Z, and the action will be exactly as in the previous example: the parameter Z will effectively be replaced by the name TEMP in the assignment statement for Z, and TEMP will be assigned the value of the expression on the right-hand side, with R0 replacing C and VAL1 replacing X, before return to statement S2. In this case, the value has been communicated from the subroutine through a parameter.

The above two examples illustrate how a single value obtained in a subroutine can be communicated back to the invoking procedure. The action of a subroutine will generally be more complex than this; many communicated variables may be involved, whether scalar, array, structure, or statement-label variables; input/output operations may be specified, etc. In contrast, the usual purpose of a function procedure is to return a scalar value.

THE ARGUMENTS IN A PROCEDURE REFERENCE

In general, an argument in a procedure reference may be any valid PL/I expression. An exception to this rule concerns the built-in functions: the only built-in functions that may be passed as arguments in a procedure reference are the Float Arithmetic Generic Functions (see Appendix 1).

The attributes of each argument in a procedure reference must, in general, match the attributes of the corresponding parameter at the named entry point. (An exception in case of a data argument is described below.)

For example, assume that the procedure SUB in a program is defined by:

```
SUB: PROCEDURE (X, Y, Z);
DECLARE X FIXED, Y ENTRY, Z LABEL;
.
.
.
.
END SUB;
```

This implies that the formal parameter X is used as a fixed-point variable with certain default data attributes, Y is used as an entry name, and Z is a statement label designator in the body of the procedure. Then if SUB is invoked in the program by the statement:

```
CALL SUB (R*S, CALC, L5);
```

it is then necessary that:

1. The expression R*S have all the data attributes of the parameter X (unless SUB is described by an ENTRY attribute, see below).
2. CALC be an entry name.
3. L5 be a statement-label designator.

THE USE OF THE ENTRY ATTRIBUTE

(The ENTRY attribute is completely described in Chapter 4.)

An identifier is contextually declared to be an entry name in a block if it appears as a label to a PROCEDURE or ENTRY statement or if it appears in the block following the keyword CALL or as the function name in a function reference whose argument list is non-empty. If it is desired to use the identifier as an entry name in a block where it is not so declared, the identifier must be given the ENTRY attribute explicitly in a DECLARE statement for the block.

As an illustration, in the above example, the CALL statement:

```
CALL SUB (R*S, CALC, L5);
```

has the entry name CALC as its second argument. This appearance of CALC is not recognizable as an entry name by context. It must previously have been declared (either contextually, or explicitly in a DECLARE statement) to have the attribute ENTRY.

A more general form of the ENTRY attribute allows the programmer to enumerate the attributes of the parameters for the named entry point.

As an illustration, in the above CALL statement example, the three parameters corresponding to the three arguments of the CALL statement might be described in the invoking procedure by the statement:

```
DECLARE SUB ENTRY (FIXED, ENTRY, LABEL);
```

This statement specifies that:

1. SUB is an entry name.
2. The entry point SUB has three parameters.
3. The first parameter has the FIXED attribute with certain default data attributes.
4. The second parameter has the ENTRY attribute.
5. The third parameter has the LABEL attribute.

The number of parameters and the attributes of each, as described in the ENTRY attribute specification, must always agree with the number of parameters and their attributes, as defined for the described entry point within the invoked procedure.

One of the applications of the extended form of the ENTRY attribute is mentioned in the immediately following description. (A detailed discussion of the various uses for the ENTRY attribute, including the ABNORMAL, USES, SETS, and GENERIC attributes, can be found in Chapter 4.)

PASSING ARGUMENTS TO THE ENTRY POINT

When a procedure is invoked at a given entry point by a procedure reference and each argument is associated with its corresponding formal parameter, the arguments are said to be passed to the entry point.

The action involved in passing the arguments generally will assume that the attributes of each argument match the attributes of its corresponding formal parameter, as described above. However, if the argument is an expression whose attributes do not correspond to those declared for the parameter associated with that argument, the expression will be evaluated and converted, before the argument is passed, to conform to the attributes described by the corresponding member of the ENTRY attribute list.

As an illustration, in the preceding example, the first argument in the CALL statement, which invokes the procedure SUB, is the expression R*S. Assume that R*S has the FLOAT attribute with certain default attributes. These do not match the attributes of the first parameter at the entry point SUB. Then the ENTRY attribute must be used in the invoking procedure to specify the same attributes for the first parameter as specified in the invoked procedure SUB. (The preceding illustration shows one way of doing this.) Thus, on execution of the CALL statement, the expression R*S is

evaluated according to the FLOAT attribute and then converted to a fixed-point value with the other required attributes, before being passed to the entry point SUB.

(A detailed description of the action involved in passing arguments to the invoked entry point can be found in Chapter 10.)

In certain circumstances, the preparatory action includes the construction of a dummy argument. For example, a dummy argument is constructed when the argument must be converted, as in the example of R*S just discussed, or when the argument is an expression involving constants or operators (R*S is again an example of this circumstance).

In each of its appearances as a reference in the procedure, the formal parameter corresponding to the argument effectively is replaced by the argument name. Thus, all appearances of the parameter during execution of the procedure are treated as appearances of the argument name. However, in the cases where a dummy argument is constructed, it is the dummy argument name that replaces the parameter. Passing an argument does not always imply a true logical substitution of the argument name for the parameter in the procedure. However, in the important case where the argument is an arithmetic, string, or label variable having identical attributes with

the corresponding parameter, a logical substitution does occur. Thus, parameters can be used to communicate values from the invoked procedure back to the invoking procedure. Example 2 of "Subroutine References," above, is an illustration of this.

In the above example, the appearance of CALC as the second argument when SUB is called does not imply that the identifier CALC is contextually declared as an entry name, even though the above ENTRY attribute for SUB has been given.

THE SPECIAL PROCEDURE ATTRIBUTE RECURSIVE

In the PROCEDURE statement or ENTRY statement for a given procedure, certain special attributes that characterize the procedure itself may be specified. (For a complete discussion of these attributes, see "The PROCEDURE Statement.") One of these, which has particular significance, is the attribute RECURSIVE. When a procedure of a program is re-activated in a task while it is still active in the same task (see "Activation and Termination of Blocks"), the procedure is said to be used recursively. Any procedure used recursively during program execution must be specified with the RECURSIVE attribute.

PROGRAM CONTROL

Every program, when it is being executed, has a control that determines the order of execution of the statements. For a discussion of their order see "Sequence of Control," in Chapter 8.

Execution of the program is initiated by the operating system, which invokes the initial procedure. This initial procedure must be an external procedure that has been specified with the MAIN attribute (see "The PROCEDURE Statement," in Chapter 8). This procedure cannot have CONTROLLED parameters (see "Storage Classes," in this chapter).

ACTIVATION AND TERMINATION OF BLOCKS

A begin block is said to be activated when control passes through the BEGIN statement for the block. A procedure block is said to be activated when the procedure is invoked at any one of its entry points.

During certain time intervals of the execution of a program, a block may be active. A block is active if it has been activated and is not yet terminated.

There are a number of ways in which a block may be terminated. These are implied by the following rules:

1. A begin block is terminated when control passes through the END statement for the block.
2. A procedure block is terminated on execution of a RETURN statement or an END statement for the block. (The END statement implies a RETURN statement; see Chapter 8.)
3. A block is terminated on execution of a GO TO statement contained in the block which transfers control to a point not contained in the block.
4. The execution of a STOP statement causes termination of the major task.
5. The execution of an EXIT statement causes termination of the task containing the statement and all tasks attached by this task. Thus, all blocks corresponding to these tasks are terminated.
6. When a block B is terminated, all of the dynamic descendants of B also are terminated.

DYNAMIC DESCENDANCE

If a block B is activated and control stays at points internal to B until B is terminated, no other blocks can be activated while B is active. (This discussion is not applicable to the multi-task, or asynchronous, mode of operation, which implies more than a single control; see "Asynchronous Operations and Tasks.")

However, another block, B1, may be activated from a point internal to block B while B still remains active. This is possible only in the following cases:

1. B1 is a procedure block immediately contained in B (the label of B1 is internal to B) and reached through a procedure reference.
2. B1 is a begin block immediately contained in B and reached through normal flow.
3. B1 is a procedure block not contained in B and reached through a procedure reference. (B1, in this case, may be identical to B (i.e., a recursive call, but conceptually it is to be regarded dynamically as a different block.)
4. B1 is a begin block or a statement specified by an ON statement (see "The ON Statement"), and reached through an interrupt. (For present purposes, even if B1 is a statement, it can be regarded as a block, and this case is dynamically similar to case 1 or case 3 above.)

In any of the above cases, while B1 is active, it is said to be an immediate dynamic descendant of B.

Block B1 may itself have an immediate dynamic descendant B2, etc., so that a chain of blocks (B, B1, B2, ...) is created, where, by definition, all of the blocks are active. In this chain, each of the blocks B1, B2, etc., is said to be a dynamic descendant of B.

It is important for the programmer to note that the termination of a given block may automatically imply the termination of other blocks and that these blocks need not necessarily be contained in the given block; storage for all AUTOMATIC variables declared in these blocks will be released at the time of termination (see "Storage Classes").

DYNAMIC ENCOMPASSING

Block A dynamically encompasses block B, or block B is dynamically encompassed by block A, if B is a dynamic descendant of A.

ALLOCATION OF DATA AND STORAGE CLASSES

Because the internal storage of any computer is limited in size, the efficient use of this storage during the execution of a program is frequently a crucial consideration. The simple static process of data allocation used by many compilers -- the assignment of a distinct storage region for each distinct variable used in the source program -- may be wasteful. Multiple use of a storage region for different data during program execution can reduce the total amount of storage required.

Provisions are included in the language to give the programmer virtually any degree of control over the allocation of storage for the data variables in a program. On the other hand, the entire problem of allocation can be ignored completely by the programmer, if storage economization is of little significance in his situation, and a reasonably efficient use of storage usually will still be obtained automatically.

DEFINITIONS AND RULES

Storage is said to be allocated for a variable when a certain region of storage is associated with, or assigned to, the variable. Allocation for a given variable may take place statically, before execution of the program, or dynamically, during execution.

Storage may be allocated dynamically for a variable and subsequently released. Thus, this storage is freed for possible use in later allocations. If storage has been allocated for a variable and not subsequently released, the variable is said to be in an allocated state.

When a variable appears in a statement of a source program, the appearance is called a reference if it corresponds either to the assignment of a value to the variable (e.g., an appearance on the left side of an assignment statement) or to a use of the value of the variable (e.g., appearance in an expression to be evaluated).

At any point where a variable appears as a reference, it must be in an allocated state.

Note: An unallocated variable may appear as an argument to a procedure with a corresponding CONTROLLED parameter, as an argument to the ALLOCATION function, or in an ALLOCATE statement.

STORAGE CLASSES

Every variable in a program must have a storage class, which specifies the manner of storage allocation for the variable.

There are three storage classes. The storage class is specified by declaring the variable with one of the three storage class attributes STATIC, AUTOMATIC, or CONTROLLED. The storage class may be declared explicitly or by default.

The Static Storage Class

Storage for a variable with attribute STATIC is allocated before execution of the program and is never released during execution.

The scope attribute (see Chapter 4) of a STATIC variable may be INTERNAL or EXTERNAL. An EXTERNAL variable with unspecified storage class has, by default, the STATIC storage class attribute.

The Automatic Storage Class

If a variable has the attribute AUTOMATIC, the status of the block containing this variable (see "Data Description") determines dynamic allocation for the variable. Whenever this block is activated during execution of a program, storage will be allocated for the variable, and the variable will remain in an allocated state until termination of the block. At the time of termination, the storage will be released. Thus, the time interval during which the variable is in an allocated state will necessarily include the intervals when the variable is known (see "Scope of Declarations").

Termination of a block by means of a GO TO statement may imply simultaneous termination of other blocks and, consequently, simultaneous release of storage for all AUTOMATIC variables declared in these blocks (see "The GO TO Statement").

If the block is a procedure and is called recursively (reactivated one or more times before return), previously allocated storage for the AUTOMATIC variable is "pushed down" on each entrance and "popped up" on each return to yield the proper generation of storage for the variable after each return, until the final return out of the procedure.

Note: The terms "pushed down" and "popped up" refer to the notion of a push-down stack. A push-down stack is a logical device S, similar in behavior to a physical stacking process. When an element is placed in S, it is conceptually placed on top of the elements already in S, which are "pushed down." At any time, if S is not empty, the top element -- the element most recently placed in S -- can be removed from S, and the remaining elements are "popped up."

The scope attribute (see Chapter 4) of an AUTOMATIC variable must be INTERNAL. An INTERNAL variable with unspecified storage class has, by default, AUTOMATIC storage class attribute.

The Controlled Storage Class

If a variable has the attribute CONTROLLED, storage allocation must be explicitly specified for the variable by the ALLOCATE and FREE statements.

The ALLOCATE statement (see Chapter 8) specifies one or more variables, each with certain optional attributes. Execution of the statement causes the allocation of storage for the variable specified.

The FREE statement specifies one or more variables, and execution of the statement causes the storage most recently allocated for the variables to be released.

At some point in a program, it may not be known whether a variable X is in an allocated state. The built-in function ALLOCATION (see Appendix 1) is provided to test this state. The function reference ALLOCATION (X) will return the value '1'B if X is in an allocated state, and the value '0'B if not.

The scope attribute of a CONTROLLED variable may be INTERNAL or EXTERNAL.

Example:

```
A: PROCEDURE;
   DECLARE X STATIC;
```

```

.
.
.
B: PROCEDURE;
   DECLARE Y (100) CONTROLLED, Z CHARACTER (1000);
   .
   .
   ALLOCATE Y;
   .
   .
   FREE Y;
   .
   .
C: BEGIN;
   DECLARE Z (100);
   .
   .
   END C;
   .
   .
   RETURN;
   .
   .
   END B;
   .
   .
   END A;
```

Assume in the above example that the termination of procedure A occurs on the return implied by END A, the termination of procedure B occurs on the RETURN statement, and the termination of block C occurs at END C. Then in this example:

Storage for the static variable X is allocated before execution and is never released.

The character-string variable Z is AUTOMATIC by default. Storage is allocated for this Z on entrance to procedure B and is released on execution of the RETURN statement.

The array-variable Z is AUTOMATIC by default. Storage is allocated for this Z at the beginning of execution of block C and is released at END C.

Storage for the CONTROLLED variable Y is allocated on execution of the ALLOCATE statement and is released on execution of the FREE statement. After execution of the FREE statement, the variable Y presumably is not used, but the character-string variable Z can be used, since storage is not released for this variable until the termination of procedure B.

ASYNCHRONOUS OPERATIONS AND TASKS

PL/I allows tasks to be created by the programmer and provides facilities for the following:

1. Synchronizing tasks
2. Testing whether or not a task is complete
3. Changing the priority of tasks.

SYNCHRONOUS AND ASYNCHRONOUS OPERATIONS


Unless the program specifies the creation of tasks, the execution of the statements of the program will proceed serially in time, according to the sequence designated by the order of the statements and the control statements (see "Sequence of Control" in chapter 8). Such operation is said to be synchronous.

In addition to full facilities for conventional synchronous processing, means are provided for performing operations asynchronously.

Some reasons for considering the use of asynchronous operations are:

1. The programmer may wish to make use of computer facilities which can operate simultaneously, e.g., input/output channels, multiple central processing units.
2. A program may be written in which input/output units initiate or complete transmission at unpredictable times, e.g., disc seeks, terminals.

The following two diagrams distinguish between synchronous and asynchronous operations. The first diagram depicts the serial action of synchronous operations, and the second diagram depicts the parallel action of asynchronous operations. (The circles represent statements.)




In asynchronous operation, once a new line has been started, the statements on that line are executed in sequence, but independently of the statements on any other line. Statements on any two lines need not necessarily be executed simultaneously -- whether this occurs depends on the resources and state of the system.

SYNCHRONIZING TWO ASYNCHRONOUS OPERATIONS

In order that the result of an asynchronous operation may be made available to other procedures, means are provided to synchronize two or more asynchronous operations.

The following diagram illustrates this:



Assume that before statement N can be executed, both M and E must have been executed. M therefore issues a WAIT statement which will suspend operation on that line until E has completed. After N, the statements O, P, ..., are executed synchronously, as are the statements F, G, ...,.

TASK AND EVENTS

In PL/I, asynchronous operations result from the creation, by the programmer, of tasks. The synchronizing of operations is obtained by waiting on events.

A task is an identifiable execution of a set of instructions. A task is dynamic, and only exists during the execution of a program or part of a program.

A task is not a set of instructions, but an execution of a set of instructions. The instructions themselves, as written by the programmer, may in fact be executed several times in different tasks.

It is necessary for at least one task to exist when a PL/I program is executed. Thus when an external procedure is first entered, its execution is part of a task. This particular task is called the major task; it is created by the operating envi-

ronment and its creation does not necessarily concern the PL/I programmer. If the programmer is concerned with only synchronous operations, then the major task will be the program itself.

In order to initiate asynchronous operations, the programmer has to create new tasks, as described below. All tasks created by the programmer are called sub-tasks.

With each task, except the major task, it is possible to associate a task name. The task name may be used to refer to and set the priority of the task.

A task may be suspended by the programmer until some point in the execution of another task has been reached. The specified point is known as an event and the record of its completion is contained in an event name. (See the EVENT built-in function and the EVENT pseudo-variable.)

An event name may be associated with the completion of a task. It is necessary to specify such an event name if the programmer wishes to synchronize a point in one task with the completion of another task, by means of the WAIT statement.

Other event names may be defined by the programmer and used in WAIT statements. In this way, the programmer can synchronize a task with events other than the completion of another task. Event names may be set by referring to them in assignment statement by means of the EVENT pseudo-variable.

THE CREATION OF TASKS

In PL/I tasks are created by writing:

- A TASK option
- An EVENT option
- A PRIORITY option

or any combination of these options in a CALL statement (see "The CALL Statement" in Chapter 8). The called procedure will then be executed asynchronously with the calling procedure. The CALL statement itself is not part of the newly-created task. The execution of the calling procedure is known as the attaching task. The execution of the called procedure is known as the attached task.

The TASK option is given in order to name the task created by the CALL. This is necessary if the programmer wishes to examine or change the priority of the called procedure, since the PRIORITY function and pseudo-variable have a task name as an argument.

The EVENT option is given if the programmer wishes to issue a WAIT statement which will wait on the completion of the task created by the CALL.

The PRIORITY option is given if the task created by the CALL is to be given a different priority from that of the task in which the CALL statement appears.

The term "task option" will be used in all later discussions to denote any one of the three options TASK, EVENT, or PRIORITY, or any part of these options, or all three.

TERMINATION OF TASKS

A task may be terminated in one of the four following ways:

1. Control for the task reaches an EXIT statement (see Chapter 8 for a discussion of each of the statements mentioned here).
2. Control for any task reaches a STOP statement.
3. Control for the task reaches a RETURN statement for the procedure defining the task.
4. Control for the task reaches an END statement for the procedure defining the task.

ALLOCATION OF DATA IN TASKS

The rules of scope and storage allocation hold across task boundaries. If storage is allocated for a variable in the attaching task, this allocation may apply to the attached task, so that the variable may appear as a reference in the attached task. It is the responsibility of the programmer to be certain that storage for such a variable is not released too early in the attaching task. (Normally, this is done by synchronizing by use of the WAIT statement.)

(Further details concerning tasks as related to storage allocation and other special considerations can be found in Chapter 10; also see "The WAIT Statement" for other information and examples.)

INTERRUPT OPERATIONS

During the course of program execution any one of a certain set of conditions may occur that can result in an interrupt. An

interrupt operation causes the suspension of normal program activities, in order to perform a special action; after the special action, program activities may or may not resume at the point where they were suspended. The time point of an interrupt is, in general, unpredictable.

For most conditions that can cause an interrupt, the special action to be taken may be specified by the programmer. To do this, he may specify the condition in an ON statement; therefore these conditions are known as the ON-conditions. A complete list and description of the ON-conditions can be found in Appendix 3. With two exceptions (see "Programmer Defined ON-Conditions," in this chapter), each ON-condition is named with a unique identifier suggestive of the condition (e.g., ZERODIVIDE names the condition obtaining whenever an attempt is made to divide by zero). This collection of names, like the built-in function names, is an intrinsic part of the language, but the names are not reserved; the programmer may use them for other purposes, so long as no ambiguity exists.

PURPOSE OF THE CONDITION PREFIX

The conditions named in the prefix to a statement may occur during program execution of a statement lying in the scope of the prefix (see below). If one of these conditions actually does occur, the appearance in the prefix of the corresponding condition name -- or its negation with the word NO -- determines whether or not an interrupt operation will then take place.

Any condition whose occurrence will cause an interrupt is said to be enabled. Enabling of the first five conditions listed above, namely, UNDERFLOW, OVERFLOW, ZERODIVIDE, FIXEDOVERFLOW, and CONVERSION, is provided automatically by PL/I; any occurrence of one of these conditions will cause an interrupt unless the enabling has been negated through the use of a prefix containing the condition name preceded by the word NO. The programmer must himself enable the other conditions through the use of a prefix. For example, no interrupt will occur for a SIZE error (see SIZE condition in Appendix 3), unless the error occurs in a calculation within the scope of a SIZE prefix. For further details, see "The ON Statement" in Chapter 8 and "System Interrupt Action" below.

SCOPE OF THE CONDITION PREFIX

The scope of the prefix depends upon the statement to which it is attached.

If the statement is a PROCEDURE or BEGIN statement, the scope of the prefix is the block or group defined by this statement, including all nested blocks, except those for which the condition is re-specified. The scope does not include procedures that lie outside the scope as defined above but which may be invoked by the execution of statements in this scope.

If the statement is an IF statement or an ON statement, the scope of the prefix does not include the blocks or groups that are part of the statement. Any such block may also have an attached prefix, whose scope rules are implied by the other rules given here.

For any other statement, the scope of the prefix is that of the statement itself, including any expressions appearing in the statement but not any procedure explicitly called by the statement.

USE OF THE ON STATEMENT

In order to define the action to be taken when an interrupt occurs, the programmer may write an ON statement, which has the general form:

ON condition-specification action-specification

The "condition specification" either is an ON-condition name or denotes a programmer-defined condition, and the "action specification" is a single simple statement or begin block, optionally preceded by the keyword SNAP (see "The ON Statement" for complete syntax and details). If the single statement is null, control is given back to the point of interrupt.

When an ON statement that is internal to a given block (for example, a block B) is executed, it causes a preparatory action with the following effect:

If, during the execution of any statement after the execution of the ON statement and before the termination of block B (including the execution of statements in all dynamic descendants of block B), the condition specified

in the ON statement ever occurs and an interrupt results, the statement or begin block specified in the ON statement will be executed as though it were invoked as a procedure block. (If SNAP also has been specified, a standard action providing program checkout information will precede this pseudo-invocation.) Control normally will be returned to the point following the interrupt.

When an ON statement specifying a given condition is executed, the action to be taken is established by the execution. The time interval during which this action specification is effective is defined above in the description of the effect of an ON statement. There are two qualifications to this description:

1. If, after a given action is established by execution of an ON statement, and while this action specification is still effective, another ON statement specifying the same condition is executed, then this latter ON statement will take effect as described above, so that its specified action will determine the interrupt action for the given condition. (The effect of the old ON statement is either temporarily suspended or completely nullified, depending upon whether or not the new ON statement is in a block dynamically descendant from the block to which the old ON statement is internal; see "The ON Statement" and "The REVERT Statement" for more details.)
2. There are eight ON-conditions whose names (possibly preceded by the word "NO") may appear in a prefix to a statement. Even when one of these conditions appears in an ON statement, occurrence of the condition will not necessarily result in an interrupt. For an interrupt to occur, there are certain additional requirements, which are described in the following paragraph.

There are three of these eight ON-conditions, SIZE, SUBSCRIPTRANGE, and CHECK (identifier list), for which an interrupt will not take place when the condition occurs unless the programmer specifically designates that the interrupt is to take place. He may enable this condition by explicitly specifying the condition in a prefix whose scope will cover the calculation where the condition may occur. If a calculation resulting in the occurrence of either of these conditions does not lie within the scope of such a prefix, no interrupts will occur. The other five of these eight special

ON-conditions, namely OVERFLOW, UNDERFLOW, ZERODIVIDE, CONVERSION, and FIXE-D-OVERFLOW, are always enabled, but the programmer may specifically designate that an interrupt is not to take place. An interrupt for any one of these conditions will always take place when the condition occurs unless the occurrence is in a calculation lying within the scope of a prefix specifying NOOVERFLOW, NOUNDERFLOW, NOZERODIVIDE, NOCONVERSION, or NOFIXE-D-OVERFLOW, respectively.

All other conditions, whose names cannot be used in a prefix, are always enabled.

SYSTEM INTERRUPT ACTION

Each of the ON-conditions has a standard action defined for it if an interrupt should occur. If there has been no previous execution of an ON statement (in which the programmer specifies the interrupt action), any interrupt caused by the occurrence of the condition during program execution will result in a standard system action, which is dependent upon the nature of the condition. If the programmer does not want the system action in the case where one of these conditions may occur and cause an interrupt, he must specify an alternative action for the condition through use of the ON statement.

In some situations, the programmer may want to specify his own action for a given condition, to have it hold for part of the execution of the program, and then to have this specification nullified and allow the standard system action. In this case, he may use the special action-specification SYSTEM, as follows:

ON condition-name SYSTEM;

Example 1:

```
A: PROCEDURE;
    .
    .
    .
    ON OVERFLOW
        BEGIN;
        DECLARE NUMBOV STATIC
            INITIAL (0);
        NUMBOV=NUMBOV + 1
        IF NUMBOV = 100 THEN GO
            TO OVERR;
        END;
    .
    .
    .
    ON OVERFLOW;
    .
```

```

.
.
ON OVERFLOW SYSTEM;
.
.
END A;

```

In the above example, assume that the program consists only of procedure A, that the three ON statements are the only ON statements involving the OVERFLOW condition, that they are internal to procedure A, and that they are executed in their physical order.

When program execution begins, the OVERFLOW condition is enabled by the system; any floating-point overflow condition that occurs before the first ON OVERFLOW statement is executed will result in an interrupt, with standard system action. However, the execution of the first ON OVERFLOW statement establishes the action specified in the BEGIN block. (The number of overflows is counted and if this number has not reached 100, the action is finished.) Any OVERFLOW interrupts will receive this action until the second ON OVERFLOW statement is executed. The action specified here is a null statement; any subsequent OVERFLOW interrupts will effectively be ignored until control reaches the third ON OVERFLOW statement, which reestablishes the standard system action.

Example 2:

```

(SIZE) : A:  PROCEDURE;
.
.
ON SIZE GO TO AERR;
.
.
CALL B;
.
.
END A;

(SIZE, NOOVERFLOW) : B: PROCEDURE;
.
.
ON SIZE GO TO BERR;
.
.
RETURN;
END B;

```

In the above example, the prefix (SIZE) enables that condition for procedure A and specifies that if a SIZE error (see Appendix 3) occurs during any calculation in procedure A, an interrupt is to take place.

The prefix (SIZE, NOOVERFLOW) for procedure B specifies the same requirement with respect to a SIZE error for procedure B; in addition, it specifies for procedure B that any interrupt that might be caused by an OVERFLOW condition is to be suppressed.

After the beginning of execution of procedure A, and before the execution of the first ON statement, any SIZE error will result in an interrupt with standard system action. After execution of this ON statement, and before execution of the ON statement in the invoked procedure B, any SIZE error will result in an interrupt with the action GO TO AERR. After execution of the ON statement in procedure B, the action GO TO BERR becomes established for the SIZE condition, but the effect of the previous ON statement is suspended only temporarily. After the RETURN statement in procedure B is executed, the effect of this previous ON statement is reinstated, so that SIZE errors occurring after this point again result in the action GO TO AERR.

If any floating-point overflow condition occurs during the execution of procedure A, an interrupt will result with the standard system action for the OVERFLOW condition. However, for any occurrence of an OVERFLOW condition during the execution of procedure B, the interrupt will be suppressed.

Example 3:

```

(NOOVERFLOW) : A:  PROCEDURE;
.
.
(OVERFLOW) : B:  BEGIN;
.
.
END B;
.
.
END A;

```

In the above example, interrupts will be suppressed for OVERFLOW conditions occurring during execution of that part of procedure A that is not included in block B. OVERFLOW conditions occurring during execution of block B will result in an interrupt.

USE OF THE REVERT STATEMENT

The REVERT statement may be used, following an ON statement, to reinstate an action specification that existed in the immediate, dynamically encompassing block

without having to return control to that block (see "The REVERT Statement," in Chapter 8. for format and rules).

Example:

```
(SIZE) :   A:  PROCEDURE;
           ON SIZE GO TO AERR;
           .
           .
           .
           CALL B;
           .
           .
           .
           END A;
(SIZE) :   B:  PROCEDURE;
           ON SIZE GO TO BERR;
           .
           .
           .
           REVERT SIZE;
           .
           .
           .
           END B;
```

In the above example, if a SIZE error occurs in procedure B after execution of the ON statement, an interrupt will take place with the resulting action GO TO BERR. After execution of the REVERT statement, the condition as specified by the ON statement in procedure A is reinstated. Program control remains in procedure B, but any subsequent SIZE error that occurs in procedure B will cause an interrupt with the action GO TO AERR.

PROGRAMMER-DEFINED ON-CONDITIONS

There are two kinds of ON-conditions the programmer may construct:

1. An arbitrary identifier can be used to create a condition name by means of the keyword CONDITION used in the ON statement, as follows:

```
ON      CONDITION(identifier)  action-
specification
```

Such a statement contextually declares the "identifier" to be a condition-name and the execution of the statement enables the named condition. The condition can be caused to "occur" only by the execution of a SIGNAL statement (see "The SIGNAL Statement").

For example, if the following statement is executed:

ON CONDITION(KEY) block

and later the following statement is executed:

```
SIGNAL CONDITION(KEY);
```

then the latter execution will (by definition of the SIGNAL statement) cause an interrupt, with the action defined by the block in the ON statement.

2. The CHECK (identifier list), where "identifier list" represents variables or labels declared in the program, can appear as the condition specification in the ON statement. Whenever one of the variables in the list is assigned a value, or one of the procedures or statements whose label appears in the list is executed and if the condition is enabled, the condition defined by this specification is regarded as occurring, and an interrupt will take place. (For a precise explanation of this kind of condition, see Appendix 3, "ON Conditions.")

