


F/6 9/2

F30602-78-C-0016

NL

Q# 7:3

08774



NAVAL RESEARCH LABORATORY
WASHINGTON, D. C. 20340

LEVEL 4

12

RADC-TR-80-66
Final Technical Report
February 1980



ADVANCED SMITE REFERENCE MANUAL

TRW/Defense & Space Systems Group

AD A 087743

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

DTIC
ELECTRONIC
S
AUG 11 1980

C.

ROME AIR DEVELOPMENT CENTER
Air Force Systems Command
Griffiss Air Force Base, New York 13441

DDC FILE COPY.

80 8 8 060

This report has been reviewed by the RADC Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RADC-TR-80-66 has been reviewed and is approved for publication.

APPROVED:

Frederick A. Normand

FREDERICK A. NORMAND
Project Engineer

APPROVED:

Wendall C. Bauman

WENDALL C. BAUMAN, Colonel, USAF
Chief, Information Sciences Division

FOR THE COMMANDER:

John P. Huss

JOHN P. HUSS
Acting Chief, Plans Office

If your address has changed or if you wish to be removed from the RADC mailing list, or if the addressee is no longer employed by your organization, please notify RADC (ISCA) Griffiss AFB NY 13441. This will assist us in maintaining a current mailing list.

Do not return this copy. Retain or destroy.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER RADC-TR-80-66	2. GOVT ACCESSION NO. AD-A087 743	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) ADVANCED SMITE REFERENCE MANUAL	5. TYPE OF REPORT & PERIOD COVERED Reference Manual	
7. AUTHOR(s) TRW Defense and Space Systems Group	6. PERFORMING ORG. REPORT NUMBER 32584-6015-RU-00	
9. PERFORMING ORGANIZATION NAME AND ADDRESS TRW/Defense and Space Systems Group One Space Park Redondo Beach CA 90278	8. CONTRACT OR GRANT NUMBER(s) F30602-78-C-0016	
11. CONTROLLING OFFICE NAME AND ADDRESS Rome Air Development Center (ISCA) Griffiss AFB NY 13441	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 63728F /25290103	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Same	12. REPORT DATE February 1980	
	13. NUMBER OF PAGES 272	
	15. SECURITY CLASS. (of this report) UNCLASSIFIED	
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE N/A	
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) Same		
18. SUPPLEMENTARY NOTES RADC Project Engineer: Frederick A. Normand (ISCA)		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Hardware Description Language Emulation Simulation Computer Architecture		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This is a reference manual for the use of Advanced SMITE which is a higher level hardware description language used to describe various computer hardware architectures and then to generate an emulation of these architectures on the Nanodata QM-1 Microprogrammable Computer. Advanced SMITE is written in a subset of PL-1 and runs on RADC's MULTICS computer system.		

DD FORM 1 JAN 73 1473


EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)



UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

Table of Contents

	page
Preface.....	1
1. The Fundamentals of SMITE.....	3
1.1 The Applications of SMITE.....	3
1.2 The Steps in Using SMITE.....	9
1.3 A Smite Computer Description.....	14
1.4 Data Definition.....	16
1.5 Operations and Expressions.....	26
1.6 Processors.....	30
1.7 External Interfaces.....	37
1.8 General Rules for Coding SMITE.....	38
1.9 Case Study 1: The Mark 1.....	39
2. Control Statements.....	43
2.1 Introduction.....	43
2.2 The IF Statement.....	44
2.3 The CASE Statement.....	47
2.4 The DO Statement.....	52
2.5 The ESCAPE Statement.....	58
2.6 Case Study 2: FTSC Floating Point Unit.....	62

Accession For
TIS GRAZI
DC TAB
Unauthorized
Classification

3.	Storage Sub-Structure and Operation Widths.....	71
3.1	Introduction.....	71
3.2	Data Item Attributes.....	71
3.3	The DEFINED Attribute: Storage Overlays.....	74
3.4	Defaults.....	77
3.5	Case Study 3: Intel 8080 Storage Declarations.....	78
4.	Timing Specification and Control.....	83
4.1	Introduction.....	83
4.2	The IN Statement.....	84
4.3	The Amount of Timing Detail in a Computer Description..	85
4.4	Case Study 4: A Shift Unit.....	86
5.	Parallelism.....	89
5.1	Introduction.....	89
5.2	The PARALLEL-BEGIN and PARALLEL-END Statements.....	89
5.3	Parallel Timing and Control Flow.....	89
5.4	Case Study 5: Instruction Pre-Fetch.....	90
6.	Input/Output and Operator Interface.....	93
6.1	Introduction.....	93
6.2	EASY Interface to SMITE.....	93
6.3	External Functions.....	94

6.4	Ports.....	96
6.5	Lights and Switches.....	96
7.	SASS: SMITE Application Support Software.....	97
7.1	Introduction.....	97
7.2	SMITE Computer Description Development Considerations..	98
7.3	Preparation of SMITE Load Tapes.....	102
7.4	SASS Commands.....	103
7.5	Performance Measurement.....	112
7.6	SMITE Concurrency Implementation.....	121
7.7	SASS Modification.....	123
8.	Advanced Capabilities	130
8.1	Syntax Macros.....	130
8.2	Direct Code.....	136
8.3	The OPDEF Statement.....	138
8.4	Examples.....	139
	References.....	141
	Appendix A: SMITE Syntax Diagrams.....	A-1
	Appendix B: SMITE Compiler Operation Procedures.....	B-1

32584-6C15-RU-CC

Appendix C: Intel 8080 Microprocessor Definition.....C-1

Appendix D: Augmented MULTI Micromachine Definition.....D-1

Appendix E: Advanced SMITE Compiler Error Messages.....E-1

Appendix F: FTSC Emulator Requirements Specification.....F-1

Appendix G: Recommended SMITE Coding Conventions and StandardsG-1

Appendix H: FTSC Description.....H-1

Appendix I. SMITE Keywords.....I-1

Appendix J: Performance Measurement Error Conditions.....J-1

EVALUATION

RADC is currently building a computer emulation facility to assist in the evaluation of hardware/software/firmware trade-offs necessary in the development of system architectures. As part of this effort, RADC has purchased a QM-1 microprogrammable computer which is designed for computer emulation. However, the programming of emulations on a microprogrammable computer at the microcode level is a difficult and time-consuming task.

The objective of this effort was to obtain SMITE which is a Higher Order Language for describing computer architecture emulations and a compiler which produces code to emulate said architectures on the QM-1 microprogrammable computer.

The Advanced SMITE Reference Manual describes the Higher Order Language, how it is used, and how an emulation can be produced.

Fredrick A. Normand
FREDERICK A. NORMAND
Project Engineer

Preface

This book is written for the person who wants to understand the use of the SMITE computer description language in the solution of problems of computer architecture and emulation in computer information systems development. The book has three distinct roles:

1. It is the primary material used in formal training classes on SMITE.
2. It is suitable for self-study use by persons familiar with basic computer architecture and emulation concepts.
3. It is suitable for use as a reference manual for users of the SMITE language, compiler, and associated support software.

This version of the book is an update to the original Basic SMITE description. It provides information about new SMITE capabilities such as syntax macros, new operation definitions, a new DECODE statement, direct code, and performance monitoring. The new implementation of the SMITE Application Support Package (SASS) is also described.

The work of Daniel D. McCracken in advancing the art of language guide development was invaluable in the writing of this book. In particular, his work "A Guide to FORTRAN IV Programming" served as the example we tried to follow in setting the outline and style we used. Any errors, omissions, or faults, however, are solely our own.

Redondo Beach, California
1979

1. The Fundamentals of SMITE

1.1 The Applications of SMITE

The consideration of computer architecture has been of concern to computer and system designers since the days of Von Neumann and earlier. Our understanding of computer architecture has become considerably more complex than the early partitioning by Von Neumann shown in Figure 1.

Today, we recognize that several levels of interest are crucial to computer architects. These levels are the gate level, the register transfer level, and device level. (Bell and Newell [1] distinguish six levels; although important from the total systems design viewpoint, the other three levels of Bell and Newell are irrelevant to an understanding of SMITE and are therefore not included in this presentation.)

At the highest level, a computer system is the interconnection of devices, such as processors, memories, switches, interfaces, busses, modems, and the like. Commercial computer installations are almost invariably configured at this level. A wide variety of analysis tools have been evolved, of which ECSS [2] and PMSL [3] are representative.

The register transfer level describes a computer system as the interconnection of circuits and components, such as ALUs (arithmetic logical units), registers, data paths, etc. The register transfer level is that hardware representation typically encountered by the assembly language programmer, and is presented in a "hardware reference" or "principles of operation" manual.

The gate level introduces even more detail and complexity into the computer description. A gate level computer description is in terms of AND, OR, and NOT gates, decoders, counters, flip-flops, etc. Languages have been defined for this level, including AHPL [4] and CDL [5], traditionally this level has been expressed with boolean equations.

Clear distinctions between these three levels do not exist.

For example, memory is found as an entity in both device and register transfer descriptions, and registers are found in both register transfer and gate descriptions. The entire spectrum has become increasingly blurred as circuit integration technology has increased the sophistication of the




FIGURE 1: CLASSICAL VON NEUMAN CPU ARCHITECTURE
(CIRCA 1942)

individual component. Bubble memories and complete microprocessors on one or two chips bring complete device level components to the physical scale of the small scale integration flip-flop.

In spite of the difficulty in specifying the exact domain of each description level, the division remains useful. Specification of processor architecture in a formal way, using languages such as SMITE or ISP [6], permits the investigation of performance in a more disciplined, in-depth manner than simple benchmarking, and also permits the development of computer emulations from the formal description [7]. Other efforts to apply register transfer descriptions are also under way, including automatic hardware design, automatic compiler targeting, and program proof.

The SMITE computer description language is intended for analysis of computer architecture and the development of microprogrammed emulations of computer architectures. The context of application of a SMITE-developed computer emulation is diagrammed in Figure 2. (We present only a brief discussion here. Fuller exposition may be found in [7], [8], and [9].)

The complete emulation facility, which we call the Flexible Analysis, Simulation, and Test (FAST) facility, provides a total simulation and analysis environment, including both the emulated hardware and the simulated external environment, as well as test and performance measurement tools. The computers shown in Figure 2 are the Nanodata QM-1 and a large, general purpose computer. Current SMITE software utilizes Honeywell 6180 equipment as the general purpose host.

Along with the SMITE emulators, programs executing in the QM-1 also include target (emulated) machine software, environmental simulations (or interfaces to environmental simulations resident on the general purpose host), performance measurement software, and the QM-1 resident operating system. The general purpose host executes compilers for target software, such as JOVIAL, meta-compilers for development and maintenance of JOVIAL, SMITE, and other compilers, and the SMITE compiler itself.

In Figure 2, the upper half of the diagram (labeled "Software Loop") uses tools to support software development, validation, and verification. The lower half (labeled "Hardware Loop") uses tools for the specification and evaluation of hardware designs. Interaction between the two loops reflecting instruction set changes into the target software is presently supported by compiler modification through a meta-compiler.




Figure 2 SMITE Emulation Application Context

32584-6015-RU-00

Future technology development is planned to make this process automatic.

The role of the SMITE emulator in this context is to provide the capability for executing programs in the target machine instruction set, and for timing the execution of those programs. The QM-1 resident operating system supports emulator execution, and the environmental simulation provides inputs to and processes outputs from the emulation. The SMITE compiler generates emulators from computer descriptions.

This book is about the latter step: the development of SMITE computer descriptions and the processing of those descriptions by the SMITE compiler to form QM-1 microcode emulations. We will at all times discuss SMITE as a problem solving tool. SMITE provides a convenient, expressive notation for describing computer designs at the register transfer level, and a means for evaluating those designs. SMITE is a means for exploring alternatives, not a substitute for invention and creativity on the part of the computer systems architect.

1.2 The Steps in Using SMITE

Using SMITE to help solve architecture problems or to provide a software development testbed involves more steps than simply coding and compiling a description. Instead, using SMITE is a software development process, and done properly entails all the steps required for good software engineering discipline. In this section we describe generically the steps we recommend for the successful use of SMITE.

Problem Formulation and Requirements Definition

The first step in the SMITE architecture methodology is to specify the problem to be solved using SMITE. Within the total problem of "Will a Motorola 6800 process fire control data fast enough to fuze and launch a salvo of 12 air-to-ground rockets in 2 minutes?" for example, SMITE can provide answers by experiment to questions such as "How long does fuzing a single rocket take?" or "What data rate of input to the fire control system exceeds a limit of 50K instructions per second for input data processing?" In general, this first step entails understanding how the software execution capability and quantitative data available from SMITE emulators contributes to solving the total system problem, and then defining the required tasks and experiments to utilize the emulator.

Once the emulator usage is analyzed, a requirements

statement can be written for development of the emulator. In addition to execution of the computer instruction set, the requirements specification should also specify:

- 1) The requirements on accuracy and fidelity of timing for instruction interpretation;
- 2) The functionality of all the external interfaces of the emulator, including input, output, and interrupts;
- 3) The program load and dump mechanisms, and the operator console interfaces,
- 4) The data collection and reduction requirements.

An example emulator requirements specification for an emulation of the Raytheon Fault Tolerant Spaceborne Computer (FTSC) is provided as Appendix F. The emulation was intended for software development, integration, and test.

Design Methodology and Structure

The next step in the application of SMITE is the development of the methodology and structure for the design of the emulator. For straightforward emulator development these concepts are well understood and are described in this section. For descriptions intended for architectural evaluation the concepts are not yet generalized; some of the problems and pitfalls will be described.

The initial step in the design sequence is the definition of the data base, including the machine memories, registers, I/O ports, and addressing conventions. The data base for emulator descriptions is primarily resident in the main processor of the description, and is referenced globally by processors nested within the main one.

The next design step is the analysis of the instruction set to determine the decoding that the emulator will perform. In the case of the Intel 8080 microprocessor (Reference the complete description in Appendix C), the primary decode is performed on the leftmost 2 bits of the 8-bit instruction. In some cases a secondary decode is performed on the rightmost three bits, such as to determine the ALU function, certain cases of these secondary decodes further require a third level of decode on the remaining 3 bits. By way of contrast, a description of the CDC-6000 series

central processor only requires a single decode on the leftmost 6 bits of the instruction.

Some design freedom exists in the way decode is performed. Returning to the Intel 8080, for example, an alternative to the three-level decode described above is to perform a single 256-way decode on all 8 bits of the instruction. Another alternative is to concatenate the leftmost 2 bits with the rightmost three bits and decode the resulting 5-bit field, using second level decodes as necessary. The consequences of exercising the possible options are to trade execution speed of the emulator against the size of the microcode generated during compilation. The 3-level Intel 8080 decode, for example, will result in a smaller but slower emulation than one performing a 1-level decode.

Once the decode is analyzed, the classes of instructions resulting from decode can be examined to determine the processors (e.g. procedures, functions, or subroutines in other languages) which will be useful in describing the computer. In the Intel 8080, for example, processors to perform an arithmetic operation, to fetch an address from memory, to push or pop the stack, and others were found useful. The external characteristics of these processors, their inputs, outputs, and function, can now be designed. Virtually every emulation will have a processor corresponding to the ALU (arithmetic-logic-unit), which performs a specified operation and sets condition bits such as overflow. In the Intel 8080 description (Appendix C), this processor is named PERFORM-ADD.

The final design step is to analyze and specify the design of the external interfaces: Input, output, interrupts, and operator console. The requirements for I/O bear heavily here, as do the facilities available at run time to support the emulation. Conventions and procedures for using the SMITE Application Support Software (SASS) to aid external interface are defined in Chapter 6.

Development Methodology

The emulator may then be coded. The goals of emulator coding should be clarity, readability, and maintainability foremost - individual coding (as opposed to design) practices usually have little impact on emulator performance or size. Some suggested coding practices are listed in Appendix G. In general, good coding practice as applied to top-down, structured programming carries over well into the SMITE domain. Consistency of style is

particularly a virtue, so that internal clarity of the emulation and non-astonishment of the reader is obtained. SMITE is well suited for top-down development and integration, in the coding of the FTSC emulation (Appendix H) for example, the data declarations and instruction decode were first coded and debugged. As individual instruction descriptions and necessary processors were written they were integrated into the growing emulator skeleton.

At some point the external interfaces and emulator-specific support software must be developed also. This software is largely coded in MULTI, the standard QM-1 microinstruction set used by SMITE. The conventions and techniques involved are discussed in chapters 6 and 7. Information about MULTI and the QM-1 may be found in references [10] and [11], respectively. Eleven new microinstructions have been added to MULTI for SMITE, and two of the standard microinstructions have been changed. These additions and changes to MULTI are defined in Appendix D.

The final step in emulator development is to integrate and test the software. The target computer program used to drive the emulator during final testing is important. If a diagnostic program for the emulated computer is available, it is the logical choice. The use of a hardware diagnostic program may be difficult, however, if the program executes illegal opcodes assuming a particular response from the hardware, or otherwise depends on particular construction of the hardware not modeled by the emulation.

Architectural Analysis Methodology

The course of events to follow is less clear for architectural analysis. If the method of analysis is timing the execution of benchmarks or complete software systems on an emulator, as might be done to evaluate the effect of replacing core memory with faster semiconductor memory, a basic emulator including modeling of the instruction timings would be required, and could be developed as outlined above. At the other extreme, if an architecture description is wanted as input to a static analysis to determine the effect of changing the number of ALUs in the implementation, an analysis which is the subject of current research, the methodology for developing the appropriate description is unknown. The problems in developing a methodology begin with requirements. The requirements on a description to be used for static analysis are not completely known. For example, the SMITE

code to describe the INTEL 8080 add (reference Appendix C), is

```
(AUX-CARRY//A<3:0>) <- CARRY-IN + A<3:0> + OP-REG<3:0>,
CARRY-A <- AUX-CARRY + A<7:4> + OP-REG<7:4>,
SIGN <- A<7>,
ZERO <- A = 0,
PARITY <- A<7> XOR A<6> XOR A<5> XOR A<4>
        XOR A<3> XOR A<2> XOR A<1> XOR 1;
```

This code is conceivably implemented with two cycles of a 4-bit ALU plus some combinatorial logic, with an 8-bit ALU and combinatorial logic, or with a special purpose device to implement the entire function. The descriptive problem is to describe the structure and operation of the architecture so as to permit identification of possible alternatives, and to highlight the opportunities inherent in the architecture design. To date, one principle has been proposed as key to the methodology: All operations within any given statement should be of the same level of complexity, as should all statements in any given sequence. It is still an open research topic, however, how to measure and compare complexity, and how to construct descriptions having this property.

The application of this "leveling" property to the short Intel 8080 addition example above helps to identify some problems. For example,

```
(AUX-CARRY//A<3:0>) <- CARRY-IN + A<3:0> + OP-REG<3:0>;
```

and

```
SIGN <- A<7>,
```

are intuitively not at the same level of complexity. Even

```
(AUX-CARRY//A<3:0>) <- CARRY-IN + A<3:0> + OP-REG<3:0>,
```

and

```
CARRY-A <- AUX-CARRY + A<7:4> + OP-REG<7:4>,
```

are of differing levels of complexity, since in the former statement the concatenation of AUX-CARRY and A<3:0> is described explicitly, while in the latter the concatenation is implicit by reference to CARRY-A, which is defined elsewhere in the description as the concatenation of the CARRY bit with the A register.

1.3 A Smite Computer Description

An example of a complete architecture is shown in Figure 3. In this example the data processing system under consideration inputs high-data-rate radar returns, performs tracking, identification and discrimination, and outputs airspace data to user consoles. Within the data processing system the architecture to be analyzed consists of a custom array processor for return correlation and track assembly, a Digital Equipment Corporation PDP 11/45 processor equipped with floating point hardware used for system control and mass data base management, and a dual-processor Honeywell H-6180 for the bulk of the computational workload analyzing tracks assembled into the data base. The processors in the system are connected together and to a high bandwidth mass store by high bandwidth busses.

A description of this example would be hierarchical. At the highest level would be the declaration of the three computers and the system mass storage device as entities, and the definitions of the capacity, protocol, and endpoints for each of the bussed interconnections. This level of description is the PMS level discussed earlier, and is a higher level than the register transfer level of a SMITE computer description.

The next hierarchical level down in the system description expands the description of each system entity. The array processor and the dual-processor H-6180 would be described in SMITE. These SMITE descriptions extend out in scope to the boundaries of the processor: everything inside the I/O and interrupt ports that the outside world connects to is described. Connecting the outside world to the emulation requires techniques discussed in chapters 6 and 7.

The PDP-11 however, presents some options. Either the PDP-11 itself can be described as a bus-connected system using PMS concepts, in which case individual SMITE descriptions are required for the CPU, I/O, and memory, or else it can be described as a single computer entirely in SMITE. The latter approach has the disadvantages that changes to the memory or I/O addressing sequences a re-write of the CPU description to reflect the new addressing conventions, and that the bus-oriented nature of the computer is completely obscured from static architectural analysis.

The monolithic approach has the advantage, however, that the description is self-contained and does not require external microcode in the implementation of the CPU.




Figure 3 Emulation Application Context

In general, SMITE is intended to describe a computer, that is a set of electronics including a CPU, memory, operator interface, interrupt system, and I/O ports. More specialized devices than complete computers may be described. For example, consider a memory device. A typical memory might be described in SMITE as follows:

```
MEMORY-DEVICE: PROCESSOR,
  DECLARE CORE[0:4095]<0:15> MEMORY,
    MAR<0:11> PORT,
    MDR<0:15> PORT,
    STROBE<0> PORT,
  DO FOREVER;
    'IF STROBE IS ZERO, READ'
    'IF ONE, WRITE. THIS MODEL'
    'DOESN'T INCLUDE SPLIT-CYCLE'
    'WRITES (READ/MODIFY/WRITE)'
    CASE STROBE;
      MDR <- CORE[MAR];
      CORE[MAR] <- MDR;
    END CASE;
  MEMORY-DEVICE: END;
```

Floating point boxes, encryption devices, and other computer sub-units may be described similarly.

The remainder of this chapter introduces the basic syntactic and semantic concepts of the SMITE language. Data, including constants and variables, is introduced, followed by expressions and processors.

1.4 Data Definition

Data is the term used to describe all quantities that are operated upon by a computer. The data representation used in SMITE assumes that all data can be and are encoded as bit strings. Therefore, the only data type provided in SMITE is the bit string. The full range of data found in "conventional" architecture description applications can be accommodated using this method: integer, fixed point, floating point and character data all have this method of representation inherent in their formats. Extended data types such as decimal data are abstractly viewed by the user as non-binary data, yet are represented within the computer and the corresponding SMITE computer description as bit strings.

This section presents the basic concepts of data employed in the SMITE language. The discussion begins with the various forms of constants recognized in SMITE, and then proceeds to the names of data items. The types of data items are discussed next, followed by the statement for declaring data items such as registers, memories, I/O ports, and the like.

1.4.1 Constants

All constants or numbers are assumed to be decimal in SMITE unless indicated otherwise. Binary constants are represented as B'binary string', octal constants as O'octal string', and hexadecimal constants as X'hexadecimal string'.

The following are valid SMITE constants:

125	(a decimal constant)
X'AB12'	(a hexadecimal constant)
B'1101100'	(a binary constant)
O'7734'	(an octal constant)

while the following are invalid SMITE constants:

A14A	(A is not a valid decimal digit)
X'Q12'	(Q is not a valid hexadecimal digit)
B'411001'	(4 is not a valid binary digit)

All data in SMITE have the characteristic of width, i.e. the number of bits being used to represent a datum is its width in that particular context. The width of a constant is the minimum number of bits needed to represent the number in binary form. The number zero is defined to have a width of one bit. The width of a constant may be increased in expression evaluation (section 1.5). All constants are positive in SMITE. For example,

5

is a constant with a width of 3 bits, and

-5

is an expression which computes the complement of the

number five. The width of this expression is 4 bits to allow space for the required sign bit. Some other examples will illustrate the computation of the width of constants:

B'11111'

has width 5,

X'FF'

has width 8,

O'7046'

has width 12,

B'000000000001'

has width 1, as does

X'0001'

and

O'0001'

1.4.2 Storage Organization: The Word

The central element of storage organization in SMITE is the word. A word is any group of one or more bits that is logically referred to as a single entity. The SMITE concept of word therefore encompasses a range of elements commonly found in modern computer structures, from the bit and byte through the halfword and word, to the doubleword. For example, in the IBM System/370, bits, 8-bit bytes, 16-bit halfwords, 32-bit words, and 64-bit doublewords are all supported data types. This word structure could be represented in SMITE as

```
DECLARE
  DOUBLEWORD<0:63>,
  WORD[0:1]<0:31> DEFINED DOUBLEWORD,
  HALFWORD[0:3]<0:15> DEFINED WORD,
  BYTE[0:7]<0:7> DEFINED HALFWORD,
  BITS[0:63]<0> DEFINED BYTE;
```

The storage structure thus represented is shown in Figure 4.




Figure 4. An Example Storage Structure

1.4.3 Identifiers

An identifier is the name by which the computer description refers to a data item. An identifier is formed from any combination of letters, numbers, and hyphens ('-'), using the following composition rules:

1. The first character must be a letter,
2. The last character must not be a hyphen, and
3. The identifier must not be identical to a SMITE reserved word (See Appendix I).

The following are valid SMITE identifiers:

ZERO-FLAG

JABBERWOCKY

R2-D2

A-FAIRLY-LONG-IDENTIFIER

The following are invalid SMITE identifiers:

O-FLAG (first character is a digit)

DECLARE (duplicates a SMITE reserved word)

TEST-STATUS- (last character is a hyphen)

1.4.4 Data Item Width

The width attribute specifies the width of each word of the data item in bits, and defines the bit numbering scheme used within the word. The width definition is a list enclosed by angle brackets ('<' and '>'), and has the following form:

<CTE>

or

<CTE : CTE>

where CTE stands for "compile time expression". Compile time expressions are discussed in section 1.5, essentially, they are any value computable at compile time (e.g.,

constants, data items, operations on data items, attributes and some specifications).

The single entry width definition specifies a one bit wide data item. The bit is assigned the number specified by the expression. For example,

```
DECLARE CARRY-STATUS<7> REGISTER;
```

declares a one bit wide register named CARRY-STATUS. The bit number is (perhaps arbitrarily) assigned the ordinal 7.

The two-entry width definition specifies the width of a multiple-bit word. The leftmost expression in the width definition corresponds to the number of the leftmost bit, and the rightmost expression corresponds to the number of the rightmost bit. For example,

```
DECLARE ACCUMULATOR<7:0> REGISTER;
```

declares an 8-bit register named ACCUMULATOR, with the leftmost bit numbered 7, and the rightmost bit numbered 0. Similarly,

```
DECLARE STATUS-WORD<0:15> REGISTER;
```

declares a 16-bit register named STATUS-WORD, with the leftmost bit numbered 0, and the rightmost bit numbered 15. The bit numbers are not required to begin at zero at either end of the word. Thus,

```
DECLARE LARGE-BIT-NUMBERS<X'FFE':X'FFF'> REGISTER;
```

declares a two bit register.

1.4.5 Storage Arrays

Many storage devices consist of an ordered set of words with an addressing mechanism to select a particular word from the entire set. SMITE provides an array definition facility to describe such structures. Arrays are sequentially ordered sets of words, each of the same width, accessed by an address. The address space is defined for an array by specification of the lower and upper address bounds.

Length definition attributes are used to describe array data items. The definition consists of two compile time

expressions separated by a colon and enclosed by square brackets, as follows:

[CTE : CTE]

The length attribute defines a block of words (the block length is defined to be one if the two address bounds are equal), with the leftmost entry as the index of the first word in the block and the rightmost entry as the index of the last word in the block. The first word address is constrained to be less than or equal to the last word address. The statement

DECLARE TO-BE[1:20];

declares a contiguous block of 20 words, with addresses ranging consecutively from 1 to 20, inclusive. Similarly, the statement

DECLARE NOT-TO-BE[0:X'FF'];

declares an array of 256 words with consecutive addresses from 0 to 255, inclusive.

The declaration of an array implies the existence of address decode and selection logic. Thus, there is a subtle difference between the declaration

DECLARE OP-HOLD[0:0];

and the declaration

DECLARE OP-HOLD;

The first declaration defines an array one word long, including the (unnecessary) address selection logic, while the second defines a word. The former statement is liable to confuse the reader of a computer description, and should only be written if such hardware actually exists.

1.4.6 Storage Classes

Several different types of storage elements are generally found in computers, such as registers, memories, input/output ports, and others. SMITE supports the declaration of different types of storage elements with the storage class attribute, which is supplied as part of the declaration for a data item. Storage class attributes are discussed further in chapter 3. Two of the storage class

attributes available in SMITE are for declaring registers and memories. For example,

```
DECLARE ACCUMULATOR<0:35> REGISTER;
```

defines a 36-bit register, and

```
DECLARE MEM[0:4095]<0:15> MEMORY;
```

declares a 16-bit by 4096-word memory as might be found in a PDP-11.

1.4.7 Data Item Declaration

Every data item used in a SMITE computer description must be declared before it may be used. Data items are defined in SMITE by the use of the DECLARE statement, which has the following form:

```
DECLARE identifier attribute(s);
```

The DECLARE statement also allows the declaration of more than one data item. Individual item declarations in a single statement are separated by commas. For example,

```
DECLARE
```

```
    identifier attribute(s),
```

```
    identifier attribute(s);
```

Within a processor, any number of DECLARE statements may appear, however they must all be grouped at the beginning. The declarations define all registers, memories, and other data definitions needed for the description.

Each individual data item declaration within the DECLARE statement has the following form:

```
    identifier length-attribute width-attribute
    class-attribute
```

The length attribute need only be present if an array is being declared.

An additional attribute, DEFINED, is presented in chapter 3. The following are examples of valid SMITE data item declarations:

```

DECLARE
  STATUS-WORD<3:8> REGISTER;
DECLARE
  ACCUMULATOR<0:7> REGISTER,
  INDEX<0:15> REGISTER,
  MEM[0:0'7777']<0:15> MEMORY;

```

1.4.8 Data Item References

References to storage elements are similar in form to the storage element declarations, but are interpreted differently. The possible forms for a data item reference are

identifier

identifier word-selector

identifier subfield-selector

or

identifier word-selector subfield-selector

A word selector looks like a length attribute, and is of the form

[expression]

Expressions are discussed in section 1.5; essentially, an expression is any computable value (e.g., constants, data items, operations on data items, processor function calls, etc).

The expression in a word selector acts as a subscript, and is input to the address selection logic for the array to determine the specific word to be referenced. Word selectors may be included in a data item reference only if the data item was defined as an array when it was declared, and may not be omitted if the item being referenced is an array. Given the declaration

```

DECLARE
  ARRAY[0:15]<0:63> MEMORY,
  MAR<0:3> REGISTER;

```

then

```

  ARRAY[0]

```

references the first word of the array, and

ARRAY[MAR]

references the word of the array addressed by the value of data item MAR.

A subfield selector looks like a width attribute, and is of the form

< constant >

or

<constant : constant>

If the width definition is omitted the entire word is referenced; otherwise the subfield identified by the width definition is referenced. Given the declaration

DECLARE ACCUM<7:0> REGISTER;

then

ACCUM

references the entire eight bits, while

ACCUM<5:4>

references a two bit subfield of ACCUM.

The word selector and width selector may be used in conjunction for array references. Using the earlier example given for word selectors,

DECLARE
ARRAY[0:15]<0:63> MEMORY,
MAR<0:3> REGISTER;

then

ARRAY[MAR]<4>

references a one bit subfield of the word addressed by MAR, and

ARRAY[1]<2:3>

references a two bit subfield of the second word of the array.

1.4.9 Limits

Certain limits on data item length and width, and on word selector expression width, are imposed on the SMITE language by the SMITE compiler and the Nanodata QM-1 computer. These limits are as follows:

1. No data item may have a word width of greater than 72 bits. In the case of array declarations, this means that the width of each individual word of the array has a maximum of 72 bits.
2. The total main store allocation for an emulator is limited to an 18-bit address by the architecture of the QM-1, and may be further limited by the amount of main store installed on a particular QM-1. Thus while there is no specific limit on the number of data items that may be declared, or on the array length of any one data item, the maximum available main store address space may not be exceeded by the aggregate data item declarations.
3. Due the means of storage allocation adopted in the SMITE compiler, there is a maximum width for any word selector. For reference to data items of from 1 to 36 bits in width, this maximum is 17 bits. For reference to data items of from 37 to 72 bits, this word selector maximum width is 16 bits.
4. The SMITE compiler does not process constants of width greater than 72 bits.

1.5 Operations and Expressions

All data manipulation (i.e. moving and/or operating on data) is achieved in SMITE by the use of expressions. Processor invocation is also accomplished through expressions. An expression is a combination of data item references, constants, operators and parentheses producing a single value upon evaluation.

An expression is one type of SMITE executable statement. The expression statement has the form:

expression ;

The binary operators defined in SMITE are:

arithmetic (two's complement)

+	addition
-	subtraction

relational (inequality comparisons are two's complement.
The relationals are defined for both signed and unsigned
quantities)

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

boolean

OR	logical sum
AND	logical product
XOR	logical difference

miscellaneous

<-	data transfer
//	concatenation

The unary operators defined in SMITE are:

+	(unary plus is ignored)
-	arithmetic negation (two's complement)
NOT	logical negation (one's complement)

Multiple character operators must be coded without spaces between the characters forming the operator.