FACILITIES FOR PROGRAM CHECKOUT

The programmer-specified condition described above is a powerful tool for program checkout. As an example of its use, suppose that a block contains the prefix (CHECK (A, SUB1, ST5)) and that the following statement is executed:

```
ON CHECK (A, SUB1, ST5) SYSTEM
```

In the example, A is a data variable, SUB1 is a procedure name, and ST5 is a statement label. Then, whenever a value is assigned to A (or to any part of A, if A is an array or structure name), an interrupt occurs, and A is printed out on a debugging file with its new value. If the statement labeled ST5 or the procedure SUB1 is executed, the label is printed out.

Another useful ON-condition is the condition named SUBSCRIPTRANGE. Parts of the program can be designated by the programmer, using the keyword SUBSCRIPTRANGE in appropriate prefixes, to receive constant monitoring of subscript values. Whenever the value of some subscript in some array goes out of its designated range, an interrupt will occur, and action, specified by a previously executed ON statement, will take place to correct the error.

The SIGNAL statement also will be found useful for checkout, since it can be used to simulate the occurrence of any ON-condition (see "The SIGNAL Statement").

CHAPTER 7: INPUT/OUTPUT

All input/output activity is performed with named collections of data called files. The name of a file is a file name. Files may be subdivided into smaller collections of data called records. Furthermore, records may be ordered within a file so that the data conceptually constitutes a single stream upon which the record structure has been superimposed. Records in such ordered files may be collected into groups.

The natural record structure of a data stream may be inappropriate to some applications. For such applications, significant divisions of data may be indicated by arbitrary symbols called segment marks. Such divisions are called segments. The programmer can, by specifying segments, provide for the crossing of record boundaries during data transmission.

Data is transmitted between the external medium and internal storage as a record. The record may be considered a continuous string of characters or bits, with the string subdivided into contiguous substrings called fields. A field may be empty, or it may contain one and only one item of data. Fields, segment marks, and item delimiters cannot span record boundaries.

The number of fields in a record, the size of those fields, the nature of the data item in each field, and the segment marks, if any, is called the format of the record.

The order of items to be transmitted is specified by a list of elements. On input, the elements are variables or pseudo-variables to which are assigned the values of the corresponding fields of data. On output, the elements are expressions whose values are given to the corresponding fields of output data. As data is transmitted, a field pointer moves across the record in synchronization with the processing of the list elements. The positioning of the pointer is governed by format items and positioning statements given for the record. The list element and the format may be specified in the record or may be specified in a list of elements and a format specification in the program.

DATA TRANSMISSION

There are four modes of data transmission: list-directed, data-directed, format-directed, and procedure-directed.

LIST-DIRECTED TRANSMISSION

List-directed transmission permits the user to specify the storage area to which data is assigned or from which data is transmitted without specifying the format.

Input: The data on the external medium is in the form of optionally signed valid constants or expressions to represent complex constants. The program storage areas to which the data is to be assigned is specified by a data list. (See "List-Directed Input" below.)

Output: The data values to be transmitted are specified by a data list. The form of the data on the external medium is a function of the data value. (See "List-Directed Output" below.)

DATA-DIRECTED TRANSMISSION

Data-directed transmission permits the user to read or write self-identifying data. (See "Data-Directed Data Specification" below.)

Input: The data on the external medium is in the form of optionally signed valid constants and includes information defining the program storage areas to which the data is to be assigned.

Output: The data values to be transmitted are specified by a data list. The data on the external medium has the form of constants and includes the name of the data being transmitted.

FORMAT-DIRECTED TRANSMISSION

Format-directed transmission permits the user to specify both the storage area to which data is to be assigned or from which data is to be transmitted and the form of fields in the record.

Input: The form of the data on the external medium is defined by a format list. The program storage areas to which the data is to be assigned is specified by a data list.

Output: The data values to be transmitted are defined by a data list. The form that the data is to have on the external medium is defined by a format list.

PROCEDURE-DIRECTED TRANSMISSION

Procedure-directed transmission provides the ability to process data fields during transmission by invoking a procedure with a CALL option.

When a statement specifying a file is executed, the specified file becomes the current file. If this statement then invokes a procedure which also contains a statement specifying a file, etc., a stack of current files is created, with the actual current file being the most recent entry into the stack. When a statement specifying a file is completed, that file is released from the stack and the next file in the stack, if any, becomes the current file.

This concept of a current file makes possible the use of GET and PUT and various positioning statements without having to specify the file name in the statements.

A GET or PUT statement that is executed when there is no current file causes the ERROR condition to be raised.

Examples:

```

1. Z1: PROCEDURE;
    DECLARE FILEX FILE;
    .
    .
    .
    READ FILE (FILEX) , LIST (A,B,C) ,
        CALL Z2;
    .
    .
    .
    END Z1;
Z2: PROCEDURE;
    DECLARE FILEY FILE;
    .
    .
    .
    READ FILE (FILEY) , DATA, CALL
        Z3;
    .
    .
    .
    END Z2;
Z3: PROCEDURE;
```

```

DECLARE FILEZ FILE;
.
.
.
WRITE FILE (FILEZ) , LIST (P,Q) ;
.
.
.
END Z3;
```

Upon execution of the READ statement in procedure Z1, FILEX becomes the current file. Since this READ statement invokes procedure Z2, the READ statement in Z2 is executed. This causes FILEY to become the current file. Now there are two files in the stack of current files (FILEX and FILEY). Similarly, the READ statement in Z2 invokes procedure Z3 in which FILEZ is added to the stack of current files. This stack is ordered from most current to least current as follows:

```

FILEZ
FILEY
FILEX
```

When the WRITE statement in Z3 is completed, FILEZ is removed from the stack of current files. Similarly, when the READ statement in Z2 is completed, FILEY is removed from the stack, and upon completion of the READ statement in Z1, there are no more current files.

```

2. A: PROCEDURE;
    DECLARE P FILE;
    .
    .
    .
    READ FILE (P) , (D,E) (2F(8,3)) ,
        CALL B;
    .
    .
    .
    END A;
B: PROCEDURE;
    SKIP (4) ;
    GET LIST (F) ;
    .
    .
    .
    END B;
```

The READ statement in procedure A causes file P to become the current file, data items D and E to be transmitted, and procedure B to be invoked. Procedure B causes a skip to the fourth record of the current group on the current file (P), and causes data item F to be read.

DATA SPECIFICATIONS

Data specifications are given in READ, WRITE, GET, and PUT statements to describe the data to be transmitted. The data specifications correspond to the modes of transmission.

DATA LISTS

List-directed, data-directed, and format-directed data specifications require a data list to specify the data items to be transmitted.

General format:

(element [, element]...)

Syntax rules:

The character of the elements depends upon whether the data list is used for input or for output. The rules for each are as follows:

1. On input, each data-list element for format- and list-directed data may be one of the following: a scalar name, an array name, a structure name, a pseudo-variable, a pseudo-array, a pseudo-structure, or a repetitive specification involving any of these elements. For a data-directed data specification, each data-list element may be an unsubscripted scalar, array, or structure name.
2. On output, each data-list element for format- and list-directed data specifications may be one of the following: a scalar expression, an array expression, a structure expression, or a repetitive specification involving any of these elements. For a data-directed data specification, each data-list element may be a scalar,

array, or structure name, or a repetitive specification involving any of these elements.

Repetitive Specification

General format is shown in Figure 1.

Syntax rules:

1. Each element in the element list of the repetitive specification is as described for data-list elements above.
2. The expressions in the specification are described as follows:
 - a. Each expression in the specification is a scalar expression.
 - b. In the specification, expression 1 represents the starting value of the control variable or pseudo-variable. Expression 3 represents the increment to be added to the control variable after each repetition of data-list elements in the repetitive specification. Expression 2 represents the terminating value of the control variable. The exact meaning of the specification is identical to that of a DO statement with the same specification. When the last specification is completed, control passes to the next element in the data list.
3. Repetitive specification may be nested to any depth. That is, each element in the element list may be a repetitive specification. A repetitive specification involving m elements repeated n times is equivalent to $m * n$ elements. For example, consider the following list element:

((A (I,J) I=1 TO 2) J=3 TO 4)

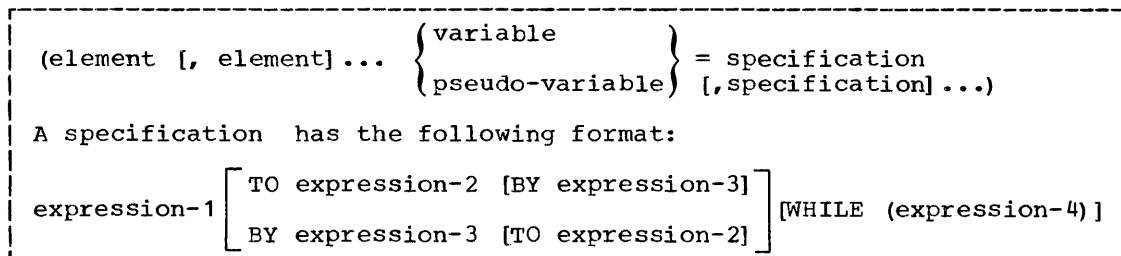


Figure 1. General Format for Repetitive Specification

This gives the elements of the array A in the following order:

A(1,3), A(2,3), A(1,4), A(2,4)

Transmission of Data-List Elements

If a data-list element is of complex mode, the real part is transmitted before the imaginary part.

If a data-list element is an array name, the elements of the array are transmitted in row-major order, that is, with the rightmost subscript of the array varying most frequently.

If a data-list element is a structure name, the elements of the structure are transmitted in the order in which they are stored.

If, within a data list used in an input statement, a variable is assigned a value, this new value is used in all later references in the data list.

Example:

In the following statement, B is a structure, XSTRING is a character string, and C is an array:

```
DECLARE A FLOAT, 1 B, 2 P, 2 E, 3 F,
        XSTRING
        CHARACTER (6), C(10) FIXED;
```

The following data list, involving these data items, and the scalar variable A, may be used for input or output.

(A,B, SUBSTR (XSTRING,2), (C(I) I=1 TO 10))

The data-list elements are transmitted in the following order:

A The scalar variable is transmitted.
P, F The elements of the structure B are transmitted.

SUBSTR (XSTRING, 2) The second through sixth characters of the string are transmitted.

C(1), C(2) ..., C(10) The ten elements of the array are transmitted.

LIST-DIRECTED DATA SPECIFICATION

General format:

LIST data-list [(scalar-expression)]

Syntax rules:

1. The "data list" is described in the preceding discussion.
2. The "scalar expression," which may optionally follow the data list, represents a separator of data items on the external medium.

List-Directed Input

When the data item is an array name and the data consists of scalar constants, the first constant is assigned to the first element of the array, the following constant to the second element, etc., in row-major order.

A structure name in the data list represents a list of the contained scalar variables and arrays in the order specified in the structure description.

List-Directed Input Format

Data on the external medium has the following general form:

$$\left\{ \begin{array}{l} [+|-] \text{ arithmetic-constant} \\ \text{'string-constant' [B]} \\ [+|-] \text{ real-constant} [+|-] \text{ imaginary-constant} \end{array} \right\}$$

Sterling constants may not be used. A string constant must be one of the four permitted forms listed above. The string value is obtained by deleting surrounding quotes and replacing contained double quotes by single quotes. If the string represents an optionally signed arithmetic constant or complex expression, this value represents the converted arithmetic value. If the string represents a bit constant, this bit string value is converted to an arithmetic value as described in Chapter 3. If the string represents a character constant, the above process is repeated.

Redundant blanks are permitted as in PL/I programs. However, if there is no specified separator, or if the separator is specified as a blank, no blanks may precede the central + or - in complex expressions.

A null field on input implies the data list item is to be skipped.

A scalar expression, enclosed in parentheses, optionally may follow the list-directed data list. If the expression is not present, data items on the external medium must be separated by a comma or a blank. If present, the expression is evaluated and converted, if necessary, to a character string; the resultant character string is recognized as the separator of data items on the external medium.

The transmission of the list of constants on input is terminated by an attempt to read the segment mark. List-directed input implies the SEGMENT option.

List-Directed Output

The values of the scalar variables in the data list are converted to a character representation of the data value, as described below, and transmitted to the external medium.

A scalar expression, enclosed in parentheses, optionally may follow the list-directed data list. If the expression is not present, a blank is used to separate data items to be transmitted. If present, the expression is evaluated and converted, if necessary, to a character string; the resultant character string is used to separate data items to be transmitted.

If the SEGMENT option appears in the WRITE statement, the final data item transmitted will be followed by the segment mark (a semicolon will not appear as a default for list-directed output).

List-Directed Output Format

Data fields written under list direction are aligned in vertical columns. This alignment is called tabbing, and is separately defined for each implementation of the language. System tabbing may be overridden by a TAB option in a LAYOUT statement for the appropriate file.

Each item of a list-directed data specification, except the last, is immediately followed by the separating character indicated in the list-directed data specification by a scalar expression enclosed in parentheses. If the data is to be reread under list direction, care should be taken to avoid including character data or numeric fields in the output and to avoid ambiguity resulting from choice of separating or terminating characters.

Data items are written on successive tab positions unless an item extends into a

succeeding position. (In that case, the next free position is used.) A data item may span several tabs, but must not span record boundaries.

Length of List-Directed Output Fields

The length of the data field on the external medium is a function of the internal precision and value of the data item.

CODED ARITHMETIC DATA: The external form of coded arithmetic data is a possibly signed valid decimal constant whose field width, w, is a function of the internal precision declared for the data item and the value of the data item. In the discussion below, the following abbreviations are used:

1. The letter w represents the field width, which is defined as the length of the data field on the external medium.
2. The letter d represents the number of positions in the external data field to the right of the decimal point.
3. The letter p represents the total number of significant digits in the data field.
4. The letter q represents the number of digits to the right of the decimal point.
5. The letter z represents the total number of leading zeros to be suppressed. A zero immediately before the decimal point, whether the decimal point is printed or not, is not suppressed.
6. The letter t represents the existence or nonexistence of a decimal point; t takes on the values 0 and 1 depending upon the value of q.

if $q = 0$, then $t = 0$
if $q \neq 0$, then $t = 1$
7. The letter u represents the existence of a minus sign; u takes on the values 0 and 1 depending on the value of the data.

if data value < 0 , then $u = 1$
if data value ≥ 0 , then $u = 0$
8. The letter s represents a scale factor as described for floating-point data.
9. The pair (r,s) represents the declared internal precision of a coded real decimal fixed point data item.

There are five kinds of coded arithmetic data to consider: coded real fixed-point decimal data, coded real fixed-point binary data, coded real floating-point decimal data, coded real floating-point binary data, and coded complex data.

Coded Real Fixed-Point Decimal Data: The data item is converted to precision (p,q) and transmitted to a field of width w where

$$\begin{aligned}q &= \text{MAX}(s, 0) \\p &= r-s+q \\w &= p-z+t+u\end{aligned}$$

Coded Real Fixed-Point Binary Data: The data is converted to fixed-point decimal and is transmitted as coded real fixed-point decimal data.

Coded Real Floating-Point Decimal Data: The data is converted according to the rules for fixed-point format items, F(w,d). For F-conversion, if p is the declared precision of the data item, $w = p + 2$ and $d = p - 1$. For list-directed output, the field containing the converted data item is padded on the right with four blank characters, such that the total field width is $w + 4$. The effect is similar to the pair of format items F(w,d), X(4).

If this conversion causes either a digit overflow into the sign position or a significant zero digit in the position immediately to the right of the decimal point, the data item is converted according to the rules for floating-point format items, E(w,d,s). For E-conversion, $w = p + 6$, $d = p - 1$, and $s = p$.

Coded Real Floating-Point Binary Data: The data is converted to floating-point decimal with a precision (p) and transmitted as coded real floating-point decimal data.

Coded Complex Data: The data is represented as two immediately adjacent real data fields, the left-hand field being the real part of the data and the right-hand field being the imaginary part of the data.

A sign always precedes the imaginary part. If the value of the imaginary part is greater than, or equal to, zero, the sign is plus; if the value of the imaginary part is less than zero, the sign is minus. The imaginary part is always followed by the letter I.

The field width of the real part of the complex data is as described in the items above, whichever applies. The field width of the imaginary part is as described above plus 2 (one position for the sign and one for the letter I) if $u = 0$, or plus 1 (for the letter I), if $u = 1$. Therefore, the field width of the complex data field is the sum of the widths of the subfields.

NUMERIC FIELD DATA: The base of numeric field data is either decimal or binary.

Numeric Decimal Data: The external format and field width of the numeric decimal data

item is that described by the associated picture specification.

Numeric Binary Data: The external format and field width of the numeric binary data item is that described by the associated picture specification. The binary digits 0 and 1 are represented by the characters 0 and 1.

CHARACTER-STRING DATA: The contents of the character string are written out. Enclosing quotation marks are not supplied, and contained quotation marks are unmodified. The field width is the current length of the string.

BIT-STRING DATA: The format of the data on the external medium is that of a bit-string constant, that is, the value is enclosed in quotation marks and followed by the letter B. The field width is $(p + 3)$, where p is the current length of the string and the three additional positions are for the two quotation marks and the letter B.

Examples of list-directed data specifications:

```
READ LIST (CARD.RATE,  
           DYNAMIC_FLOW);  
READ LIST ((THICKNESS (DISTANCE)  
           DISTANCE = 1 TO 1000));  
WRITE LIST (P, Z, M, R);  
WRITE LIST (A * B / C, (X + Y) ** 2  
           (','));
```

The first two examples are list-directed input specifications and the latter are output specifications.

DATA-DIRECTED DATA SPECIFICATION

General format:

- Option 1
DATA
- Option 2
DATA data-list

Syntax rules:

1. The data list is described in "Data Lists," in this chapter. It may not include formal parameters.
2. Option 1 implies that all of the data items to be transmitted are known to the block containing the READ or GET statement and are declared with a SYMBOL attribute. Option 1 may be used for data-directed input only.
3. Option 2 may be used for both data-directed input and output.

Data-Directed Data On External Medium

The data on the external medium associated with data-directed transmission is in the form of a list of scalar assignments having the following general format:

```
scalar-variable = [+|-] constant  
[,scalar-variable=[+|-] constant] ...;
```

Sterling constants may not be used.

The scalar variable must be unsubscripted and cannot be a pseudo-variable.

Syntax rules:

1. The "scalar variable" may be a subscripted name with decimal integer constant subscripts, but may not be declared with the DEFINED attribute.
2. On input, the scalar assignments may be separated by either a blank or a comma. On output, the assignments are separated by blanks.
3. The terminating semicolon may be replaced by a segment-mark if so specified in the READ or WRITE statement.
4. The constant in the general format above may be a complex expression used to represent a complex constant.

General rules for data-directed input:

1. If the data specification in Option 1 is used, the names on the external medium may be any unqualified name known at the point of transmission and declared with the SYMBOL attribute.
2. If Option 2 is used, each element of the data list must be an unsubscripted scalar, array, or structure name. The names on the external medium must appear in the data list; however, the order of the names need not be the same, and the data list may include names that do not appear on the external medium.

For example, consider the following data list, where A, B, C, and D are names of scalar variables:

```
DATA (B,A,C,D)
```

This data list may be associated with the following input data stream:

```
A = 2.5, B = .00476, D = 125;
```

Note that C appears in the data list but not on the external medium.

3. If the data list in Option 2 includes the name of an array, subscripted references to that array may appear on the external medium. The entire array

need not appear on the external medium.

Let X be the name of a two-dimensional array declared as follows:

```
DECLARE X (2,3);
```

Consider the following data list and input data stream:

<u>Data List</u>	<u>Input Data Stream</u>
DATA (X)	X(1,1) = 7.95, X(2,1) = 8085, X(1,3) = 73;

Although the data list has only the name of the array, the associated input stream may contain values for individual elements of the array.

4. If the data list includes qualified names, then qualified names of identical form may appear on the external medium. Consider the following structures:

```
DECLARE 1 CARDIN, 2 PARTNO, 2 DESCRP  
        SYMBOL(X), 2 PRICE SYMBOL(X1),  
1 CARDOUT, 2 PARTNO, 2 DESCRP  
        SYMBOL(X2), 2 PRICE SYMBOL(X3);
```

If it is desired to read a value for CARDIN.PARTNO, then the data list and input data stream have the following forms:

<u>Data List</u>	<u>Input Data Stream</u>
DATA (CARDIN.PARTNO)	CARDIN.PARTNO = 737314;

5. When a structure name is given in a data list, elementary structure elements, which have been declared with the SYMBOL (identifier) attribute, may be transmitted. In this case, the scalar variable in the assignment list on the external medium is the identifier given in the SYMBOL attribute. The equivalent qualified name may not appear in the assignment list unless the qualified name also appears in the data list.

Consider the structures described in item 4. If it is desired to read values for DESCRP in CARDIN and for PRICE in CARDOUT, the following data list and input stream may be used:

<u>Data List</u>	<u>Input Stream</u>
DATA (CARDIN,CARDOUT)	X = 65, X3 = 38;

Both of these data items may have been transmitted by using identical qualified names in both the data list and input stream as described in item 4, above.

General rules for data-directed output:

1. The elements of the data list may be a scalar name, an array name, a structure name, a repetitive specification involving any of these elements or further repetitive specifications. The data with names appearing in the data list is transmitted, in the form of a list of scalar assignments separated by blanks and terminated by a semicolon.
2. Array names in the data list are treated as a list of the contained subscripted elements in row-major order.

Let X be an array declared as follows:

```
DECLARE X(2,4);
```

Let X appear in a data list as follows:

```
DATA (X)
```

Then, on output, the output data stream is as follows:

```
X(1,1)=1 X(1,2)=2 X(1,3)=3 X(1,4)=4  
X(2,1)=5 X(2,2)=6 X(2,3)=7 X(2,4)=8;
```

3. Subscript expressions in a data name are evaluated and replaced by integer constants.
4. Qualified names appearing in the data list are transmitted with the same qualification, but subscripts follow the qualified name rather than being interleaved. If a data list is specified for a structure element transmitted under data-directed output as follows:

```
DATA (Y(1,3).Q)
```

Then the associated data field in the output stream is as follows:

```
Y.Q(1,3)=3.756;
```

5. Structure names in the data list are interpreted as a list of the contained scalar or array elements, and arrays are treated as above.

Consider the following structure:

```
1 A, 2 B, 2 C, 3 D
```

If a data list for data-directed output is as follows:

```
DATA (A)
```

Then the associated data fields in the output stream are as follows:

```
B=2 D=17;
```

6. Data-directed output is suitable for data-directed input only if it includes no numeric fields of binary base or numeric fields of decimal base that do not have the form of valid arithmetic constants.
7. If a data-directed output list contains an item declared with the attribute, SYMBOL (identifier), the synonym is used rather than the name in the data list.

Data-Directed Output Format

Data fields written under data direction are tabbed, that is, aligned in vertical columns. This tabbing may be overridden by a TAB option in a LAYOUT statement for the appropriate file.

Data items are written on successive tab positions unless an item extends into a succeeding position. (In that case, the next free position is used.) A data item may span several tabs; however, data items must not span records.

Length of Data-Directed Data Fields

The length of the data field on the external medium is a function of the internal precision, the value of the data item being written, and the length of the data identifier and its associated subscript list. The field length for coded arithmetic data, numeric field data, and bit-string data is the same as described for list-directed output (see "Length of List-Directed Output Fields").

For character-string data the contents of the character string are written out enclosed in quotation marks. Each quotation mark contained within the character string is represented by two successive quotation marks.

Example:

Assume that A is declared as a one-dimensional array of six elements; B is a one-dimensional array of seven elements. If it is desired to calculate values, the procedure in Figure 2 calculates and writes out values for $A(I) = B(I+1) + B(I)$.

```

AB:  PROCEDURE;
      Input Stream on External Medium
      DECLARE A (6) , B (7) ;
      READ DATA (B) .
      DO I = 1 TO 6;
      A (I) = B (I+1) + B (I) ;
      Output Stream on External Medium
      END;
      WRITE DATA (A) ;
      END AB;

```

Figure 2. Example of Data-Directed Transmission, Both Input and Output

FORMAT-DIRECTED DATA SPECIFICATION

General format:

data-list format-list

General rules:

1. The data list is described in "Data Lists," the format list in "Format Lists." Only this form of transmission can be used for sterling data.
2. On output, the value of each data item in the data list is converted to a format specified by the associated format item in the format list. The first scalar data item is associated with the first format item, the second scalar data item with the second format item, etc. Suppose the format list effectively contains *j* format items, and the data list effectively contains *k* data items. Then if *j*<*k* after *j* scalar data items have been transmitted, the format list is re-used, the (*j*+1)th scalar item being associated with the first format item, etc. This re-use is performed as many times as required. If *j*>*k*, redundant format items are ignored.
3. An array or a structure in a data list is equivalent to *n* data items, where *n* is the number of scalar elements in the array or structure.
4. If a data list item is associated with a control format item, that control action is executed and the data list item is paired with the next format item.
5. The specified transmission is complete when the last data item has been

processed using the corresponding format item. Subsequent format items, including control format items, are ignored.

Examples:

The first of the following examples is a format-directed input specification, and the second is an output specification:

1. (NAME,DATE,SALARY) (A (COLA_COLB) , X (2) , A (6) , F (M+2,2))
2. ('RESULT(' || I || ')=' , A (I) I=1 TO 20) (A,F (2) ,A,F (8,33))

FORMAT LISTS

The format-directed data specification and the FORMAT and POSITION statements require a format list.

General format of a format list:

$$\left(\begin{matrix} \text{item} \\ n \text{ item} \\ n \text{ format-list} \end{matrix} \right) \left[\begin{matrix} ,\text{item} \\ ,n \text{ item} \\ ,n \text{ format-list} \end{matrix} \right] \dots$$

Syntax rules:

1. Each "item" represents a format item as described below.
2. The letter *n* represents an iteration factor, which is either an expression enclosed in parentheses, or a decimal integer constant. The iteration factor specifies that the associated format item is to be used *n* successive times. A zero or negative iteration

factor specifies that the associated format item is to be skipped and not used.

If an expression is used to represent the iteration factor, it is evaluated and converted to an integer once for each set of iterations. The associated format item is that item or list of items to the right of the iteration factor.

General rule:

There are two types of format items: data format items and control format items. Data format items specify the form of data on an external medium. Control format items specify control over records and groups being read or constructed.

Data Format Items

Data format items describe data representation in two modes: external and internal. The external mode is designed to be readable and uses character representation. The internal mode is in coded form, which is individually defined for each implementation of PL/I, and is primarily used for compact intermediate storage. Arithmetic internal mode format items, other than the numeric picture item (P), specify coded internal form. The P format item specifies that the internal form is numeric field or character string.

External Mode Format Items

The discussion of external mode format items requires the following definitions:

1. The letter *w* represents the length of the field in characters used by the external representation (including signs, decimal or binary points, blanks, editing characters, and the letters E and B, as used in representation of constants).
2. The letter *d* represents the number of positions after the decimal or binary point.
3. The letter *s* represents the number of significant digits (binary or decimal) to appear.
4. The letter *p* represents a scale factor, which may be positive or negative.

The quantities *w*, *d*, *s* may be specified by an expression. When the format item is used, the expression is evaluated and converted to an integer. If $w \leq 0$ in a format

specification, then, on input, the associated list item is skipped, unless it is a string, in which case the data value is taken as the null string. On output, the FIELD_OVERFLOW condition is raised unless the format item is A or B and the associated list item has the null string as its value.

On input, the data item in the external data field is converted to the characteristics of the list item. Rules for the conversion are given in Chapter 3.

There are six format items associated with data in external mode: fixed-point (F), floating-point (E), complex (C), Picture specification (P), character-string (A,P), and general (G).

FIXED-POINT FORMAT ITEMS: Numeric data may be described by a fixed-point format item.

General format:

Option 1.

F (*w*)

Option 2.

F (*w*, *d*)

Option 3.

F (*w*, *d*, *p*)

General rules:

1. On input, the data items in the external data field are the character representation of decimal fixed-point numbers anywhere in a field of width *w*.

In Option 2, if no decimal point appears in the number, it is assumed to appear immediately before the last *d* digits (if trailing blanks are treated as zeros, they are included in the count of *d* digits). If a decimal point does appear, it overrides the *d* specification. Option 1 is treated as option 2, with *d* equal to zero.

In Option 3, the scale factor effectively multiplies the external data value by 10 raised to the value of *p*. If *p* is positive, the number is treated as though the decimal point appeared *p* places to the right of its given position. If *p* is negative, the data is treated as though the decimal point appeared *p* places to the left of its given position. The given position of the decimal point is that indicated either by an actual point, if it is given, or by *d*, in the

- absence of an actual point.
- On output, the external data is a decimal fixed-point number, right-adjusted in a field of width w .

In Option 1, only the integer portion of the number is written; no decimal point appears.

In Option 2, both the integer and fractional parts of the number are written. If d is specified, a decimal point is inserted before the last d digits, and the value is appropriately positioned. Trailing zeros are supplied if the number of fractional digits is less than \underline{d} (where \underline{d} must be less than w).

In Option 3, the scale factor effectively multiplies the internal data value by ten raised to the power of p before it is edited into its external character representation. If d is omitted, only the integer portion of the number is considered.

For all options, if the value of the number is less than zero, a minus sign will be prefixed to the external character representation; if it is greater than or equal to zero, no sign will appear. Therefore, for negative values, w must encompass both sign and decimal point.

FLOATING-POINT FORMAT ITEMS: Numeric data may be described by a floating-point format item.

General format:

E (w, d [, s])

General rules:

- On input, the data item in the external data field is an optionally signed character representation of a decimal floating-point number anywhere within a field of width w .

The external form of the number is as follows:

$$\left[\begin{array}{c} + \\ - \end{array} \right] \text{ fixed-point-} \left[\begin{array}{c} \left\{ \begin{array}{cc} [E] & \pm \\ E & [\pm] \end{array} \right\} \text{ exponent} \\ \text{number} \end{array} \right]$$

- If there is no decimal point in the external data field, the decimal point is assumed to be before the last d digits. If there is a decimal point in the external data field, it overrides the decimal point placement specified by d .
- The "exponent" is a decimal integer. If the exponent and the

preceding E or sign are omitted, a zero exponent is assumed.

- On output, the data item in the external data field has the following general form:

$$[-] s-d \text{ digits.} d \text{ digits E } \left\{ \begin{array}{c} + \\ - \end{array} \right\} \text{ exponent}$$

- The "exponent" is a decimal integer of n digits, where n is defined individually for each implementation. The exponent is adjusted so that the leading digit of the characteristic is nonzero.
- If the above form does not fill the field of width w , it is right-adjusted, and blanks are inserted on the left. If s is omitted it is taken as equal to d . The field width w must be greater than or equal to $(s + n + 3)$ for non-negative values and $(s + n + 4)$ for negative values of the data item.