Expressions are composed of one or more terms (operands) separated by binary operators. A term can be a storage element reference, processor call, constant, an expression surrounded by parentheses, or a unary operator with any of the previous four items as the operand. The following are expressions:

AFB

(A + C) // D

A // B + D = Q

A + B > (Q + C) OR S

A <- B + (-C + 5) // D <- E<3:4>

Expressions are evaluated from right to left, with the exception that unary operators are evaluated before binary operators, and the special binary operator '//' (concatenation) is evaluated even before unary operators since it is a part of basic operand formation. The normal order of evaluation may be altered by parenthesization. The right-to-left convention, although a departure from most procedural languages, has the benefit that all evaluation and operator definition has is a simple rule to define the order of computation, rather than the complicated tables of operator precedence found in many other languages. In addition, the evaluation from right to left makes the embedding of transfer operators in an expression natural. (i.e. B <- C + D <- MEMORY[A] is interpreted as: fetch MEMORY[A], store that value into D, add the same value to C and store the result in B).

The operands allowed on the left hand side of the transfer operator ('<-') are restricted to storage elements and the names of function processors. For example,

A <- 4

and

A//B <- C

are a valid assignments, but

4 <- A

and

$(C + D) < - E$

are not.

Explicit rules have been defined in SMITE for the width of a result after the evaluation of an expression. These rules are as follows:

1. The width of a data item reference is the declared width of the data item unless a subfield extraction has been included in the data item reference, in which case the width of the reference is established by the subfield reference. The width of a constant is the minimum number of bits required to hold the constant.
2. The width of an expression after the evaluation of the unary plus and NOT operators is the same as the width of the operand. The width of an expression after the evaluation of a unary minus operator is also the same as the width of the operand unless the operand is a constant, in which case the width of the expression is one greater than the width of the constant.
3. The width of an expression after the evaluation of any binary arithmetic or logical operator is the same as the width of the two operands if these widths are identical. If the widths of the two operands are not equal, then the smaller-width operand is expanded to the width of the wider operand, and the operation performed. The width of the result will then be the width of the wider operand. When the smaller operand is expanded, it will be zero-padded on the left unless the sign extension (SE) function is used (see section 1.5).
4. The width of an expression after evaluation of any relational operator is one bit. The operands are adjusted to conforming widths as in rule 3 above before comparison if their widths are not equal. The relational operation is performed assuming both operands are unsigned magnitudes unless the sign extension (SE) function is used on one of the operands. In the latter case, an arithmetic comparison assuming both operands are signed is performed.
5. The width of an expression after evaluation of a concatenation operator is the sum of the widths of the two operands.

6. If a transfer of data is being performed, the source (right hand) operand is constrained to be no wider than the receiving (left hand) operand. If the source operand is smaller than the destination operand, padding is performed as in rule 3 above. The width of the expression is equal to the width of the destination operand.

Compile time expressions (CTE) are a special class of expressions which are evaluated at compile time instead of at execution time. The result of the evaluation is applied as a constant to the SMITE description, e.g. a width attribute. The terms of a CTE must be capable of evaluation at compile time to produce valid results. In addition, a CTE cannot contain a '>' operator. The concept of compile time expressions gives a broader meaning to a constant than that of a simple number. A constant can now be represented by an expression as long as all the terms are defined at compile time. Such terms would be the length or width attribute of a data item, e.g. LENGTH(MEM), or WIDTH(ACC). The MIN and MAX of a set of CTE's is also recognized as a valid constant. For example, the width of a data item could be declared:

```
DECLARE A<MAX(WIDTH(C), WIDTH(D))>REGISTER;
```

1.6 Processors

A SMITE computer description conceivably could be written as one self contained program. However, the description is much more readable if it is separated into smaller modules. In SMITE, these modules are called processors. The SMITE processor is analogous to the function, subroutine, or procedure capability found in conventional higher order procedural languages. The SMITE processor is typically used to isolate common processing steps and parameter dependent processes, or to modularize the description in order to highlight features of the computer architecture.

SMITE processors may be called re-entrantly, as may occur from within parallel contexts, but may not be used recursively.

1.6.1 Processor Structure

The basic module of a SMITE computer description is the processor. The processor is begun with a processor header, and is terminated by an END statement. The processor header and END statement must both be labeled with an

identifier, the name of the processor, and the name must be the same on both statements. The processor returns to the point from which it was called when control passes to the END statement. A processor, in its most basic form, has the framework:

```

identifier: PROCESSOR;
  data declarations (if any)
  executable SMITE statements (at least one required)
identifier: END;

```

For example, the processor to fetch an address from memory in the Intel 8080 description is:

```

GETADD: PROCESSOR;
  OP-PAIR <- MEM[PC+1] // MEM[PC];
  PC <- PC + 2;
GETADD: END;

```

This example contains no data declaration statements. An example of a simple processor containing data declarations as well as executable statements is the memory processor defined earlier:

```

MEMORY-DEVICE: PROCESSOR;
  DECLARE CORE[0:4095]<0:15> MEMORY,
    MAR<0:11> PORT,
    MDR<0:15> PORT,
    STROBE<0> PORT;
  DO FOREVER;
    'IF STROBE IS ZERO, READ''
    'IF ONE, WRITE. THIS MODEL''
    'DOESN'T INCLUDE SPLIT-CYCLE''
    'WRITES (READ/MODIFY/WRITE)''
    CASE STROBE;
      MDR <- CORE[MAR];
      CORE[MAR] <- MDR;
    END CASE;
  MEMORY-DEVICE: END;

```

In addition to data declarations and executable statements, processors may contain definitions of subprocessors. Subprocessor definitions are placed between the last data declaration and the first executable statement of the surrounding processor. Subprocessor definitions themselves have the exact same form as the processor definition just introduced. For example, the memory device example could be re-coded as follows to use subprocessors:

```

MEMORY-DEVICE: PROCESSOR;
  DECLARE CORE[0:4095]<0:15> MEMORY,
    MAR<0:11> PORT,
    MDR<0:15> PORT,
    STROBE<0> PORT;
  READ: PROCESSOR;
    MDR <- CORE[MAR];
    READ: END;
  WRITE: PROCESSOR;
    CORE[MAR] <- MDR;
    WRITE: END;
  DO FOREVER;
    'IF STROBE IS ZERO, READ'
    'IF ONE, WRITE. THIS MODEL'
    'DOESN'T INCLUDE SPLIT-CYCLE'
    'WRITES (READ/MODIFY/WRITE)'
    CASE STROBE;
      READ;
      WRITE;
    END CASE;
  MEMORY-DEVICE: END;

```

Subprocessors may themselves, in turn, contain other subprocessors. Data declared within a subprocessor is temporary data which is lost after the processor is exited.

1.6.2 Scope of Names Recognition

The last example illustrates the use of global data by subprocessors. In that example, the ports MAR and MDR, and the array CORE, were accessed globally by the subprocessors. This usage is an example of the general concept in SMITE of scope of name recognition. SMITE allows data items to be referenced globally from within subprocessors, and also allows subprocessors to redefine data item and processor names. The rules for scope of names recognition are as follows:

1. Within and inclusive of the main processor (the outermost processor of a SMITE computer description), the names known are only those data items and subprocessors declared in the main processor. These names are local to the main processor.
2. Within and inclusive of a subprocessor (whether the subprocessor was itself declared in the main processor or another subprocessor), data items and processors declared within the subprocessor are known as local names. All processors and data items declared in

processors surrounding this subprocessor are known as global names.

3. Any globally known name may be re-declared locally, superseding the global definition.

For example, consider the following skeleton of a SMITE computer description.

```
A: PROCESSOR;
  DECLARE
    U<0:1> REGISTER,
    V<0> MEMORY;
  B: PROCESSOR;
    DECLARE
      W<0:15> REGISTER,
      U<0:15> MEMORY;
  T: PROCESSOR;
    DECLARE
      V <0:15> REGISTER;
      : END;
  B: END;
  D: PROCESSOR;
    DECLARE
      X<0:15> MEMORY;
  D: END;
A: END;
```

The names known within the text of processor A are U (a 2-bit register), V (a 1-bit memory), B (a processor), and D (a processor). The names W, C, and X are not known to A. The names known within the text of processor B are V (a 1-bit memory known globally from the definition in A), W (a 16-bit register declared locally in B), U (a 16-bit memory declared locally in B redefining the definition of U in processor A), C (a processor declared in B), and D (a processor known from the global definition in A). The names known within the text of processor C include everything known in B, plus a redefinition of the global name V (now a 16-bit register). The names known within the text of D include the names declared in A (i.e. U as a 2-bit register, V as a 1-bit memory, and B as a processor), plus X (a 16-bit memory).

1.6.3 Function Processors

In the preceding examples, processors have been used exclusively as subroutines, that is they have been defined and called strictly for their effect, and have had no parameters. SMITE provides the capability for function processors as well as subroutine processors, and permits both types of processors to be defined with parameters. The type and parameters of a processor are specified in the processor header statement.

The header for a simple subroutine processor with no parameters, as used in the previous examples, has the form

```
name: PROCESSOR;
```

If a value (a word) is to be returned as a function result from invocation of the processor, then a width attribute is added to the processor header:

```
name: PROCESSOR width-attribute;
```

For example, a processor to compute an arithmetic/logical function of 8 bits from global data (so no parameters are required) might be defined with the header:

```
ALU: PROCESSOR<0:7>;
```

Processors are invoked in expressions by referencing the processor name. For example, the ALU processor defined above could be invoked and the returned value stored by the statement

```
ACCUMULATOR <- ALU;
```

The ALU processor could also be invoked without the function result being used, such as might be done if only the setting of status bits was desired. A statement to perform that task would be

```
ALU,
```

1.6.4 Processor Parameters

Both subroutine processors and function processors may be defined to receive or return parameters. A subroutine processor header including parameters has the form:

```
name: PROCESSOR ( argument-list );
```

A function processor header including parameters has the following form:

```
name: PROCESSOR width-attribute ( argument-list );
```

The argument list is a list of local data items (formal parameters) which are replaced by the calling (actual) parameters at the time of processor invocation. The number of actual parameters must agree with the number of entries in the argument list.

Processor parameters are passed to and from the processor using either call-by-value (CBV), if the parameter is only referenced in the processor and never assigned a value, or call-by-value-returned (CBR), for parameters which are assigned a value within the scope of the processor. The operation of these parameter passing techniques closely models the operation of many hardware units: the values input to the hardware are gated off of a bus (call-by-value), and the unit performs its function. The returned results, if any, are then gated back onto a bus for transmission to the caller (call-by-value-returned).

For CBV parameters, the actual parameter may be any expression. For CBR parameters, the actual parameter may be any expression allowed on the left hand side of an assignment operator (e.g. data items, concatenated data items, function processor names (for defining the function value), etc.).

If a particular SMITE variable is a calling parameter, and is referenced as a global variable in the processor, the global references and the parametric references use separate data values. For example:

```
A: PROCESSOR;
  DECLARE
    ACTUAL<0:15> REGISTER;
  B: PROCESSOR(FORMAL);
    DECLARE
      FORMAL<0:15> REGISTER;
      FORMAL <- 1;
      IF FORMAL /= ACTUAL
        THEN ...
      END IF;
    B: END;
  ACTUAL <- 0;
  B(ACTUAL);
  IF ACTUAL = 1
```

```

      THEN ...
      END IF,
A: END;

```

Both IF statements will evaluate the relational test as true. The operation of the example is as follows. The main processor (A) sets ACTUAL to 0. The processor B is then invoked, and the value 0 is bound to FORMAL. B then assigns the value 1 to FORMAL, but that value will not be gated back into ACTUAL until B terminates execution and returns to the point of the call. The IF test in B will therefore find that ACTUAL is in fact not equal to FORMAL. B eventually returns to A, at which time the value assigned to FORMAL within B is gated into ACTUAL. The IF test in A therefore finds ACTUAL to have a value of 1.

Every formal parameter in the processor header must be declared in the subprocessor. If the actual parameter is a simple data item or a subfield of a simple data item, then the storage class attribute of the two data items (formal and actual) must be identical (MEMORY and REGISTER are considered identical for this test). The width of the formal parameter must in all cases be greater than or equal to the width of the actual parameter; if the subprocessor assigns a value to the formal parameter, then the widths of the formal and actual parameters must be identical. Parameters may not have a storage class of EXTERNAL, DATA, or PORT, and may not be declared as arrays in the subprocessor.

1.6.5 Intrinsic Processors

SMITE has the following built-in functions to perform shifts and sign extensions:

Function	Operation
SE(D)	Sign extend D
SLL(D,S)	Shift D left logical S bits
SRL(D,S)	Shift D right logical S bits
SLA(D,S)	Shift D left arithmetic S bits
SRA(D,S)	Shift D right arithmetic S bits
SLC(D,S)	Shift D left circular S bits
SRC(D,S)	Shift D right circular S bits

In the above definitions, both S and D may be any arbitrary expressions.

The sign extension operator is useful when the width of an

operand is to be expanded: the most significant bit of the data word is extended until a data word of the appropriate width is formed. The sign extension operator is also used to indicate that an operand of a relational operator is signed, and therefore that an arithmetic comparison should be made.

The shift operators have no predefined width. The width of the result is equal to the width of the word being shifted. The word is shifted as if the shift register is exactly as wide as the data word. For example,

SLC(ACCUMULATOR<0:3>, 2)

performs a 2-bit left circular shift of the 4-bit quantity in a 4-bit register.

The shift operators follow standard shift definitions. For logical shifts, each vacated bit position is replaced by zero. For circular shifts, each vacated bit is replaced by the bit shifted out (wraparound). For arithmetic shifts, vacated bits are replaced by zero (left shift) or the sign bit (right shift).

1.7 External Interfaces

A SMITE computer description may be linked to external processors or to input/output devices of the host computer through data items of storage class EXTERNAL or PORT. These external interfaces support description of a computer's connections to the outside world. In the context of emulation, external interfaces permit routines implemented in QM-1 MULTI microcode to be linked to the emulator.

A data item declared with the attribute EXTERNAL is defined to be an external processor. If the symbol is also declared with an explicit width, the external processor is thus defined to be a function rather than a subroutine. A data item declared as EXTERNAL may not be declared as an array, and may have no parameters. An external processor is invoked in the same manner that any internal SMITE processor is invoked. The differences between external processors and the internal SMITE processors may be summarized:

1. External processors are declared by the appearance of the EXTERNAL storage class attribute in a data declaration

statement. Internal processors are declared by their definition as a PROCESSOR. External processors may have no parameters.

2. For emulation, the microcode for an external processor is not produced by the SMITE compiler. Chapters 6 and 7 discuss the methods used to add external processors to the emulation system. SMITE internal processors are compiled to microcode as part of the computer description.

3. The external and internal processor are invoked in the same manner. However, an external processor may not have arguments.

A data item declared as a PORT is an externally accessible register used as an interface cell. The PORT storage class attribute implies that a value stored by the computer into the port register is transferred as a data or control word to an external device. Similarly, a value read from a port is a status or data word transferred to the computer from the external device. Unlike other data declarations, PORT provides a direct interface with elements outside of the SMITE computer data base. PORT is functionally identical to REGISTER in the emulation microcode, except that an external assembly language routine is automatically activated to process the implied I/O function.

The implementation of PORT and EXTERNAL, the methods of adding external processors to the emulator, and the rules for using predefined processors provided with SMITE are described in Chapters 6 and 7.

1.8 General Rules for Coding SMITE

SMITE computer descriptions may be coded completely free field within columns 1 through 72 of the line image. Blanks may be used freely to improve readability. The following rules define the correct usage requirements for blanks in SMITE:

1. At least one blank is required to separate alphanumeric symbols. For example,

DECLARE ACCUM<0:7>REGISTER;

is illegal because no space separates the symbol DECLARE

from the symbol ACCUM. The correct coding of this statement would be

```
DECLARE ACCUM<0:7>REGISTER;
```

Note that no space is required before the symbol REGISTER, since it is not adjacent to another alphanumeric symbol.

2. Due to the lack of an underscore character in some character sets, the hyphen has been forced to do double duty. Because of this, incorrect interpretations of expressions involving the minus operator are possible. For example,

```
A-B;
```

is a reference to the data item or processor named 'A-B', while

```
A - B;
```

is an expression which will subtract A from B.

3. Blanks may not appear within a symbol. For example,

```
A: PROC ESSOR;
```

is illegal since a space appears in the middle of the symbol PROCESSOR. Symbols may not be split across line boundaries.

4. Statements in SMITE are invariably terminated by a semicolon. This includes all the control statements (chapter 2), expression statements, declaration statements, processor headers, and END-type statements (e.g. END IF;).