COMPLEX FORMAT ITEMS: Complex numeric data may be described by a complex format item.

General format:

C (real-format-item
[, real-format-item])

General rules:

- Each "real format item" is specified by F, E, G, I, or P formats. P can specify a numeric field only; it cannot specify a sterling field.
- On input, the external data is the real and imaginary parts of the complex number in adjacent fields described by the two contained format items. If the second real format item is omitted, it is assumed to be the same as the first.
- On output, the form of the real and imaginary parts is specified by enclosed real format items. If the second is omitted, it is assumed to be the same as the first.

PICTURE FORMAT ITEM: Numeric data may be described by a numeric picture using the P format item.

General format:

P 'numeric-picture-specification'

The "numeric picture specification" is described in "The PICTURE Attribute," in Chapter 4.

On input, the picture specification describes the form of the data on the external

medium and how it is to be interpreted numerically. The external representation of binary numeric fields uses the characters 0 and 1.

On output, the value of the list item is edited to the form specified by the picture before it is transmitted. Binary numeric fields will have a character representation after transmission.

BIT-STRING FORMAT ITEMS: The bit-string item describes the external representation of a bit string using the characters 0 and 1.

General format:

A [(w)]

General rules:

1. If w is omitted, it is taken to be the maximum length of the associated data list element on input or the current length on output; it must be specified if conversion is to be performed.
2. On input, the external data is a character representation of bit string anywhere within the field of width w.
3. On output, the character representation of bit string is left-adjusted in the field of width w. Truncation, if necessary, is performed on the right. Blanks are used for padding.

CHARACTER-STRING FORMAT ITEMS: Character data may be described by a character-string format item.

General format:

{ A [(w)]
P 'character-picture-specification' }

General rules:

1. The "character picture specification" is described in "The String Attributes," in Chapter 4.
2. The external representation is a string of w characters.
3. On output, truncation, if necessary, is performed on the right. If the associated list element is too short, it is extended on the right with blanks. If the picture form is used, w is implied. Checking and editing are performed.
4. w can be omitted on output, in which case the associated data list element must be a character string, and w is taken to be the maximum length of that string.

GENERAL FORMAT ITEMS: Both character data and numeric data may be described by the general format item.

General format:

G [(w) | (w,d) | (w,d,s)]

General rules:

1. The type of the external character representation of the data is assumed to be that of the associated data-list element.
2. In the case of strings and numeric fields, the effect of the general format item is identical to A (w); d and s, if specified, are ignored. Coded bit-string external representation may not be described by a general format item.
3. On input for arithmetic data, the scale of the external character representation is deduced. The effect of the general format item is then identical to F (w), F (w,d) for fixed-point numbers and E (w,d,s) for floating-point numbers.

On output for arithmetic data, the data is analyzed with respect to the specified field width w.

4. If the data item may be represented without loss of accuracy as a fixed-point number, the external form is that specified by F (w), or F (w,d) if d is specified. If the data item cannot be suitably represented by an F format item, it is necessary that d be specified in the general format item. The effect is then identical to E (w,d) or E (w,d,s) if s is specified.

Internal Mode Format Items

Internal mode format items may specify precision and length. This is given in exactly the same way as with the precision attribute. The base or the precision is that of the format item, or, where this is indeterminate, that of the associated data item. If size or precision is omitted, it is assumed to be that of the associated data item. The type, base, scale, mode, and precision of a data item may differ from its associated format item. Wherever this occurs, conversion is performed. Precision must be specified in this case.

There are six format items associated with data in internal mode: fixed-point (IF), floating-point (IE), picture (P), bit-string (B), character-string (A,P), and general (IG).

INTERNAL FIXED-POINT ITEMS: Numeric data in internal mode may be described by the internal fixed-point format item.

General format:

$$\left\{ \begin{array}{l} \text{IF } [(\text{precision})] \\ \text{IFB } [(\text{precision})] \end{array} \right\}$$

The first form is used for decimal data, the second form for binary data.

INTERNAL FLOATING-POINT FORMAT ITEMS: Numeric data in internal mode may be described by the internal floating-point format item.

General format:

$$\left\{ \begin{array}{l} \text{IE } [(\text{precision})] \\ \text{IEB } [(\text{precision})] \end{array} \right\}$$

The first form is used for decimal data, the second form for binary data.

INTERNAL PICTURE FORMAT ITEMS: Decimal numeric data in internal mode may be described by an internal picture format item.

General format:

P 'picture-specification'

The "picture specification" is described in "The PICTURE Attribute," in Chapter 4.

INTERNAL BIT-STRING FORMAT ITEMS: The internal bit-string format item describes the internal representation of a bit string. It may also be used for the output of binary numeric fields in internal mode, since the conversion from binary numeric to bit string is a null conversion.

General format:

B [(length)]

General rules:

1. "Length" is the length of the string in bits. If omitted, it is taken as the current length of the associated bit-string list item.
2. The external representation is the coded form for a bit string. If s bits are encoded in one character, the width of the external field may be represented as follows:

$\text{CEIL}(\text{length}/s)$

3. On input, the coded string is interpreted as bit string and is truncated, if necessary, to the specification length.
4. On output, the string is extended with zeros to length $s \cdot \text{CEIL}(\text{length}/s)$ and the external form is this string.

INTERNAL CHARACTER-STRING FORMAT ITEMS: Character data in internal mode may be

described by an internal character-string format item.

General format:

$$\left\{ \begin{array}{l} \text{A } [w] \\ \text{P 'character-picture-specification'} \end{array} \right\}$$

The "character picture specification" is described in "The Picture Attribute," in Chapter 4.

INTERNAL GENERAL FORMAT ITEMS: Both character data and numeric data in internal mode may be described by an internal general format item.

General format:

IG

The IG format item specifies that the format of the data on the external medium is to be identical to its internal form.

Control Format Items

There are two types of control format items: the spacing format item, X, and the positioning format items, which effect transmission in exactly the same way as do the statements of the same names. All of these format items except POSITION, are for use only with data in the external mode.

Spacing Format Item

General format:

X (w)

General rules:

1. On input, the format item specifies that the next w characters of the external data are to be ignored.
2. On output, the format specifies that w characters of blanks are to be inserted into the external data.

Positioning Format Items

The positioning format items are:

SPACE [(expression)]

SKIP [(expression)]

GROUP [(expression)]

POSITION (format list)

TAB [(expression)]

General rules:

1. The effect of each of these format items is identical to the statements of the same names (see the individual statements, described in Chapter 8, for a description of the action taken).
2. Only the POSITION item of the above, may be used in lists intended for internal string editing.

Remote Format Item

If it is desired to locate format items remotely from a format list, the remote format item, R, may be used.

General format:

R (statement-label-designator)

General rules:

1. The "statement-label-designator" is a label constant or a label variable that is the statement label of a FORMAT statement. The FORMAT statement includes a format list that is taken to replace the remote format item.
2. The R format item and the specified FORMAT statement must be internal to the same block.

PROCEDURE-DIRECTED DATA SPECIFICATION

The data specification for procedure-directed transmission has the following general format:

CALL entry-name [(argument [,argument] ...)]

The CALL option causes the procedure whose name is "entry name" to be invoked in the same manner as the CALL statement.

The invoked procedure may perform further action on the data to be transmitted by using GET, PUT, and the positioning statements.

INPUT/OUTPUT STATEMENTS

The input/output statements may be classified as follows: file preparation, data specification, data transmission, positioning, report generation, SAVE and RESTORE

statements, and DISPLAY statements. A description of each statement is given in Chapter 8.

FILE PREPARATION STATEMENTS

The OPEN statement causes certain checking and allocation of facilities in preparation for input/output on a file. The CLOSE statement causes disposition of a file and release of facilities upon completion of input/output. Both statements are optional.

DATA SPECIFICATION STATEMENT

The format of a record to be transmitted may be specified by the FORMAT statement or in the data transmission statements.

DATA TRANSMISSION STATEMENTS

The READ and WRITE statements cause the transmission of data between storage and external media. The GET and PUT statements cause data to be moved between the current record and specified variables in conjunction with procedure-directed data transmission. The DISPLAY statement causes messages to be transmitted between the program and the machine operator.

POSITIONING STATEMENTS

Positioning within and between records may be accomplished with the POSITION, TAB, SKIP, SPACE, and GROUP statements. The first two of these apply only to current files. The remainder may apply either to the current file or an explicitly designated file. The SEGMENT statement is used to position between segments. The REPOSITION statement has a special use with the ON statement (see "The REPOSITION Statement").

REPORT GENERATION STATEMENTS

The PAGE and LAYOUT statements are provided to facilitate preparation of files for printing. The statements may, however, be used for nonprint files. The statements refer explicitly (or in the case of procedure-directed transmission, implic-

itly) to a particular file. Each statement applies to that file until overridden by another statement of the same type. Until such statements are encountered, system standards are assumed to apply.

The execution of a PAGE or LAYOUT statement for a file destroys all options established by previously executed PAGE or LAYOUT statements for the same file. Execution of a CLOSE statement releases the PAGE and LAYOUT for the indicated file.

RECORD IDENTIFICATION OPTIONS

THE KEY OPTION

Function:

The KEY option is used when direct access to a particular record is required. The file containing this record must have the REGIONAL or INDEXED organization (see "File Organization Attributes," in Chapter 4).

General format:

KEY (expression)

General rules:

1. The "expression," which, if necessary, is converted to characters, is the key value used to locate the particular record in the file. This expression is evaluated whenever transmission of another record is required.
2. If the KEY option appears in an output context and if the key already exists within the file, that record is replaced by the record being written; if the key does not exist, the record is added to the file.
3. The ACCESS condition is raised (see "ON-Conditions, in Appendix 3) if the KEY option is used in an input context and the key does not exist.
4. The KEY option may appear only in the READ or WRITE statement.

THE NEWKEY OPTION

Function:

The NEWKEY option serves the same purpose as the KEY option except that the key of the record being written must not already exist within the file. If such a duplicate key exists, an error condition is

raised. When a file is being created with keys, the NEWKEY option must be employed.

General format:

NEWKEY (expression)

General rules:

1. See "The KEY Option" for a discussion of the purpose of the KEY and NEWKEY options.
2. The key of the record is represented by the "expression," which is converted to characters, if necessary. The NEWKEY option can be used in an output context only.
3. The NEWKEY option may appear only in a WRITE statement.

THE REGION OPTION

Function:

The REGION option is used when direct access to a file organized in the REGIONAL mode is required (see "File Organization Attributes," in Chapter 4, for a definition of the REGIONAL file organization).

General format:

REGION (expression)

General rules:

1. The "expression" is converted to integer; this integer n represents the n th region of the file (relative to the beginning of the file) to which the specified input/output device is to be positioned. The value of the expression must be within the limits of the file.
2. The REGION option may be used in conjunction with the KEY option if the region is defined to contain more than one record. In this case, the specified device first is positioned to the specified region and the region is then searched for a record with the specified key value.
3. The REGION option may appear only in the READ or WRITE statement.

This section includes a description of each statement in the language. These descriptions are presented in alphabetic order.

To show the relationships among these statements, they are also classified into logical groups.

RELATIONSHIP OF STATEMENTS

CLASSIFICATION

Statements may be classified into the following logical groups: assignment, control, data declaration, error control and debug, input/output, program structure, sorting, and storage allocation.

Assignment Statement

The assignment statement is used to evaluate expressions and to assign values to scalars, arrays, and structures.

Control Statements

The control statements alter the normal sequential flow of control through a program. The control statements are GO TO, IF, DO, CALL, RETURN, WAIT, STOP, EXIT, DELAY.

Data Declaration Statement

The data declaration statement, DECLARE, specifies attributes for names and identifiers. This statement is described in Chapter 4.

Error Control and Debug Statements

When an interrupt occurs during program execution, standard operating system action is taken; however, the language provides the facility to override system action on

these interrupts. By using the ON statement, a programmer may specify the action to be taken when an interrupt occurs and can record the status of the program at the point of the interrupt. By using the SIGNAL statement, the programmer may initiate programmed interrupts and may simulate machine interrupts to facilitate debugging.

Input/Output Statements

See "Input/Output Statements," in Chapter 7, for a classification and discussion of statements used in input/output operations.

Program Structure Statements

The program structure statements are: PROCEDURE, BEGIN, END, DO, and ENTRY. The first three statements delimit the scope of declarations within a program. The ENTRY statement provides a secondary entry point for a procedure.

Sorting Statement

The SORT statement sorts and, optionally, merges records on a file.

Storage Allocation Statements

The storage allocation statements are: ALLOCATE, FREE, FETCH, and DELETE. The ALLOCATE and FREE statements allocate and free storage for variables. The FETCH and DELETE statements allocate and free storage for programs.

SEQUENCE OF CONTROL

Within a block, control normally passes sequentially from one statement to the next. If a DECLARE, FORMAT, or ENTRY is encountered, control passes to the next statement. If a PROCEDURE statement is

encountered, control passes to the statement following the end of the procedure. Control passes to the statement following an IF statement when control reaches the end of the THEN-unit. Sequential operation is modified by the following statements: CALL, END, EXIT, GO TO, PROCEDURE, RETURN, SIGNAL, and STOP.

A CALL statement passes control to the specified entry point.

An END statement, logically terminating a procedure, acts as a RETURN statement, causing control to return to the invoking procedure.

The EXIT statement causes control to leave a task; the STOP statement causes control to leave a program.

A GO TO statement causes control to transfer to the specified statement label.

A PROCEDURE statement heads a procedure. Procedures may be considered as independent blocks and are placed anywhere within an external procedure, consistent with desired identifier scopes. However, a procedure may be invoked only by a CALL statement, a statement with a CALL option, or a function reference. Thus, control passes around a nested procedure, from the statement before a PROCEDURE statement to the statement after the appropriate END statement for the procedure.

A RETURN statement returns control from a procedure to the invoking procedure.

A SIGNAL statement specifying an enabled condition causes control to pass to the on-unit of the associated ON statement. If there is no associated ON statement, control is passed to the appropriate system routine.

The following conditions may cause sequential operation to be modified:

1. A function reference in any expression causes control to pass to the specified function procedure.
2. The occurrence of an enabled condition specified in an ON statement causes control to pass to the statement or

block contained in the statement. If there is no ON statement, control is passed to the appropriate system routine.

3. The flow of control through the IF and ON statements and through a DO group may or may not be sequential.
4. In an appropriate environment, the asynchronous execution of several operations may involve transfer of control under the influence of external occurrences.

The following example illustrates sequence of control:

```
A: PROCEDURE;
B: X = Y + Z;
C: CALL D;
E: W = P*Q;
  D: PROCEDURE;
  G: S = T/P;
  H: RETURN;
  I: END D;
J: U = V**W;
K: GO TO N;
.
.
.
N: END;
```

Control flows in the following order: A, B, C, D, G, H, E, J, K, N.

PSEUDO-VARIABLES

The following built-in functions (see Appendix 1 for a more complete description) may be used as pseudo-variables on the left side of an equal sign in an assignment statement, or a DO statement, or in a data list in a READ statement or a GET statement. In the definitions below, the item in the data list of a GET or READ statement may be considered to correspond to the item on the left side of the equal sign in an assignment statement; the value being transmitted may be considered to correspond to the expression on the right side.

COMPLEX (a,b) The letters a and b represent variables that need not have the same characteristics. During execution of an assignment statement, the real part of the expression on the right is assigned to a, the imaginary part to b.

REAL (c) The letter c represents a complex variable. During execution of an assignment statement, the real value of the expression is assigned to the real part of c.

IMAG (c) The letter c represents a complex variable. During execution of an

assignment statement, the real value of the expression is assigned to the imaginary part of *c*.

ONCHAR The expression on the right is converted to a character string of length 1. On assignment, the character that caused the input EDIT error interrupt is replaced by the value assigned. This pseudo-variable is defined only while such an interrupt is being processed.

ONFIELD The expression on the right is evaluated and converted to a character string. On assignment, the field that was being processed when the input interrupt occurred is replaced by the value assigned. This pseudo-variable is defined only while an interrupt is being processed.

SUBSTR (*s*,*i* [,*k*]) The letter *s* represents a string. During execution of an assignment statement, the expression is assigned to the substring of *s* defined by the built-in function SUBSTR (see Appendix 1). This substring is always treated as a fixed length string.

EVENT (*v*) The letter *v* represents a scalar event name. When used in an assignment statement, the expression on the right-hand side is evaluated and converted to a bit string of length 1. The value of this bit string is used in an assignment to the named event (see "Asynchronous Operations and Tasks" in Chapter 6).

PRIORITY [*V*] The letter *v* represents a scalar task name. When used in an assignment statement, the expression on the right-hand side is evaluated and converted to FIXED (*m*,*o*) where *m* is implementation defined. The priority of *v*, the named task, is adjusted to be *n*, relative to that of the task in which the assignment is performed, prior to that assignment. If *v* is not specified, this is the task in which the assignment statement is executed (see "Asynchronous Operations and Tasks" in Chapter 6).

UNSPEC (*v*) The letter *v* represents a scalar variable. The expression on the right is evaluated and converted to a bit string (whose length is an implementation defined function of the characteristics of *v*), and assigned to *v* without conversion to the type of *v*.

ALPHABETIC LIST OF STATEMENTS

The ALLOCATE Statement

Function:

The ALLOCATE statement causes storage to be allocated for specified controlled data.

General format:

```
ALLOCATE [level] identifier
          [dimension] [attribute] ...
          [, [level] identifier [dimension]
          [attribute] ...] ...;
```

Syntax rules:

1. Each identifier must represent data of the controlled storage class or be an element of a controlled major structure.
2. "Dimension" indicates a dimension attribute. "Attribute" indicates a BIT, CHARACTER, or INITIAL attribute. "Level" indicates a level number.
3. A dimension attribute, if present, must specify the same number of dimensions as that declared for the associated identifier.
4. The attribute BIT may appear only with a BIT identifier; CHARACTER may appear only with a CHARACTER identifier.
5. A structure or structure element name may appear only if the entire major structure including all level numbers and identifiers appear as in the DECLARE statement for that structure.
6. The length specification may be dropped from BIT or CHARACTER attributes.

General rules:

1. An ALLOCATE statement for an identifier for which storage was allocated and not freed causes storage for the identifier to be "pushed down" or stacked. This pushing down creates a new generation of data for the identifier. When storage for this identifier is freed, using the FREE statement, storage is "popped up" or removed from the stack.
2. Bounds for arrays and lengths of strings are fixed at the execution of an ALLOCATE statement.
 - a. If a bound or length is explicitly specified in an ALLOCATE statement, that bound or length overrides any bound or length given in the DECLARE statement.
 - b. If a bound or length is specified by an asterisk in an ALLOCATE statement, that bound or length is taken from the most recent genera-

tion of data for the identifier. In case no such generation exists, the bound or length is undefined.

c. If a bound or length is not specified in an ALLOCATE statement, it must be explicitly specified in the DECLARE statement. The scope of this declaration must include the ALLOCATE statement. The expression from the DECLARE statement is evaluated at the point of allocation.

3. Upon allocation of an identifier, initial values are assigned to it if the identifier has an INITIAL attribute in either the ALLOCATE statement or DECLARE statement. Expressions or a CALL option in the INITIAL attribute are executed at the point of allocation. If an INITIAL attribute appears in both DECLARE and ALLOCATE statements, only the INITIAL attribute in the ALLOCATE statement is used.
4. To determine whether or not storage has been allocated for an identifier the built-in function ALLOCATION may be used.
5. A parameter that is declared CONTROLLED may be specified in an ALLOCATE statement if the associated argument is given the CONTROLLED attribute. (See "Relationship of Arguments and Parameters," in Chapter 10).

Examples:

1. The following examples illustrate the use of the ALLOCATE statement when the DECLARE statement contains explicit bounds for an array A:

```
DECLARE A (N1,N2) CONTROLLED;
The values of N1 and N2 are assumed to
be known at this point.
```

```
N1, N2 = 10;
ALLOCATE A;           The bounds are 10 and
                      10
```

```
ALLOCATE A           The bounds are K1 and
  (K1,K2);           K2 which override N1
                      and N2.
```

```
N1 = N1 + 1;
ALLOCATE A;           The bounds are 11 and
                      10.
```

```
ALLOCATE A           The bounds are 11 and
  (*, *);           10.
```

```
ALLOCATE A           The bounds are J1 and
  (J1, J2);         J2.
N1, N2 = 20;
```

2. The following example illustrates the use of the ALLOCATE statement when the DECLARE statement contains asterisks for the length of a bit string B:

```
DECLARE B BIT (*) VARYING CONTROLLED;
ALLOCATE B         Illegal; violates rule
  BIT (*);         2b.
```

```
ALLOCATE B;        Illegal; violates rule
                   2b.
ALLOCATE B         The length is N.
  BIT (N);
ALLOCATE B CHAR-   Illegal; violates syn-
  ACTER (4);       tax rule 4.
ALLOCATE B         The length is 8.
  BIT (8);
```

4. The following example illustrates the use of the built-in function ALLOCATION and of the INITIAL attribute for an identifier in an ALLOCATE statement:

```
DECLARE A (N,N) CONTROLLED INITIAL
  ((N*N) 0);
.
.
.
IF 1 ALLOCATION (A) THEN ALLOCATE A
  INITIAL (1, (N-1) ((N) 0, 1));
.
.
.
ALLOCATE A;
```

The Assignment Statement

Function:

The assignment statement is used to evaluate expressions and to assign values to scalars, arrays, and structures.

General format:

Option 1. (Scalar Assignment)

$$\left\{ \begin{array}{l} \text{scalar-} \\ \text{variable} \\ \text{pseudo-} \\ \text{variable} \end{array} \right\} \left[\begin{array}{l} , \text{ scalar-} \\ \text{variable} \\ , \text{ pseudo-} \\ \text{variable} \end{array} \right] \dots = \text{scalar-} \\ \text{expression};$$

Option 2. (Array Assignment)

$$\left\{ \begin{array}{l} \text{array} \\ \text{pseudo-array} \end{array} \right\} \left[\begin{array}{l} , \text{ array} \\ , \text{ pseudo-array} \end{array} \right] \dots \\ = \text{array-expression [BY NAME]};$$

Option 3. (Structure Assignment)

$$\left\{ \begin{array}{l} \text{structure} \\ \text{pseudo-structure} \end{array} \right\} \\ = \text{structure-expression [, BY NAME]};$$

Option 4. (Statement Label Assignment)

$$\text{scalar-label-variable} \\ [, \text{scalar-label-variable}] \dots = \\ \left\{ \begin{array}{l} \text{label-constant;} \\ \text{scalar-label-variable;} \end{array} \right\}$$

array-label-variable [,array-label-variable]...=
 { label-constant;
 scalar-label-variable;
 array-label-variable; }

Syntax rules:

1. In Option 1, each variable on the left of the equal sign may be of arithmetic, bit, or character data type.
2. In Option 2, each array referred to on the left of the equal sign may be an array variable name or a pseudo-array. If the BY NAME option is present, those arrays must be arrays of structures. A pseudo-array is a pseudo-variable with array arguments whose arguments are array variable names. (In the case of the pseudo-variable SUBSTR (s,i,k), this requirement applies only to the argument s; see "Pseudo-Variables.")

All of the arrays on the left and the arrays in the array expression must have the same number of dimensions and identical dimension bounds.

3. In Option 3, in the absence of the BY NAME option, the structure indicated on the left must have structuring identical to the structures indicated in the structure expression. Actual level numbers of the structures involved need not be the same; only the structuring described need be the same.

General rules:

1. The assignment statement is evaluated as follows:
 - a. In Options 1 and 4, if any expressions appear on the left of the equal sign, either in subscripts or in pseudo-variables, these expressions are evaluated exactly once from left to right. The expression on the right of the equal sign is evaluated. The value of the expression on the right of the equal sign is assigned to the variables on the left of the equal sign, from left to right.
 - b. In Options 2 and 3, the assignment statement is treated as if it were a sequence of scalar assignment statements applied on an element-by-element basis. See Rules 3 and 4 below for a discussion of the evaluation of a structure or array assignment BY NAME.
 - c. In the following definition of order of assignment, A is an array of dimensionality n:

L1: DO I1 = LBOUND (A,1) TO HBOUND (A,1) ;

```
DO I2 = LBOUND (A,2) TO HBOUND (A,2) ;
.
.
DO In = LBOUND (A,n) TO HBOUND (A,n) ;
A(I1, I2,...,In) = array-expression;
```

Subscripts (I1,..., In) are inserted for the appropriate arrays on the righthand side, thus yielding a sequence of scalar assignments.

The result of the evaluation for a later position in an array or structure may be affected by the evaluation and assignment to an earlier position (see Example 1, below).

- d. When necessary, the expression value, or values, is converted to the characteristics of the variable on the left according to the rules in "Expressions," in Chapter 3.
2. When a variable on the left is a bit or character string or the UNSPEC pseudo-variable, the expression is evaluated as above, and the assignment is performed from left to right, starting with the leftmost position.
 - a. If the string has a fixed length and the value of the expression is longer than the string, the value is truncated at the right.
 - b. If the string has a fixed length and the value of the expression is shorter than the string, the value is extended on the right with zeros for bit strings or with blanks for character strings.
 - c. If the string has a varying length and the value of the expression is longer than the maximum length of the string, the value is truncated; the assigned string is of the maximum length.
 - d. If the string has a varying length and the value of the expression is shorter than the maximum length of the string, the value is assigned; the new length of the string is the length of the value.
 - e. If the variable on the left is the pseudo-variable SUBSTR with an argument that is a varying-length string, the assignment is performed to this substring in precisely the same way as it would be if the argument were of fixed length, where this fixed length is the current length of the string.
3. If the BY NAME option is used for arrays of structures in option 3, the assignment statement is treated as a sequence of BY NAME structure assignments applied on an element-by-element basis.

4. If the BY NAME option is used in Option 3, the assignment statement causes the following activity:
- Subscript expressions on the left are evaluated.
 - The names of all contained arrays and scalars are extracted from each structure operand appearing on both left and right.
 - These names are qualified by all the minor structure names that contain them, up to but not including the structure names specified in the structure expression BY NAME.

For example, suppose there are three structures:

```

1 ONE                1 TWO
2 PART1              2 PART1
3 RED                3 RED
3 WHITE              3 GREEN
3 BLUE               3 WHITE
2 PART2              2 PART2
3 GREEN              3 BLUE
3 YELLOW             3 YELLOW
3 ORANGE (3)         3ORANGE (3)
2 PART3
3 BLACK
3 WHITE

1 THREE
3 PART1
5 BLACK
5 WHITE
5 RED
3 PART2
5 YELLOW
5 WHITE
5 ORANGE (3)
5 PURPLE

```

Note that the structures contain array names.

The elements of these structures are as follows:

```

Structure ONE      Structure TWO
PART1.RED          PART1.RED
PART1.WHITE        PART1.GREEN
PART1.BLUE         PART1.WHITE
PART2.GREEN        PART2.BLUE
PART2.YELLOW       PART2.YELLOW
PART2.ORANGE       PART2.ORANGE
PART3.BLACK
PART3.WHITE

Structure THREE
PART1.BLACK
PART1.WHITE
PART1.RED
PART2.YELLOW
PART2.WHITE
PART2.ORANGE
PART2.PURPLE

```

- The largest subset of qualified names is selected such that each selected name is contained in all structures involved in the assignment statement.

From the above example, this produces:

```

PART1.RED
PART1.WHITE
PART2.YELLOW
PART2.ORANGE

```

- All expressions involving the selected names are evaluated and values from the right are assigned to items on the left for identical qualified names. The order of the left hand structure is used. These assignments must be valid; for example, arrays may not be assigned to arrays of different dimensions or bounds.

For example, the statement ONE=TWO-2 * THREE, BY NAME is then equivalent to:

```
ONE.PART1.RED =TWO.PART1.RED-2*THREE.
                PART1.RED
```

```
ONE.PART1.WHITE =TWO.PART1.WHITE-2*THREE.
                PART1.WHITE;
```

```
ONE.PART2.YELLOW=TWO.PART2.YELLOW-2*THREE.
                PART2.YELLOW;
```

```
ONE.PART2.ORANGE=PART2.ORANGE-2*THREE.
                PART2.ORANGE;
```

- In BY NAME structure assignment, it is unnecessary for the structuring of all participating structures to be identical. Names of variables that are defined on structures appearing in BY NAME assignment take no part in name matching (see "The DEFINED Attribute").

In Option 4, the value of the label constant or the label variable is qualified by an identification of the current invocation of the block containing the label and by the current task.

This qualification information is used when a GO TO statement specifies the label variable to make the identified invocation current and to check that control does not cross task boundaries.

Examples:

- The following example illustrates array assignment (Option 2):

Given the array A

2	4
3	6
1	7
4	8

```
C=COMPLEX (U,V) +REAL (Q) ;
U1=REAL (C) ;
U2=IMAG (C) ;
```

and the array B

1	5
7	8
3	4
6	3

4. The following example illustrates structure assignment (Option 3):

```
DECLARE 1X, 2Y, 2Z, 2R, 3S, 3P, 1A,
        2B, 2C, 2D, 3E, 3Q;
X = X*A;
```

Consider the assignment statement:

```
A = (A+B) **2 - A(1,1) ;
```

The second statement is equivalent to the following statements:

After execution, A has the value

7	74
93	189
9	114
93	114

```
Y = Y*B;
Z = Z*C;
S = S*E;
P = P*Q;
```

Note that the new value for A(1,1), which is 7, is used in evaluating the expression for all other elements.

2. The following example illustrates string assignment:

Given:

A is a fixed-length string whose value is 'XZ/BQ'.
 B is a varying-length string of maximum length 8 whose value is 'MAFY'.
 C is a fixed-length string of length 3.
 D is a varying-length string of maximum length 5.

5. The following example illustrates statement label assignment (Option 4):

```
DECLARE P LABEL;
P = A;
GO TO P;
.
.
.
A: X = Y**2;
```

This set of statements causes control to transfer to A when the GO TO P statement is executed.