1.9 Case Study 1: The Mark 1

In this section we present a SMITE description of a complete computer, the Mark 1. The SMITE description was adapted from a description of the Mark 1 found in Bell and Newell [1]. The complete Mark 1 is as follows:

```
1. MARK1: PROCESSOR;
2.   DECLARE
3.     M[0:8191]<0:31> MEMORY,
```

```

4.      PI<0:15> REGISTER,
5.      F<0:2> DEFINED PI<0:2>,
6.      S<0:12> DEFINED PI<3:15>,
7.      CR<0:12> REGISTER,
8.      ACC<0:31> REGISTER;
9.  DECLARE
10.     STOP EXTERNAL;
11.  DO FOREVER;
12.     BEGIN;
13.         PI <- M[CR]<0:15>;
14.         CASE F;
15.             ' 0 '
16.                 CR <- M[S]<19:31>;
17.             ' 1 '
18.                 CR <- CR + M[S]<19:31>;
19.             ' 2 '
20.                 ACC <- - M[S];
21.             ' 3 '
22.                 M[S] <- ACC;
23.             ' 4 '
24.                 ACC <- ACC + M[S];
25.             ' 5 '
26.                 ACC <- ACC - M[S];
27.             ' 6 '
28.                 IF SE(ACC) < 0
29.                     THEN CR <- CR + 1;
30.                 END IF;
31.             ' 7 '
32.                 STOP;
33.         END CASE;
34.         CR <- CR + 1;
35.     END;
36.  MARK1: END;

```

Line 1 is the processor header. It specifies that a processor is being defined, and that the name of the processor is 'MARK1'. The processor is the main processor since it is not embedded in any other processor.

Lines 2 through 8 are a declaration statement. The registers and memories of the Mark 1 will be defined in this statement.

Line 3 is the declaration of the primary memory. It is specified as an 8192-word array of 32-bit words, and is given the name 'M'.

Line 4 is the declaration of the instruction register. It is 16 bits wide, and will be referred to by the name 'PI'. Its

bits are numbered sequentially from 0 to 15, starting with the leftmost bit of the register.

Line 5 is the declaration of the opcode, or function, subfield of the instruction register. The F subfield is specified to be three bits wide, and is the leftmost three bits of the PI register. The DEFINED attribute, used here to overlay F onto PI, will be explained in Chapter 3.

Lines 6 through 8 complete the register declarations. The remaining subfield of the instruction register, S, is defined, as are the program address register (CR) and the accumulator (ACC).

Lines 9 and 10 declare an external processor, STOP. In this description, no attempt was made to define the clocking or timing of instructions in the Mark 1. Rather than include lines of SMITE to show the manner in which the Mark 1 was halted, an external processor was defined to stop the computer (in some unspecified manner).

Line 11 is a SMITE control statement. Control statements will be explained in Chapter 2; the DO FOREVER statement in this example specifies that the statements in lines 12 through 35 are to be executed repetitively without end.

Line 13 is a SMITE expression statement (an expression statement is any expression terminated by semicolon, the statement termination character). The statement specifies that the upper 16 bits of the memory word at the address specified by the contents of register CR are to be transferred into register PI, thus fetching the next instruction.

Line 14 is a CASE statement, which specifies that one of the following executable statements (lines 16, 18, 20, 22, 24, 26, 28, or 32) is to be executed. The particular statement chosen is determined by the value of F.

Lines 15, 17, 19, 21, 23, 25, 27, and 31 are SMITE comments. A comment in SMITE is any string of text (except two quotes) enclosed in two single-quotes (i.e. '). Comments may be placed anywhere that blanks may be placed.

Line 16 is the specification of the Mark 1 branch instruction, opcode 0. The contents of memory at the address defined by the contents of the S subfield of the current instruction are transferred to the program address register. Note that the program address register is incremented later in the

instruction execution cycle (line 34), and so the address of the next instruction is actually $M[S] + 1$.

Line 18 is a relative jump, opcode 1. The sum of the program address register and the memory word at the address contained in the S subfield of the current instruction is transferred to the program address register. As with opcode 0, the next instruction executed will be at the address one greater than the value left in CR at line 18.

Line 20 is a load complement instruction, opcode 2. The arithmetic complement of the memory word addressed by the contents of S is transferred to the accumulator.

Line 22 is a load accumulator instruction, opcode 3. The memory word at the address specified in the S subfield of the current instruction is transferred to the accumulator.

Lines 24 and 26 are the arithmetic instructions, opcodes 4 and 5. The content of the memory word addressed by S is added or subtracted from the value in the accumulator, and the result is transferred back into the accumulator.

Line 28 is the conditional skip instruction, opcode 6. If the content of the accumulator is negative, then one is added to the program address register, and that result is transferred back into the program address register. Note the use of the SE built-in function in the IF statement to specify that the accumulator contains a signed quantity. The expression in the IF statement could be written in other ways, such as

```
IF ACC<0>
  THEN ...
```

Line 32 is the halt instruction, opcode 7. The computer is halted by a call on the STOP external processor.

Line 34 increments the program address register.

Line 36 is the end of the Mark 1 processor definition.

2. Control Statements

2.1 Introduction

Control structures are provided in SMITE to form groups of statements (context blocks) and to alter the normal sequential flow of statement execution. The control structures in SMITE are designed to provide the capabilities needed to produce computer descriptions, and to adhere to good programming practices.

The need for control statements in computer description is clear even from the short case study of the Mark 1. The processing logic of a computer contains branches based on the presence/absence of a signal, the value of one or more bits, or the value of a status line. The decoding logic extracts several bits of the instruction register to determine the opcode, source register, addressing mode, or destination register. The processing logic may contain repetitive loops for shift operations, normalization, division, and multiplication. The presence of an illegal condition, or of special case logic, may cause the normal processing to be bypassed. SMITE control statements provide the capability for describing these hardware structures.

The term 'context block' refers to a group of SMITE statements forming a unit. Every SMITE control statement (except ESCAPE) forms a context block around its scope. For example, the IF statement forms a context block encompassing the THEN and ELSE statements. The DO statement forms two context blocks, one for the entire loop structure including both the loop controls and the controlled statement, and one just for the controlled statement. The CASE statement forms a variable number of context blocks, one for the entire CASE structure, and one for each substatement within the CASE statement.

SMITE also provides the capability to group multiple statements together into a context block which is treated as a single statement. This may be done with the BEGIN and END statements, e.g.

```
BEGIN;  
  .  
  . (statements)  
  .  
END;
```

in which case statements in the sequence are to be executed in sequential order, first to last, or with the PARALLEL-BEGIN AND PARALLEL-END statements, e.g.

```
PARALLEL-BEGIN;
.
.
.
PARALLEL-END;
```

in which case statements in the group are all to be executed simultaneously in asynchronous parallel.

A context block may be labeled or unlabeled. If a label is used, both the statement beginning the block and the statement ending the block must be labeled with the context block name. For example,

```
A: BEGIN;
.
.
A: END;
```

or

```
B: IF ACC<0>
  THEN ...
  ELSE ...
B: END IF;
```

or

```
C: CASE F;
.
.
.
C: END CASE;
```

2.2 The IF Statement

The IF statement selects a single statement from its component statements for execution. It specifies that the statement following the symbol THEN be executed only if a certain condition (expression) is true. If the condition is false, then either no statement, or the statement following the

symbol ELSE, is executed. The SMITE IF statement has the following two forms:

```
IF expression
  THEN statement
  END IF;
```

or

```
IF expression
  THEN statement
  ELSE statement
  END IF;
```

A context block is associated with the IF statement and may be labeled. If the label option is exercised, then both labels must be present and identical. In addition to labeling the context block these labels allow the compiler to perform verification of correct block nesting. The labeled IF statement has the forms

```
label: IF expression
  THEN statement
  label: END IF;
```

or

```
label: IF expression
  THEN statement
  ELSE statement
  label: END IF;
```

The expression is evaluated according to the standard rules described in section 1.5, and must evaluate to a result that is one bit wide. A value of one is defined as true, and a value of zero as false.

Control passes to the statement immediately following the THEN if the expression evaluates as true (one). Control passes to the statement immediately following the ELSE if the expression evaluates as false (zero). IF no ELSE clause is present and the expression is false, then the statement immediately following the END IF statement receives control.

In either case, at most one of the statements associated with the IF statement is executed. When that statement has completed execution, control passes to the statement immediately following the END IF, unless the statement

selected for execution by the IF is an ESCAPE (which causes an explicit transfer of control).

IF statements may incorporate any SMITE executable statement as the controlled statements, including BEGIN/END or another IF. For example,

```

IF expression1
  THEN statement1
  ELSE IF expression2
    THEN statement 2
    ELSE IF expression 3
      .
      .
      .
      ELSE IF expression
        THEN statement
        END IF;
      .
    END IF;

```

To avoid costly execution by an emulator of the above statement, it should be ordered so that the probability of expression 1 occurring is greater than the probability of expression 2 occurring, and so forth. Using this technique ensures that the least number of expression evaluations and tests is performed for each execution of the IF.

An unnecessary use of the IF statement is the following:

```

IF A = 0
  THEN ZERO <- 1;
  ELSE ZERO <- 0;
  END IF;

```

This entire IF statement can be replaced with the simpler statement:

```

ZERO <- A = 0;

```

2.3 The CASE Statement

The CASE statement provides a limited and highly disciplined branching capability. It consists of an expression, called the selector, and a list or collection of statements. The CASE statement selects one and only one statement from the collection for execution. The CASE statement is similar in usage to an indexed jump, or to a FORTRAN computed GO TO statement, and has the following form:

```
CASE expression;  
    statement  
    statement  
    .  
    .  
    .  
    statement  
END CASE;
```

The CASE statement forms a context block, which may or may not be labeled. If the label option is exercised, then both the CASE and the END CASE statements must be labeled with the identical context block name. In addition to labeling the context block, labels are used by the SMITE compiler for verification of correct context block nesting. A labeled CASE statement has the form:

```
label: CASE expression;  
  
    statement  
  
    statement  
  
    .  
  
    .  
  
    .  
  
    statement  
  
label: END CASE;
```

The selector expression in the CASE statement is evaluated according to the standard expression evaluation rules described in section 1.5. The value of the expression is interpreted as a positive number n bits wide, where n is the width of the result. This number is used as an index into the

collection of statements that follows, and the appropriate statement receives control. The first statement in the sequence is executed if the number is zero, the next one if the number is one, and so forth.

Since n bits can select one of 2^n statements, 2^n statements are required to occur between the CASE and END CASE statements. One and only one statement in the collection receives control; when that statement has finished executing, control proceeds to the END CASE, unless the selected statement contained an ESCAPE to alter the normal flow of control.

Often situations will occur when one or more of the statements of the 2^n statement collection are to evoke no action, i.e. control is to proceed directly to the END CASE statement with no operation performed. For these instances, the NULL statement is provided. It may be coded as a statement in a CASE sequence, and will act as a do-nothing statement.

There are various alternatives to using the CASE statement. One is the nested IF:

```

IF expression1
  THEN statement1
  ELSE IF expression2
    THEN statement 2
    ELSE IF expression 3
      .
      .
      .
      ELSE IF expression
        THEN statement
        END IF;
      .
    END IF;

```

As discussed in section 2.2, the nested IF statement can be very inefficient in an emulation if the order that the expressions appear in the statement is not chosen carefully.

Another alternative to the CASE statement is the parallel IF construct, an example of which is:

```

PARALLEL-BEGIN;
  IF expression1 THEN statement1;
  IF expression2 THEN statement2;

```

```

.
.
IF expression m THEN statement m;
PARALLEL-END;

```

In this form, all decodes are intended to occur simultaneously, and the assumption is usually made that only one of the expressions will ever be true. The order of the expressions in this method has no impact on the efficiency of execution, since the IF statements are not nested.

When deciding on which method of decoding to use, two considerations must be made. The first is to determine whether all the conditions to be decoded are mutually exclusive. If they are not, the parallel IF is required. If the conditions are mutually exclusive, then two choices are available in addition to the parallel IF, namely the nested IF and the CASE. The CASE statement is usually preferable to the nested IF statement on grounds of clarity of the resulting description.

Instruction operation code decoding with the CASE statement is often performed in a single field within the instruction. For example:

```

DECOD: CASE INSTRUCTION <7:6>;
.
.

```

```

DECOD: END CASE;

```

The selector can also be more complex:

```

DECODE: CASE OP-FIELD(INSTRUCTION);
.
.

```

```

DECOD: END CASE;

```

where OP-FIELD is a processor defined as:

```

OP-FIELD: PROCESSOR<3:0> (IN-VALUE);
DECLARE IN-VALUE REGISTER;
OP-FIELD <- ACCUM-STATE // IN-VALUE<5:4> // IN-VALUE<1>;
OP-FIELD: END;

```

In this case the selector is a processor that returns a value

four bits wide. The value is a concatenation of two fields within the instruction and the one bit wide ACCUM-STATE.

2.3.1 The DECODE Statement.

The decode statement provides still another alternative to the CASE structure. It is used primarily where the number of statements selected by the case selector expression is not a simple relationship. This happens in many processors where all values of the opcode are not implemented as valid instructions, or where the operation is based on other factors in addition to the opcode. The decode statement allows the user to specify only the statements which will actually be selected and to specify the selection logic in nanocode. The selected statements are specified in SMITE; the selection processing must be done in nanocode which is separately loaded on the QM-1.

The form of the DECODE statement is:

```
label DECODE length expression;
      statement-list
label END DECODE;
```

where:

label is optional but if present it must be present in both places and match.

length is optional and specifies the number of statements in the statement-list.

expression is any valid SMITE expression which is used to determine which of the statements in the statement-list is to be selected.

statement-list is one or more SMITE statements to be selected by the decode operation. A NULL statement may also be used in the statement list.

END DECODE may also be ENDDECODE.

The DECODE statement generates a context block, the same as a CASE statement, which encompasses all statements in the statement-list. After execution of a selected statement control will pass to the END DECODE statement. The DECODE statement will cause a micro instruction developed specifically for the emulation, called DECODE (opcode = 176 octal) to be executed. This instruction will contain the

address of the selection expression result and a pointer to an address table giving the beginning address of all statements in the statement-list. The format of the DECODE instruction is:

```

bit 17-11  opcode
    10- 6   QM-1 register with selector value, right
            justified
    - 5- 0   QM-1 register which points to the jump table
  
```

The format of the address table is:

```

word 0      address of first SMITE statement in selector
list
word n-1    address of Nth SMITE statement in selector
list
  
```

In the example of the MARK 1 from Chapter 1, the DECODE statement could be used instead of the CASE by changing the following statements:

```

CASE F;      to    DECODE F;

and

END CASE;    to    END DECODE;
  
```

Although the form is the same for this simple example, the user must now develop the appropriate nanocode to select one of the eight statements based on the value in "F". While this example is not significant, it can be used to demonstrate the "open endedness" of the DECODE statement. Once control is transferred to the user's nanocode, by the DECODE statement, he can implement very complex selection/decode algorithms in high speed nanocode and may provide additional processing in the central decode loop. For instance, on the decode nanoinstruction, the Mark I could load the instruction register, separate it into fields, and perform the statement selection. Sophisticated application and interface to nanocode can be accomplished with the OPDEF, syntax macro and direct code capabilities described in Chapter 8.

2.4 The DO Statement

The SMITE DO statement provides the means for repetitive execution of statements or groups of statements. All forms of the DO statement are based on the concept of a controlling statement and a controlled statement. The controlled statement may be compound (e.g. BEGIN/END) to form complex loop bodies. The general form of the DO statement is:

```
DO terminator-clauses;
    controlled statement;
```

The DO statement may be labeled:

```
label: DO terminator-clauses;
    controlled statement;
```

Two basic DO statements are defined in SMITE: DO with controlled termination clause(s), and DO FOREVER (i.e. DO with an uncontrolled terminator clause). The form with controlled terminator clause(s) provides controlled repetition (with optional iteration), while the DO FOREVER form provides uncontrolled repetition and must be used with the ESCAPE (section 2.5) statement if loop termination is desired.

2.4.1 The Uncontrolled DO Statement

The DO FOREVER statement has the following forms:

```
DO FOREVER;
    statement;
```

or

```
label: DO FOREVER;

    statement;
```

The DO FOREVER form of the DO is equivalent to the DO WHILE 1. This form specifies an unlimited repetition of the controlled statement. Dependent upon the application, an infinite number of repetitions may be desired, such as the main loop of an emulator. In other cases, some set of circumstances may imply termination of the loop, in which case the ESCAPE statement (section 2.5) provides the method of termination.

In certain cases where a DO FOREVER has an ESCAPE, the

FOREVER may be replaced by a terminator clause, and the
ESCAPE eliminated:

```

LOOP: DO FOREVER;
      BEGIN;
        IF ERROR THEN ESCAPE LOOP;
        .
        .
      END;

```

may be changed to

```

DO WHILE NOT ERROR;
  BEGIN;
  .
  .
  END;

```

2.4.2 The Controlled DO Statement

There are three forms of controlled repetition DO
statements:

- 1) Termination when a given condition becomes true (DO
UNTIL).
- 2) Termination after a given condition becomes false
(DO WHILE).
- 3) Termination when a specified iteration sequence is
completed (DO FOR).