6. The example below illustrates assignment to an array of structures (Options 2 and 3).

Then in the statement:

```
C=A, the value of C is 'XZ/'.
C='X', the value of C is 'Xbb'.
D=B, the value of D is 'MAFY'.
D=SUBSTR(A,2,3) || SUBSTR(A,2,3),
the value of D is 'Z/BZ/'.
SUBSTR(A,2,4)=B, the value of A is
'XMAFY'.
SUBSTR(B,2,2)='R', the value of B
is 'MRbY'.
SUBSTR(B,2)='R', the value of B is
'MRbb'.
```

In the following statement, A is an array of structures, and R is a structure:

```
DECLARE 1A(2,2), 2B, 2C, 2D, 3E, 3F,
        1R, 3S, 3T, 3U, 5V, 5W;
```

The following is an array assignment statement:

```
A=R;
```

The above assignment statement is equivalent to the following four structure assignment statements:

```
A(1,1)=R;
A(1,2)=R;
A(2,1)=R;
A(2,2)=R;
```

The four statements above are, in turn, equal to the following:

```
A(1,1) . B, A(1,2) . B, A(2,1) . B,
A(2,2) . B=S;
```

```
A(1,1) . C, A(1,2) . C, A(2,1) . C, A(2,2) .
C = T;
```

3. The following examples illustrate scalar assignment (Option 1):

a. A,B,C = A+SIN(B) + C**2; provided X has the characteristics of the expression, this is the same as
 X = A+SIN(B) + C**2;
 A = X;
 B = X;
 C = X;

b. COMPLEX (U1, V1) = COMPLEX (U, V) + REAL (Q) ;

This is the same as

```
A (1,1) .E, A (1,2) .E, A (2,1) ,E, A (2,2) .
E = V;

A (1,1) .F, A (1,2) .F, A (2,1) .F, A (2,2) .F
= W;
```

(If R is ABNORMAL, sixteen statements are actually generated.)

7. The following example illustrates conversion of data defined by a picture description assigned to floating-point data, and vice versa:

```
DECLARE A FLOAT, B PICTURE '999V99';

A = B; (B is converted from fixed-
point to floating-point.)

B = A; (A is converted from floating-
point to fixed-point.)
```

The BEGIN Statement

Function:

The BEGIN statement is the heading statement of a begin block.

General format:

```
BEGIN;
```

General rules:

1. A BEGIN statement is used in conjunction with an END statement.
2. See Chapter 1 for a discussion of blocks.

Examples:

1. ON OVERFLOW BEGIN;


```
      .
      .
      .
      END;
```
2. (SIZE) : PROCEDURE;


```
      .
      .
      .
      (NOSIZE) : A: BEGIN;
      .
      .
      .
      END;
      .
      .
      .
      END;
```

The SIZE condition is enabled with the prefix to the PROCEDURE statement. This enabling is negated throughout the begin

block with the prefix NOSIZE. On exit from the begin block, SIZE errors are again enabled because statements again are in the scope of the SIZE prefix.

The CALL Statement

Function:

The CALL statement invokes a procedure and causes control to be transferred to a specified entry point of the procedure.

General format:

```
CALL { entry-name
      } (scalar-expression)
      [(argument [,argument] . . .)]
      [TASK-option] [EVENT-option]
      [PRIORITY-option];
```

where the three options have the format:

```
TASK (scalar-task-name)
EVENT (scalar-event-name)
PRIORITY (expression)
```

Syntax rules:

1. The entry name or the value of the scalar expression represents the entry point of the procedure invoked. When necessary, the value of the expression is converted to a character string.
2. Each argument may be any of the following: any type of expression, a statement label constant, a statement label variable, a statement label array, an entry name, an entry parameter, a file name, a file parameter, a task name, a task parameter, an event name, or an event parameter.
3. The TASK, EVENT, and PRIORITY options can appear in any order, and are separated from each other and the initial part of the CALL statement by a blank.
4. The scalar event and task names may be subscripted references to event or task arrays.

General rules:

1. When the scalar expression is used to designate the entry point and invoked procedure, the scalar expression is evaluated to give a character string, whose length is implementation-defined. This string specifies a program name that must appear in an active FETCH statement logically prior to the call.

- An active `FETCH` statement is one whose function has not been voided by a subsequent `DELETE` statement.
- When the procedure name is represented by a scalar expression, no conversion is performed for the arguments (see "Relationship of Arguments and Parameters," in Chapter 10), and the arguments may not be entry names, statement label designators, or built-in function names.
 - The `TASK`, `EVENT`, and `PRIORITY` options, when used alone or in any combination, specify that the invoked and invoking procedures are to be executed asynchronously. Note that if either the `EVENT` option or the `PRIORITY` option, or both, are used without the `TASK` option, the created task will have no name (see "Asynchronous Operations and Tasks" in Chapter 6).
 - When the `TASK` option is used, the task name is associated with the task created by the `CALL`. Reference to this name enables the priority of the task to be controlled.
 - When the `EVENT` option is used, the event name is associated with the completion of the task created by the `CALL` statement. Another task can then wait for completion of this created task by specifying the event name in a `WAIT` statement. The value of the completion status for the event name (i.e., the value of `EVENT` (event name)) is set to '0'B on execution of the `CALL` statement and to '1'B on completion of the created task. (see "Event Names" in Chapter 2 and "The `WAIT` Statement" in this chapter.)
 - If the `PRIORITY` option is used, the expression in the above form is evaluated when the `CALL` statement is executed. The result of this evaluation is converted to `FIXED` (m,o) where m is implementation defined. The priority of the named task is then made m relative to the task in which the `CALL` is executed.
 - See "Relationship of Arguments and Parameters" for a detailed description of the interaction of `CALL` arguments and invoked entry parameters.

Examples:

- ```
CALL CRITICAL_PATH (A,B*C,D);
.
.
.
CRITICAL_PATH: PROCEDURE (ALPHA,BETA,
 GAMMA) FLOAT;
.
.
.
END;
```
- `FETCH (A||B);`

- ```
.
.
.
CALL (A||B) (C,D,E);
3. CALL PAYROLL (NAME, DATE, HRRATE);
4. CALL PRINT (A,B) TASK (T2) EVENT (ET2)
    PRIORITY (-2);
```

The CLOSE Statement

Function:

The `CLOSE` statement releases facilities that were allocated during the opening of a file and causes proper disposition of the file.

General format is shown in Figure 3.

Following is the format of "ident":

```
IDENT { data-list format-list
      { CALL entry-name [(argument)
                        [,argument]]...}
```

Syntax rules:

- The "file name" is the name of a file to be closed.
- Each file name is separated from its option by a blank.
- An option that is common to two or more file names may be factored in the same way that attributes in a `DECLARE` statement may be factored (see "Factoring of Attributes" in Chapter 4).

General rules:

- The `CLOSE` statement causes certain actions to be performed on the file, whose name is one of the file names of the `CLOSE` statement. The file is repositioned to its logical beginning, and the facilities allocated to it are released.

If a `CLOSE` statement is encountered and the file has not been opened, or has already been closed, the statement is ignored.

If, however, the file is not closed by a `CLOSE` statement, the file is closed at the completion of the task in which the file was opened.

- The `IDENT` option specifying a data list and format list yields a character string that is compared with the file label for an input file or is written as the file label for an




Figure 3. General Format for the CLOSE Statement

output file. For an INOUT file, INPUT is assumed. The data list and format list are described in Chapter 7.

The IDENT option specifying an entry name and argument list causes the specified file to become the current file and the designated procedure to be invoked for reading or writing the file label through GET or PUT. The form of file labels is implementation defined.

If the IDENT option is not specified, no special label operations are performed.

Examples:

1. CLOSE MASTER;

The file, MASTER, is closed, and the facilities allocated to it are released. The file is repositioned to its logical beginning.

2. CLOSE TABLEA, TABLEB, TABLEC;

The three files, TABLEA, TABLEB, and TABLEC, are closed in the same way as MASTER, in the preceding example.

The DECLARE Statement

See "The DECLARE Statement", in Chapter 4, for a discussion of the DECLARE statement.

The DELAY Statement

Function:

The DELAY statement causes execution of the controlling task to be suspended for a specified period of time.

General format:

DELAY (scalar-expression);

General rule:

Execution of the DELAY statement caus-

es the scalar expression to be evaluated and converted to an integer n and execution to be suspended for n milliseconds.

Execution resumes after n milliseconds only if the controlling task is of sufficiently high priority to be selected in preference to all other ready tasks.

Example:

DELAY (10);

Execution of the controlling task is suspended for ten milliseconds.

The DELETE Statement

Function:

The DELETE statement causes a specified program to be made inaccessible.

General format:

DELETE (scalar-expression);

General rules:

1. The scalar expression is evaluated and, where necessary, converted to a character string, whose length is implementation-defined. This string represents the name of the program to be deleted.

The DELETE statement makes the specified program inaccessible and also deletes the STATIC data areas of the deleted program.

2. The specified program must have appeared in a previously executed FETCH statement.
3. After execution of a DELETE statement, the program name may not be specified in a CALL statement before it appears in another FETCH statement.

Examples:

1. DELETE ('PROCTL');
2. DELETE (A||B);

The DISPLAY Statement

Function:

The DISPLAY statement causes a message to be displayed to the machine operator. A response may be requested.

General format:

Option 1.

```
DISPLAY (scalar-expression) [task]
```

Option 2.

```
DISPLAY (scalar-expression)      REPLY  
(character-variable)
```

General rules:

1. Execution of the DISPLAY statement causes the scalar expression to be evaluated and, where necessary, converted to a character string. This character string is the message to be displayed.
2. In Option 2, the character variable receives a string that is a message to be supplied by the operator.
3. In option 2, execution of the program is suspended until the operator's message is received. In option 1, execution continues uninterrupted.

Example:

```
DISPLAY ('END OF JOB');
```

This statement causes the message, "END OF JOB" to be displayed.

The DO Statement

Function:

The DO statement delimits the start of a DO group (see "Groups") and may specify iteration of the statements within the group.

General format is shown in Figure 4.

Syntax rules:

1. The "variable" in Option 3 is a subscripted or unsubscripted scalar variable. Label variables, string variables, complex variables are allowed provided the expansions given below result in valid PL/I programs.
2. Each "expression" in the specification list is a scalar expression.
3. If BY expression3 is omitted from the specification list, expression3 is assumed to be one (1).
4. If TO expression2 is omitted from the specification list, the iteration is performed indefinitely until terminated by some other statement within the scope of the DO or the WHILE clause.
5. If both TO expression2 and BY expression3 are omitted, this form of the specification list implies a single execution of the DO group with the

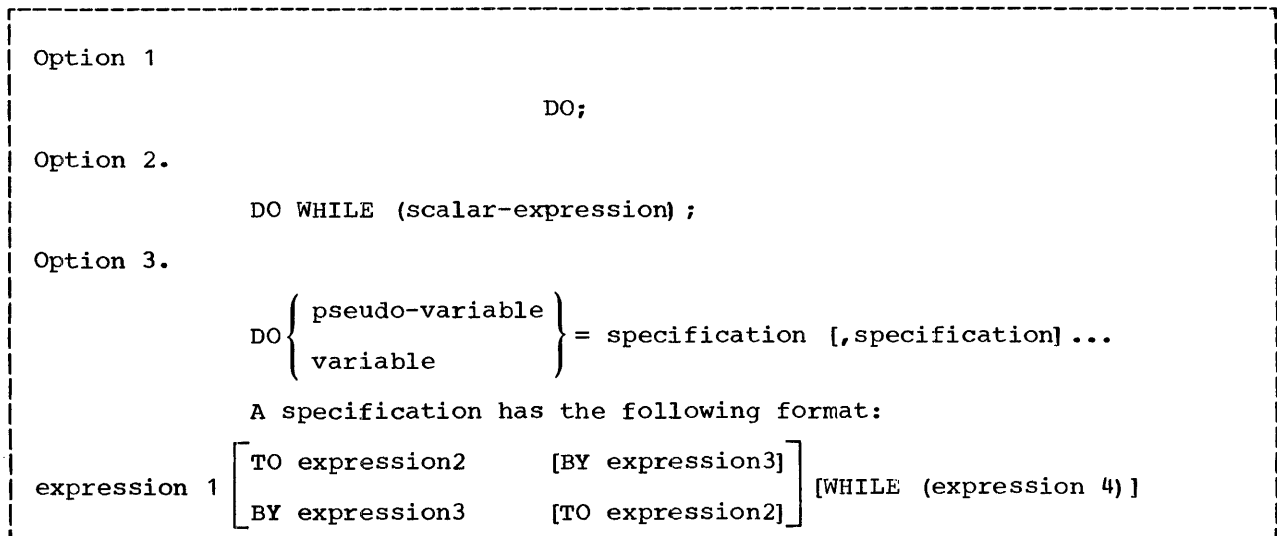


Figure 4. General Format for the DO Statement

control variable having the value of expression 1.

General rules:

1. In Option 1, the DO statement delimits the start of a DO group.
2. In Option 2, the DO statement delimits the start of a DO group and specifies an iteration defined by the following expansion:

```
LABEL: DO WHILE (expression) ;
        statement 1
        .
        .
        .
        statement n
        END;
NEXT: statement
```

The above expansion is exactly equivalent to the following expansion:

```
LABEL: IF (expression) THEN; ELSE GO TO
        NEXT;
        statement 1
        .
        .
        .
        statement n
        GO TO LABEL;
NEXT: statement
```

3. In Option 3, the DO statement delimits the start of a DO group and specifies controlled iteration defined by the following expansion:

```
LABEL: DO variable = expression1
        TO expression2 BY expression3
        WHILE expression4;
        statement 1
        .
        .
        .
        statement n
        END;
NEXT: statement
```

The above expansion is exactly equivalent to the following expansion:

```
LABEL:t1=sexp1; t2=sexp2;...; tm=sexpm;
        e1=expression1; e2=expression2;
        e3=expression3;
        v=e1;
LABEL1: IF (e3>=0) & (v>e2) | (e3<0) & (v<e2)
        THEN GO TO NEXT;
IF (expression 4) THEN; ELSE GO TO NEXT;
        statement 1
        .
        .
        .
        statement n
        v=v+e3;
        GO TO LABEL1;
NEXT: statement
```

In this expansion $s_{exp1}, \dots, s_{expm}$ are the expressions which appear in subscripts of the controlled variable or pseudo-variable, followed by the second and third argument positions if the SUBSTR pseudo-variable is being used. The letter v denotes the controlled variable with all exp_i replaced by t_i . In the simplest cases, $m=0$ and the first statement is $e1=expression1$. The variables t_1, \dots, t_m , are BINARY FIXED integer variables of default precision, inserted by the compiler. The variables e_1, e_2 , and e_3 have the characteristics of the corresponding expressions.

- a. If more than one specification is given, the statement labeled NEXT refers to the initialization for the next specification; for example:

```
NEXT: e5 = expression 5;
```

The t_i variables are computed only once in each DO statement.

- b. If the WHILE clause is omitted, the IF statement involving expression4 is replaced by a null statement.
 - c. If the TO clause is omitted, the IF statement and the assignment statement involving e_2 are omitted.
 - d. If both the TO clause and the BY clause are omitted, all statements involving e_2 and e_3 are omitted as well as the statement "GO TO LABEL1;".
4. The WHILE clause in Options 2 and 3 specifies that before each associated execution of the DO group, the expression is evaluated and, if necessary, converted to give a bit-string value. If any bit in the resulting string has the value '1', the iterations continue uninterrupted. If all bits have the value '0', the iterations associated with the current specification are terminated.
 5. In the specification list, in Option 3, expression1 represents the starting value of the control variable. Expression3 represents the increment to be added to the control variable after each iteration of the statements in the DO group. Expression2 represents the terminating value of the control variable. Iteration terminates as soon as the value of the control variable passes its terminating value. When the last specification is completed, control passes to the statement following the DO group.
 6. Control may transfer into a DO group from outside the DO group only if the DO group is delimited by the DO state-

ment in Option 1; that is, iteration is not specified.

Examples:

1. DO INDEX = CTR WHILE (A>B), 5 TO 10
WHILE (A = B), 100;
2. DO I = J TO K BY I, I+1 TO N BY 1;
3. DO WHILE (P);
4. DO;
5. DO WHILE (TAX-DEDCT < ESTTAX * 4);
6. DO COMPLEX(X,Y) = 0 BY 1+1I WHILE
(X<10);

The END Statement

Function:

The END statement terminates blocks and groups.

General format:

END [label];

General rules:

1. If a label follows END, the END statement terminates that group or block having that label.
2. If a label does not follow END, the END statement terminates that group or block headed by the nearest preceding DO, BEGIN, or PROCEDURE statement for which there is no other corresponding END statement.
3. An END statement may be used to terminate more than one group or block (see "Use of the END Statement," in Chapter 1).
4. The END statement may have a label preceding the END. This label may be referred to anywhere in the program where the label is known.
5. If control reaches an END statement, terminating a procedure, it is treated as a RETURN statement.
6. If control reaches an END statement which terminates a BEGIN block that is an on-unit, either control is returned to the point following the interrupt location, or an appropriate system action is taken.

For examples, see "Use of the END Statement," in Chapter 1.

The ENTRY Statement

Function:

The ENTRY statement specifies a secondary entry point to a procedure.

General format:

entry-name: ... ENTRY [(parameter
[,parameter] ...)]
[data-attributes];

General rules:

1. The parameters are names that specify the parameters of the entry point. When the procedure is invoked, a relationship is established between the arguments of the invocation and the parameters of the invoked entry point (see "Relationship of Arguments and Parameters").
2. The data attributes specify the characteristics of the value returned by the procedure when invoked as a function at this entry point. The value specified in the RETURN statement of the invoked entry is converted to the specified data attribute.

If insufficient data attributes are specified at the entry point, default attributes are applied, as determined by the name of the entry point.

If an ENTRY statement has more than one label, each label is interpreted as if it were a single entry name for a separate ENTRY statement.

Consider the statement:

A:I: ENTRY;

This statement is equivalent to:

A: ENTRY;
I: ENTRY;

The ENTRY statement must be internal to the procedure block for which it defines a secondary entry point. The ENTRY statement may not be internal to any block contained in this procedure; nor may it be within a DO group that specifies iteration.

Example:

NAME: PROCEDURE (N) CHARACTER (15);
DECLARE TABLE (100) CHARACTER (15)
EXTERNAL;

INITIAL: ENTRY (N) CHARACTER (1);
RETURN (TABLE (N));
END;

The EXIT Statement

Function:

The EXIT statement causes immediate termination of the task that contains the statement and all tasks attached by this task (see "Asynchronous Operations and Tasks," in Chapter 6). If the EXIT statement is executed in a major task, it is equivalent to a STOP statement (see this chapter).

General format:

```
EXIT;
```

The FETCH Statement

Function:

The FETCH statement causes a program to be fetched and made available for invocation by a CALL statement, with the entry name specified by an expression.

General format:

```
FETCH (scalar-expression);
```

General rules:

1. On execution of the FETCH statement, the scalar expression is evaluated and, where necessary, converted to a character string whose length is implementation-defined. This string specifies the name of a program to be fetched.

It is assumed that the specified program was not available before the FETCH.

2. After execution of the FETCH statement, the fetched program may be invoked by a CALL statement, with the entry name specified by an expression (see "The CALL Statement").
3. Data declared EXTERNAL, task identifiers, and file names may be shared only among procedures within a program. Consequently, any program which is made available by a FETCH statement may not share externals with any other program.
4. Initial values for data in static storage are established at the time of fetching.

Examples:

1. FETCH ('PROCTL');
2. FETCH ('PROG' || BETA);

```
CALL ('PROG' || BETA) (ALPHA);
```

The FORMAT Statement

Function:

The FORMAT statement specifies a format list for use with data transmitted under format direction.

General format:

```
label:...FORMAT format-list;
```

Syntax rules:

1. The "format list" is as described for use with a format-directed data specification (see "Format Lists" in Chapter 7).
2. At least one "label" is required. It is the name of a statement label appearing in a remote format item.

General rules:

1. A READ, WRITE, GET, or PUT statement, or an IDENT option may include a remote format specification, R, in the format list of a format-directed data specification. That portion of the format list covered by the R (statement label designator) format item must be specified in a FORMAT statement with a corresponding statement label.
2. The remote format item and the FORMAT statement must be internal to the same block.

The FREE Statement

Function:

The FREE statement causes the storage most recently allocated for specified controlled variables to be freed. The next most recent allocation is made available, and subsequent references to the identifier refer to that allocation.

General format:

```
FREE identifier [,identifier] ...;
```

Syntax rule:

Each identifier is a scalar, array, or major structure name of the controlled storage class.

General rules:

1. Controlled storage allocated in a task cannot be freed by a task which it attaches.
2. If a specified identifier has no allocated storage at the time the FREE statement is executed, no action is taken.

Examples:

1. FREE X,Y,Z;
2. The following excerpt from a procedure illustrates the FREE statement in conjunction with an ALLOCATE statement:

```

DECLARE A(100) INITIAL ((100) 0)
        CONTROLLED, C(100), X(100);
        .
        .
        .
ALLOCATE A;
        .
        .
        .
C=A;
        .
        .
        .
FREE A;
        .
        .
        .
X=SIN(C**2 + X/Y);

```

The GET Statement

Function:

The GET statement causes data to be fetched from the current record, converted from external data form, if necessary, and assigned to variables as specified. The GET statement has meaning only when there is a current file; if there is no current file, the ERROR condition is raised.

General format:

```
GET data-specification [, data-specification]...;
```

General rules:

1. The "data specifications" are discussed in Chapter 7. Only those forms specified for input may be used; a CALL option may not be used.
2. As data is fetched from the record, the action that occurs is as if a pointer moved across the record as demanded by the data specifications. This pointer may be repositioned within the record by use of the POSITION statement or the REPOSITION statement.

Example:

```

READAB: PROCEDURE;
        READ (A,B) (2F(7,3)), CALL GETC;
        .
        .
        .
        END READAB;

GETC:   PROCEDURE;
        GET (C) (G(8,5));
        .
        .
        .
        END GETC;

```

The GO TO Statement

Function:

The GO TO statement causes control to be transferred to the specified statement.

General format:

$$\left\{ \begin{array}{l} \text{GO TO} \\ \text{GOTO} \end{array} \right\} \left\{ \begin{array}{l} \text{label-constant} \\ \text{scalar-label-variable} \end{array} \right\};$$

General rules:

1. If a label variable is specified, the GO TO statement has the effect of a multi-way switch. The value of the label variable is the label of the statement to which control is transferred.

Since the label variable may have different values at each execution of the GO TO statement, control may not always pass to the same statement. (Example 2 illustrates a GO TO statement used as a multi-way switch.)

2. A GO TO statement may not pass control to an inactive block (see "Activation and Termination of Blocks," in Chapter 6, for a discussion of active and inactive blocks).

A GO TO statement may not transfer control from outside a DO group to a statement inside the DO group if the DO group specifies iteration unless the GO TO terminates a procedure invoked from within the DO group.

3. A GO TO statement that transfers control from one block (D) to a dynamically encompassing block (A) has the effect of terminating block D, as well as all other blocks that are dynamically descendant from block A. Conditions are reinstated, and automatic variables are freed in the same way as if the blocks terminated normally. When a GO

TO statement transfers control out of a procedure invoked as a function, the evaluation of the expression that contained the corresponding function reference is discontinued, and control is transferred to the specified statement.

4. A GO TO may not terminate any procedure invoked within an input/output statement, unless the GO TO is encountered in an ON unit.
5. A GO TO may not terminate any procedure invoked during a prologue (see "Prologues" in Chapter 10), or an ALLOCATE statement.
6. A GO TO statement may not be used to transfer control from a task to its attaching task or to any of its descendant tasks.

Examples:

```
1. GO TO A234;
   .
   .
   .
   A234: ...
```

2. The following example illustrates a GO TO statement that effectively is a multi-way switch.

```

   .
   .
   .
   DECLARE L LABEL (L1, L2) INITIAL
   (L2);
   GO TO MEET;
L1: X = Y - 1;
   L = L2;
   GO TO MEET;
L2: Y = X - 1;
   L = L1;
MEET: CALL FUDGE (X, Y, Z);
      IF Z = LIMIT GO TO L;
   .
   .
   .
```

3. The following procedure illustrates use of the GO TO statement with a subscripted label variable to effect a multi-way switch:

```

CALC: PROCEDURE (N1, N2);
      DECLARE SWITCH(3) LABEL INITIAL
      (CALC1, CALC2, CALC3);
      I=MOD(N1+N2,3)+1;
      GO TO SWITCH(I);
CALC1: ...
   .
   .
   .
      RETURN;
CALC2: ...
   .
   .
   .
      RETURN;
CALC3: ...
   .
```

```

.
.
END CALC;
```

The GROUP Statement

Function:

The GROUP statement causes a group mark to be inserted in the file on output, or positioning of the file to the next group mark on input.

General format:

```
GROUP [(expression)] [FILE
      (file-name)];
```

Syntax rule:

The "expression", if specified, is a scalar expression and it is evaluated and converted, where necessary, to an integer *n*. If the expression is not specified, it is assumed to be 1.

General rules:

1. A group is defined as a sequence of records delimited by a group-delimiter. A group is created (1) by the GROUP format item specified in a WRITE or PUT statement specifying format-directed transmission, or (2) by a GROUP statement.
2. In a GROUP statement, input records are skipped until a group-delimiter is encountered, with synchronization occurring at the next group, or, if *n* is the value of the expression, at the *n*th subsequent group. Output records are followed by a separate record containing a group-delimiter and released. If *n* is negative or zero, the group statement or format item is ignored.
3. The FILE option specifies that the action is to be taken on the named file. In the absence of a FILE option, the current file is assumed (see "Procedure-Directed Transmission" for a discussion of current files).
4. The techniques for marking a group are implementation defined.

Examples:

1. GROUP FILE (X);

If X is an input file, records are skipped until a group-delimiter is encountered. The file is then positioned immediately following the group-delimiter.

2. GROUP;

Since no file is specified, the GROUP statement positions the current file. If the current file is an output file a group-delimiter is placed on the file, where it is currently positioned, and the group is released from the program. If the current file is an input file, records are spaced until a group mark is encountered.

The IF Statement

Function:

The IF statement causes program flow to depend on the value of an expression.

General format:

```
IF scalar-expression THEN unit-1 [ELSE
  unit-2]
```

Syntax rules:

1. Each "unit" is either a group or a begin block, either of which would be terminated by a semicolon.
2. The IF statement is not itself terminated by a semicolon.

General rules:

1. When the ELSE clause -- ELSE, and its following unit -- is not specified, the scalar expression is evaluated and, if necessary, converted to a bit string. If any bit in the resulting string has the value 1, the unit is executed, and control passes to the statement following the IF statement. If all bits have the value 0, the unit is not executed, and control passes to the next statement. When the ELSE clause is specified, the expression is similarly evaluated. If any bit is 1, unit-1 is executed, and control passes to the statement following the IF statement. If all bits have the value 0, unit-2 is executed, and control passes to the next statement. The units may contain statements that specify transfer of control (see "Sequence of Control"), and so override these normal sequencing rules.
2. IF statements may be nested, that is, either unit-1 or unit-2, or both, may themselves be IF statements. Since each ELSE clause is always associated with the innermost preceding IF, an ELSE with a null statement may be required to specify the desired effect.

Examples:

1. IF QUEUE = EMPTY THEN CALL COMPILE;
ELSE GO TO MULTIPROCESS;
2. A: IF X > Y THEN
IF Z = W THEN
IF W < P THEN Y = 1;
ELSE P = Q;
ELSE;
ELSE X = 4;
J: Z = 5;

The LAYOUT Statement

Function:

The LAYOUT statement specifies the horizontal layout of data on input and output.

General format:

```
LAYOUT
[FILE (file-name [,file-name] ...)]
[MARGIN (expression-1, expression-2)]
[TAB (expression [,expression] ...)];
```

Syntax rules:

1. The options may appear in any order.
2. The "expressions" are scalar expressions.

General rules:

1. The FILE option specifies the files to be operated upon. In the absence of a FILE option, the current file is used, or if there is no current file, the standard output file is assumed.
2. The MARGIN option specifies left and right margins. The values of both expressions are converted to integers when the LAYOUT statement is executed. These values are interpreted as the positions of the left and right margins of the record and line, respectively, relative to the beginning of the record or line. On input, data before the left margin or after the right margin is ignored. On output, the first data item of a record or line is aligned on the left margin, with blanks before it; data is not placed beyond the right margin. If the left margin is specified to the right of, or equal to the right margin, or if either margin is negative, the ERROR condition is raised.
3. The TAB option specifies tabbing. The expressions are converted to integers when the LAYOUT statement is executed. The values are used to indicate character positions from the left end of the line or record. These values need

not be in ascending order. Column 1 is the leftmost column. During list-directed and data-directed output, successive items are aligned on successive free tabs. There is a default tabbing for list- and data-directed output which is implementation defined. There is no automatic tabbing on input. During format-directed transmission, alignment on a tab can be achieved by use of the TAB format item. In other cases, alignment on a tab can be achieved by using the TAB statement.

4. In the absence of a LAYOUT statement, system standards apply.
5. Execution of a LAYOUT statement destroys all options established by a previously executed LAYOUT statement for the same file.

The Null Statement

Function:

The null statement causes no action and does not modify sequential operation.

General format:

[label:]...;

Example:

```
.
.
.
CN OVERFLOW;
.
.
.
```

The on-unit (see "The ON Statement") is a null statement.

The ON Statement

Function:

The ON statement specifies the action to be taken when an interrupt occurs for the named condition. For a discussion of "enable" and "interrupt," see "Interrupt Operations" in Chapter 6.

General format:

Option 1

ON condition [SNAP] on-unit

Option 2

ON condition SYSTEM;

Syntax rules:

1. The "condition" may be any one of those described in Appendix 3.
2. The "on-unit" is an action specification and it is either an unlabeled single statement (other than BEGIN, DO, END, RETURN, or DECLARE) or an unlabeled begin block. Since the on-unit itself requires a semi-colon, no semi-colon appears in Option 1.
3. The on-unit may not be a RETURN statement, nor may a RETURN statement appear within the begin block.

General rules:

1. The standard action to be taken for all ON-conditions is established by the language. When an interrupt takes place before an ON statement for that condition has been executed, standard system action is taken. This standard system action is described in Appendix 3. The ON statement in Option 2 specifies that standard system action is to be taken when an interrupt results from the occurrence of the specified condition.
2. The ON statement in Option 1 is a means for the programmer to specify a special action, that is, execution of the on-unit, to take place when an interrupt occurs for the specified condition.
3. In Option 1, if SNAP is specified, then when the given condition occurs, a calling trace is listed.
4. Control can reach an on-unit only when an interrupt occurs for the condition associated with this on-unit in an ON statement.
5. If an action specification is established by an ON statement in a given block, it remains in effect throughout this block and throughout all dynamic descendants of this block (see "Activation and Termination of Blocks," in Chapter 6, for a discussion of blocks and generations of blocks).

If an action is specified more than once in a given block, the effect of the old (or prior) ON statement is either temporarily suspended or completely nullified by the new (or later) ON statement, as follows:

- a. If the new (or later) ON statement is in a block dynamically descended from the block containing the old (or prior) ON statement, the

effect of the old ON statement is temporarily suspended or stacked. The effect of the old ON statement is restored upon termination of the block containing the new ON statement.

- b. If the new (or later) ON statement and the old (or prior) ON statement are internal to the same block, the effect of the old ON statement is completely nullified.
- 6. If an action is specified by an ON statement in a particular task, the effect of this ON statement is inherited by each attached task and by each task attached by the attached task, etc. (see "Asynchronous Operations and Tasks," in Chapter 6, for a discussion of attached and attaching tasks).
- 7. A condition raised during execution results in an interrupt if and only if the condition is enabled at the point where it is raised.
 - a. The conditions OVERFLOW, UNDERFLOW, ZERODIVIDE, CONVERSION, the input output conditions, and the conditions CONDITION, FINISH, and ERROR are enabled by default.
 - b. The conditions SIZE, SUBSCRIPTRANGE, and CHECK are disabled by default.
 - c. The enabling status of OVERFLOW, UNDERFLOW, ZERODIVIDE, CONVERSION, SIZE, SUBSCRIPTRANGE, and CHECK are controlled by the condition prefix (see "Prefixes" in Chapter 1).

Examples:

```

1. IOPR: PROCEDURE;
   .
   .
   .
   R1: READ FILE (FILEX) (A,B)
       (2F(7,3));
       ON CONVERSION (FILEX)
       CONVQ = 9999;
   .
   .
   .
   R2: READ FILE (FILEX) (X)
       (A(6));
       END IOPR;

```

Assume that program execution begins with procedure IOPR. At the beginning of execution, all conditions are enabled.

If an illegal character is read from FILEX during the execution of statement R1, the standard system action occurs.

The ON statement specifies that the execution of the statement CONVQ =

9999 is to occur in the event that a conversion error causes an interrupt subsequent to execution of the ON statement. Thus, if a conversion error occurs during the transmission of X in statement R2, the normal sequence of control is interrupted, and the statement CONVQ = 9999 is executed.

```

2. ZCHK: PROCEDURE;
   .
   .
   .
   S1: ON OVERFLOW OVSWCH = 1;
   .
   .
   .
   CALL Q;
   .
   .
   .
   Q: PROCEDURE;
   .
   .
   .
   S2: ON OVERFLOW OVSWCH = 2;
   .
   .
   .
   S3: ON OVERFLOW SYSTEM;
   .
   .
   .
   END Q;
   END ZCHK;

```

Assume that program execution begins with procedure ZCHK.

If an overflow occurs prior to execution of the S1 : ON statement, an interrupt with standard system action occurs. If an overflow occurs subsequent to execution of the S1 : ON statement, an interrupt occurs, and the statement OVSWCH = 1 is executed.

When procedure Q is invoked, the S1 : ON statement remains in effect until the S2 : ON statement is executed. At this point, the effect of the S1 : ON is temporarily suspended, and the S2 : ON goes into effect.

If an overflow occurs between the S2 : ON and the S3 : ON, an interrupt occurs, and the statement OVSWCH = 2 is executed.

When the S3 : ON is executed, it completely replaces the S2 : ON (the S1 : ON is still stacked). If an overflow occurs after the S3 : ON is executed and before the end of procedure Q, it causes the standard system action to take place.

After control is returned from Q to ZCHK, the S3: ON is completely replaced by the S1: ON, whose effect is restored. Any overflows occurring from this point to the end of procedure ZCHK cause the statement OVSCH = 1 in S1: ON to be executed.

```

3. SBCHK: PROCEDURE;
    DECLARE A (9) ;
B1: . . .A (I) ...;
    ON SUBSCRIPTRANGE BEGIN;
        IF I>9 THEN
            GO TO BIGER
        ELSE GO TO
            LITLER;
        .
        .
        .
        BIGER: ...;
        .
        .
        .
        LITLER: ...;
        END;
(SUBSCRIPTRANGE) : B2:...A (I) ...;
    B3:...;
    END SBCHK;

```

Assume that procedure SBCHK is the only procedure in the program.

At the beginning of execution, any occurrence of the condition SUBSCRIPTRANGE will not give an interrupt; it is not enabled, since the condition name does not appear in a prefix in the PROCEDURE statement. However, the occurrence of any other ON condition, except SIZE and CHECK (identifier), will give an interrupt. If in statement B1, the value of I is greater

than 9 or less than 1, no interrupt action is taken.

When the ON statement for the condition SUBSCRIPTRANGE is executed, any interrupt that results from a subsequent occurrence of the SUBSCRIPTRANGE condition will result in the action specified by the begin block in the ON statement.

The prefix for statement B2 specifies that the condition SUBSCRIPTRANGE is enabled and should cause an interrupt if it occurs during the execution of statement B2. In this case, the begin block in the ON statement is executed.

In the execution of B3 and subsequent statements, the occurrence of a subscript that is not within the specified range does not cause an interrupt action to occur.

For further examples, see "Interrupt Operations" in Chapter 6.

The OPEN Statement

Function:

The OPEN statement acquires and prepares files for subsequent transmission.

General Format is shown in Figure 5.




Figure 5. General Format for the OPEN Statement




Figure 6. Format of "option" Allowed in the OPEN Statement

Syntax rules:

1. The options may appear in any order.
2. The "file name" may be described in a DECLARE statement with the file attributes discussed in Chapter 4.
3. Options that are common to two or more file names may be factored in the same way that attributes in a DECLARE statement may be factored (see "Factoring of Attributes" in Chapter 4). Only one level of factoring is permitted.

General rules:

1. The OPEN statement causes certain actions to be performed upon the file, whose name is one of the file names of the OPEN statement. These actions are specified by the options.

If, however, a file is not opened by an OPEN statement, the file is opened during the first READ or WRITE statement that refers to it.

If an OPEN statement is encountered for a file already opened, the statement is ignored.

2. The following options may be given to a file in the OPEN statement:

- a.

```
[ INPUT
  OUTPUT
  INOUT ]
```

One of these options may be given to specify the direction of data transmission that is permitted for the file. INOUT may be given for both direct and sequential access files, stating that both INPUT and OUTPUT are permitted.

Either the OPEN statement or the file declaration must specify the direction of data transmission. References to the file in PAGE and LAYOUT statements before INPUT or OUTPUT is established, forces no assumptions. References to the file in GROUP, SPACE, SKIP, or SEGMENT statements before INPUT

or OUTPUT is established, forces the default assumption INPUT.

It should be noted that INPUT files may not be written upon and OUTPUT files may not be read.

- b. [TITLE (expression)]

A file name may be associated with more than one set of data. The choice of the desired set may be delayed until the OPEN statement is executed. At this point, the "expression" in the TITLE option is evaluated, converted to a character string, and used to identify the data set. The original file name is remembered, such that TITLE does not permanently override it. If the TITLE option is omitted, the file name is taken as the data set name. The TITLE option can be used to let the file name refer to more than one actual file.

- c. [ACTIVITY (expression)]

The ACTIVITY option causes the "expression" to be evaluated and converted to an integer that indicates the relative activity of the file. This relative activity is represented in units defined individually for each implementation of the language.

- d. The format of the ident option is as follows:

```
IDENT { data-list format list
      CALL entry-name [(argument
                        [,argument]) ...]
```

The ident option in an OPEN statement for an output file specifies that a label is to be placed on the external medium. For an input file, IDENT provides information for label checking.

The ident option specifying a data list and format list yields a character string that is compared

```

PAGE  [FILE (file-name [,file-name] ...)]
      [NUMBER [(expression)]] [HEAD (expression)]
      [FOOT (expression)] [SIZE (expression)]
      [SPACE (expression)]
      [AT (expression-1) { (expression-2)
                          { CALL entry-name (argument [,argument] ...) } } ] ;

```

Figure 7. General Format for the PAGE Statement

with the file label for an input file or is written as the file label for an output file. For an INOUT file, INPUT is assumed. The data list and format list are described in Chapter 7.

The ident option specifying an entry name and argument list causes the specified file to become the current file and the designated procedure to be invoked for reading or writing the file label through GET or PUT. The form of file labels is implementation defined.

If the ident option is not specified, no special label operations are performed.

- e. The ENVIRONMENT option specifies various characteristics of the FILE being opened. See "The ENVIRONMENT Attribute," in Chapter 4.

Examples:

1. OPEN MATRIX INPUT.

The file MATRIX is made available for use as an input file.

2. OPEN WORKFILE OUTPUT ENVIRONMENT (CREATE, BUFFER (2));

The PAGE Statement

Function:

The PAGE statement specifies the vertical format of files.

General format is shown in Figure 7.

Syntax rules:

1. The options may appear in any order.
2. "Expression" is a scalar expression.
3. The AT option may appear more than once.

General rules:

1. The PAGE statement is used only with CONSECUTIVE SEQUENTIAL OUTPUT files; it causes a skip to the start of the next page.
2. The FILE option specifies the files to be operated upon. If the FILE option is omitted, the current file is assumed (if the current file is an INPUT file, the ERROR condition is raised). If there is no current file, the standard system output file is assumed.
3. The NUMBER option specifies that the pages are to be numbered on the right of the heading, starting at the number that is the integer value of the expression. If the NUMBER option is not specified, numbering is not generated. If the expression is omitted, numbering starts at one (1).
4. The HEAD option provides a page title, left adjusted on every page. The character string, which is the value of the expression, is the page title. The expression is evaluated when the PAGE statement is executed.
5. The FOOT option provides a left-adjusted line at the foot of each page. The character string, which is the value of the expression, is the footing line. The expression is evaluated when the PAGE statement is executed.
6. The SIZE option specifies the number of lines per page, including heading, footing, and blank lines. The integer value of the expression provides this information. If this option is unspecified, system standards apply. Lines here mean actual page lines, not print lines.
7. The SPACE option specifies the line spacing. If the integer value of the expression is *n*, then (n-1) blank lines are effectively generated between each two lines explicitly specified. In the absence of this option, SPACE (1) is implied. The SPACE option implies spacing before printing.

8. The AT option specifies that certain action is to occur at a specified location on every page. Expression 1 is evaluated and converted to an integer *n* when the PAGE statement is executed. (If the value of *n* is zero or negative, the AT option is ignored.) Subsequently, when the *n*th line of each page is reached, the following occurs before continuing with the output:
- If expression 2 is specified, it is evaluated and converted to a character string that is generated as the *n*th line.
 - If a CALL option is specified, the arguments are evaluated, where necessary, and the procedure entry, specified by "entry name," is invoked. This procedure may cause special page handling.

The scope of the arguments includes the block containing the PAGE statement. Since the arguments are evaluated at each invocation, the block containing the PAGE statement must still be active at each invocation, that is, when the *n*th line is reached on each page. A procedure invoked by the CALL option can contain only GET and PUT statements for data transmission. The file which caused the procedure to be invoked becomes the current file during entry to the invoked procedure.

Note:

The *n*th line is said to be "reached", in the sense used above, when:

- A WRITE statement is encountered for which the previous WRITE for that file released line (*n*-1).
- A WRITE statement with the CROSS option is being executed and a SPACE or SKIP releases line (*n*-1), and the first item to be output for line *n* has just been evaluated (or if a further SPACE or SKIP attempts to pass line *n*).

(In the above definition, when *n*=1, *n*-1 refers to the last line of the previous page.)

- In the absence of a PAGE statement, system standards apply.
- AT options override spacing and skipping, and the SPACE factor is still applied to the data line which caused the AT to be activated. AT lines do not have the SPACE factor applied to them, as only enough blank lines are generated to cause the AT line to appear on the specified line number. The AT procedure must have a SPACE

format item or statement in order to emit the line; otherwise the data being output from the WRITE which caused the procedure to be invoked will be written on the same line.

- The SEGMENT option may not be applied to a file controlled by a PAGE statement.

Example:

```
PAGE FILE (FILEX), NUMBER (100), HEAD
('PAGE HEADING')
FOOT ('BOTTOM OF PAGE'), SIZE (34);
```

The POSITION Statement

Function:

During data transmission, the action that occurs is as if a pointer moved across the records as demanded by the data specifications. The POSITION statement manipulates this pointer.

General format:

```
POSITION format-list;
```

Syntax rules:

- The format list is as described for format-directed data specification in "Format Lists," in Chapter 7.
- The following format items are not allowed in the format list of a POSITION statement: GROUP, SEGMENT, SKIP, SPACE, and the remote format item, R.

General rules:

- When the POSITION statement is executed, the pointer is first reset to the beginning of the current record. The format elements are then used to determine the movement of the pointer as if there were associated data list elements corresponding to the format items. Since no data list exists, all format items must have an explicit or implicit field width (precision) specification.
- If the POSITION statement moves the pointer across parts of an output record that have no information edited into them, the record is assumed to be initially blank.

Examples:

```

GETAB: PROCEDURE;

      GET (A,B) (2F (5,2) , X (6)) ;

      IF (A>0 & B = 0) THEN POSITION
        (X (25)) ;
      RETURN;
      END GETAB;

READY : PROCEDURE;

      READ (Y) (F (8,3)) , CALL GETAB;

      END READY;

```

The PROCEDURE Statement

Function:

The PROCEDURE statement has the following functions:

1. Heads a procedure
2. Defines the primary entry point to a procedure
3. Specifies the parameters for the primary entry point
4. Defines any special attributes of the procedure
5. Specifies the attributes of the value that is returned if the procedure is invoked as a function at the primary entry point

General format:

```

entry-name: ...PROCEDURE
            [ (parameter [, parameter] ...)]
            [OPTIONS (option-list)]
            [RECURSIVE] [data-attributes];

```

Syntax rules:

1. The data attributes and the OPTIONS and RECURSIVE attributes may appear in any order and are separated by blanks.
2. The attributes in the OPTIONS list are separated by commas, where necessary.

General rules:

1. The "parameters" are names that specify the parameters of the entry point. When the procedure is invoked, a relationship is established between the arguments of the invocation and the parameters of the invoked entry point (see "Relationship of Arguments and Parameters," in Chapter 10).
2. The OPTIONS attribute specifies a list of options, separated by commas where necessary. The list, depending upon implementation, includes the options

MAIN and REentrant. The OPTIONS attribute may be specified only for an external procedure.

3. The RECURSIVE attribute specifies that this procedure may be invoked recursively. It does not apply to contained procedures which, if recursive, must also have the attribute.
4. The data attributes permitted with a PROCEDURE statement are the arithmetic and string attributes. The data attributes specify the characteristics of the value returned by the procedure when invoked as a function at the primary entry point. The value specified in the RETURN statement of the invoked procedure is converted to the specified data attributes.

If insufficient data attributes are specified at the entry point, default attributes are applied, as determined by the name of the entry point.

If a procedure has more than one label and no data attributes, there is, because of default attributes, potential ambiguity in the characteristics of the value to be returned (see "Assignment of Attributes to Identifiers" in Chapter 4). To avoid this ambiguity, the first label is interpreted as if it were a single entry name for a separate PROCEDURE statement, and each subsequent label is interpreted as if it were a separate ENTRY statement.

For example, the statement:

```

A:I: PROCEDURE;

is equivalent to:

A: PROCEDURE;
I: ENTRY;

```

Examples:

```

B: PROCEDURE;
  .
  .
  .
  C=A (X, Y) ;
  END B;
A: PROCEDURE (B,C) FIXED;
  .
  .
  .
  RETURN (B*C + SIN (P))
  END A;

```

If procedure A is invoked as a function, as it is in procedure B, then when control is returned to B, the expression (B*C + SIN (P)) is evaluated, converted to fixed point, and the value assigned to C in procedure B.

The PUT Statement

Function:

The PUT statement has meaning only when there is a current file; if there is no current file, the ERROR condition is raised. The PUT statement then causes data to be fetched from variables as specified and to be moved to the record being constructed for the current file (see "Procedure-Directed Transmission," in Chapter 7).

General format:

```
PUT data-specification ...;
```

General rules:

1. The "data specifications" are discussed in "Data Specifications," in Chapter 7. Only those forms specified for output may be used; a CALL option may not be used.
2. As the data record is being formed, the action that occurs is as if a pointer moved across the record as demanded by the data specifications. This pointer may be repositioned by use of the POSITION statement or the REPOSITION statement.

The character count of varying-length records depends upon the rightmost sweep of the pointer. The character count of fixed-length records is predetermined.

The READ Statement

Function:

The READ statement is normally used to transmit data from an external storage

medium to internal storage. However, if the STRING option is specified, the READ statement causes the movement of data from an internal storage area to other internal storage areas.

General format is shown in Figure 8.

Syntax rules:

1. The options may appear in any order.
2. At least one "data specification" must appear, but more than one is permissible. The CALL option may appear in conjunction with other data specifications.
3. When the STRING option is used, only the data specifications may be used; the other options must not appear, nor may the CALL option.
4. Each "expression" is a scalar expression.

General rules:

1. The FILE option specifies the name of the file from which the data is to be acquired.

The STRING option provides for the internal editing and moving of strings. It specifies the name of a string variable or the name of an element in a string array from which data is transmitted to the data list.

In the absence of a FILE or STRING option, the standard system input file is assumed.

2. The data specifications are discussed in "Data Specifications," in Chapter 7. Only those forms specified for input may be used. All modes of transmission may be arbitrarily specified together. The transmissions associated with each data specification and edit procedure are performed in the order that the options appear.

```
READ [ FILE (file-name) ] [ STRING (name) ] {data-specification} ...  
      [ CROSS [(expression)] [HOLD] ]  
      [ SEGMENT (expression) ]  
[PRINT] [KEY (expression)] [REGION (expression)]  
[ZERO];
```

Figure 8. General Format for the READ Statement

3. Each READ statement normally processes one record; an error condition is normally produced if the data specification causes the record boundary to be crossed. However, the CROSS option permits data acquisition to proceed through any number of records in order to satisfy the specified data requirements. The number of records read may be limited by the integer value of the expression in the CROSS option. If no expression is specified, unlimited crossing is allowed. The margin qualifications for the data file, if specified by a LAYOUT statement, remain valid while under control of the CROSS option. Record boundary crossing due to LIST, DATA, or SEGMENT does not require the presence of the CROSS option. It may be specified, however, to limit the number of records crossed. Crossing due to SPACE, SKIP, or GROUP does require its presence. Data items may not span record boundaries.

A HOLD option permits part of one record to be processed. This HOLD option causes the position of the record pointer to be set on completion, so that the next READ begins its data scan at the point where the previous operation ceased scanning. If HOLD is not specified, the remaining part of the record is skipped. HOLD may not be specified for a file which is accessed in more than one task.

The SEGMENT option implies both the CROSS and HOLD options. The expression in the SEGMENT option is converted, if necessary, to a character string. This string serves as a segment mark. If it is the null string, the ERROR condition is raised. This option permits the data input stream to be synchronized, not at the record boundary, but at the mark, effectively causing the segment to be operated upon as a record. Before data items are transmitted from the input stream, a scan for the segment mark is made in order to delimit the segment. The segment mark is not part of the segment. A subsequent READ will begin after the mark. Should the end of the segment be encountered while transmitting data, transmission ends for that data specification.

4. The PRINT option specifies that data being read is, at the same time, to be written, in the same format, on the standard output file.
5. The KEY and REGION options may be used when direct access to a particular record is required (see "The KEY Option" and "The REGION Option" in

Chapter 7). If a file is declared with the access attribute DIRECT, then the KEY or REGION option must be used with each READ for that file, unless the immediately preceding READ of the file had the HOLD option.

6. The ZERO option specifies that trailing blanks in numeric data input fields are to be treated as zeros when read under F, G, or E format.
7. A count is kept of the number of scalar data items transmitted. The COUNT (file-name) built-in function may be used to determine this number of transmitted data items.
8. If a group mark is encountered during a read operation, the END GROUP condition is raised.

Examples:

1. READ FILE (INVENTORY) , (ITEM.NAME, ITEM.COST) (A(20) , F(5,2)) ;

The file name INVENTORY is read under format-directed transmission for one record. The first 20 characters of the record are placed in the character-string variable ITEM.NAME, the next 5 are converted from fixed-point decimal external format to the internal form of the variable, ITEM.COST. A subsequent READ of the data file is synchronized to the next record boundary.

2. READ FILE (TABLES) , (TABLE.POOL) (F(5)) , KEY (Q) ;

The file named TABLES is read for the record composed of five-digit, fixed-point integers. The record is converted to integer representation, and each item is assigned to the array TABLE.POOL.

3. READ FILE (FILEZ) , (AB) (A(10)) , SEGMENT ('*') ;

The file-FILEZ is read for alphabetic data items, each ten characters in length, that are assigned to the character-string array AB. Assignment ceases when either the complete array is satisfied or the SEGMENT mark, the asterisk, is encountered (in the former case, the input data stream is subsequently synchronized to the segment mark).

4. READ DATA;

This statement under data-directed transmission specifies that data is to be read under data-directed transmission from the standard system input tape. It is assumed that the records to be read are composed of scalar

assignments giving the names of the data items to be read and the values of these data items.

The REPOSITION Statement

Function:

During data transmission, the action that occurs is as if a pointer moved across the records being processed. If an error condition occurs during this activity, then, under the control of an ON statement, a REPOSITION statement may reset the pointer to the start of the data item that caused the error condition.

General format:

```
REPOSITION;
```

Examples:

```
READX:  PROCEDURE;
        READ FILE (FILEX) , (X) (F(7,2));
        CALL GETY;
        .
        .
        END READX;

GETY:   PROCEDURE;
        ON CONVERSION (FILEX) REPOSITION;
        GET (Y) (X(17) , F(7,2));
        RETURN;
        .
        .
        END GETY;
```

The RESTORE Statement

Function:

The RESTORE statement causes data previously saved by name in auxiliary storage to be restored (see "The SAVE statement").

General format:

```
RESTORE (item, [,item]...)
        [, (expression)] ;
```

Syntax rules:

1. Each "item" may be an array, major structure name, or a scalar which is not part of an array or structure.
2. Each item must have appeared in a previously executed SAVE statement.
3. The "expression" is a scalar expression.

General rules:

1. The RESTORE statement without an expression is equivalent to a series of simple RESTORE statements as follows:

```
RESTORE (item1);
RESTORE (item2);
.
.
.
```

The RESTORE statement with an expression is equivalent to the following statements:

```
temp=expression
RESTORE (item1) (temp);
RESTORE (item2) (temp);
.
.
.
```

- Each simple RESTORE statement causes the specified data to be identified by the data name qualified by the integer value of the expression (if an expression is specified) converted to BINARY FIXED (s,0), where s is implementation defined.
2. Once an item has been restored, it may not be restored again. If the same item has been saved repeatedly with no qualifying expression, the action of restoring the data causes the top item of the stacked information to be deleted. Therefore, the stacked information is treated in "first-in last-out" manner.
3. An item may be saved in one external procedure and restored in another if the data name is declared EXTERNAL.
4. One SAVE statement may be used to save more than one scalar, array, or structure; however, these items may be restored separately.
5. The extents of the data restored must be the same as the data saved.

Examples:

1. RESTORE (A,B,C);

Assume that the scalar data items A, B, and C were previously saved by using the SAVE statement. The RESTORE statement then causes A, B, and C to be made available for computation.

2. SAVERM: PROCEDURE
 DECLARE TABLE (10) , 1 RAINBOW,
 2 RED, 3 CRIMSON, 3 PINK, 2
 BLUE, 3 NAVY, 3 TEAL, 2
 YELLOW;
 .
 .
 .

```

SAVE (TABLE, RAINBOW);
.
.
RESTORE (TABLE);
.
.
RESTORE (RAINBOW);
.
.
END SAVERM;

```

Since TABLE is an array and RAINBOW is a structure, the SAVE statement causes all 10 data items in TABLE to be saved and the elementary items (CRIMSON, PINK, NAVY, TEAL, and YELLOW) of the structure to be saved.

The first RESTORE statement causes the entire array to be restored; the second RESTORE statement causes the elementary items of the structure to be restored.

The RETURN Statement

Function:

The RETURN statement terminates execution of the procedure that contains the RETURN statement and returns control to the invoking procedure. It may also return a value.

General format:

Option 1.

```
RETURN;
```

Option 2.

```
RETURN (expression);
```

General rules:

1. Only the RETURN statement in Option 1 can be used to terminate procedures not invoked as function procedures; control is returned to the point logically following the invocation.

Option 1 represents the only form of the RETURN statement that may be used to terminate a procedure invoked with the task option. If the task option involved an EVENT option (see "The CALL Statement," in this Chapter), then the execution of the RETURN statement will cause the completion status of the associated event name to be set to '1'B.

2. The RETURN statement in Option 2 is used to terminate a procedure invoked as a function procedure. Control is returned to the point of invocation, and the value returned to the function reference is the value of the expression specified.

If the entry point at which the procedure is invoked specifies data attributes, the value of the expression is converted to the implicit or explicit data attributes specified at the entry point, before it is returned.

3. If control reaches an END statement corresponding to the end of a procedure, this END statement is treated as a RETURN statement (of the Option 1 form) for the procedure.

Example:

```

A: PROCEDURE (X,Y) FIXED;
   DECLARE (X,Y) FLOAT;
   .
   .
   RETURN (X**2+Y**2);
   END;
B: PROCEDURE;
   DECLARE A ENTRY FIXED;
   .
   .
   R = A (P,Q);
   .
   .
   END;

```

In the assignment statement (R = A(P,Q)), procedure B invokes procedure A as a function. Procedure B specifies that the scalar expression in the RETURN statement is to be evaluated; since X and Y are floating-point variables and the PROCEDURE statement specifies that the value returned is to be fixed point, the value of the expression is converted to fixed point, and this value is returned to B.

The REVERT Statement

Function:

A REVERT statement specifying a given ON-condition is used to nullify the effect of the most recent previously executed ON statement for that condition and to cause the action specification to be reestablished as it was in the immediate, dynamically encompassing block (see "Activation and Termination of Blocks," in Chapter 6).

General format:

```
REVERT condition;
```

Syntax rule:

The "condition" is any ON-condition (see Appendix 3).

General rules:

The execution of a given REVERT statement, specifying a given condition and internal to a given block, has the effect described above only if an ON statement, specifying the same condition and internal to the same block, was executed after the block was activated. If such an ON statement was executed, and if the execution of no other similar REVERT statement has intervened, then the execution of the given REVERT statement does have the effect described above. Otherwise, the REVERT statement is effectively treated as a null statement. Thus, a repeated REVERT statement results in no operation.

Examples:

```
A: PROCEDURE;
.
.
.
ON1: ON ZERODIVIDE GO TO ERRSPEC;
.
.
.
CALL B;
.
.
.
B: PROCEDURE;
.
.
.
ON2: ON ZERODIVIDE;
.
.
.
REVERT ZERODIVIDE;
.
.
.
END B;
.
.
.
ON3: ON ZERODIVIDE SYSTEM;
.
.
.
END A;
```

Unless it is stated otherwise, the condition ZERODIVIDE always is enabled. If division by zero occurs prior to execution

of statement ON1, an interrupt with standard system action takes place.

If division by zero occurs after execution of ON1 and prior to execution of statement ON2, an interrupt takes place and control transfers to the statement GO TO ERRSPEC.

If division by zero occurs after execution of ON2 and prior to the REVERT statement, an interrupt takes place effectively with no action.

When the REVERT statement is executed, the effect of the statement ON2 is nullified, and statement ON1 again becomes effective. If division by zero occurs after execution of the REVERT statement and prior to the execution of statement ON3, an interrupt takes place, and control transfers to the statement GO TO ERRSPEC.

After the execution of statement ON3, division by zero causes standard system action to take place.

The SAVE Statement

Function:

The SAVE statement causes data to be placed in auxiliary storage, identified by its name, and if the name is INTERNAL, by the block in which it is declared.

General format:

```
SAVE (item-1 [,item-2] ...)
      [(expression)];
```

Syntax rule:

Each "item" may be an array, major structure name, or a scalar which is not part of an array or structure. The "expression" is a scalar expression.

General rules:

1. The SAVE statement without an expression is equivalent to a series of simple SAVE statements as follows:

```
SAVE (item 1);
SAVE (item 2);
.
.
.
```

The SAVE statement with an expression is equivalent to the following statement:

```
temp=expression;
SAVE (item 1) (temp);
SAVE (item 2) (temp);
.
.
.
```

Each simple SAVE statement causes the specified data to be identified by the data name, qualified by the integer value of the expression (if an expression is specified) converted to BINARY FIXED (s,0) where s is implementation defined.

2. If no expression is specified, and items of the same name are repeatedly saved, the values are stacked, and restored in a last-in, first-out basis.

If an expression is specified and items of the same name are repeatedly stored, only one value for a given name and given expression value is saved at any one time. Subsequent execution of a SAVE statement with matching identification causes the previously saved value to be overridden.

3. The extents of the data, when restored, must be the same as they were when the data was saved.

Example:

```
SAVE (A, B , C);
```

The scalar data items A, B, and C are saved in auxiliary storage.

```
.
.
.
DO N=1 TO 10;
X=Y**N;
SAVE (X) (N);
END;
.
.
.
I=J+M/2;
RESTORE (X) (I);
Q = SIN (X);
.
.
```

Each execution of the SAVE statement causes the current value of X to be identified by the current value of N and to be saved.

The RESTORE statement causes one of the previously saved values of X to be restored. In particular, that value of X identified by an integer whose value is J+M/2 is restored.

The assignment statement Q=SIN(X); uses this restored value of X.

The SEGMENT Statement

Function:

The SEGMENT statement specifies positioning within a segmented file. Segment marks may also be written on a file by a WRITE statement that specifies a SEGMENT option (see "The WRITE Statement").

General format:

```
SEGMENT [ (expression) ] [ FILE (file-
name) ];
```

General rules:

1. The scalar "expression," if specified, is converted to character, if necessary; this character string is the segment mark. If the SEGMENT statement appears within a procedure invoked by a CALL option in a READ or WRITE statement that specifies a SEGMENT option, the expression may be omitted; in this case, the SEGMENT option defines the segment mark.
2. The SEGMENT statement causes the pointer to be positioned at the next segment mark after the current position. On input, sufficient records may be skipped to effect this positioning. On output, the segment mark is added to the data stream. Segments positioning need not, however, cross record boundaries.
3. The FILE option specifies that the action is to be taken on the named file. In the absence of a FILE option, the current file is assumed (see "Procedure-Directed Transmission" for a discussion of "current" files).