2.4.2.1 The DO WHILE Statement

The form of the DO WHILE statement is:

```

DO WHILE expression;
  controlled-statement;

```

or

```

label: DO WHILE expression;
  controlled-statement;

```

The expression must evaluate to a one bit wide result.
The WHILE clause executes the controlled statement as
long as the specified expression is true (evaluates to a

one), or until control is transferred out of the context of the loop (i.e. by an ESCAPE statement). The evaluation and test of the expression is performed once each time through the loop, and is made before the controlled statement is executed.

2.4.2.2 The DO UNTIL Statement

The form of the DO UNTIL statement is:

```
DO UNTIL expression;
    controlled-statement;
```

or

```
label: DO UNTIL expression;
    controlled-statement;
```

The expression must evaluate to a one bit wide result. The UNTIL clause executes the controlled statement as long as the expression is false, or until control is transferred out of the context of the loop. The evaluation and test of the expression is performed once each time through the loop, and is made after execution of the controlled statement.

2.4.2.3 The DO FOR Statement

The form of the statement is:

```
DO FOR expression UP TO expression STEP expression;
```

or

```
DO FOR expression DOWN TO expression STEP expression;
```

A context block label may be used (e.g. 'label: DO').

The words UP and DOWN may be omitted, in which case UP is assumed. The phrase 'TO expression' may be omitted as well, in which case loop termination is not provided by the DO FOR statement, and if required must be separately provided using an ESCAPE statement. The STEP phrase may be omitted when the keyword UP (or no UP/DOWN keyword) is used; if omitted a value of one is used. The keyword STEP may not be omitted if the keyword DOWN is used, as a negative step is required. The phrases after 'FOR expression' may occur in any order:

DO FOR expression STEP expression UP TO expression;

or

DO FOR expression TO expression STEP expression DOWN;

or any of the other possible orderings.

The FOR clause provides capabilities similar to the FORTRAN or PL/I iterative DO loops. The UP/DOWN keyword only affects the form of the termination test, and not the incrementation of the loop counter. The action followed by a FOR clause loop is as follows:

- 1) The initial-value expression (the one immediately following the FOR) is evaluated. If the last evaluated operator in the expression is a data transfer ('<-'), then the evaluation is stored in the indicated data item, and all subsequent incrementations of the initial expression value will also be stored in the same data item. Note that the initial-value expression is evaluated only once, before the first execution of the controlled statement.
- 2) The body of the loop (the controlled statement) is executed.
- 3) The current value of the loop control variable is incremented by the value of the STEP expression, or by one if no step was specified. The STEP expression is only evaluated once, before the initial execution of the loop, and will not be evaluated again. If DOWN is specified, the STEP expression should evaluate to a negative number. Otherwise, loop execution occurs until the loop counter overflows and becomes negative. The loop counter is maintained in an 18-bit register by the compiler if not explicitly named, i.e.

DO FOR 0 TO 7;

causes the compiler to maintain the counter in an 18-bit register, while

DO FOR COUNT<0:3> <- 0 TO 7;

maintains and tests a 4-bit counter in the data item COUNT.

4) The resulting value (stored if so specified in the DO statement) is compared against the (optional) loop termination value. No testing is performed if no termination value is given, and the loop becomes a form of the DO FOREVER with automatic incrementation of the counter. If a termination value is given, then the test is made based on the UP or DOWN direction, and if the limit has not been exceeded then control returns to step 2. Otherwise, control passes to the succeeding statement.

2.4.3 DO Statement Context Blocks

A context block is associated with the DO statement, and may be optionally labeled. This context block includes the controlled statement, no matter how complex. The end of the context block can be thought of as being just after the controlled statement, and just before the statement immediately following the range of the DO. The controlled statement will define a context block of its own, unless it is a simple statement. A few examples will explain this more clearly.

```
DO WHILE SP < PC;
  PUSH(PC);
```

The first example consists of a DO statement that has a simple statement as its controlled statement. In this case there is only one context block, which is the context block associated with the entire DO statement.

```
DO-BLOCK: DO FOREVER;
  CONTROLLED-BLOCK: BEGIN;
    P <- PC + 1;
    DECODE(PC);
  CONTROLLED-BLOCK: END;
```

The above example is composed of two context blocks. The block named DO-BLOCK is the context block for the DO statement, while the block named CONTROLLED-BLOCK is the context block of the compound statement controlled by the DO statement.

```
DO FOR PC UP TO PC+5 STEP 1;
  BEGIN;
    PUSH(PC);
    DECODE(PC);
  END;
```

This example also consists of two context blocks, with the exception that this time they are unnamed. The outer block is the context block for the DO statement, while the BEGIN-END block is the context block associated with the controlled statement.

2.5 The ESCAPE Statement

Occasionally within a computer description, situations occur in which it is desirable to immediately transfer control from the statement currently being executed to the end of some surrounding context block. The reason for the transfer may be an error condition, a "found-it" condition in a search, etc. The method of causing this transfer is the SMITE statement ESCAPE.

ESCAPE causes control to be immediately transferred to the end of a surrounding context block, i.e. a context block that lexically contains the ESCAPE statement. Escapes may not be made to a point outside the current processor. Using the ESCAPE statement with parallel context blocks and timing statements imposes additional restrictions which are described in sections 5 and 4.2, respectively.

The two forms of the ESCAPE statement are:

ESCAPE;

or

ESCAPE label;

The use of the form without the label causes control to be transferred to the end of the context block immediately surrounding the ESCAPE. For example:

```
BEGIN;  
  .  
  . (statement set 1)  
  .  
  ESCAPE,  
  .  
  . (statement set 2)  
  .  
  END;
```

Execution of statement set 2 is bypassed. The ESCAPE statement transfers control to the statement following the END statement.

The more common form of ESCAPE statement includes the label of the context block to be escaped, because the ESCAPE without a label is restricted in application. For example,

```

IF DECODE-ERROR
  THEN ESCAPE;
END IF,

```

is a no-operation statement. The context block surrounding the ESCAPE is the IF statement, and so the action of the ESCAPE is to transfer control to the statement following the IF.

The use of the form of ESCAPE with a label allows control to be transferred to the end of any surrounding context block within the current processor. This form of ESCAPE is useful for terminating the current iteration of a DO statement, for terminating loops, and for transfer of control out of nested context blocks upon the detection of an error condition.

```

BLOCK1: BEGIN;
.
.
.
  BLOCK2: BEGIN;
  .
  .
  .
  IF OP-BAD
    THEN ESCAPE BLOCK1;
  END IF;
  .
  .
  .
  IF BAD-DATA
    THEN ESCAPE BLOCK2;
  END IF;
  .
  .
  BLOCK2: END;
  . (statement A)
  .
  .
  BLOCK1: END;
  . (statement B)
  .
  .

```

If OP-BAD is a one, an escape is made to the end of context block BLOCK1. The next statement to be executed is that at (statement B). If BAD-DATA is a one, an escape is made to the end of context block BLOCK2. The next statement to be executed is that at (statement A).

.
.
ESCAPE; (1)
.
.
.

ESCAPE A-BLOCK; (2)
.
.
.

END;

In this example, the ESCAPE marked (1) causes control to be passed to the end of the controlled statement. The controlled statement will be executed again (from the top, of course) by virtue of the DO FOREVER. The ESCAPE marked (2) causes control to be transferred to the end of the DO statement, and thereby terminates the action of the infinite loop defined by the DO FOREVER statement.

2.6 Case Study 2: FTSC Floating Point Unit

The Raytheon Fault Tolerant Spaceborne Computer (FTSC) is a computer designed for use in future satellite applications. An emulation-oriented description developed of the FTSC computer provides an illustration of the use of SMITE for complex arithmetic operations. In this case study, we will examine the description of the floating point unit of the FTSC.

The FTSC is a 32 bit computer with the following floating point operations:

Opcode	Definition
LDNF	Complement and load operand
LDAF	Load absolute value of operand
ADDF	Floating add
SUBF	Floating subtract
MPYF	Floating multiply
DIVF	Floating divide
SRTF	Floating square root
VADDF	Vector floating add
VSUBF	Vector subtract
VMPYF	Vector multiply
VIPF	Vector inner product
VSMF	Vector scalar multiply
CFX	Floating to fixed conversion
UPF	Unpack floating point
PKF	Pack fixed point to floating
CFL	Fixed to floating conversion
JZEF	Jump if floating point 0
JNZF	Jump if not floating point 0
JPSF	Jump if positive, non-zero floating point
JMZF	Jump if negative or zero floating point

The exact operation of each floating point instruction need not be described at this time. There are a few considerations, though, which are important for understanding the FTSC description.

- 1) The 32 bit FTSC floating point data word is organized as a 24 bit two's complement fractional mantissa in the upper bits and an 8 bit exponent in the lower bits. The exponent is biased by $X'80'$. Hence, 1.0 is represented by $x'40000081'$ (i.e. $0.5 \times 2^{+1}$).
- 2) Any number which is an exact power of two is normalized

to have a mantissa of X'800000' and an exponent of the power of two plus one. Numbers with only the sign bit set, such as are obtained after extracting the mantissa from floating point powers of two, are singular for certain operations, such as two's complement, and therefore precautions are sometimes required to detect these special cases.

3) Floating point 0 is represented as 00000080, a mantissa of 0 and an exponent of 0. This presents special problems for the floating point addition and subtraction operations, where the operands must be aligned.

4) The MPYF (multiply floating) instruction produces a 48 bit result. The Version 1 SMITE compiler could not manipulate words longer than 36 bits, and therefore the description was written to separate the result into two smaller words.

5) The floating point instructions must set an error flag if floating point exponent overflow or underflow occurs.

2.6.1 Overall Design

The first task is to identify the instructions or functions which are used as processing steps in other instructions. This division produces a natural set of floating point processors. These low level floating point instructions and functions may then be examined to identify additional common processes.

Figure 5 shows the organization of the FTSC floating point emulation in terms of SMITE processors. The top-down ordering of the floating point instructions is straightforward. In this design, the VIPF instruction could also have been implemented as a series of MPYF and ADDF instructions. The SRTF instruction has an algorithmic design very similar to the manual 'two digits at a time' method for determining decimal square roots, which precluded the use MPYF and ADDF subprocessors.

Once the basic floating point instructions were identified, common processes within these instructions were also readily apparent.

1) The results of the floating point operation are normalized. Common normalize processors perform this function.




Figure 5 FTSC Floating Point Processors

2) The mantissa and exponent operations in the FTSC ALU require more than 24 or 8 bits of data to determine carry and overflow. The FLOATING-PREP processor sign extends the floating point data, a process patterned after the actual hardware operation. By extending the data to 32 bits, commonality with fixed point arithmetic processors is achieved.

3) The FLOAT processor converts fixed point numbers to floating point for the PKF instruction (variable exponent pack floating) and CFL instruction (fixed exponent convert to floating).

2.6.2 Top Level Processor Coding

The top level floating point processors illustrate simple DO statements and processor invocations. The vector scalar multiply (VSMF) processor was designed and coded under the following ground rules:

1. The instruction operand is contained in the variable OPERAND, which will be destroyed by subsequent processors.
2. The processor FLOATING-MULTIPLY (no parameters) multiplies REG-OP by OPERAND, and returns the value in REG-OP.
3. The result REG-OP must be stored in GPXR[RB] after each multiply, and RB must then be incremented. RB is a 3 bit field of the instruction register. The new value of REG-OP is loaded from the GPXR array addressed by the incremented value of RB.
4. Three multiplies are performed in the scalar multiply instruction. The index RB may wrap around the address space of register array GPXR.

The FLOATING-ADD processor was designed and coded under the following constraints:

1. The two mantissas are contained in the variables OPERAND and REG-OP. The two exponents are the variables OPERAND-EXP and REG-OP-EXP.
2. The result of the floating point add operation, prior to normalization, will be stored back in REG-OP and REG-OP-EXP.

3. The operand values are moved to REG-OP and REG-OP-EXP if REG-OP equals 0. This test is necessary due to the representation of 0. The values in REG-OP and REG-OP-EXP are the correct result of the floating point add if OPERAND equals 0.
4. The OPERAND value replaces the value in REG-OP if OPERAND-EXP exceeds REG-OP-EXP by more than 23. OPERAND-EXP also replaces REG-OP-EXP in this case. This test covers the case where REG-OP adds no significance to the result.
5. The values in REG-OP and REG-OP-EXP are unchanged if REG-OP-EXP exceeds OPERAND-EXP by more than 23 (i.e. OPERAND adds no significance to the result).
6. If the exponents differ by less than 23, invoke the processor UNNORMALIZE, which aligns the mantissas REG-OP and OPERAND and forms the resulting exponent. Add the aligned mantissas.
7. Invoke the NORMALIZE processor to normalize the results.

The description written of the scalar multiply was the following:

1.	WORK3 <- OPERAND;	'SAVE OPERAND'
2.	DO FOR 1 TO 3;	'LOOP 3 TIMES'
3.	BEGIN;	
4.	FLOATING-MULTIPLY;	'INVOKE
	PROCESSOR'	
5.	GPXR[RB] <- REG-OP;	'STORE RESULT'
6.	REG-OP <- GPXR[RB<-RB+1];	'LOAD NEXT REG'
7.	OPERAND <- WORK3;	'RESTORE
	OPERAND'	
8.	END;	

Line 1 is a simple transfer operation. Line 2 is an example of a simple DO statement. The loop index is not used by the loop code, and is therefore not maintained in an explicit loop variable. Lines 5 and 6 could be recoded to use the loop index. In that case, line 2 would be recoded as

DO FOR I <- 0 TO 2;

and lines 5 and 6 would become

```

GPXR[RB+I] <- REG-OP;

REG-OP <- GPXR[RB+I+1];

```

Line 3 starts the block for the DO statement. The body of the loop exceeds on statement, and therefore is surrounded by a BEGIN/END context block. Line 8 is the END statement which closes the block.

Line 4 invokes the processor FLOATING-MULTIPLY. Line 5 moves a variable onto an array entry. (The range of address values for the array GPXR was defined in the data declaration as 0 to 7.) The contents of the variable RB supply the address.

Line 6 performs a compound operation. Following the right to left order of operations, the contents of RB plus 1 replace RB. The new value of RB is used as the address for indexing the array GPXR, and the corresponding entry is loaded into REG-OP. Instruction 7 is a simple data transfer.

The description written for the FLOATING-ADD processor was the following:

```

FLOATING-ADD:  PROCESSOR;
1.  IF (OPERAND-EXP > REG-OP-EXP + 23)
    OR (REG-OP = 0)
2.  THEN BEGIN;
3.  REG-OP <- OPERAND;
    REG-OP-EXP <- OPERAND-EXP;
4.  END;
5.  ELSE
6.  IF (OPERAND-EXP > REG-OP-EXP - 23)
    AND (OPERAND /= 0)
7.  THEN BEGIN;
8.  UNNORMALIZE;
9.  REG <- OP-REG-OP + OPERAND;
10. END;
11. END IF;
12. END IF;
13. NORMALIZE;
FLOATING-ADD:  END;

```

This example could have been coded with 4 simple IF tests, instead of the 2 compound IF tests actually used. With either method, the branches of the IF statements should converge only at the concluding NORMALIZE operation.

2.6.3 VSMF Example

One of the examples in the previous section was the vector scalar multiply (VSMF) instruction description. In this section, the related coding is followed down to the low-level processors.

The VSMF processor invoked only the FLOATING-MULTIPLY processor. It supplied two 32-bit registers (REG-OP and OPERAND) to the FLOATING-MULTIPLY processor, and received a 32 bit result in REG-OP. The processor FLOATING-MULTIPLY was constructed under the following design constraints:

1. The variables REG-OP and OPERAND are operated upon, and the result is stored back in REG-OP.
2. The MULTIPLY processor used for the fixed point multiply instruction is used to multiply the mantissas. MULTIPLY performs a 32-bit multiplication of the variables REG-OP and OPERAND and returns the result in REG-OP and EX. The actual multiplication algorithm uses the absolute values of the variables.
3. The result is normalized and stored in REG-OP and EX.
4. A zero result is stored as 00000080.
5. Exponent overflow and underflow must be tested.

The description of FLOATING-MULTIPLY is:

```

FLOATING-MULTIPLY:  PROCESSOR;
  FLOATING-PREP;
  REG-OP <- SLL(REG-OP, 8);
  MULTIPLY;
  IF REG-OP = 0
    THEN REG-OP <- X'80';
  ELSE BEGIN;
    REG-OP-EXP <- REG-OP-EXP + OPERAND-EXP;
    DBL-NORMALIZE;
  END;
  END IF;
FLOATING-MULTIPLY:  END;

```

As a consequence of requirements 2 and 5, the first operation, performed by an invocation of FLOATING-PREP, is to separate the mantissa and exponent of each input into

two variables. The mantissas are sign extended, and the characteristic is padded with zeros.

The call to FLOATING-PREP is the first code in the FLOATING-MULTIPLY routine. The value of REG-OP is now tested for zero after a left shift of 8 bits, and a floating point zero value is stored if required. Given unnormalized multipliers, however, it is conceivable that the upper 24 bits of the product are zero and the lower bits non-zero. This process, although intuitively inaccurate, is the one described because it is a modeling of the actual process of the computer.

In the case of a non-zero product, the mantissa has been calculated and is stored in the appropriate variables. The pre-normalization exponent is calculated. The 48 bit result must then be normalized. Although this appears to be a cumbersome operation, only the more significant data word is operated upon. If it is all zeroes or ones, the less significant data word is moved up and the normalization count is incremented by 24. The normalization count of the upper word is found. The corresponding number of bits is shifted out of the lower word. The exponent of the result is adjusted by the normalization count and tested for overflow/underflow. The mantissa and exponent are then merged back into the result variable.

The important concept in this case study is that complex arithmetic processes may be described with SMITE. Analysis of the operations performed in the actual hardware will usually lead the SMITE programmer to a natural modularization and flow of processors which closely models the operations performed in the hardware.

3. Storage Sub-Structure and Operation Widths

3.1 Introduction

Many classes and types of memory devices and structures exist in digital computer systems. Memories vary widely with respect to access time, storage capability, read and/or write capabilities, interfaces, error detection, etc. SMITE provides a means of declaring data items, the attributes in the data item declaration statement, that describes the differing structural aspects of different memory devices. The active characteristics of some memory devices, such as error detection and correction or associative access, are properly described with SMITE processors. This chapter presents the full storage attribute facility provided in the SMITE language.

In addition, SMITE provides a basic abstract data format capability to define data patterns which are independent of their physical location. For example, an instruction format may successfully be defined as subfields of the instruction register. Floating point data, however, is more properly defined using the abstract data format technique, since floating point data may occur in many different locations of the machine, and is not bound to a specific register. The abstract data format definition attribute is also presented in this chapter.

3.2 Data Item Attributes

As described in section 1.4, a data item declaration within a DECLARE statement consists of the data item name followed by its attributes. The complete form is as follows:

identifier array width class defined

where 'array' represents an array attribute ('[...]'), 'width' represents a width attribute ('<...>'), 'class' represents a storage class attribute (e.g. REGISTER or MEMORY), and 'defined' represents a defined attribute (section 3.3).

Some (or all, when the DEFAULT facility of section 3.4 is used) of the attributes may be omitted. Those attributes which do appear must appear in the order shown above. The

defined and array attributes are always optional, while the width and class attributes are optional only if provided by a DEFAULT. The width and array attributes are described in sections 1.4.4 and 1.4.5 respectively. The defined attribute is described in section 3.3; the remainder of this section is devoted to the description of the storage attribute.

The various types of storage attributes defined in SMITE are used to describe the physical characteristics of the storage elements (register, memory, I/O port, etc.) and to define abstract data formats. The following storages attributes are provided in SMITE:

REGISTER - indicates a storage element likely to be used with high frequency. It is otherwise equivalent to MEMORY, and therefore corresponds to the intuitive notion that a register is a fast memory.

MEMORY - indicates an undistinguished storage element in the sense that the use of the data item does not constrain the element to have any special properties (e.g. I/O port, speed, etc.). Mass storage devices could (potentially) be represented using huge arrays having the MEMORY attribute; however, SMITE does not support any virtual memory or secondary storage capability to implement this feature if the limits of the QM-1 are exceeded.

SWITCH - corresponds intuitively to switches on the operators console of the target computer. It represents a read-only binary storage device which is to be connected to the operator interface mechanism.

LIGHT - is the display analog of SWITCH. It is a write-only binary storage device which is to be output through the operator interface mechanism.

PORT - denotes a register which is an I/O interface. PORT functions identically to REGISTER in the resultant microcode, except that an external microcode routine is automatically activated to process the implied I/O function.

FLAG - denotes a 1 bit wide REGISTER.

EXTERNAL - indicates that the named data item is an external microcode closed procedure. If a width definition is attached to the name, the procedure is assumed to be a function; no array attribute is allowed with a storage element that has an EXTERNAL attribute.

CLOCK - defines the storage element name to be used to reference the internal clock (section 4). The clock is kept in a SMITE internal location which imposes the following restrictions:

1. only one clock declaration is permitted,
2. the declaration must not have an array attribute,
3. the clock may not be DEFINED (section 3.3) onto another storage element, and
4. the clock must have a width attribute 36 bits wide.

One other class attribute, DATA, is available to define abstract data formats which are referenced in several different physical locations. One typical use of DATA would be to define the floating point number format. No storage is allocated for DATA items; reference to these items is always qualified by a simultaneous reference to a register or other storage elements.

The interpretation of a complete qualified reference (of the form STORAGE.DATA) is as follows. For a parent DATA item (i.e. one which is not DEFINED onto another item), if the bit width of the storage and DATA items correspond, then they are directly overlapped. If the bit widths differ, however, then if the DATA item bit definitions are numbered in the same direction as the storage item, and if the subfield indicated by the DATA item bit definition exists in the storage element, then the corresponding overlay is made. For DATA fields which are DEFINED onto other DATA items, the parent item is first overlapped onto the storage item as explained above. The DEFINED subfield as positioned by the parent is then used. For example, given the declarations

```
DECLARE
  REG<15:0> REGISTER,
  FULL-WORD<0:15> DATA,
    LOW-BYTE<0:7> DATA DEFINED FULL-WORD<8:15>,
    HIGH-BYTE<0:7> DATA DEFINED FULL-WORD<0:7>,
  MIDDLE-BYTE<11:4> DATA;
```

then

```
REG.FULL-WORD
```

is equivalent to the entire register,

REG.LOW-BYTE

selects the rightmost 8 bits of the register,

REG.HIGH-BYTE

selects the leftmost 8 bits of the register, and

REG.MIDDLE-BYTE

selects the middle 8 bits of the register.

3.3 The DEFINED Attribute: Storage Overlays

SMITE allows a tree structured data definition technique, whereby subfields may be defined as overlays on a parent data item. Attributes of the parent data item (except an array attribute) are passed to the subfield unless overridden by explicit declarations for the subfield. The overlay of one data item on another is achieved through the use of the DEFINED attribute in the data declaration statement for the subfield. A DEFINED attribute has the following form:

DEFINED parent-subfield-reference

In the following example,

```
DECLARE
  ACCUM<7:0> REGISTER,
  SIGN FLAG DEFINED ACCUM<7>;
```

the data item ACCUM is the parent, and SIGN is the subfield.

The parent subfield reference is a description of the portion of the parent to be overlaid by the subfield. It consists of the parent data item name, an array attribute, and a width attribute. If either or both of the attributes are omitted, the entire scope of the parent is assumed.

The two data items must be compatible, i.e. the total the number of bits referenced in the parent must be equal to the total number of bits in the subfield, an element of the subfield must not cross a word boundary of the parent, and if multiple words of the parent are referenced then the entire width of each parent word must be referenced. The following overlay operations are permitted:

1. A word of the parent may be broken up into one or more subfields, or into an array of subfields. For example, the following declarations are permitted:

```
DECLARE
  DATA-REGISTER<0:31> REGISTER,
  ADDRESS-FIELD<0:23> DEFINED DATA-REGISTER<8:31>,
  BYTES[0:3]<0:7> DEFINED DATA-REGISTER;
```

2. A parent array may be broken up into an array of smaller sized words. For example,

```
DECLARE
  PARENT-ARRAY[0:7]<0:31> REGISTER,
  SUBFIELD-ARRAY[0:15]<0:15> DEFINED PARENT-ARRAY;
```

3. A portion of a parent array's address space may be overlaid by a subfield. For example,

```
DECLARE
  PARENT-ARRAY[0:0'77777']<0:47> MEMORY,
  GENERAL-REGISTERS[0:7]<0:47> DEFINED
  PARENT-ARRAY[0:7];
```

Rather complicated structures may be constructed, if required, with the DEFINED attribute. For example,

```
DECLARE
  DWM[0:X'FFFF']<0:63> MEMORY,
  WM[0:X'1FFFF']<0:31> DEFINED DWM,
  HWM[0:X'3FFFF']<0:15> DEFINED WM,
  BM[0:X'7FFFF']<0:7> DEFINED HWM,
  BTM[0:X'3FFFFF']<0> DEFINED BM,
  INTERRUPT-VECTOR[0:7]<0:31> DEFINED
  WM[X'300':X'307'];
```

The larger-to-smaller requirement for the structure hierarchy must be strictly observed. For example,

```
DECLARE
  MAIN[0:1023]<0:15> MEMORY,
  SUBFIELD<0:31> DEFINED MAIN[0:1];
```

is illegal since the item SUBFIELD is required to cross a word boundary. One way to correctly define this structure is as follows:

```
DECLARE DOUBLEWORD[0:511]<0:31> MEMORY,
  MAIN[0:1023]<0:15> DEFINED DOUBLEWORD,
```

SUBFIELD DEFINED DOUBLEWORD[0];

Note the use of omitted attributes in the DEFINED attribute. The array and width subfield reference attributes are missing from the parent description for the MAIN subfield, and therefore are obtained from the data declaration of the parent item. The attributes obtained are [0:511] and <0:31> respectively. The width subfield reference attribute is also omitted from the SUBFIELD declaration, and so the full width of the parent is referenced.

The above example also illustrates the inheritance of attributes from the parent to the subfield itself. In the declarations of both MAIN and SUBFIELD, the storage class attribute is omitted, and therefore MEMORY is inherited from the parent DOUBLEWORD. In the declaration of SUBFIELD, the width attribute is omitted, and so a width of <0:31> is inherited from its parent, DOUBLEWORD. Note that the length attribute, '[0:511]' is not inherited by SUBFIELD from DOUBLEWORD.

Class attributes declared in subfield data items must be identical with the class attribute of the parent. For purposes of comparison, REGISTER and MEMORY are considered equivalent.

An interaction occurs between the block structuring produced by nested processors and global data references, and the DEFINED attribute. In the following example:

```
A: PROCESSOR;
  DECLARE Q <15:0> REGISTER;
  : PROCESSOR;
    DECLARE B FLAG DEFINED Q <15>;
    DECLARE Q <7:0>;
    .
    .
    J: END;
  .
  .
  A: END;
```

The item B is a subfield of the global data item Q instead of the data item Q local to processor J, because the declaration of the local data item Q has not been seen at the time B is declared.

```

A: PROCESSOR;
  DECLARE Q<15:0> REGISTER;
J: PROCESSOR;
  DECLARE Q<7:0>;
  DECLARE B FLAG DEFINED Q<15>;
  .
  .
  J: END;
  .
  .
A: END;

```

The only difference between this and the last example is the order of the DECLARE statements in processor J. The definition of the subfield B is now in error, since it refers to the local Q as the parent and no bit numbered 15 exists in the Q local to J. To avoid confusion, a local parent should always be declared immediately prior to its subfields.

3.4 Defaults

Frequently, a number of data items will be declared in a SMITE computer description with identical attributes. For example:

```

DECLARE
  A<7:0> REGISTER,
  B<7:0> REGISTER,
  C<7:0> REGISTER,
  D<7:0> REGISTER,
  E<7:0> REGISTER,
  F<7:0> REGISTER;

```

SMITE allows the declaration of common attributes by the declaration of the special data item DEFAULT. Any time a data declaration is made and an attribute is missing, the DEFAULT declaration is searched for the missing attribute. The use of defaults simplifies the computer description, and permits concentration on the unusual characteristics of a data item declaration.

A declaration of defaults is made by declaring the special identifier DEFAULT, using the same data declaration statement form used for all other data item definitions. The DEFAULT declaration may appear any place that a data declaration is

valid. The name DEFAULT is a SMITE reserved word, however, and may never be used as an actual data item, nor is it ever allocated space by the compiler. The following example

```
DECLARE
  DEFAULT <31:0> REGISTER;
```

defines the default width attribute to be a field 32 bits wide numbered from 0 to 31 and from right to left. It also defines the default storage class to be REGISTER.

There exist no pre-set defaults in SMITE, and therefore in the absence of a default, the following example

```
DECLARE
  ACCUM <7:0>,
  XTENSION REGISTER;
```

is in error because both of the data declarations are incomplete. However, this example

```
DECLARE DEFAULT <7:0> REGISTER,
  ACCUM,
  XTENSION <0:3>;
```

is acceptable. The data item ACCUM is declared to be type REGISTER, eight bits wide, and has bit numbering 0 to 7 from right to left. The data item XTENSION is declared to be type REGISTER, four bits wide, and has bit numbering 0 to 3 from left to right.

Only one DEFAULT declaration may exist in a SMITE description. That DEFAULT should be in the main processor, preferably as the first declaration.

3.5 Case Study 3: Intel 8080 Storage Declarations

The storage declaration section of the Intel 8080 computer description provides several illustrative examples of data item definition. The 8080 storage declaration is the following:

1. DECLARE DEFAULT<7:0> REGISTER,
2. PROD-FLAGS<35:0>,
3. STEP-FLAG FLAG DEFINED PROD-FLAGS <35>,
4. INTERRUPT-IR<7:0> DEFINED PROD-FLAGS<25:18>,

```

5.      INTE FLAG DEFINED PROD-FLAGS<17>,
6.      INTERRUPT FLAG DEFINED PROD-FLAGS<16>,
7.      PC<15:0> DEFINED PROD-FLAG <15:0>,
8.      REGS[0:3]<15:0>,
9.      REGISTERS[0:7]<7:0> DEFINED REGS,
10.     B DEFINED REGISTERS[0],
11.     C DEFINED REGISTERS[1],
12.     D DEFINED REGISTERS[2],
13.     E DEFINED REGISTERS[3],
14.     H DEFINED REGISTERS[4],
15.     L DEFINED REGISTERS[5],
16.     STATUS DEFINED REGISTERS[6],
17.     CARRY FLAG DEFINED STATUS<0>,
18.     PARITY FLAG DEFINED STATUS<2>,
19.     AUX-CARRY FLAG DEFINED STATUS<4>,
20.     ZERO FLAG DEFINED STATUS<6>,
21.     SIGN FLAG DEFINED STATUS<7>,
22.     A DEFINED REGISTERS[7],
23.     CARRY-A<8:0> DEFINED REGS[3]<8:0>,
24.     SP<15:0>,
25.     B-PAIR<15:0> DEFINED REGS[0],
26.     D-PAIR<15:0> DEFINED REGS[1],
27.     H-PAIR<15:0> DEFINED REGS[2],
28.     IR,
29.     IR-OPERAND-SELECT<2:0> DEFINED IR<2:0>,
30.     IR-DEST-SELECT<2:0> DEFINED IR<5:3>,
31.     IR-SUB-FUNCTION<2:0> DEFINED IR<2:0>,
32.     IR-TEST-SELECT<2:0> DEFINED IR<5:3>,
33.     OP-REG,
34.     HOLD-A,
35.     OP-PAIR<5:0>,
36.     TEST-FLAG,
37.     HOLD-REG,
38.     MEM[0:4095] MEMORY;
39.     DECLARE OP-STEP EXTERNAL,
40.             OP-HALT EXTERNAL,
41.             OP-ERROR EXTERNAL,
42.             OP-NOTSIM EXTERNAL
43.             IN-PORT PORT,
44.             OUT-PORT PORT;

```

In line 1, the default width and data attributes are declared. As a result of this statement, when no width is declared for a variable, the default width <7:0> is used. When no data attribute is used, the default attribute of REGISTER is used.

Lines 2 through 7 define the variable PROD-FLAGS and the substructures of that word. PROD-FLAGS is defined as a 36 bit

word, the sign bit being represented as bit 35 and the least significant bit as bit 0. The statement

```
DECLARE
  PROD-FLAGS<0:35>;
```

would also have defined a 36-bit word, but with a different bit representation (right to left in increasing order). Three variables overlayed on PROD-FLAGS are defined with the FLAG attribute, each data item thereby being specified as being one bit wide. Line 3 could be recoded equivalently as

```
DECLARE
  STEP-FLAG<0> DEFINED PROD-FLAGS<35>;
```

The result of these definitions is to form 5 subfields on the parent word. When any of these variables is read, the specified bits are extracted from the parent word for use in the expression. When the variables are stored into, only the specified bits in the parent word are altered. When the parent word is used in an expression, the entire word is used regardless of the substructure defined on the word.

Line 8 defines a register array of 4 words length and 16 bits width. The first element of the array is designated as REGS[0], and the last element as REGS[3].

Line 9 defines a substructure of the array REGS. REGISTERS is an array of 8 words length and 8 bits width. REGISTERS is a valid substructure of REGS since it does not cross word boundaries.

The registers could have been defined as separate independent variables. However, the instruction decoding algorithm often uses a 3 bit field to select the register. Grouping the registers in an 8 word array permits the description to access the proper emulated register by an access to the register file REGISTERS. Similarly, the registers are allocated within 16 bit register pairs because that structure allows the computer description to adhere to the architecture more closely.

Lines 10-23 define alternate names for the registers to improve the clarity of the computer description. For example, the register C is a renaming of REGISTER[1]. The C register can also be referred to as a subfield of REGS[0] or B-PAIR.