Examples:

1. SEGMENT FILE (FILE PROC);
2. SEGMENT (A+B**3);
3. DECLARE FILX FILE STANDIN;


```
.
.
.
SEGMENT FILE (FILX);
```

In Example 3, the SEGMENT statement positions the system standard input file; the name FILX is declared as a pseudonym for this tape.

The SIGNAL Statement

Function:

The SIGNAL statement simulates the occurrence of an interrupt (see "Interrupt Operations," in Chapter 6, and "The ON Statement"). It may be used to test the action specification of the current ON statement.

General format:

```
SIGNAL condition;
```

Syntax rule:

The condition may be any one of those described in "ON-Conditions," in Appendix 3.

General rules:

1. When a SIGNAL statement is executed, it is as if the specified condition had actually occurred. The sequence of control through the program is interrupted, and control is transferred to the current ON statement for the specified condition. After execution of the on-unit, control normally returns to the statement immediately following the SIGNAL statement.
2. If an ON statement specifies the CONDITION condition, the condition can cause an interrupt only if a SIGNAL statement, specifying this condition, is given.
3. If the condition specified in the SIGNAL statement is disabled, no interrupt occurs, and the statement is equivalent to a null statement.
4. If the condition has no current ON statement, then the normal system action for the condition is performed.

Examples:

- ```
1. X: PROCEDURE;
 .
 .
 .
ON1: ON ENDFILE (DATIN) Y,Z = 0;
 .
 .
 .
S1: SIGNAL ENDFILE (DATIN) ;
 .
 .
 .
ON2: ON ENDFILE (DATIN) SYSTEM;
 .
 .
 .
S2: SIGNAL ENDFILE (DATIN) ;
 .
 .
 .
END X;
```

The S1: SIGNAL statement causes an interrupt in the same way as if an attempt to read past a file delimiter, had actually occurred. Control is transferred to the statement Y,Z = 0 in the ON1: ON statement.

When the S2: SIGNAL statement causes an interrupt, control is transferred to the ON2: ON statement, and standard system action is taken.

- ```
2. ON CONDITION (TAX) TAXCT = TAXCT+1;
  .
  .
  .
SIGNAL CONDITION (TAX) ;
```

The ON statement establishes an action for the programmer-specified condition TAX. This condition can occur only when a SIGNAL statement causes the condition to occur.

The SKIP Statement

Function:

The SKIP statement causes records (or lines) to be skipped.

General format:

```
SKIP [ (expression) ] [ (FILE file-
name) ] ;
```

General rules:

1. The scalar "expression," if specified when the SKIP statement is executed, is evaluated and converted, where necessary, to an integer n. If the expression is not specified, n is assumed to be 1.
2. When used with print files, lines and pages are considered, otherwise, records and groups are indicated.

On input, the SKIP statement causes a skip to the nth record of the group. If the current record is greater than n, a skip to the nth record of the next group occurs.

On output, the SKIP statement causes the creation of a sufficient number of empty records to cause alignment on the record as described for input.

3. The FILE option specifies that the action is to be taken upon the named file. In the absence of a FILE option, the current file is assumed.

```

SORT FILE (file-name [,file-name]...) [RECORD (format-list)
{UP }
{DOWN } (integer [,integer]...)... [GIVING (file-name)];

```

Figure 9. General Format for the SORT Statement

Examples:

1. OPEN PQR INPUT;
SKIP (N) FILE (PQR);
2. OPEN FILEA OUTPUT;
.
.
.
SKIP FILEA;

The SORT Statement

Function:

The SORT statement specifies that records on a particular file are to be sorted and, optionally, merged. The sorting is performed on specified fields in ascending or descending order.

The general format is shown in Figure 9.

Syntax rules:

1. The options may appear in any order.
2. The UP and DOWN options may be repeated as required, to specify an ascending sort on some fields and a descending sort on others, in the required order.

General rules:

1. The size of the records to be sorted either must be specified within the ENVIRONMENT attribute for the file name or must be implied by a record description using the RECORD specification.
2. The FILE specification specifies the files to be sorted. If more than one file name is specified, a merge is also performed.
3. The RECORD specification describes either the format of the whole record

or merely an initial portion of the record. When only an initial portion of the record is to be described, the ENVIRONMENT attribute must be declared for the file name, giving the actual length of the record or the maximum length for varying-length records.

The format list (see Chapter 7) defines fields on the record; the nth format item describes the nth field. If a format item has an iteration factor of m, this constitutes m fields. Of the positioning format items listed in Chapter 7, only POSITION is permitted; this item does not constitute a field.

4. The UP/DOWN specification indicates the sorting order. UP specifies an ascending sort; DOWN specifies a descending sort.

The integers in the specification are decimal integer constants that specify the fields to be sorted with respect to the record description. The fields to be sorted are taken from the UP/DOWN specification in left-to-right order. The file is ordered on the leftmost specified field first, and within that ordering, on the next field, and so on.

The sort comparisons are performed using the character collating sequence for character-string fields, bit comparison for bit-string fields, and algebraic comparison for arithmetic fields.

5. The GIVING option specifies the file on which the sorted output is to be written. If omitted, the standard output file is used.

If an output file is specified that differs from the standard output file and that differs from any of the files to be sorted, the file must not be currently open. Rather the operating system opens it for output, produces the sorted file, and closes it.

If an output file is specified that differs from the standard output file but is identical to one of the files being sorted, then, after reading the file, the operating system closes it, opens it for output, produces the sorted file, and closes it.

Example:

```
    SORT FILE (MASTER) , RECORD (A (1) ,  
      (6) P'99999') , UP (1,3,5) ;
```

This SORT statement specifies that the file MASTER is to be sorted in ascending sequence. The RECORD specification indicates that each record in MASTER is composed of 7 fields; the first field contains one character; subsequent fields each contain five digits.

The records are sorted on the first, third, and fifth fields, in that order.

The sorted file is written on the standard output file.

The SPACE Statement

Function:

The SPACE statement causes spacing over records on input files and the completion and release of records on output files.

General format:

```
    SPACE [(scalar-expression)] [FILE  
      (file-name)] ;
```

General rules:

1. The expression, if specified, is evaluated and converted, where necessary, to an integer *n*. If the expression is not specified, it is assumed to be 1.
2. The FILE option specifies that the action is to be taken on the named file. In the absence of a FILE option, the current file is assumed (see "Procedure-Directed Transmission" for a discussion of current files).
3. On input, this statement causes the file to be positioned at the start of the *n*th record following the current record. On output, this statement causes the completion and release of the current record and the emission of *n*-1 empty records.
4. The use of SPACE on a REGIONAL or INDEXED file accessed in the DIRECT mode is limited in that an expression may not be specified. Use of SPACE in

such instances causes the current record to be released. Subsequent data transmission to or from the current file will cause the KEY and REGION values to be reevaluated in order to determine the next record to be accessed.

The STOP Statement

Function:

The STOP statement causes immediate termination of the major task and all sub-tasks (see "Asynchronous Operations and Tasks," in Chapter 6).

General format:

```
    STOP ;
```

The TAB Statement

Function:

During data transmission, the action that occurs is as if a pointer moved across the records being processed. The TAB statement causes this pointer to be aligned on the *n*th tab of the record or line (see "The LAYOUT Statement"). Intervening data is skipped.

General format:

```
    TAB [(scalar-expression)] ;
```

General rule:

The "scalar expression," if specified, is evaluated and converted to an integer *n*. This integer represents the *n*th tab for alignment. If the scalar expression is omitted, the smallest tab value to the right of the current position in the record is used.

Example:

```
    TAB (3) ;
```

Suppose the LAYOUT for the current file contains the specification TAB (10, 50, 20, 60); then the above statement causes a TAB to column 20 to occur, since 20 is the third tab. This may involve backing up certain character positions. If backing up has occurred, it is meaningful to use a GET statement even on an output file; and a PUT statement may be used on input files followed by TAB or POSITION to cause re-reading of a field.

If the TAB statement moves the pointer across a part of an output record which has no data edited into it, the record is assumed to be initially blank.

The WAIT Statement

Function:

The WAIT statement is used to cause the suspension of operations in the task where it appears until certain events have been completed.

General format:

```
WAIT (event-name [,event-name]...)  
[(scalar-expression)];
```

Syntax rule:

The event name is as described in "Event Names," Chapter 2.

General rules:

1. The execution of this statement causes the task in which it is executed to be suspended until, for some or all of the event names in the list above, the condition

```
EVENT (event-name) = '1'B
```

is satisfied. (See "Asynchronous Operations and Tasks," in Chapter 6, "Event Names," in Chapter 2, "Pseudo-Variables," in this chapter, and the description of the EVENT built-in function in Appendix 1.)

2. If the optional expression does not appear, all the event names in the list must satisfy the above condition before the task issuing the WAIT statement can resume.
3. If the optional expression appears, the expression is evaluated when the WAIT statement is executed and converted to an integer. This integer specifies the number of events that must satisfy the above condition

before the task issuing the WAIT statement can resume. If the value of the expression is zero or negative, the WAIT statement is treated as a null statement. If the value of the expression is greater than the number, n, of event names in the list, the value is taken to be n. If the statement refers to an array event name, then each of the array elements may contribute to the count.

Example:

```
PI: PROCEDURE;  
. . .  
CALL P2 EVENT (EP2);  
. . .  
WAIT (EP2);  
. . .  
END;
```

The CALL statement, when executed, attaches a task whose completion status is associated with the event name EP2. When the WAIT statement is encountered, the execution of the task corresponding to PI is suspended until the value of EVENT(EP2) is '1'B, i.e., until the attached task is completed.

The WRITE Statement

Function:

The WRITE statement is normally used to transmit data from internal storage to an external storage medium. However, if the STRING option is specified, the write statement causes the movement of data from one or more internal storage areas to another internal storage area.


```

WRITE [ FILE (file-name) ] {data-specification} ...
      [ STRING (name) ]

      [ CROSS [(expression)] [HOLD] ]
      [ SEGMENT (expression) ]

      [KEY (expression)] [NEWKEY (expression)] [REGION (expression)]

      [FROM (file-name)];

```

Figure 10. General Format for the WRITE Statement

Syntax rules:

1. The options may appear in any order.
2. At least one "data specification" must appear, but more than one is permissible. However, only one procedure-directed data specification can be given, but it may appear in conjunction with other data specifications.
3. When the STRING option is used, only the data specification may appear; the other options must not appear.
4. Either a "data specification" or a FROM option must appear.

General rules:

1. The FILE option specifies the name of the file to be associated with the data transmitted.

The STRING option provides for the internal editing and moving of strings. It specifies the name of a string variable or the name of an element in a string array into which data is transmitted from the data list. A PICTURE attribute given for the string is ignored.

In the absence of a FILE or STRING option, the standard system output file is assumed.

2. The data specifications are discussed in Chapter 7. Only those forms specified for output may be used. All modes of transmission may be arbitrarily specified together. The transmissions associated with each data specification and edit procedure are performed in the order that the options appear.
3. Each WRITE statement normally processes one record; an error condition is produced if the data specification causes the record boundary to be crossed. However, the CROSS option permits data transmission to proceed through any number of records in order to satisfy the specified data require-

ments. The number of records written may be limited by the integer value of the expression in the CROSS option. If no expression is specified, unlimited crossing is allowed. The margin qualifications for the data file, if specified by a LAYOUT statement, remain valid while under control of the CROSS option. Record boundary crossing due to LIST, DATA, or SEGMENT does not require the presence of the CROSS option. It may be specified, however, to limit the number of records crossed. Crossing due to SPACE, SKIP, or GROUP does require its presence.

A HOLD option permits part of one record to be processed. The HOLD option causes the position of the record pointer to be remembered on completion, so that the next WRITE begins its data scan at the point where the previous one ceased scanning. If a HOLD option is not specified and the record is of fixed length, the remaining part of the record is padded. HOLD may not be specified for a file which is accessed in more than one task.

The SEGMENT option implies both the CROSS and HOLD options. The expression in the SEGMENT option is converted, if necessary, to a character string. This string serves as a segment mark.

At completion of record construction for the WRITE operation (through one or possibly more records), the specified mark is added to the last record. Emission of the record is not thereby implied. Subsequent data is appended to this record until the maximum record length is met. A maximum SEGMENT length is defined by each implementation.

4. The KEY, NEWKEY, and REGION option may be used when direct access to a particular record is required (see "The KEY

Option," "The NEWKEY Option," and "The REGION Option," in Chapter 7. If a file is declared with the access attribute DIRECT, then the KEY, NEWKEY, or REGION option must be used with each WRITE for that file unless the previous WRITE for that file had the HOLD option.

5. The FROM option specifies that the last record read from the file, the name of which is specified in the FROM option, is to be written on the output file. Data specifications given in a WRITE with the FROM option replace the initial portion of the record read from. If both the FILE option and the FROM option specify the same file, the file must be a direct access file. In this case, the KEY option may be dropped to signify that the new and old keys are the same. FROM may not be used with files accessing in two simultaneous asynchronous procedures.
6. A count is kept of the number of scalar data items transmitted. The COUNT (file-name) built-in function may be used to determine this number of transmitted data items.

Examples:

1. WRITE FILE (INVENTORY), (ITEM.NAME, ITEM.COST) (A (20), F (5,2));

One record of the file named INVENTORY is written under format-directed transmission. The character-string variable ITEM.NAME is written in the first 20 characters of the record. The variable ITEM.COST is converted to

fixed decimal and written in the next 5 characters of the record.

2. WRITE FILE (TABLES), (TABLE.POOL) (F (5)), KEY (Q);

A record is constructed from the contents of the TABLE.POOL array, and a key, the value of the variable Q, is appended. Should the key be identical to an existent key within the data file, the WRITE causes replacement of that keyed record.

3. WRITE FILE (FILEZ), ('FINAL_DATA', X, Y) (A, F (3,2), E (5,2));

Three data items are transmitted to FILEZ (assuming X and Y are scalar variables). The first is the character string of length 10, FINAL_DATA, then the fixed-point form of the value of X, then the floating-point form of the value of Y.

4. WRITE DATA (X, Y, Z);

The values of the three variables X, Y, and Z are transmitted to the standard output file in the data-directed format. If X is floating-point value 3.1141593, Y is fixed-point value 347, and Z is character-string value MATH then the output data stream would appear as the following:

X=3.1141593, Y=347, Z='MATH'

It would be in a form suitable for data-directed input.

A PL/I source program may contain compile-time variables and statements. These may be used to effect program parameterization and modification and iterative program generation. The name "macro" is used to describe such compile-time activity. To differentiate the PL/I macro language from the rest of PL/I, the latter is called the "execution-time" language.

When a source program involves macro activity, the compiler is supplied with text consisting of a skeleton execution-time program with interspersed macro statements. This text is called the "source text." Compile-time activity is performed by a part of the compiler called the "macro processor." The macro processor converts the source text into text representing an execution-time PL/I program. This created text is called "program text."

For ease of learning, the PL/I macro statements and the variables and labels have been made very similar to the corresponding execution-time entities. With one exception, the macro language forms a subset of the execution-time language. (The exception is the CHARACTER attribute in the DECLARE statement.)

Macro statements and skeleton execution-time text can, in general, be interspersed in any way in the source text. Each macro statement begins with a percent sign. Macro statements are recognized by the occurrence of a percent sign other than within a comment or character string. The restrictions on the form of source text are as follows:

If the source program involves any macro activity, the first statement must be the macro DECLARE statement.

Quotation marks and comments delimiters must be matched within the skeleton execution-time program.

Beyond the requirements above there is no requirement that the skeleton execution-time program be syntactically correct.

MACRO VARIABLES

If compile-time activity is to take place, the source text must contain one and only one macro DECLARE statement. The first character of the source text must be

a percent sign. (A comment or blank may appear just after the initial percent sign, but not before.) The percent sign is the beginning of the macro DECLARE statement. This statement is used to declare the macro variables. All macro variables must be declared. The scope of macro variables is the entire source text. Each macro variable and macro label must have a unique name.

THE MACRO DECLARE STATEMENT

General format:

```
% DECLARE macro-declaration
    [, macro-declaration] ...;
```

A "macro declaration" may be either of the following:

1. macro-variable-specification
2. (macro-variable-specification
[,macro-variable-specification]...)
attribute...

A "macro variable specification" is of form:

```
identifier[attribute]...
```

The "identifier" is the name of the macro variable to be declared. Attributes are associated with it in the same way as for the execution time DECLARE statement. In particular, the same rules apply for conflicting attributes. However, at most, one level of factoring is permitted. The following are the permissible attributes on a macro DECLARE statement:

FIXED

This specifies that the macro variable is an integer of implementation defined precision.

CHARACTER (decimal-integer-constant)

This specifies that the variable is a fixed-length character string. The length is given by the decimal integer constant.

CHARACTER VARYING

No length may be given. This attribute specifies that the macro variable is a varying-length character string. The length of the variable is defined as the

length of the last value assigned to it. Initially the length is undefined. (Storage is reserved for the variable as it is needed.)

INITIAL
(optionally-signed-decimal-integer-constant)

This attribute may be given only to macro variables with the FIXED attribute. The specified constant is assigned to the macro variable(s) to which the attribute applies. The assignment is performed at the start of compile-time activity.

INITIAL (character-string-constant)

This attribute may be given only to macro variables of character data type. The specified constant is assigned to the macro variable(s) to which the attribute applies. The assignment is performed at the start of compile-time activity.

The first three attributes in the above list are data attributes. One and only one of these must apply to each macro variable.

MACRO EXPRESSIONS

Macro expressions may take one of the following forms:

1. [+ | -] operand [{ + | - | / | * } operand]
2. operand || operand [|| operand] ...

An operand appearing alone in form 1 or the operands in form 2 may be any of the permitted operands. Operands used with a sign or an arithmetic operator must be decimal integer constants or fixed macro variables. Expressions of form 2 may involve conversion from type integer to type character. The conversion is performed according to list-directed transmission rules (see "List-Directed Specifications," in Chapter 7).

Macro expressions are subdivided into fixed and character expressions. Macro expressions involving macro variables that have not been assigned a value are in error.

Macro expressions are evaluated in exactly the same way as execution-time expressions.

EXECUTABLE MACRO STATEMENTS

The executable macro statements are enumerated below. Any executable macro statement may optionally be preceded by one or more labels, each consisting of an identifier.

THE MACRO ASSIGNMENT STATEMENT

General format:

```
% [ label : ] ... macro-variable =  
macro-expression;
```

The statement causes the value of the macro expression to be assigned to the macro variable. If the expression is of fixed type, the variable on the left may be either of fixed or character type. If the expression is of character type, the variable on the left must also be of character type unless the value of the expression is a string that contains a decimal integer constant, optionally signed, and with optional surrounding blanks. In this latter case, the variable may be of fixed type.

All conversions implied in assignment are performed according to list-directed transmission rules (see "List-Directed Data Specifications").

THE MACRO NULL STATEMENT

General format:

```
% [label:]...;
```

Macro null statements are used for placing macro labels in the text.

THE MACRO GO TO STATEMENT

General format:

```
% [label:]... GO TO label;
```

The macro label following the GO TO must appear on another macro statement in the source text. The execution of the macro GO TO statement is described below.

THE MACRO IF STATEMENT

General format:

```
% [label:]... IF macro-expression  
comparison-operator macro-expression  
THEN GO TO label;
```

The six "comparison operators" are listed in "Scalar Expressions," in Chapter 3. The relationship is evaluated to yield a true value or a false value in the same way as for the execution-time IF statement. If the relationship is true, the statement acts as a macro GO TO statement. If the relationship is false, the statement acts as a null statement.

The two macro expressions to be compared may be of differing data types. In this case, the expression of character type is converted to an integer; it must therefore represent a decimal integer, optionally signed and with optional surrounding blanks.

ACTION OF THE MACRO PROCESSOR

The way in which the macro processor creates program text from source text is described below. Initially program text is null.

Initially the macro DECLARE statement is processed to form a list of macro variables. INITIAL values are assigned. The DECLARE statement is then deleted from the source text.

After this, the source text is scanned sequentially, starting from the beginning of the text. The scan acts as follows:

If the name of a macro variable occurs in the source text delimited at each end by any of the special characters listed in Chapter 1, other than the break character, and is not within a character string, a comment, or a macro statement, then the scan behaves exactly as if the name had

been replaced by the current value of the macro variable it represents, converted to a character string if necessary. The source text, however, remains unchanged. The conceptually inserted value is not enclosed in quote marks. The value must not contain unmatched quotation marks or comment delimiters. If the name of a macro variable occurs within conceptually inserted text, then this name is, in turn, conceptually replaced by its value. (Thus the replacement of JIM by JIM+1 would cause the macro processor to go into an infinite loop.)

If the macro variable A currently contains the null string, it may appear between the two characters of a composite operator, e.g., *A* or ,A=. However, the combinations /A* and *A/ are not allowable when A has null value.

If the scan encounters a percent sign, other than in a comment or character constant, a syntactically correct macro statement must follow. (Comments may be embedded in a macro statement as usual.) This macro statement is executed. If the statement involves a GO TO, the scan resumes at the designated statement. Otherwise, the scan resumes with the text following the statement just executed; the macro statement itself is replaced by a blank space.

All text passed over by the scan is added to program text with the following provisions:

1. If a macro statement is encountered by the scan, none of the text from the opening percent sign up to, and including, the closing semicolon is added to program text.
2. The circumstances under which redundant blanks or comments appear in program text are individually defined for each implementation of PL/I.

Macro activity ends when the scan encounters the end of the source text. The program text then is compiled normally. It is impossible for the program text to contain any macro statements.

RELATIONSHIP OF ARGUMENTS AND PARAMETERS

When a procedure is invoked, a relationship is established between the arguments of the invoking statement and the parameters of the invoked entry point.

A parameter may be a scalar, array, or structure name (including a label variable name, a task name, or an event name) that is unqualified and unsubscripted, or it may be a file parameter or an entry parameter.

A file parameter may be used within a procedure wherever a file name may be used; an entry parameter may be used wherever an entry name may be used.

A parameter is accessible in the procedure only if the parameter is in the parameter list of the entry point at which the procedure is invoked.

Parameters must be declared in the invoked procedure; they cannot be declared in outer containing blocks. If no explicit declaration is given, an implicit or contextual declaration is assumed, internal to the invoked procedure.

Parameters cannot be declared with the storage class attributes `STATIC` or `AUTOMATIC`, with scope attributes, or with the `DEFINED` attribute.

A parameter may have the `CONTROLLED` storage class attribute. In this case, the associated argument must also have the `CONTROLLED` attribute.

EVALUATION OF ARGUMENT SUBSCRIPTS

When an argument is a subscripted variable, the subscripts are evaluated before invocation. The specified element is then passed as the argument. Subsequent changes in the subscript during the execution of the invoked procedure have no effect upon the corresponding parameter.

USE OF DUMMY ARGUMENTS

A constructed dummy argument containing the argument value is passed to a procedure if the argument is one of the following:

- a constant,
- an entry name,
- an expression other than a single unparenthesized scalar variable, array variable, or structure variable, or
- an expression whose data attributes may disagree with the declared data attributes of the parameter.

In all other cases the argument as it appears is passed. The parameter becomes identical with the passed argument; thus, changes to a dummy will be reflected in the original argument only if a dummy is not passed.

USE OF THE ENTRY ATTRIBUTE

An `ENTRY` attribute may be specified for the invoked entry name; this `ENTRY` attribute appears in a `DECLARE` statement whose scope includes the invoking block. If an `ENTRY` attribute is not specified in the invoking procedure for the invoked entry name, the attributes of the arguments must agree with those of the corresponding parameters of the invoked entry.

If an `ENTRY` attribute without parameter attribute lists is specified for an identifier, it indicates that the identifier is an entry name. In this case also, the argument and parameter attributes are assumed to agree.

However, if an `ENTRY` attribute with parameter attribute lists is specified for the invoked entry name, then the attributes of the parameter of the invoked entry are assumed to be the same as those specified for it in the `ENTRY` attribute specification. If an argument has data attributes that differ from the corresponding set of attributes defined in the `ENTRY` attribute specification (string lengths are considered to match only if they have the same decimal integer constant as length), then a dummy argument, with the value of the given argument, is constructed by converting the argument to the data attributes defined for the corresponding parameter in the `ENTRY` attribute specification. If conversion is impossible, then the program is in error (e.g., conversion of file name to bit). The dummy argument is then passed to the invoked entry. Dummy arguments have `CONTROLLED` storage class in the invoking procedure. They are allocated immediately

before invocation of the procedure and freed upon return, unless the invocation has a task option, in which case they are freed upon exit from the invoking block.

The asterisk notation may be used in the ENTRY attribute to specify that for varying length strings, or arrays of adjustable dimensions, the current argument bounds or length are to be assumed for the parameter.

Example:

```
A: PROCEDURE;
   DECLARE B ENTRY (FIXED, FLOAT),
           (C,D) FLOAT;
       .
       .
       .
   CALL B(C,D);
       .
       .
       .
   END A;

B: PROCEDURE (P,Q);
   DECLARE P FIXED, Q FLOAT;
       .
       .
       .
   END B;
```

The specification of the ENTRY attribute in procedure A indicates that B has two parameters, the first with attribute FIXED and the second with attribute FLOAT. However, the arguments C and D both have the FLOAT attribute. Since C is to be fixed-point when it is passed to procedure B, a dummy argument is constructed by converting C from floating-point to fixed-point. This dummy argument is then passed to B.

CORRESPONDENCE OF PARAMETERS AND ARGUMENTS

If a parameter of an invoked entry is a scalar, the argument must be a scalar expression. The data attributes of the argument must agree with the corresponding attributes of the parameter.

If a parameter of an invoked entry is an array, the argument must be an array expression. The argument may also be a scalar expression so long as an ENTRY attribute is given for the invoked entry, specifying the dimension attribute for the relevant parameter. Asterisks may not be given in the dimension attribute if the argument is a scalar. In this case, a dummy array argument will be constructed where the value of each element of the array is the value of the scalar expression. The data attributes of the argument must agree with those of the parameter. If

the asterisk notation is not used to specify the dimensions of the parameter in the invoked procedure, the values of the bounds of the array argument must agree with the values of the bounds specified for the parameter in the invoked procedure.

If a parameter is a structure, the argument must be a structure expression. When a structure description is given for a parameter in an ENTRY attribute specification, a scalar expression may be specified as the corresponding argument. A dummy structure argument will then be constructed where the value of each element of the structure is the value of the scalar expression. The data attributes of the elements of the structure argument must match those of the associated parameter as specified in the invoked procedure. The relative structuring of the argument and the parameter must be the same, although the level numbers need not be identical.

If a parameter is a scalar-label variable, the argument must be a scalar-label variable or constant. If a parameter is an array-label variable, the argument must be an array-label variable. If an ENTRY attribute is given for the invoked entry in the invoking procedure, and if the appropriate parameter attribute list specifies that the parameter is a label array, then the argument may be a scalar-label variable or constant; a dummy label array argument will be suitably constructed. A dummy argument is always constructed when the argument is a label constant or label array.

If the argument is a statement label constant, this statement label constant is qualified by an identification of the current invocation of the block containing the label; this information is passed as a dummy argument to the invoked entry.

If a parameter is an entry parameter, the argument must be an entry name or entry parameter. When a parameter is specified as an entry parameter in the parameter description of an ENTRY attribute and is not given data attributes, no default data attributes are assumed. If it is necessary that the entry parameter have data attributes, they may be specified in the parameter description and a check will be made to insure that a correct argument is provided.

If a parameter is a file parameter, the argument must be a file name or file parameter.

An argument passed to a parameter that is a fixed-length string variable or an array must be of fixed length. An argument passed to a parameter that is a varying-

length string variable or an array must be of varying length.

Example:

```
M1: PROCEDURE;
  DECLARE A (10) , AA (10) , AAA (10) ,
    N EXTERNAL;
    .
    .
    .
  N=10; CALL S1 (A,AA,AAA) ;
    .
    .
    .
  END M1;

S1: PROCEDURE (P,PP,PPP) ;
  DECLARE P (10) ,PP (*), PPP (N) ,
    N EXTERNAL;
    .
    .
    .
  END S1;
```

In the above example, P, PP, and PPP are parameters. Procedures M1 and S1 are both external procedures. P is declared with constant bounds; thus, the bounds of any argument associated with P must be 10. PP is declared with the asterisk notation; thus, any one-dimensional argument of the same type may be associated with it. PPP is declared with an adjustable bound; thus, the bound of any argument associated with PPP must be equal to the value of N when S1 is activated. Note that a similar effect would result if S1 were internal to M1 and N were an internal variable declared in M1.

ALLOCATION OF PARAMETERS

A parameter that has no storage class may correspond to an argument of any storage class; if more than one generation of the argument exists, however, the parameter is synonymous only with the generation existing at the point of invocation. A CONTROLLED parameter, however, always must be presented with a CONTROLLED argument; the argument must be an unsubscripted name of CONTROLLED data that is not an element of a structure. The parameter is synonymous with the entire allocation stack of the controlled variable. Thus each reference to the parameter is a reference to the current generation of the associated argument. A controlled parameter may be allocated and/or freed in the invoked procedure, thus manipulating the allocation stack of the associated argument.

Parameters, Bounds and Length

If an argument is a string or an array, the length of the string or the bounds of the array must be declared in the invoked procedure by using the asterisk notation, by giving explicit bounds or length or by declaring the bounds or length as an expression that, when evaluated, gives the appropriate value. The expressions specified for the bounds or length must be formulated according to the rules stated in "Evaluation of Expressions," in Chapter 3.

The number of dimensions and the bounds of the array argument or the length of the string argument must be the same as those of the corresponding parameters. However, the actual bounds or length may not be known at the time the invoked procedure is written; the invoked procedure may assume either that storage has been allocated prior to the invocation or that storage will be allocated explicitly in the procedure for those parameters declared CONTROLLED.

Asterisk Notation for Bounds or Length

The correspondence between argument and parameter in the invoked procedure can be achieved by specifying the length by an asterisk or by specifying each and every bound by an asterisk, thus indicating that the length or bounds are the same as those for the corresponding argument.

If storage has been allocated for an argument, the corresponding parameter in the invoked procedure is assumed to have the same length or bounds as the argument. If the parameter is controlled, further allocations of the data will use these same bounds or length unless different length or bounds are specified in the ALLOCATE statement.