Lines 17-21 define the substructure of the STATUS register. The variable CARRY, for example, is equivalent to STATUS<0>, REGISTERS[6]<0>, and REGS[3]<8>. Referring to the status

bits by the individual name improves the readability of the description.

Lines 22 and 23 provide an example of an extended variable. REGS[3] is allocated in one parent word. REGISTERS [6] and REGISTERS [7] form a substructure of the parent word, and so the registers STATUS and A occupy contiguous subfields in the parent word. The bit immediately adjacent to the sign bit of A is CARRY. Arithmetic operations which reference the subfield A//CARRY may instead be written to reference CARRY-A, which holds both CARRY and A. In addition to offering a certain compactness of expression, this technique improves the speed of emulation when referencing the concatenated field.

Lines 25-27 provide an alternative representation of REGS[0], REGS[1], and REGS[2] for improved readability.

Lines 28-32 define IR, the instruction register, and its substructure. Notice that IR-OPERAND-SELECT and IR-SUB-FUNCTION define the same bit positions. The use of two names reflects the different usages of those bits during the instruction decode. For certain opcodes, those bits select the destination register; in other cases, they provide a secondary opcode.

The Intel 8080 computer description does not use the DATA attribute to define substructure. The definitions in lines 28-32 could be recoded as follows:

```
IR,
INSTRUCTION<7:0> DATA,
  OPERAND-SELECT<2:0> DATA DEFINED INSTRUCTION<2:0>,
  DEST-SELECT<2:0> DATA DEFINED INSTRUCTION<5:3>,
  SUB-FUNCTION<2:0> DATA DEFINED INSTRUCTION<2:0>,
  TEST-SELECT<2:0> DATA DEFINED INSTRUCTION<5:3>;
```

With this form of instruction format specification, for example, all uses of the variable IR-OPERAND-SELECT in the computer description would be replaced by IR.OPERAND-SELECT. The advantage to this technique is clearer, for example, if a data structure called SIGN-BIT is defined as INSTRUCTION<7>. Then the sign bit of an 8 bit variable, say register A, may be referenced by A.SIGN-BIT rather than the less descriptive A<7>.

Line 38 defines the primary memory array.

Lines 39-44 define the external processors and I/O port registers.

4. Timing Specification and Control

4.1 Introduction

For many applications, it is sufficient to describe the data transfer and control operations of a computer. There is an inherent value in the computer description, and an emulator compiled from such a description can be used to debug and execute a large class of software.

In other applications, it is necessary to consider timing relationships in the computer description. Each instruction requires a certain amount of time to execute on the target computer. Interrupts may occur based upon the real-time clock. Input operations require a given amount of time to complete, and status lines may change within certain intervals after the initiation or completion of input. Output operations may require a fixed minimum interval between successive commands.

When a complete data processing system is to be simulated, the need for an emulator clock is even more apparent. For many command and control applications, the computer is used to control specialized hardware. An accurate simulation of the hardware requires knowledge of the exact time that the computer issues various commands.

A timing-free description of the instruction operation of a computer specifies the architecture of a general family of computers. For a particular computer in the architecture family, there are many possible timing options, such as the instruction clock frequency, memory cycle time, I/O device access times, I/O transfer rates, and clock interrupt frequency.

One application of an emulator is the prediction of the effect on system performance of a particular timing change. For example, an emulation could be used to answer questions such as "What is the effect of doubling the instruction clock frequency for a computer when a particular program is run?" The execution time for a particular instruction may depend on both the instruction clock frequency, the memory cycle time, and I/O access delays and transfer rates. The new execution time of each instruction can be calculated. The degree of change in instruction execution times probably vary for different instructions, depending on the operations performed within each instruction. Any given program has a specific

instruction mix, and will not necessarily benefit proportionately from the increase in CPU speed. The use of an emulation provides the means to experiment to determine the effect of the proposed change. Similar tradeoffs may be performed for memory cycle time or various I/O options.

4.2 The IN Statement

The SMITE construct used to express timing relationships is the IN statement. IN defines the amount of time required to perform the next SMITE statement. The syntax of IN is the following:

IN expression statement

The statement so prefixed with an IN statement is specified to require 'expression' time units to execute. The time increment, of course, may be a constant or some variable expression. The units of time are left to the discretion of the programmer. The units may be unspecified ('clock counts'), or else the expression may be followed with one of the following noise words:

Unit	Abbreviation
SECONDS	S
MILLISECONDS	MS
MICROSECONDS	US
NANOSECONDS	NS

These time units are noise words. No form of units checking, conversion, or other significance is attached to the word by the compiler. It is the programmers responsibility to ensure that all IN statements in a description utilize the same units.

The current clock value, as updated by execution of timed statements, may be referenced through a 36-bit data item declared with the CLOCK attribute. The clock is initialized to zero when emulation execution is initiated.

The IN statement forms a context block surrounding the timed statement, and may be labeled:

label: IN expression statement

No escape may be made past the IN context block; an escape to the label on the IN statement will result in the clock being updated as specified by the expression.

The clock is updated at the completion of the statement. If one or more IN statements are nested within another IN statement, the inner statements define the subintervals in the interval defined by the outer statement. The expression should not evaluate to a negative result. The clock will remain unaltered should such an attempt be made.

Should nested IN statements advance the clock past the time specified for completion of an outer statement (through incorrect description of timing), then this later time will remain in the clock after execution of the outer statement completes. This procedure is employed to prevent the occurrence of negative time intervals.

4.3 The Amount of Timing Detail in a Computer Description

Several options exist as to the level of timing detail present in a SMITE computer description. For example, an IN statement may surround the invocation of an entire processor which executes the current instruction, in which case the timing expression would be an average execution time, or, at the other extreme, individual timings may be specified for every lowest level operation in the description.

In general, lower level timing specification requires more information about the timings of the computer due to the increased level of detail present. Very low level timing specifications tend to be constants, while as the level of timing specification increases more complicated expressions tend to occur in the attempt to accurately model the aggregate effect of conditionals or loops nested under the IN statement. There is no 'correct' level of detail for inserting timing into a computer description. The level of detail employed depends on the intended use of the description and the amount of detailed timing information available.

4.4 Case Study 4: A Shift Unit

For a very large number of computers, the execution time of a shift instruction depends upon the number of bits shifted. In this case study, we consider three cases of description of a shifter unit. In all cases, the shift performed will be left logical, shifting the data in a register A by a shift count in a register N.

1. The execution time is 4 units plus 1 for every bit shifted ($4 + N$).
2. The execution time is 12 units plus 3 for every bit shifted ($12 + 3N$).
3. The execution time is 4 units when $N=0$, 5 units for $N=1$ or 2, and 6 units when $N=3$ or 4 (two bit at a time shifts).

For the first case, there are two possible alternatives:

```
IN 4+N A <- SLL(A,N),
```

or

```
IN 4 SHIFT-COUNT <- N;
DO WHILE SHIFT-COUNT /= 0;
  IN 1 PARALLEL-BEGIN;
    A <- SLL(A,1)
    SHIFT-COUNT <- SHIFT-COUNT - 1;
  PARALLEL-END;
```

In the second case, the two methods become

```
IN 12+N+N+N A <- SLL(A,N);
```

and

```
IN 12 SHIFT-COUNT <- N;
DO WHILE SHIFT-COUNT /= 0;
  IN 3 PARALLEL-BEGIN;
    A <- SLL(A,1);
    SHIFT-COUNT <- SHIFT-COUNT - 1;
  PARALLEL-END;
```

In the third case, the two methods become

```
IN 4 + SRL(N+1,1) A <- SLL(A,N);
```

and

```
IN 4 SHIFT-COUNT <- N;  
DO WHILE SHIFT-COUNT > 1;  
  IN 1 PARALLEL-BEGIN;  
    A <- SLL(A,2);  
    SHIFT-COUNT <- SHIFT-COUNT - 2;  
  PARALLEL-END;  
IF SHIFT-COUNT /= 0  
  THEN IN 1 A <- SLL(A,1);  
END IF;
```

5. Parallelism

5.1 Introduction

Within a computer there are a number of operations which occur in parallel or overlap each other. Oftentimes by making simplifying assumptions about the operation or relative independence of processes, it is satisfactory to describe a computer in terms of serial steps. In the case of advanced processors or computer systems, it is more natural and efficient to describe the system in terms of parallel processes. SMITE provides the parallel context block to support the description of parallel processes.

5.2 The PARALLEL-BEGIN and PARALLEL-END Statements

The PARALLEL-BEGIN and PARALLEL-END statements define a context of parallel processes in the same way that BEGIN and END define a serial context. Each statement in a parallel context is executed simultaneously and asynchronously with all the other statements in the parallel context. The parallel processes start at the same time; the parallel context completes execution when the last process finishes. The implementation of the parallelism is discussed in Chapter 7.

The syntax of the PARALLEL-BEGIN and PARALLEL-END statements is similar to BEGIN and END. A label may be used with the parallel statements. The same label must be used for both PARALLEL-BEGIN and PARALLEL-END, if a label is used. The full set of executable SMITE statements is allowed within the parallel context.

5.3 Parallel Timing and Control Flow

The interval required to execute the statements in a parallel context is the maximum execution time for any one statement, not the sum of the processor execution times. Within a parallel context, an IN statement causes a temporary clock to be incremented. At the conclusion of the parallel context, the maximum temporary clock value updates the aggregate clock. This reflects the fact that the parallel processes which

complete first are suspended until the last process is finished.

An ESCAPE statement within a serial context causes subsequent statements to be bypassed. In a parallel context, the ESCAPE statement can only bypass statements within the parallel process in which it lies. The remaining parallel processes are not affected by execution of the ESCAPE. It is not possible to escape from a parallel context out of that context. This constraint coincides with the intuitive understanding that the parallel processes operate independently of each other and an ESCAPE should be a local condition.

5.4 Case Study 5: Instruction Pre-Fetch

A simple example of parallelism is the instruction prefetch logic of a computer. The emulation of each target computer instruction can be regarded as four consecutive steps: The instruction fetch, the instruction decode, the instruction operand fetch, and the instruction execution. In many computers, the instruction fetch overlaps the other steps to produce faster execution times. The technique of reading the next instruction while executing the current one is referred to as instruction prefetching.

One common prefetch convention, sometimes called instruction pipelining, is to load the instruction register with the next instruction, and to begin the fetch of the following instruction location at the end of the current instruction. Ignoring complicating factors, such as multi-word instructions or branch instructions, the computer may be described in terms of its basic steps as follows:

```

COMPUTER-WITH-PREFETCH:  PROCESSOR;
    INSTRUCTION-FETCH;
    DO FOREVER;
        PARALLEL-BEGIN;
            BEGIN;
                DECODE;
                OPERAND-FETCH;
                INSTRUCTION-EXECUTE;
            END;
        INSTRUCTION-FETCH;
    PARALLEL-END;
COMPUTER-WITH-PREFETCH:  END;

```

The first complicating factor is multi-word instructions. At the time of the instruction fetch, reading the next word in memory is begun on the assumption that this will be the next instruction. For certain computers, this memory word (current instruction location + 1) is part of the instruction or the instruction operand.

The computer description may then be recoded in the following manner:

```

COMPUTER-WITH-PREFETCH:  PROCESSOR;
  INSTRUCTION-FETCH;
  DO FOREVER;
  BEGIN;
    DECODE;      ''MAY SET NEED-NEXT-WORD''
    IF (NEED-NEXT-WORD)
      THEN PARALLEL-BEGIN;
        NEED-NEXT-WORD <- 0;
        DECODE2;    ''DECODE NEXT WORD''
        INSTRUCTION-FETCH;
        PARALLEL-END;
      END IF;
    IF (OPERAND-MEMORY-NEEDED)
      THEN PARALLEL-BEGIN;
        OPERAND-MEMORY-NEEDED <- 0;
        OPERAND-FETCH;
        INSTRUCTION-FETCH;
        PARALLEL-END;
      ELSE OPERAND-FETCH2;
      END IF;
    PARALLEL-BEGIN;
    INSTRUCTION-EXECUTE;
    INSTRUCTION-FETCH;
    PARALLEL-END;
  END;
COMPUTER-WITH-PREFETCH:  END;

```

The third case occurs when the current instruction is a branch instruction which forces the program out of sequence or when an interrupt occurs. The instruction available in the pipeline is then not the next instruction to execute. For the computer description to reflect this capability, we may recode the first lines of code as follows:

```

COMPUTER-WITH-PREFETCH:  PROCESSOR;
  INSTRUCTION-FETCH;
  DO FOREVER;
  BEGIN;
    IF BRANCH OR INTERRUPT

```

32584-6015-RU-00

```
      THEN INSTRUCTION-FETCH;  
      END IF;  
DECODE;  
.  
.  
.
```

AD-A087 743

TRW DEFENSE AND SPACE SYSTEMS GROUP REDONDO BEACH CA
ADVANCED SMITE REFERENCE MANUAL.(U)

F/G 9/2

FEB 80

F30602-78-C-0016


UNCLASSIFIED

TRW-32584-6015-RU-00

RADC-TR-80-66

NL

2 3
AD-A087 743



1.0 1.1 1.25 1.4 1.6 1.8 2.0 2.2 2.5 2.8 3.2 3.6 4.0 4.5

6. Input/Output and Operator Interface

6.1 Introduction

For many intended applications, the SMITE programmer is not concerned with input/output operations of the emulated computer. From the standpoint of a CPU description, the I/O channels or ports are very similar to memory. The CPU provides data to a port, and sometimes also provides an enabling signal to the device. A black-box operation transmits the data to the particular output device. Status registers or ports may be read by the CPU, but the changes in the status lines occur independently of the CPU. Similarly, from the CPU viewpoint, input data appears on a channel from an external source which itself determines the values and the time of arrival.

There are a number of significant advantages in partitioning the CPU description at the CPU-to-port interface. First, it is a realistic description of the CPU operation. The CPU relinquishes control to the I/O devices through these ports. Second, the computer description is then independent of the I/O devices present in the computer system. Finally, interface software may be required between the SMITE code and an I/O device driver, and so it is not unreasonable for the interface software to simulate the operation of the external device.

For an emulation of the CPU and some I/O devices, software linking the SMITE microcode to device drivers or the operating system is required. This linkage is also important for operator interfaces and the invocation of microcoded routines. This section will discuss the interface of the SMITE object code to other software executing on the QM-1.

6.2 EASY Interface to SMITE

EASY and TASK are the two basic elements of the QM-1 operating system. TASK provides the overall system control, processes interrupts, schedules tasks, handles inter-task linkages, and communicates with I/O devices. EASY is the primary task running under TASK. It is written in a higher order language called SIMPLQ [14]. The intermediate language generated by its compiler is executed by the EASY interpreter [16]. EASY

allows operator communication with an emulator, provides display and control commands, and provides the emulator's interface with the outside world.

It is important to understand the method of communication between TASK, EASY, and the emulator. Figure 6 shows the relative organization of main store and control store in the system. Each of the three tasks, TASK, EASY, and the SMITE emulation, can access its own main memory and the main memory of lower priority tasks. Control store is allocated to the tasks in a predetermined manner. Each task executes with its own set of virtual QM-1 registers. Higher priority tasks control their descendant tasks. Lower priority tasks communicate with higher priority task through SYSTEM instructions.

Of particular importance to the SMITE emulator is the SYSTEM Recall instruction, which transfers control back to EASY for further action. The SMITE program requests an external process to occur by a particular SYSTEM Recall instruction. This initiates a TASK interrupt which transfers control to the interrupt handler of EASY. This routine is an EASY processor, using the EASY register set and main store addressing. This method of communication frees the SMITE compiler from providing explicit linkages with EASY and from protecting the emulator registers from inadvertent changes in EASY. The SYSTEM Recall instructions generated by the SMITE compiler provide the linkage to EASY and TASK.

6.3 External Functions

A processor declared as an external processor in the SMITE computer description is implemented as an EASY processor. A SYSTEM recall instruction is generated by the SMITE compiler to link to the processor. The first six external processors are implemented as SYSTEM recalls 0 through 5. For remaining external functions, the emulator register R.ADR is set to the ordinal assigned by the compiler to the particular external function. The emulator registers are accessed by EASY indirectly through the emulator register save area. Emulator main store is accessed through special SIMPL-Q system primitives.

MAIN STORE	0-13777	TASK Overlays
	14000-37777	EASY Region
	14000-BPS	Emulator Main Store
	BPS-37777	EASY Program and Data
CONTROL STORE	0-23677	Emulator Program
	23700-23777	Emulator Register Save
	24000-31777	EASY Machine
	32000-35777	TASK
NANOSTORE	0-240	MULTI, QM36 Instructions
	241-254	SMITE Instructions
	277	System Instruction

Figure 6 QM-1 Memory Allocation

6.4 Ports

A port is implemented as a particular type of external function. The linkage to EASY is identical to that for external functions, except that the intermediate I/O value is passed in the emulator local store. Ports of from 1 to 18 bits in width transmit data right justified in local store register R.0. Ports of from 19 to 36 bits transmit data right justified in local store registers R.0 and R.1. Ports of from 37 to 72 bits transmit data right justified in local store registers R.0, R.1, R.2, and R.3.