If storage has not been allocated for an argument passed to a parameter declared with the asterisk notation, explicit bounds or length must be declared in an ALLOCATE statement given before another reference to the parameter in the invoked procedure.

Expressions as Bounds or Length

If storage has been allocated for an argument passed to a parameter for which explicit bounds or length are specified, then upon entry to the invoked procedure, any expressions are evaluated and must give values such that the bounds or length of the parameter are the same as the argument. If the parameter is controlled and is subsequently reallocated, these expressions

are again evaluated to give new bounds or length for the new allocation, unless they are specified in the ALLOCATE statement.

If storage has not been allocated for the argument, then, at the point of entry, no requirements are made on the value of the expressions specified for the corresponding parameter bounds or length. These expressions are evaluated at a subsequent point of allocation, unless they are specified in the ALLOCATE statement.

Example:

```
M2: PROCEDURE;
    DECLARE A (10) , AA (25) CONTROLLED;
    .
    .
    .
    CALL S2 (A,AA,10) ;
    .
    .
    .
    END M2;
```

```
S2: PROCEDURE (P,PP,N) ;
    DECLARE PP (*) CONTROLLED, P (N) ,
           Q (25) , S (5) ;
    .
    .
    .
    PP = Q;
    .
    .
    .
    ALLOCATE PP (5) ;
    .
    .
    .
    PP = S;
    .
    .
    .
    END S2;
```

PROLOGUES

On entering a block, certain initial actions are performed, e.g., allocation of storage for automatic variables. These initial actions constitute the prologue.

On entry to the prologue, the following items are available for computation:

1. variables declared outside the block and known within it
2. variables declared STATIC and known within the block
3. arguments passed to the block
4. The most recent generations of controlled variables known within the block

The prologue makes available for computation all the other variables known within the block as follows:

5. automatic variables declared in the block
6. Defined variables declared within the block

In making these items available, the prologue may need to evaluate expressions defining lengths, bounds, iteration factors, and initial values. Such expressions may depend on items of 1, 2, 3 or 4. They may also be dependent on items 5 and 6 under the following circumstances: If an item is referred to in an expression and the allocation or initialization of a second item depends on that expression, then that first item must in no way be dependent on the second item for its own allocation and initialization. Further, the first item must in no way be dependent on any other item that so depends on the second item.

Example:

The following is illegal:

```
DECLARE (A (M) INITIAL (1) ,
        M INITIAL ((A (1)) ('ABC'))
        AUTO;
```

The evaluations must not invoke abnormal functions. The entry invoked with the INITIAL CALL attribute may be abnormal only in that it sets the data being initialized. The sequence in which the evaluations refer to any abnormal data is not defined.

Function calls within the evaluations must not refer to items being made available by the prologue.

DATA ALLOCATION ACROSS TASKS

The scope of an identifier declared in an attaching task may include the attached task. Thus, the WAIT statement should properly be used in the attaching task to avoid freeing storage allocated in the attaching task and used in the attached task.

An attached task has almost the same access to the attaching task's data as it would have if it were executed synchronously; however, when it is attached, only the generations of CONTROLLED variables current at the time of attachment are passed to the attached task. Subsequent allocations in the attached task are known only within the attached task; subsequent allocations in the attaching task are known only within

the attaching task. A task may only free storage that it has allocated. All storage allocated within a task is destroyed when that task is completed.

Allocation of Task and Event Names

Like variables, task names and event names have scope and storage class attributes. Storage will be allocated for task and event names in the same manner as for variables (by virtue of either an explicit or contextual declaration). If a given task is active and there is a task or event name associated with the task, then storage must not be released for the name until the task is terminated.

ABNORMALITY

The ABNORMAL, NORMAL, USES, and SETS attributes are provided in PL/I to enable the compiler to generate optimized code.

In the absence of any information, the following assumptions are made:

1. All external function references are normal.
2. All other procedure references are abnormal.
3. All variables are normal.

A variable is said to be abnormal if its value may be altered or otherwise accessed without an explicit indication. Thus, for example, the appearance of a variable name on the left side of an assignment statement, in the data list specification of a READ or GET statement, or as an argument to an abnormal function or procedure (see below) indicates a predictable situation where the variable may change its value. However, when the variable is subject to change by the occurrence of an ON-condition, or if it is subject to change in a procedure invoked with the TASK option (see "Asynchronous Operations and Tasks"), then there is no way to predict the point at which the change in value will occur or, in fact, if it will occur.

Such possibilities cannot always be recognized contextually. Furthermore, if a portion of a source program contains several references to such a variable, the order in which the indicated operations are executed becomes significant. (For example, if B is abnormal, the expression B + B is not necessarily equivalent to the expression 2 * B.)

The implication is that the programmer expects the operation to be performed in a particular order. Such variables must therefore be declared ABNORMAL, to inhibit the optimization of such portions of a source program.

A procedure may possess varying degrees of abnormality. A procedure is said to be "definitively abnormal" if it, or any procedures invoked by it, accesses, modifies, allocates, or frees external data or modifies, allocates, or frees arguments. In addition, an internal procedure is abnormal if it, or any procedures invoked by it, accesses, modifies, allocates, or frees any variables known in the invoking block. Such procedures are only definitively abnormal because the exact nature of their abnormality is described by the USES and SETS attributes, thus inhibiting some, but not all, optimization in the neighborhood of a reference to the procedure (see "The USES and SETS Attributes" in Chapter 4).

However, if a procedure is "completely abnormal," all optimization of successive references must be inhibited. A procedure is completely abnormal if it, or any procedures invoked by it, does any of the following:

1. returns inconsistent function values for identical argument values
2. maintains any kind of a history
3. performs input or output operations
4. returns control from the procedure by means of a GO TO statement

The ABNORMAL attribute (described in Chapter 4) is used to describe such a procedure. It may also, of course, be used to describe a procedure that is "definitively abnormal."

When abnormality is specified, the order of execution becomes significant. In particular, if an expression contains a reference to an abnormal function that may affect values in other parts of the expression, the value of the expression will, in general, depend upon the order in which data is accessed (see "Order of Evaluation of Expressions," in Chapter 3).

If an ABNORMAL procedure, referred to in a statement, allocates or frees controlled data that has been referred to elsewhere in the same statement, then the effect of the statement is undefined.

PROGRAMS

A program constitutes a domain of reference for external identifiers and a

domain of persistence for data. A program consists of a set of external procedures, linked to each other in the following sense. Every external identifier has constant meaning across all external procedures in a program. Values of data known in a program exist only as long as the program is available for execution.

When a program becomes available for execution, its `STATIC` data is initialized according to `INITIAL` attributes in the program; all stacks of `CONTROLLED` data are empty; no data is available to the `RESTORE` statement; these values disappear when the program is made unavailable for execution.

A program is made available for execution either by the operating environment as a major task, or by the `FETCH` statement. It is made unavailable by termination of the major task or by the `DELETE` statement.

There is complete data isolation between programs except for arguments passed from one to the other and for files with the same `TITLE`. This is true even if both programs contain apparently identical external procedures.

APPENDIX 1: BUILT-IN FUNCTIONS

ARITHMETIC GENERIC FUNCTIONS

The generic functions listed in this section return a value of type coded arithmetic. The arguments may, unless otherwise specified, be any expressions. If necessary they will be converted to type coded arithmetic before the function is invoked according to the rules stated under "Type Conversion," in Chapter 3. Also certain conversions of arithmetic characteristics will be performed before the function is invoked, where this is explicitly defined to be the case for particular functions below. Where conversion to highest characteristics is specified, these are determined by the rules for mixed characteristics, as explained in Chapter 3, applied to the arguments. Where reference is made to an argument, it should be taken to mean the converted argument when an argument that is not coded arithmetic has been specified. The magnitude of a complex number is the positive square root of the sum of the squares of the real and imaginary parts where this value has the base and scale of the complex number and the mode REAL.

Name Arguments and Function Value

ABS
 Arguments: One is given.
 Function value = absolute value of argument, i.e., positive value of real argument, positive magnitude of complex. The mode is REAL. Base, scale, and precision are those of the argument, unless the argument is fixed complex, in which case the precision is $\text{MIN}(N, p+1), q$ for an argument of precision (p, q) .

MAX
 Arguments: Two or more are given. Complex arguments are not permitted.
 Function value = value of maximum argument, converted to highest characteristics of all arguments specified. If the arguments are FIXED of precisions $(p_1, q_1), (p_2, q_2), \dots, (p_n, q_n)$, the resulting precision is $(\text{MAX}(p_1, \dots, p_n), \text{MAX}(q_1, \dots, q_n))$.

MIN
 Arguments: Two or more are given. Complex arguments are not permitted.

Function value = value of minimum argument, converted to highest characteristics of all arguments specified. If the arguments are FIXED of precisions $(p_1, q_1), (p_2, q_2), \dots, (p_n, q_n)$, the resulting precision is $(\text{MAX}(p_1, \dots, p_n), \text{MAX}(q_1, \dots, q_n))$.

MOD
 Arguments: two are given, x and y. Base and scale of the arguments are converted to the higher characteristics of the pair. Complex arguments are not permitted.
 Function value = positive remainder after division of x by y to yield an integer quotient. The mode is REAL; base and scale are those of the converted arguments. Precision for FLOAT is the higher of the precisions of the arguments, and for FIXED is defined as follows:
 Let the precision of x be (p, q) and the precision of y be (r, s) . The resulting precision is $(\text{MIN}(N, r-s+u), \text{MAX}(q, s))$.

SIGN
 Arguments: One is given. Complex arguments are not permitted.
 Function value = integer 1 if argument >0 ; = 0 if argument = 0; = -1 if argument <0 . The result is fixed binary with default precision.

FIXED
 Arguments: Three are given. The second and third are optional decimal integer constants specifying the number of digits after the decimal or binary point and the scale factor of the result. If omitted, the second argument assumes a value specified by each implementation, the third assumes zero.
 Function value = first argument converted to fixed-point scale with precision as specified but base and mode unchanged.

FLOAT
 Arguments: Two are given. The second is an optional decimal integer constant specifying the precision of the result. If omitted, a value specified by each implementation will be assumed.

Function value = first argument converted to FLOAT scale with precision as specified but base and mode unchanged.

FLOOR

Arguments: One is given, x. A complex argument is not permitted.

Function value = largest integer not exceeding x. Base, scale, and mode are those of the converted argument. Precision of result for x FIXED (p,q) is (MIN(N, MAX(p-q+1, 1)), 0).

CEIL

Arguments: One is given, x. A complex argument is not permitted.

Function value = smallest integer not exceeded by x. Base, scale, and mode are those of the converted argument. Precision of result for x FIXED (p,q) is (MIN(N, MAX(p-q+1, 1)), 0).

TRUNC

Arguments: One is given, x. A complex argument is not permitted.

Function value = FLOOR (x) if $x \geq 0$, = CEIL (x) if $x < 0$. Base, scale and mode are those of the converted argument. Precision of result for x FIXED (p,q) is (MIN(N, MAX(p-q+1, 1)), 0).

BINARY

Arguments: Three are given. The second and third are optional decimal integer constants specifying the binary precision of the result. If the scale is FIXED, all three are required; if the scale is FLOAT, the third is not required. If both the second and third arguments are omitted, the precision of the result is as defined for base conversion in Chapter 3.

Function value = first argument converted to binary base with scale and mode unchanged.

DECIMAL

Arguments: Three are given. The second and third are optional decimal integer constants specifying the decimal precision of the result. If the scale is FIXED, all three are required; if the scale is FLOAT, the third is not required. If both the second and third arguments are omitted, the precision of the result is as defined for base conversion in Chapter 3.

Function value = first argument converted to decimal base with scale and mode unchanged.

PRECISION

Arguments: Three are given. The second and third are decimal integer constants specifying the precision of the result. If the scale is FIXED, all three are required; if the scale is FLOAT, the third is not required.

Function value = first argument converted to specified precision. Base, scale, and mode are unchanged.

ADD

Arguments: Four are given. The third and fourth are decimal integer constants specifying the precision of the result. If the scale of the result is FIXED, all four are required; if the scale is FLOAT, the fourth is not required.

Function value = the sum of the first and second arguments. Base and scale of the result are the higher of those of the first two arguments. Precision is as specified.

MULTIPLY

Arguments: Four are given. The third and fourth are decimal integer constants specifying the precision of the result. If the scale of the result is FIXED, all four are required; if the scale is FLOAT, the fourth is not required.

Function value = the product of the first and second arguments. Base and scale of the result are the higher of those of the first two arguments. Precision is as specified.

DIVIDE

Arguments: Four are given. The third and fourth are decimal integer constants specifying the precision of the result. If the scale of the result is FIXED, all four are required; if the scale is FLOAT, the fourth is not required.

Function value = the result of dividing the first argument by the second. Base and scale of the result are the higher of those of the first two arguments. Precision is as specified.

COMPLEX

Arguments: Two are given. The first is the real part, the sec-

ond is the imaginary part.
 Function value = complex number formed from the two arguments. Base, scale, and precision of result are the highest characteristics of those of the arguments.

REAL

Arguments: One is given, complex value.
 Function value = real part of argument. Base, scale, and precision are unchanged.

IMAG

Arguments: One is given, complex value.
 Function value = imaginary part of argument. Base, scale, mode, and precision are unchanged.

CONJG

Arguments: One is given, complex value.
 Function value = conjugate of the argument. Base, scale, mode, and precision are unchanged.

FLOAT ARITHMETIC GENERIC FUNCTIONS

The following generic functions may have as arguments any expression. This expression will be converted to floating point before the function is invoked. The result will be of scale FLOAT with the precision and base of the converted argument. If the mode of the argument is COMPLEX, the mode of the result will be COMPLEX. The following functions are defined only for REAL arguments: LOG2, LOG10, ATAND, TAND, SIND, COSD, ERF, ERFC, and ATAN with two arguments.

The following table specifies the meaning of these functions for real arguments:

<u>Function Reference</u>	<u>Function Value</u>
EXP (x)	exp (x)
LOG (x)	ln (x). Error if $x \leq 0$.
LOG10 (x)	$\log_{10} (x)$. Error if $x \leq 0$.
LOG2 (x)	$\log_2 (x)$. Error if $x \leq 0$.
ATAND (x)	arctan (x) in degrees.
ATAN (x)	arctan (x) in radians.
	ABS (arctan (x)) $< \pi/2$.
TAND (x) degree argument	tan (x)
TAN (x) radian argument	tan (x)
SIND (x) degree argument	sin (x)
SIN (x) radian argument	sin (x)
COSD (x) degree argument	cos (x)
COS (x) radian argument	cos (x)
TANH (x) radian argument	tanh (x)
ERF (x)	Two divided by square root of pi, multiplied by the integral from 0 to x of EXP (-t ²) with respect to t.
SQRT (x)	The positive square root of x.
ERFC (x)	1 - ERF (x)
COSH (x) radian argument	cosh (x)
SINH (x) radian argument	sinh (x)
ATANH (x)	arctant (x). Error if ABS (x) ≥ 1 .
ATAN (y, x)	The arguments are converted to the highest characteristics of the pair. The value is: arctan (y/x) if $x > 0$ $\pi/2$ if $x = 0, y > 0$ error if $x = 0, y = 0$ $-\pi/2$ if $x = 0, y < 0$ $\pi + \arctan (y/x)$ if $x < 0, y \geq 0$ $-\pi + \arctan (y/x)$ if $x < 0, y < 0$
ATAND (y, x)	ATAN (y, x) in degrees, i.e. $(180/\pi) * \text{ATAN} (y, x)$

With complex mode many of these mathematical functions are formally multiple-valued, so the following table defines the principal values which are returned by the built-in functions. Here $Z = x+iy$ is the argument, and $w = u+iv$ is the value.

<u>Function Reference</u>	<u>Function Value</u>
EXP (Z)	exp (Z)
LOG (Z)	Log (Z) , where $-pi < v \leq pi$. Error if Z=0.
ATAN (Z)	(LOG ((1+Z) / (1-Z))) / 2. Error if Z= +1 or -1.
ATAN (Z)	iATANH (iZ) . Error if Z= +1i or -1i.
SIN (Z)	sin (Z) =sin (x) cosh (y) + icos (x) sinh (y)
COS (Z)	cos (Z) =cos (x) cosh (y) - isin (x) sinh (y)
SQRT (Z)	Z** (1/2) . Either REAL (Z) >0, or REAL (Z) =0 and IMAG (Z) ≥0.
COSH (Z)	cosh (Z) =cosh (x) cos (y) + isinh (x) sin (y)
SINH (Z)	sinh (Z) =sinh (x) cos (y) + icosh (x) sin (y)

when converted to integer, the third is optionally any expression having value j when converted to integer.

The function value is defined as follows:

Let k be the length of the first argument.

If $i > k$, the value is the null string.

If $i \leq k$, the value is that substring beginning at the Mth character or bit of the first argument, and extending N characters or bits, where M and N are defined by:

$$M = \max(i, 1)$$

$$N = \max(0, \min(j + \min(i, 1) - 1, k - M + 1)), \text{ if } j \text{ is specified.}$$

$$N = k - M + 1, \text{ if } j \text{ is not specified.}$$

STRING GENERIC FUNCTIONS

The generic functions listed in this section may be used for manipulation of strings. The arguments specified as strings may be any expression. If the argument is arithmetic, it will be converted to bit string (if binary base) or character string (if decimal base) before the function is invoked.

<u>Name</u>	<u>Arguments and Function Value</u>
-------------	-------------------------------------

BIT	Arguments: Two are given. The second is an optional decimal integer specifying the size of result. Function value = first argument converted to type bit string. If the size is unspecified, the size of the result will be a function of the first argument characteristics (see "Type Conversion," in Chapter 3).
CHAR	Arguments: Two are given. The second is an optional decimal integer specifying the size of result. Function value = first argument converted to type character string. If the size is unspecified, the size of the result will be a function of the first argument characteristics (see "Type Conversion," in Chapter 3).
SUBSTR	Arguments: Three are given. The first is a string, the second is any expression having value i

INDEX

Arguments: Two are given. If both arguments are bit strings, no conversion occurs, otherwise conversion to character string is performed.

Function value = decimal integer with implementation defined precision giving:

- The index of the first element of the first argument such that starting at this element the second argument appears as a substring.
- Zero, if no such index satisfying (a) exists, or if either of the arguments is of zero length.

LENGTH

Arguments: One is given, a string.
Function value = fixed binary integer of default precision giving current length of argument.

HIGH

Arguments: One is given, a decimal integer constant.
Function value = character string of the length specified and composed of the highest characters of the data character set.

LOW

Arguments: One is given, a decimal integer constant.
Function value = character string of the length specified and composed of the lowest characters of the data character set.

REPEAT

Arguments: Two are given. The first is a string and the second a decimal integer constant n.

Function value = string argument concatenated with itself n times, giving a total of $n+1$ terms in the concatenation. If n is zero or negative, the result is the argument itself.

UNSPEC

Arguments: One is given.
Function value = bit string which is the internal coded representation of the argument. The length is an implementation defined function of the argument characteristics. If the argument is a varying length string, the result is adjusted to be just large enough to hold the internal form of the argument expression after conversion to bit string.

BOOL

Arguments: Three are given, bit string X , Y , and W . W is converted if necessary, to a bit string of length $4, n^1n^2n^3n^4$. This string defines which of the 16 possible boolean functions is desired, in the manner implied below.

Function value = bit string Z where if X and Y are of different lengths, the shorter is extended with zeros, and Z is of the longer length. The following table relates the j th bit of Z to the j th bits of X and Y .

X_j	Y_j	Z_j
0	0	n^1
0	1	n^2
1	0	n^3
1	1	n^4

BUILT-IN FUNCTIONS FOR MANIPULATION OF ARRAYS

The following built-in functions have array expression arguments and return scalar values. In the following functions X is any array expression unless otherwise specified.

Function Reference	Function Value
SUM (X)	A scalar value equal to the

sum of all the elements of X . Precision, mode and base are that of argument elements. (The argument is converted to arithmetic FLOAT before the function is invoked.)

PROD (X)
ALL (X)

As above but product.
The argument is converted to bit string. The result is a bit string of the length (or max length if variable) of the elements of X . The i th bit of the result is 1, if the i th bits of all the elements of X are 1. Otherwise 0.

ANY (X)

As above, i th bit of the result is 1 if any of the i th bits of the elements of X are 1. If all 0, then the result bit is 0.

POLY (X,Y)

$POLY (A, X); A (M:N)$ and $X (P:Q)$ are vectors.
Result is

$$A (M) + \sum_{J=1}^{N-M} (A (M+J) * \prod_{I=0}^{J-1} X (P+I))$$

If $Q-P < N-M-1$, then $X (P+I) = X (Q)$ for $P+I > Q$.

A scalar second operand X is interpreted as a vector with one element, $X (1)$. The function result is then

$$\prod_{J=0}^{N-M} A (M+J) * X ** J$$

The characteristics of the result are the higher of those of the arguments (after conversion to arithmetic type) except for scale, which is always FLOAT.

LBOUND (X,S)

S is a scalar expression which is converted to a binary integer n , of default precision. The function value is an integer of default precision giving the current lower bound of the n th dimension of X .

HBOUND (X,S)
DIM (X,S)

As above but higher bound.
 S is as above. The function value is a binary integer n of default precision giving the current extent of the n th dimension of X .

NOTE: The functions LBOUND, HBOUND, and DIM are not defined if the argument X is unallocated, if it has less than n dimensions, or if n≤0.

SCAN (A,I, 'operator') A is any array expression; I is a decimal integer constant. The third argument may be any binary operator in quotes. The function value is defined by the value of TEMP on exit from the following loop:

```
TEMP = A (*,.....,*, LBOUND (A,I),
          *,.....,*) ;
DO J = LBOUND (A,I) + 1 TO HBOUND
  (A,I) ;
TEMP = TEMP operator A (*,.....,*,J,
          *,.....,*) ;
END;
```

TEMP has dimensions N-1 where A has N. The bounds of TEMP are the first (I - 1) and the last (N - I) of A. TEMP has the base, scale, mode and precision of A if arithmetic, and the length of elements of A, if string.

ARRAY AND STRUCTURE BUILT-IN FUNCTIONS

All of the built-in functions listed under "Arithmetic Generic Functions" and "String Generic Functions" in this appendix may have array or structure expressions as arguments, except where decimal integer constants are required. They yield an array or structure of the same dimension bounds or structuring as the argument--the function being performed on each element. The rules are the same as those for the scalar functions.

CONDITION BUILT-IN FUNCTIONS

The following built-in functions (with no arguments) are available to allow investigation of interrupts arising from enabled ON conditions. They may be referred to only in ON units.

<u>Function Reference</u>	<u>Function Value</u>
---------------------------	-----------------------

ONPOINT	An integer, being the value of the I/O buffer pointer when the I/O condition arose.
---------	---

ONLOC	A character string of variable length, being the name of the procedure in which the condition arose.
ONFIELD	A character string of variable length, being the contents of the field being processed when the input condition arose.
ONCHAR	A character string of length 1, being the character which caused an input conversion error.
ONCODE	A binary integer of default precision whose value depends on a detected error. Each of the following error categories has a set of contiguous code values: I/O errors Conversion errors Control program errors Built-in function errors

OTHER BUILT-IN FUNCTIONS

<u>Function Reference</u>	<u>Function Value</u>
---------------------------	-----------------------

DATE	Character string of length six of the form YYMMDD, where YY is year, MM is month, DD is day.
------	--

TIME	Character string of length nine of the form HHMMSSTTT, where HH is hours, MM is minutes, SS is seconds, TTT is milliseconds.
------	--

ALLOCATION (X)	X is a CONTROLLED major structure or unsubscripted array or scalar variable not in a structure. The function value is '1'B if storage has been allocated for X and '0'B if not.
----------------	---

POINT (Filename)	The value of this function is a decimal integer of precision (n), where n is implementation defined. It specifies the current position of the pointer relative to the start of the current logical record for the named file.
------------------	---

COUNT (Filename)	The value of this function is a binary fixed-point integer of default precision. It returns a value that is the number of scalar data items transmitted during the last
------------------	---

read or write operation on the specified file.

ROUND (Expression,
Decimal Integer Constant)

The expression may be scalar, array, or structure. The function value is the expression value rounded on the n'th digit after the point where n is the value of the integer. (Binary digits if binary base, decimal if decimal base.) If the expression is of string type, the function value is the string value unmodified. Floating point rounding is a bias removal rather than systematic rounding; the decimal point is assumed at the left. Base, scale, mode and precision of the value are those of argument. For fixed-point scale, digits after the rounded digit are set to zero.

STRING (Structure Name)

The argument must be a packed structure composed either of

all bit strings and numeric fields of binary base, or character strings and numeric field of decimal base. The function value is a string, being the concatenation of all the structure elements. The argument must not be a parameter.

EVENT (scalar event name)

This function will return the value 'O'B OR '1'B, depending on the current status of the referenced event name (see "Asynchronous Operations and Tasks," in Chapter 6 and "The WAIT Statement," in Chapter 8).

PRIORITY (scalar task name)

This function will return the priority of the named task relative to the priority of the task in which the function is evaluated (see "Asynchronous Operations and Tasks," in Chapter 6 and "The WAIT Statement," in Chapter 8).

DIGIT POINT AND SUBFIELD DELIMITING CHARACTERS

- 9 Specifies that the associated field position will contain any decimal digit.
- 1 Specifies that the associated field position contains a binary digit. This character may not appear in a picture with either 2 or 3.
- 2 Specifies that the associated field position contains a binary digit, being part of a binary value in 2's complement notation. This character may not appear in a picture with either 1, 3, or S.
- 3 Specifies that the associated field position contains a binary digit, being part of a binary value in 1's complement notation. This character may not appear in a picture with either 1, 2, or S.
- V Specifies that a decimal or binary point should be assumed to appear at this point in the associated field. It does not specify a character in the field.
- K Specifies that the exponent subfield should be assumed to follow the point in the field associated with the K. It does not specify a character in the field.
- E Specifies that the associated field position will contain the letter E, indicating the start of the exponent subfield.

ZERO SUPPRESSION CHARACTERS

A leading zero in a numeric subfield is a zero to the left of the actual occurrence of the digits 1 to 9 in the subfield. The leftmost of these latter digits and all digits in the subfield following it, are significant digits (including any zeros). Picture characters are provided for zero suppression, leading zero suppression, and the replacement of these zeros by blanks or asterisks.

- Z Specifies a conditional digit position. If the associated field position involves a leading zero it will be represented in the field by a blank, otherwise the digit will appear. The character may not appear to the right of 9 T I R or a drifting string in a

- subfield. It may not appear with * in a subfield.
- * Specifies a conditional digit position. If the associated field position involves a leading zero it will be represented in the field by *, otherwise the digits will appear. The character may not appear to the right of 9 T I R or drifting string in a subfield. It may not appear with Z in a subfield.
- Y Specifies a conditional digit position. If the associated field position involves a zero (leading or otherwise) it will be represented in the field by a blank, if it involves a digit other than zero that digit will appear.

DRIFTING EDITING SYMBOLS

The following picture characters may be static or drifting:

<u>Character</u>	<u>Name</u>
{ S }	sign characters
{ + }	
{ - }	
\$	currency symbol

The static use of these characters specifies that there is a field position where a sign, a currency symbol, or a blank always appears. The drifting use specifies that leading zeros may be suppressed, and the suppressed positions may contain blanks. In this case, the rightmost suppressed position associated with the picture character will contain a sign, a blank, or a dollar sign.

A drifting character is specified by multiple use of that character in a picture subfield. Thus, if a subfield contains one dollar sign, it is interpreted as static; if it contains more than one, as drifting. The drifting character must be specified in each position through which it may drift.

Drifting characters must appear in strings. A string is a sequence of the same drifting character, optionally containing interspersed editing characters comma (,), point (.), slash (/), or V or B. Picture characters slash, comma, point, and B following the last drifting symbol of the string are considered part of the string. However, a following V terminates the string and is not part of it. A subfield

may only contain one drifting string. The picture characters * and Z may not appear to the right of a drifting string in a subfield.

The field position associated with the character slash, comma, point, and B appearing in a drifting string will contain one of the following:

1. slash, comma, point, or blank if a significant digit has appeared to the left
2. the drifting symbol, if the next position to the right contains the leftmost significant digit of the subfield
3. blank, if the leftmost significant digit of the subfield is more than one position to the right

If a drifting string contains the drifting character n times, then the string is associated with $n - 1$ conditional digit positions. The field position associated with the leftmost drifting character may only contain the drifting character or blank, never a digit. If a drifting string is specified for a subfield, the other potentially drifting characters may only appear once to the left of the string in the subfield, i.e., the other characters represent a static sign or dollar sign.

If a drifting string contains a V, then all digit positions of the subfield following the V must also be part of the drifting string.

If one of the characters Z or * follows the V in a subfield, then all digit positions in the subfield following the V must be Z or asterisk (*).

In the case where all digit positions after the V contain suppression characters, suppression will only occur where all the fraction digits are zero. The resulting field will then be all blanks or asterisks. If there are any significant fraction digits they all will appear unsuppressed.

DRIFTING CHARACTERS

\$ If this character appears more than once in a subfield it is a drifting character, otherwise it is a static character. The static character specifies that the character \$ be placed in the associated field position. The static character must appear either to the left of all digit positions in a subfield or to the right of all digit positions in a subfield. See details above for the drifting use of the

character.

- S Specifies the sign character + if the field value is ≥ 0 , otherwise -. The character may be drifting or static. The rules are identical to those for the dollar sign.
- + Specifies the sign character + if the field value is \geq to 0, otherwise blank. The character may be drifting or static. The rules are identical to those for the dollar sign.
- Specifies the sign character - if field value is < 0 , otherwise blank. The character may be drifting or static. The rules are identical to those for the dollar sign.

EDITING CHARACTER

- B Specifies that a blank appear in the associated field position.

CONDITIONAL EDITING CHARACTERS

, If the subfields in which the comma appears involve no zero suppression, that character specifies that a comma will appear in the associated field position. If zero suppression is involved the comma will appear only if there is an unsuppressed digit to the left of the comma position in the subfield. If there is no such unsuppressed digit, the associated field position will contain a character that depends on the first digit (conditional or otherwise) picture character preceding the comma.

If the preceding character is an asterisk the field position will contain an asterisk.

If the preceding character is a drifting sign or dollar sign the action taken will be identical to that which would have occurred if the picture specification had contained the drifting character in place of the comma.

If the preceding picture character is anything other than the above, the field position associated with the comma will contain a blank.

- / Exactly as comma, but a slash will appear when indicated.
- . Exactly as comma, but a point will appear when indicated.

SIGN CHARACTERS

Digit characters in numeric fields may contain an overpunched sign. The following picture characters are used to specify overpunching:

- T Specifies that the associated field position will contain a digit overpunched with the sign of the containing subfield.
- I Specifies that the associated field position will contain a digit overpunched with + if the containing subfield is ≥ 0 ; otherwise it will contain the digit with no overpunching.
- R Specifies that the associated field position will contain a digit overpunched with - if the containing subfield is < 0 ; otherwise it will contain the digit with no overpunching.

The two character picture items CR and DB may be used to reflect the sign of REAL numeric fields.

CR Specifies that the associated field positions will contain the letters CR if the containing field value is < 0 . Otherwise the positions will contain two blanks. The characters CR may appear only to the right of all digit positions of a field.

DB As CR, except that a DB appears if the containing field value is greater, or equal to, zero.

SCALING FACTOR SPECIFICATION

- F Specifies that the optionally signed decimal integer enclosed in parentheses following the picture character F in the picture string is the scaling factor (see "The PICTURE Attribute," in Chapter 4).

STERLING PICTURES

The following additional characters are provided for use in sterling pictures.

- 8 Specifies the position of a shilling digit in BSI single-character representation.
- 7 Specifies the position of a pence digit in BSI single-character representation.
- 6 Specifies the position of a pence digit in IBM single-character representation.
- G Specifies the start of a sterling picture. It does not specify a character in the numeric field.
- H Specifies that the associated field position contains the shilling character S.
- P Specifies that the associated field position contains the pence character P.

PICTURES FOR CHARACTER STRINGS

A form of picture may be given for character strings. The following are used to indicate the form:

- A The associated field position may contain any alphabetic character or blank.
- X The associated field position may contain any character.
- 9 The associated field position may contain any decimal digit or blank. At least one X or A must appear in the picture.

APPENDIX 3: ON-CONDITIONS

The ON-conditions are those conditions that may be specified in the ON statement. These conditions are also specified in SIGNAL and REVERT statements.

For each condition name, the description in this appendix includes the circumstances under which the condition occurs, the standard system action that would be taken in the absence of programmer-specified action, and, where applicable, the result. ("Standard system action" does not refer to any operating system but to standard action prescribed for the language.)

For the conditions OVERFLOW, UNDERFLOW, ZERODIVIDE, CONVERSION, or FIXEDOVERFLOW, an interrupt action will always take place on occurrence of the condition unless the occurrence is in a calculation lying within the scope of a prefix specifying NOOVERFLOW, NOUNDERFLOW, NOZERODIVIEE, NOCONVERSION, or NOFIXEDOVERFLOW. For the conditions SIZE, SUBSCRIPTRANGE, or CHECK (identifier list), an interrupt will not take place on occurrence of the condition unless the occurrence is in a calculation lying within the scope of a prefix specifying the condition. (See "Prefixes," in Chapter 1).

For any other condition, whose name may not be used in a prefix, an interrupt always will result from the occurrence of the condition.

CLASSIFICATION OF CONDITIONS

The ON-conditions are classified as follows: computational conditions, input/output conditions, program-checkout conditions, programmer-named conditions, and system-action conditions.

The computational conditions are associated with data handling, expression evaluation, and computation.

The input/output conditions are associated with data transmission.

The program-checkout conditions facilitate debugging of programs.

The programmer-named conditions permit the programmer to use conditions of his own naming. These conditions are raised only by a SIGNAL statement.

The system-action conditions provide facilities to the programmer to extend the standard system action taken after the occurrence of a condition or at the completion of a program.

COMPUTATIONAL CONDITIONS

CONVERSION: This condition is raised whenever an illegal internal conversion is attempted on character string data. The condition will be raised for such errors as characters other than 0 or 1 in conversion to bit string, characters not permitted in conversion to numeric field, or illegal characters in conversion to arithmetic. The CONVERSION condition is analogous to the EDIT condition for input/output.

Result: Undefined.

Standard system action: Comment and raise the ERROR condition.

FIXEDOVERFLOW: This condition occurs during fixed-point arithmetic operations if the results of these operations exceed N, the maximum field width as defined by the implementation. See SIZE for a related condition that occurs on assignment.

Result: Truncation on the left to size N.

Standard system action: Comment and continue.

OVERFLOW: This condition occurs when the exponent of a floating-point number exceeds the permitted maximum, as defined by the implementation.

In some implementations, the condition may be detected by hardware interrupt, in others by special coding.

Result: Maximum positive value.

Standard system action: Comment and raise the ERROR condition.

SIZE: This condition is raised by internal conversions between data types, or between differing bases, scales, or precisions. The condition arises when a value is assigned to a data item, with a loss of high-order bits or digits. The error situations are similar to those listed under the input/output condition FIELD_OVERFLOW. The string assignments, analogous to

item 1 under "Input" and item 1 under "Output" (see "Input/Output Conditions" in this appendix) of FIELD_OVERFLOW, do not raise the SIZE condition. However, assignments analogous to the other listed situations do raise the SIZE condition.

The SIZE condition should be distinguished from FIXED_OVERFLOW that occurs during arithmetic calculations. A value too large for the field to which it is assigned will raise a SIZE condition on assignment, regardless of whether there was a FIXED_OVERFLOW in the calculation of the value. FIXED_OVERFLOW depends upon the size of fixed-point numbers allowed in the implementation. SIZE depends upon the declared size of the item of data receiving a value.

Result: Modulo assignment for fixed-point; other assignments are undefined.

Standard system action: Comment and raise the ERROR condition.

UNDERFLOW: This condition occurs when the exponent of a floating-point number is smaller than the permitted minimum, as defined by the implementation.

The condition does not occur when equal numbers are subtracted (often call significance error).

In some implementations, the condition may be detected by hardware interrupt, in others by special coding.

Result: Smallest positive non-zero value.

Standard system action: Comment and continue execution.

ZERODIVIDE: This condition occurs on an attempt to divide by zero. The condition does not distinguish between fixed-point and floating-point division; either can cause it.

In some implementations, the condition may be detected by hardware interrupt, in others by special coding.

Result: Undefined.

Standard system action: Comment and raise the ERROR condition.

INPUT/OUTPUT CONDITIONS

EDIT (filename): This condition is caused by an illegal character in the input data from a specified file, or a character that

is illegal for an output editing operation. For example, characters other than 0 or 1 when the list item is a bit string, or characters not permitted by the PICTURE for a numeric field (for both input and output), or an illegal character in an arithmetic constant.

Standard system action: Comment and raise the ERROR condition.

ENDFILE (filename): This condition is caused by an attempt to read past a file delimiter from the specified file.

Standard system action: Comment and raise the ERROR condition.

ENDGROUP (filename): This condition is caused by an attempt to read past a group delimiter from the specified file. The file is positioned past the group mark.

Standard system action: Comment and continue.

ENDRECORD (filename): This condition is caused by an illegal attempt to read past a record delimiter from the specified file. The file position is undefined.

Standard system action: Comment and raise the ERROR condition.

FIELD_OVERFLOW (filename): This condition can be raised for input and output operations in which the source value is too large for the destination.

Input: The condition is raised in the following circumstances:

1. For a string list item--if the string is too long for the maximum (varying) or actual (fixed) length of the string list item.
2. For a numeric field list item--if the input value cannot be edited according to the PICTURE for the numeric field list item, or according to the PICTURE format item. These conditions exclude those covered by the EDIT condition.
3. For an arithmetic list item--
 - a. If the input value exceeds the implementation defined floating-point number range (for a floating-point list item).
Note: The condition is not raised if the precision of the list item is insufficient to hold all given digits of the input value.
 - b. If the input value (after insignificant leading zeros are removed) exceeds the maximum value that the destination field can hold. Digits on the right end are truncated without notice.

Output: The condition is raised if:

1. The length of a string list item exceeds the field width of the format.
2. The pictured field cannot hold the value of the list item.
3. The arithmetic value in a string (e.g., 1.25 E95) exceeds the floating-point number range for an arithmetic format.
4. The arithmetic value of a list item exceeds the maximum value that can be accommodated in the field designated by the format.

Standard System Action: The system action for output is to fill the output field with asterisks and continue; for input, the system action is to comment and continue.

IDENT (filename): This condition is raised if the OPEN or CLOSE IDENT option does not compare with the label on the designated file. This applies only to the IDENT option form that specifies both data list and format list.

Standard system action: Comment and return from the ON unit. Processing will continue, with the unmatched label ignored.

NAME (filename): This condition is caused by an unrecognizable identifier on data-directed input.

Standard system action: Comment and raise the ERROR condition.

ACCESS (filename): This condition is raised whenever a programming error prevents successful access of a record from the designated file. The particular error may be determined by means of the ONCODE built-in function. If a return is made from the ON unit, the ERROR condition is raised.

Standard system action: Comment and raise the ERROR condition.

TRANSMIT (filename): This condition is caused by a transmission error on the specified file.

Standard system action: Comment and retry, and if unsuccessful after a standard number of retries (defined by the implementation), comment and raise the ERROR condition. A READ transmission error may be accepted by a return from the ON unit; a WRITE transmission error cannot be accepted (a return from the ON unit will raise the ERROR condition).

UNDEFINEDFILE (filename): This condition is raised when the specified file is not available.

Standard system action: Comment and raise the ERROR condition.

PROGRAM CHECKOUT CONDITIONS

SUBSCRIPTRANGE: This condition occurs when a subscript is evaluated and found to lie outside its specified bounds.

The condition does not distinguish between values that are too large and values that are too small.

Result: Undefined.

Standard system action: Comment and raise the ERROR condition.

CHECK (identifier list): A statement prefix specifying this condition may only be applied to PROCEDURE or BEGIN statements.

In the identifier list, each identifier is one of the following:

- a statement label
- an unsubscripted variable name representing a scalar, array, structure, or label variable
- an entry label

Each item in the list is, in effect, enabled independently.

None of the conditions that follow (up to but not including "Programmer-Named Conditions") will be raised in a prologue.

Statement Label: For a statement-label identifier the condition is raised prior to the execution of the statement to which the label is prefixed. If the label is prefixed to a non-executable statement, no condition will be raised.

Variables: For identifiers representing variables, the condition is raised whenever the value of the variable, or any generation of any part of the variable, may have been changed by any statement within the scope of the prefix.

The condition will be raised by the explicit reference to an identifier ID in the circumstances listed below, where ID is:

- an identifier in the list
- an identifier representing a structure or element contained by, or containing, an identifier in the list

The reference to ID may be subscripted or qualified.

The condition will be raised for ID if:

1. ID appears on the left hand side of an assignment statement. (This applies to assignment BY NAME even if the identifier mentioned does not appear in the final expansion of the statement.)
2. ID is set as a result of a pseudo-variable or pseudo-array appearing on the left hand side of an assignment.
3. ID appears as the controlled variable of a DO statement (or ID is set as a result of a pseudo-variable appearing as the controlled variable of a DO loop).
4. ID appears in a data list on a READ or GET statement.
5. ID has the SYMBOL attribute and a data-directed READ or GET statement is executed.
6. ID appears as the second argument of a DISPLAY statement.
7. ID appears as a STRING option on a WRITE statement.
8. ID is passed as an argument to a programmer-defined procedure, and no dummy is created.
9. ID appears on a RESTORE statement.

However, the condition is NCT raised under any of the following circumstances:

1. If the value of a variable defined upon ID or upon part of ID changes value in any of the ways described above.
2. If the value of a variable upon which ID is defined changes value.
3. If a parameter which represents ID changes value.
4. If ID appears in a GO TO or RETURN statement or any statement which involves the execution of a GO TO or RETURN statement.

Each condition is raised after the statement which caused it to be raised has been executed. If the statement has a task option, the condition is raised when the attaching task regains control. If the statement is a DO statement, the condition is raised each time control proceeds sequentially to the statement following the DO statement. If the DO specifies iteration, the condition is raised once for every iteration.

No statement other than a DO statement can cause a condition to be raised more than once for the same identifier. If a statement causes a CHECK condition to be raised for several identifiers, then the conditions will be raised in the left-to-right order of appearance of the identifiers in the statement.

Entry Labels: For an entry label, the condition is raised prior to each invocation of the entry label. The condition is raised only if the entry label is invoked by the name given in the ON list.

Result: Continue. The statement is executed normally.

Standard system action: If the identifier is a statement label or an entry label, the label will be printed on a debugging file.

If the identifier represents data, the identifier and its new value will be printed on a debugging file in the format of data-directed output.

PROGRAMMER-NAMED CONDITIONS

CONDITION (identifier): This condition is always enabled and may not appear in a condition prefix. The identifier is specified by the programmer, and is EXTERNAL. The condition is raised by the execution of a SIGNAL statement having the same identifier.

Standard system action: Comment and continue.

SYSTEM ACTION CONDITIONS

The following conditions are always enabled and may not appear in a condition prefix.

FINISH: This condition is raised immediately before the major task terminates by executing a STOP, RETURN, END, or EXIT statement. The ON unit is executed as part of the task in which the interrupt takes place.

Standard system action: Terminate the major task.

ERROR: This condition is raised when a major task is forced to terminate because of some error situation.

Standard system action: Comment and raise the FINISH condition.

APPENDIX 4: PERMISSIBLE KEYWORD ABBREVIATIONS

Abbreviations are provided for certain keywords. The abbreviations themselves are keywords and will be recognized as synonymous in every respect with the full keywords. The abbreviated keywords are shown to the right of the full keywords in the following list.

PROCEDURE	PROC
DECLARE	DCL
DECIMAL	DEC
BINARY	BIN
COMPLEX	CPLX
CHARACTER	CHAR
VARYING	VAR
POSITION	POS

INITIAL	INIT
INTERNAL	INT
EXTERNAL	EXT
AUTOMATIC	AUTO
CONTROLLED	CTL
DEFINED	DEF
ABNORMAL	ABNL
PRECISION	PREC
OVERFLOW	OFL
UNDERFLOW	UFL
FIXEDOVERFLOW	FOFL
SUBSCRIPTRANGE	SUBRG
ZERODIVIDE	ZDIV
CONVERSION	CONV
ENVIRONMENT	ENV
PICTURE	PIC

The characters that make up the 48-character set are same as those that make up 60-character set except for certain restrictions.

The following characters are not included:

Percent	%
Colon	:
Not	!
Or	
And	&
Greater Than	>
Less Than	<
Break character	-
Semicolon	;
Number sign	#
Commercial At sign	@
Question mark	?

The following three characters are replaced as indicated:

60-Character Set	48-Character Set
:	..
;	::
%	//

The two periods which replace the colon must be immediately preceded by a blank if the preceding character is a period.

The following operators, as used in the 60-character, set are replaced in the

48-character set by alphabetic operators as indicated:

60-Character Set	48-Character Set
>	GT
>=	GE
=	NE
<=	LE
<	LT
!	NOT
	OR
&	AND
	CAT

The above nine words are "reserved" in the 48-character set; that is, they must not be used as programmer-specified identifiers.

In each case, one or more blanks must immediately precede the alphabetic operator if the preceding character would otherwise be alphameric, and one or more blanks must immediately follow if the following character would otherwise be alphameric. Thus, to indicate the comparison of the variables A6 and BQ2Y for inequality, one would write A6 NE BQ2Y, but not A6NEBQ2Y, A6 NEBQ2Y, or A6NE BQ2Y. As the equal symbol is usable, however, the comparison of these two variables for equality may be written A6=BQ2Y.

The break character, commercial at-sign, and number sign are not used and consequently may not be employed in identifiers.

The default segment delimiter for data-directed transmission is a single semicolon regardless of which character set is used; if no hardware representation of a semicolon is available, the programmer must supply an explicit segment delimiter.

(If more than one page number is given, the primary discussion is listed first.)

abbreviation of keywords	154	BIT;	
ABNORMAL attribute	49,138	see string attributes	
abnormality	138,49	bit-string data	23
defaults for	50,62,63	bit-string operations	33
access attributes	59	blanks	
activation;		use of	17
see blocks, activation		with qualified names	27
ACTIVITY option	115	with structure level numbers	25
ALIGNED attribute	53	in picture specification	148
ALLOCATE statement	97	trailing, in input fields	62
allocation	72,10	blocks	19,10
also see storage class attributes		activation of	71
of parameters	136	begin	19
in tasks	75,137,138	nested	20
test for	145,98	procedure	19
ALLOCATION built-in function	145	termination of	71,109,122
arguments	68,65,69,134	bounds;	
dummy	70,134	see array	
evaluation of subscripts	134	overriding DECLARE statement	97
list	65	of parameters	136
arithmetic built-in functions	140	BUILTIN attribute	52,67
arithmetic data	22	built-in functions	140,17,66
attributes	43	BY and TO	105
arithmetic operations	31	BY NAME option	98,99
array	25,10	CALL option	19,62,65,67,116
allocation	53	CALL statement	102
assignment	98,101	with FETCH statement	108
bounds	25,48	for creating tasks	75
also see asterisks		CHARACTER;	
cross section of	27	see string attributes	
defining	55	character string	
dimensions	25,26	data	22
expressions;		pictures	149
see expressions		also see string	
manipulation	144	characters	
of statement labels	29	alphabetic	14
of structures	26	alphanumeric	14
assignment		data character set	16
array	98,101	48-character set	15
scalar	98	language character set	14
statement	98,95	60-character set	14
evaluation of	99	special	14
statement-label	98,99	CLOSE statement	103
string	99,101	coded arithmetic data	22
structure	98,100	collating sequence	16
asterisks		comment	17
for bounds or length	136,48,97,98	comparison operations	34
for cross sections of arrays	27	compile-time activity	131,11
with INITIAL attribute	56,57	COMPLEX attribute;	
with USES or SETS attributes	50	see mode	
asynchronous operations	74	COMPLEX pseudo-variable	96
AT option (with PAGE statement) ...	116,117	composite operand	37
attached task	75,11,137	compound statement	18
attaching task	75,11,137	concatenation operations	34
attributes	38,43,17,11	condition prefixes	18
also see individual attribute		conditions;	
defaults for	62,9	see ON-conditions	
also see individual attribute		CONSECUTIVE attribute	59
factoring of	39	constants	23
with macro DECLARE statement	131	bit-string	24
AUTOMATIC attributes;		character-string	24
see allocation, storage class attributes		fixed-point binary	23
base	22,43	fixed-point decimal	23
base identifier	54	floating-point binary	23
BEGIN statement	102,19	floating-point decimal	23
BINARY attribute;		imaginary	24
see scale		real arithmetic	23

statement-label	24	delimiters	15
sterling	23	descendence of blocks	71
contained in	20	dimension attribute	48
contextual declarations	40	with ALLOCATE statement	97
also see declarations		DIRECT attribute	59
control		DISPLAY statement	105
modification of	96	DO groups	19, 105
program	71	DO statement	105
return of	122, 66, 67	DOWN specification	126
sequence of	95	editing;	
statements	95	see PICTURE attribute	
CONTROLLED attribute	53, 97	symbols	147, 148
also see storage		drifting	147
conversion	33	ELSE clauses	111
arithmetic base and scale	33	nesting of	111
arithmetic mode	33	enable	76
integer	33	encompassing blocks	72
in expressions	31	END statement	107
type	34	use of	21
with RETURN statement	122	ENTRY attribute	51
COUNT built-in function	145, 120, 130	declaration of	40, 51
COUNT option	120, 130	use of	69, 134
CROSS option	120, 129	entry name	18, 20, 40
cross sections;		attributes	51
see array		default for	51
crossing of records boundaries	120, 129	with IDENT option	104, 115
current file	81	passing arguments to	69
		required for PROCEDURE statement ...	118
data		entry point	
aggregates	25	primary	20
arithmetic	22	secondary	20, 107
bit-string	23	ENTRY statement	107
character set	16	evaluation	
character-string	22	of argument subscripts	134
coded arithmetic	22	in array assignment	99
description	38	of assignment statement	99
elements	22	of expressions	37
list	82	EVENT	
numeric	22	built-in function	146
specification	82	option	75, 102
repetitive specification for	82	pseudo-variable	97
statements	93	event name	30
statement-label	23, 29	EXIT statement	108, 71
transmission	80	explicit declarations;	
statements	93	see declarations	
types	22	exponentiation	32
default for	45, 62	expressions	31
data-directed transmission	80	array	35
data specification for	85	as bounds or length	136
input	86	evaluation of	37
length of field	87	scalar	31
output	87	structure	36
DECIMAL attribute;		extended values on assignment	99
see base		EXTERNAL attribute	52, 41
declarations	38	external declarations	41
contextual	40	external names	20, 42, 62
explicit	38	scope of	41
external	41	external procedure	20
implicit	41	factoring	
multiple	41	of attributes	39
scope of	41	of options	103
DECLARE statement	38	FETCH statement	108
default;		relationship with CALL statement ..	108
see attributes		file	80
DEFINED attribute	54	attributes	58, 114
defined on	54	closing	103
DELAY statement	104		
DELETE statement	104		

conditions	151	IMAG pseudo-variable	96
current	81	imaginary numbers	24
names	80,58,63	also see mode	
opening	114	implicit declarations;	
preparation statements	93	see declarations	
specification	59	INDEXED attribute	59
FILE attribute	58	infix operators;	
FILE option .. 111,114,115,116,119,124,125,		see operators	
..... 126,127,129		INITIAL attribute	56,29
FIXED attribute;		rules with ALLOCATE statement	98
see scale		initial value for statement-label	
fixed-point;		arrays	29
see constants, precision, variables		INOUT attribute	59,114
FLOAT attribute;		INPUT attribute	59,114
see scale		input/output	80
floating-point;		statements	93
see constants, precision, variables		INTERNAL attribute	52
FOOT option	116	internal name	42
format		internal procedure	20
of data-directed output	87	internal to	20
of list-directed I/O	83,84	interleaving	28
of records	84	interrupt	75,18,114,150
format-directed data specification	88	system	77,114,150
format-directed transmission	80	iteration	105
format items	89	factor	56,24,88
control	92	KEY option	94,59,120,129,130
external mode	89	KEYLENGTH attribute	61
data	89	keyword	16
internal mode	91	abbreviations of	154
remote	93	separating	17
format list	88,103,108,115,117	known	43
FORMAT statement	108	label	18
label required for	108	also see statement label	
48-character set	155,15	required for FORMAT statement	108
FREE statement	108	LABEL attribute	47
FROM option	130	LAYOUT statement	111
function	65	length	
built-in	52,66,140	data-directed data fields	87
generic	66	format-directed data fields	89
procedure	66	identifiers	16
termination of	122	keys	61
reference	66	list-directed data fields	84
GENERIC attribute	51,66	overriding DECLARE statement	97
generic functions	66	parameters	136
arguments of the reference	51,66	strings	47
GET statement	109	level numbers	25,63
with current file	81	also see structures	
with procedure-directed transmission	81	optional blank	25
GIVING option	126	LIKE attribute	58
GO TO statement	109	with DEFINED attribute	54
groups	19	list-directed	
DO groups	19,105	data specification	83
of I/O files	80,110	input	83
single statement	19	length of field	84
GROUP statement	110	output	84
HEAD option	116	transmission	80
heading statements	19	macro	
HOLD option	120,129	assignment statement	132
IDENT option	104,115	DECLARE statement	131
identifiers	16	expressions	132
attributes of	38	GO TO statement	132
length of	16	IF statement	133
statement labels	16	null statement	132
keywords	16	processor	133
IF statement	111	statements	132
		variables	131

MAIN attribute	71,118	PACKED attribute	53
MARGIN option	111	PAGE statement	116
mode	22,44	parameters	65,134
multiple declarations	39	allocation of	136
multiple labels	18,107,118	bounds and length	136
names	16,26,42	controlled	98
external	42	with ENTRY statement	107
internal	42	with PROCEDURE statement	118
qualified	27	PICTURE attribute	45,147
scope of	41	with numeric data	45
simple	26	specification	147
subscripted	26	with string data	47
subscripted qualified	28	picture format items	
use of	43	external	90
nesting		internal	92
of blocks	20,71	picture specification tables	147
of ELSE clauses	111	POOL attribute	59
NEWKEY option	94,129	POSITION attribute	55
NORMAL attribute	49,138	POSITION statement	117
NOSYMBOL attribute	57	format items that are not	
null parameter list	51	allowed with	117
null statement	18,112	positioning statements	93
macro	132	precision	22,44
null string	22	in expressions	31
NUMBER option	116	of format items	89
numeric field	22	in picture specifications	46
ON statement	112	of real arithmetic constants	23
with REPOSITION statement	121	prefix	
use of	76	condition	18,76,112,150
ONCHAR pseudo-variable	97	label	18
ON-conditions	18,76,112,122,123,150	prefix operators;	
also see ON statement		see operators	
built-in functions	145	PRINT option	120
input/output	151	PRIORITY	
nullification of	18	built-in function	146
prefixes used with	18,76	option	75,102
program checkout	79,152	pseudo-variable	97
programmer-defined	79,153	procedure	19,65
with SIGNAL statement	125	activation of	71
ONFIELD pseudo-variable	97	attribute	70,118
on-unit	112	external	20
cannot be RETURN statement	112	internal	20
OPEN statement	114	invocation	65,66
operations		name	20
arithmetic	31	parameters	65
array-array	36	termination of	67,71
bit string	33	also see termination of blocks	
comparison	34	procedure-directed data specification ..	93
concatenation	34	procedure-directed transmission	81
scalar-array	35	PROCEDURE statement	118,19
operands		program	21,11
composite	37	control	71,95
simple	37	deletion of	104
operators		elements	14
arithmetic	15	modification	131
bit string	15	structure	17,71
comparison	15	prologues	137
infix	31	pseudo-array	98
prefix	31	pseudo-structure	98
string	15	pseudo-variables	96
options	17	PUT statement	119
also see individual options		without current file	81
OPTIONS attribute	118	with procedure-directed transmission	81
output;		qualified names	27,39
see input/output		READ statement	119
OUTPUT attribute	59,115	with procedure-directed transmission	81

REAL attribute;	
see mode	
REAL pseudo-variable	96
record	80
RECORD specification	126
RECURSIVE attribute	70,118
REENTRANT attribute	118
REGION option	94,59,120,127,129
REGIONAL attribute	59
relationship of arguments and	
parameters	134
remote format specification	93,108
report generation statements	93
REPOSITION statement	121
with ON statement	121
RESTORE statement	121
return of control	66,71,122
return of value	66,122
RETURN statement	122,66,71
cannot be an on-unit	112
returned value	
characteristics of	107,118
specifications	118
REVERT statement	122
use of	78
SAVE statement	123
scalar	23
assignment	98
constant;	
see constants	
defining	54
expression;	
see expressions	
variable	
see variables	
scale	22,44
scope	
of declarations	41
of names	41
of condition prefixes	76
scope attributes	41,52
default for	52,62,63
SECONDARY attribute	49
secondary entry point	20,107
segment	80
mark	80,124
SEGMENT option	117,120,129
SEGMENT statement	124
separators	15
sequence	
collating	16
of control	95
modification of	96
SEQUENTIAL attribute	59
SETS attribute	50,138
sign picture characters	149
SIGNAL statement	125
with programmer-defined	
ON-conditions	79
simple operand	37
60-character set	14
SIZE option	116
SKIP statement	125
SORT statement	126
SPACE option	116
SPACE statement	127
specification	106
stack, push-down	73
stacking current files	81
standard attributes	58
STANDIN attribute	58
STANDOUT attribute	58
statement label	16,18,47
array	29
initial values for	29
assignment	98,99
constant	28
data	23
designator	29
required for FORMAT statement	108
variable	28
statements	17,95
also see individual statement	
alphabetic list of	97
classification	95
compound	18
heading	19
identifiers	16
input/output	93
relationship	95
simple	18
STATIC attribute;	
see storage class attributes	
sterling	
constants	23
pictures	46,149
STOP statement	127
storage;	
also see allocation	
ALLOCATE statement	97
automatic	72,53
controlled	73,53,97
FREE statement	108
static	72,53
storage class attributes	53,72
default for	53,63
restrictions	53
with structures	64
storage equivalence attribute;	
see POOL attribute	
string	
assignment	101
attributes	47
built-in functions	143
data	22
STRING option	119,129
structure	25
assignment	98
BY NAME;	
see BY NAME	
declarations and attributes	63
with DEFINED attribute	56
with LIKE attribute	58
level numbers	25,63
storage allocation	98,53
with storage class attributes	53
subroutine	65
references	67
subscripts	26
interleaved	28
SUBSTR pseudo-variable	97
SYMBOL attribute	57
with DEFINED attribute	54
symbol table attributes	57
default for	57,63

syntactical unit	11	UNSPEC pseudo-variable	97
syntax notation	11	UP option	126
TAB option	111	USES attribute	50
TAB statement	127	variables	
task	11,29,74	array	25
attached	11,75	scalar	25
attaching	11,75	range of	25
major	74	default for range	63
synchronization of	74	statement-label	28
termination of	75	WAIT statement	128,137
task option	75	WHILE clause	106
TASK option	75,102	WRITE statement	128
termination		with procedure-directed	
blocks	71,109,122	transmission	81
function procedure	66,122	ZERO attribute	61
program	127	ZERO option	120
task	75	zero suppression	147,45
TITLE option	115		
TO and BY	105		
truncation on assignment	99		

COMMENT SHEET

IBM OPERATING SYSTEM/360
PL/I: LANGUAGE SPECIFICATIONS
Form C28-6571-1

FROM

NAME _____

OFFICE NO. _____

FOLD

CHECK ONE OF THE COMMENTS AND EXPLAIN IN THE SPACE PROVIDED

FOLD

- SUGGESTED ADDITION**
- SUGGESTED DELETION (PAGE)**
- ERROR (PAGE)**

EXPLANATION

CUT ALONG LINE

FOLD

FOLD


NO POSTAGE NECESSARY IF MAILED IN U.S.A.
FOLD ON TWO LINES, STAPLE, AND MAIL

FOLD

FOLD

BUSINESS REPLY MAIL
NO POSTAGE STAMP NECESSARY IF MAILED IN U. S. A.

FIRST CLASS
PERMIT NO. 33504
NEW YORK, N. Y.



POSTAGE WILL BE PAID BY
IBM CORPORATION
1271 AVENUE OF THE AMERICAS
NEW YORK, N. Y. 10020

**ATTN: PROGRAMMING SYSTEMS PUBLICATIONS,
DEPARTMENT D39**

FOLD

FOLD

CUT ALONG LINE

STAPLE

STAPLE

IBM

**International Business Machines Corporation
Data Processing Division
112 East Post Road, White Plains, N. Y. 10601**