6.5 Lights and Switches

A light is implemented as a 1 bit output port. A switch is a 1 bit input port. In either case, the least significant bit of R.0 contains the data.

7. SASS: SMITE Application Support Software

7.1 Introduction

Once the SMITE computer description compiles successfully, it is ready to test on the QM-1. The process of transporting the object code of the compiler and executing on the QM-1 is relatively straightforward. A special system, the SMITE Application Support Software (SASS), is used to load and test emulators developed with SMITE. SASS is an augmented version of the EASY operating system which provides:

1. Additional commands to load and test SMITE emulators,
2. Interface for EASY control and display commands,
3. A predefined set of external functions, and
4. Emulator performance measurement.

The SMITE programmer's interface with SASS does not begin the first day he tests the emulator on the QM-1. By considering the SASS capabilities during the computer description, the programmer can significantly improve his test efficiency. This section will focus on six particular aspects of SASS:

1. The requirements imposed on the SMITE computer description by the SASS implementation,
2. The preparation of magnetic tapes to load the emulator and SMITE main memory,
3. The steps used to load, run, and debug SMITE emulators,
4. The method for obtaining performance measurement data from an emulation,
5. The implementation of parallelism, and
6. The procedures necessary to modify SASS for particular emulators.

7.2 SMITE Computer Description Development Considerations

When writing the computer description, the SMITE programmer should be aware of conventions established by the implementation of SASS. SASS requires knowledge of the location of interface variables in the SMITE data base, and the order of the external function declarations in the computer description. These items are predefined in SASS, thus establishing interface conventions for the SMITE programmer to follow.

7.2.1 Step Flag

The first SMITE data item should be declared as a 36 bit wide data item, such as

```
DECLARE
  EASY-FLAG<35:0>REGISTER;
```

Bit 35 of this variable is the emulator step flag. The STEP command causes the SMITE step flag to be set. For the step to then occur, the computer description must call the SASS OP-STEP external function (which, in turn, resets the step flag). For example:

```
DECLARE
  EASY-FLAG<35:0>REGISTER,
  STEP-FLAG FLAG DEFINED EASY-FLAG<35>,
  OP-STEP EXTERNAL;
DO FOREVER,
  BEGIN;
    IF STEP-FLAG
      THEN OP-STEP;
    END IF,
    .
    . 'PROCESS EMULATED INSTRUCTION'
    .
  END;
```

The emulator will cycle through emulated instructions. Each execution of the emulator STEP command will cause the execution of one target machine instruction. After the emulator has been checked out, this emulator STEP implementation may be used to debug software executing on the emulator.

Although the name STEP and the description in EASY imply that the emulator cycles through the emulation of one

target computer instruction, SMITE STEPs may occur several times for an emulated instruction, or may be spaced several emulated target machine instructions apart. For example:

```

DECLARE
  EASY-FLAG<35:0>REGISTER,
  STEP-FLAG FLAG DEFINED EASY-FLAG<35>,
  OP-STEP EXTERNAL;
DO FOREVER,
  BEGIN;
  .
  .
  .
  IF STEP-FLAG
    THEN OP-STEP;
  END IF;
  .
  .
  .
  IF STEP-FLAG
    THEN OP-STEP;
  END IF;

```

This code may be used to step the emulation through one or more SMITE statements. The primary advantage of this technique is that the emulator may be debugged at the SMITE description level with little examination of the MULTI code produced. The drawback is the large amount of code which must be inserted and then removed from the computer description. A technique for minimizing the amount of code shuffling required is as follows:

```

Declare DBG constant 1;

```

```

:
IF DBG
  THEN IF STEP-FLAG
    THEN OP-STEP
  ENDIF,
ENDIF;
:

```

Setting the declaration of DBG to zero suppresses

compilation of the code, setting it to 1 causes it to be compiled.

7.2.2 Emulated Program Counter

The emulated program counter must be declared as the rightmost bits of the first SMITE variable. For example, to declare a 16-bit program counter,

```
DECLARE EASY-FLAGS<35:0> REGISTER,
      PC<15:0> REGISTER DEFINED EASY-FLAGS<15:0>;
```

The program counter is used for the SMITE state display, the BREAK command, and the PC command. The SMITE programmer may change the location of the emulated program counter if the corresponding SASS definition is modified.

7.2.3 Emulated Memory

Most SMITE computer descriptions will contain a data item representing the primary memory of the computer. For convenient inspection and modification of the primary memory, EASY includes the MEMDISPLAY, MEMPRINT, MEMVAL (memory inspection) and CHANGEMS (memory modification) commands. An emulation which has been fully integrated into SASS, includes support routines to properly reference primary memory.

The SMITE programmer may activate either a specific SASS system (e.g., SMITE, EMULATION=642) or the general purpose system (SMITE, EMULATION=SMITE). The general purpose version makes no assumptions about the location of emulated memory in the SMITE data base; the memory display and memory modify commands accept octal addresses and data, and refer to QM-1 addresses in the SMITE main memory data base. The memory reference commands in SASS assume memory is composed of 18 bit words beginning at main SMIE store address 0. The user must translate from emulated addresses to QM-1 addresses.

When a specific version of SASS is activated, however, the memory display and modification commands may be programmed to refer to the emulated memory. Addresses and data are input in a predetermined format, and the display only outputs the machine-specific data image.

7.2.4 Order and Operation of Predefined SASS External Functions

A number of useful external functions are defined for SMITE computer descriptions by SASS: The linkage between SMITE and SASS is maintained by ordinal. The first external function or port declared in the SMITE description is assigned the ordinal 0 by the compiler, the next one 1, and so forth. This ordinal is passed to SASS by the SMITE emulator when the function is to be called. The predefined functions and their ordinals are as follows. The specific names shown are irrelevant and are only used for mnemonic purposes.

OP-STEP

Ordinal 0. OP-STEP clears the emulator step flag, places the emulator in step mode, and returns control to EASY. Control is returned to the emulator as the completion of the external call after any of a number of EASY commands, including STEP and GO.

OP-HALT

Ordinal 1. OP-HALT places the emulator in halted mode, as after a target computer halt instruction, and returns control to EASY. Control is returned to the emulator as the completion of the external call after any of a number of EASY commands, including STEP and GO.

OP-ERROR

Ordinal 2. OP-ERROR displays an 'emulator error' message, and returns control to EASY. Control is returned to the emulator as the completion of the external call after any of a number of EASY commands, including STEP and GO.

OP-NOTSIM

Ordinal 3. OP-NOTSIM displays an 'illegal instruction' message, and returns control to EASY. Control is returned to the emulator as the completion of the external call after any of a number of EASY commands, including STEP and GO.

7.3 Preparation of SMITE Load Tapes

The object code generated by the SMITE compiler, as well as any software which runs on the emulator, is transmitted to the QM-1 through a serial line or on magnetic tape. The procedures used to generate these magnetic tapes are determined by the host QM-1 installation. This section specifies the format of these transfer files.

7.3.1 SMITE Emulator Load Tape

The SMITE load file consists of a series of variable length records, containing header information, the object program, and symbolic records. The file consists of a stream of eight bit bytes, organized into logical records. Each logical record contains a five byte prolog, specifying the type of information, and the record length. An object code record contains an additional six bytes to designate the start and step addresses.

The header record has information about the emulation size and the number of emulation symbols, labels, and statements. The object record contains the control store image generated by the compiler. The statement record lists the control store addresses corresponding to SMITE statements. The label and symbol records contain name and attributes of each SMITE label and data item. The emulator load file is created as the compiler output file LGO.

The MULTICS program LGOTRAN translates the compiler LGO file for serial data transmission. Each transmission character contains 6 bits of data, biased to form it into a legal character. The downlink process must transmit this file into a NOVA file (two characters per 18 bits). This procedure may vary at different installation, depending upon the interface equipment. The SASS utility LOADTAPE reconstructs the original compiler file and packs the data file.

7.3.2 Target Software Load Tape

At some point, the emulator is ready to execute target computer software. The user must process the output of the target computer assembler, compiler or loader and create a load tape. The command LOADTAPE is used to copy the tape onto disk. The command LOADMS is used to load the file into QM-1 main store for execution by the emulator.

The format of this tape is an absolute QM-1 core image of the SMITE emulator main store data base, including any other declared SMITE variables allocated to mainstore. The full 18 bits of data must be loaded for each QM-1 word, regardless of the SMITE definition.

The tape generator program is responsible for mapping data onto main store images. For instance, if each emulated word is 8 bits wide, and is stored in the right most 8 bits of a QM-1 36-bit memory word pair, such as is the case for the Intel 8080 emulation, the tape generator program must supply 28 leading zero bits for each 8-bit data byte so that the tape image is 36 bits of data with emulator data in the lower 8 bits.

The tape is arranged in 512 word physical records (512 x 18 = 9216 bits) and written in odd parity. A 9-track tape physical record, therefore, contains 1152 frames of data.

7.4 SASS Commands

7.4.1 EASY Commands

SASS is a subsystem of EASY which provides the capabilities to load and execute emulators developed with the SMITE compiler. EASY provides a wide range of editing, compilation, file handling, and debug features. Table 7.4.1.1 lists the primary EASY commands. Table 7.4.1.2 contains the applicable CONTROLQ debug commands.

Initially the SMITE programmer requires only a small subset of the available EASY commands. The system is initialized by following the standard disk boot procedure and entering LD2/EASY from the keyboard. The following set of commands invokes the SASS system after a bootstrap of the EASY system:

DATE

TIME

0

DEADSTART,BS=N, where N is larger than the total static + dynamic main store required

Table 7.4.1.1 Primary EASY Commands

<u>Command</u>	<u>Function</u>
BIND	Link Editor
CEXPOR	UT-200 Source Transfer
CIMPORT	UT-200 Source Transfer
COPY	Disk copy
DEADSTART	System initialization
DEBUG	Traceback control
DIRECTORY	Disk access control
DISK-SAVE	Disk backup
EDIT	Source editor
EDITLIB	Library update
EXEC	Input file of commands
IMPORT	UT-200 Binary Transfer
LIBRARY	Library Specification
LIST	Source File List
LISTCF	Command File List
LOCK	Reset privileged command access
MAKECF	Create command file
MT-TO-DISK	Binary tape to disk
PRUDMP	File dump
RESTORE-DISK	Disk file retrieval
REWIND	Tape rewind
SIMPLQ	Compile
SMITE	Activate SASS
SOURCE-TO-DISK	Tape copy
TEST	Execute module
UNLOCK	Set privileged access
<control Q>	Activate debugger

Table 7.4.1.2 Debug Commands

<u>Command</u>	<u>Function</u>
ARRAY	Display SIMPLQ ARRAY
BASE	Set main store base
CHANGE	Symbolic SIMPLQ change
CHANGECS	Patch control store
CHANGECSM	Patch control store multiple
CHAGEMS	Patch main store
CHAGEMSM	Patch main store multiple
CODE	EASY code display
CONVERT	Number conversion
DISPLAY	EASY symbolic display
DISPLAYCS	Control store display
DISPLAYMS	Main store display
EXEC	Command file execution
GO	Resume execution
GOTO	EASY branch
HALT	EASY breakpoint
INTEGER	Integer display
MAPPROGSPACE	Resident code map
MAPSTACK	Stack map
MBPPRO	Proceed from microbreakpoint
MBP	Set/clear microbreakpoint
PRINT	List file
PRINTCS	Dump control store
PRINTMS	Dump main store
REGISTERS	Display QM-1 registers
REGISTERSET	Change QM-1 registers
SCHEMA	EASY symbolic file
STEP	EASY step
STRING	Display EASY string
TRACEBACK	Program traceback
UNLOCK	Privilege access
*	Short program map
+	Next memory page
-	Previous memory page
.	Short stack map
<LF>	Print CRT display
<at>	Exit debug

DIRECTORY,01

SMITE, EMULATION=

From this point, SASS is used to load, execute, and test the SMITE emulation. The following paragraphs describe the SASS inputs.

The EASY system reference manual [15] contains a complete description of the EASY system inputs which may be used to create input data files, obtain additional debug information, or modify SASS. Input commands to the EASY system are controlled by a procedure named MASTER. MASTER provides user input cues, name expansion, and first-level format checking. The following special control characters are recognized by MASTER.

<blank>	use default value for next parameter
<CR> parameters	use default values for remaining parameters
<backspace>	delete last character
<control U>	delete command line
<control Q>	enter debug
<control Z>	Deadstart

7.4.2 SMITE

The command SMITE activates the SASS subsystem. The SMITE Application Support System contains a standard set of capabilities independent of the emulator used. Capabilities may be added for a particular emulator and the implementation of standard SASS features may be tailored to a specific emulator. An input parameter string X is present to select the particular SASS version. The nominal parameter (EMULATION=SMITE) selects SASS commands from the file CMDSMITE and SASS code from the library file SMITE. A parameter such as EMULATION=642 selects SASS commands from the file CMD642 and SASS code from the library files 642 and then SMITE. The remaining parameters specify the directory files, which the emulation system may access. The nominal case does not alter the user spaces which have been established by a previous DIRECTORY command.

7.4.3 Load Commands

7.4.3.1 LOADTAPE

The command LOADTAPE copies a magnetic tape or a downlinked file onto a QM-1 disk file which can be processed by the LOADCS or LOADMS command. The first parameter specifies the magnetic tape unit or the name of the downlinked file. The second parameter is the file number (0-N) on the magnetic tape. This parameter is ignored for a downlinked file. The third parameter specifies the type of copy (DATA or EMULATOR). The DATA copy packs the input bitstream into 18 bit QM-1 words prior to the LOADMS command. The EMULATOR copy collects similar logical records into a word addressable disk file, reformats data, and creates additional system level information.

7.4.3.2 LOADCS

The command LOADCS/Filename loads the SMITE emulator from the specified file into control store and initializes the emulator. Error checks are performed for the control store or main store requirements of the emulator exceeding the emulator allocation. The LOADCS command initializes the emulator task, the concurrent processing module, and the performance measurement subsystem. This command must precede the LOADMS, GO, and performance measurement commands.

7.4.3.3 LOADMS

The command LOADMS/Filename/Offset/Header loads SMITE main store from the specified disk file, presumably built as a result of the LOADMT command.

The data records of the file form a direct image of SMITE main store. If a header record is present on the file, the first 3 QM-1 words of the file contain the load address, length, and initial execution address. The remaining 253 words on the header record are not used.

SMITE main store is loaded from the specified disk file, starting from the input OFFSET address plus the load address on the header record. If the offset address is 0, the nominal case, all SMITE main store is loaded. The offline program to create the load tape must supply initial values for SMITE variables not in emulated

memory. To load only emulated memory, the offset address should be set to the location of emulated memory within SMITE main store, which is available from the SMITE allocation map.

The initial value for the emulated program counter is set to the initial execution address specified on the header record or to 0 if no header record is present. The program counter is assumed to be the low order 16 bits of the first 36 bit variable in SMITE main store.

7.4.4 Memory Display and Modify Commands

7.4.4.1 CHANGEMS

The command CHAGEMS allows 1 to 4 values to be changed in emulated memory. The first parameter specifies the address and the remaining parameters specify the new memory values. The format of the inputs and the mapping to emulated memory are a function of the particular emulation system. In the nominal case, the address refers to the SMITE main memory address and the values are input as 6 digit octal numbers. For specialized emulations, the input emulated memory address is mapped to the QM-1 memory location and an appropriate input format is specified.

7.4.4.2 CHANGECS

The command CHANGECS permits the contents of control store to be altered. This allows patches to be made to the emulator programs or SMITE variables allocated to control store to be changed. CHANGECS is a privileged debug command. The following procedure is used to change control store:

<ControlQ>

UNLOCK

CHANGECS, address, value

other debug commands

@

7.4.4.3 DISPLAYCS

The Debug command DISPLAYCS outputs a 64 word block of

control store on the CRT. The Debug commands "+" and "-" may be used to alter the base address of the display block. The following procedure is used to leave the SASS subsystem temporarily to display control store and then return for further SASS commands:

<control Q>

DISPLAYCS, address

other Debug commands

@

7.4.4.4 MEMDISPLAY

The command MEMDISPLAY causes SMITE main memory, starting at the specified input address, to be displayed on the CRT. In the nominal case, a block of 64 QM-1 words is output in O6 format. For specialized emulations, the input address is mapped to its QM-1 address and converted to a predetermined format.

7.4.4.5 MEMPRINT

The command MEMPRINT causes the SMITE main memory between two specified addresses to be printed. If the second parameter is not specified, a 64 word block is printed.

7.4.4.6 MEMVAL

The command MEMVAL displays one line beginning at the specified address on the CRT. This command does not erase a previous MEMDISPLAY output; so the MEMVAL command can be used to display specific memory cells referenced by the primary MEMDISPLAY block.

7.4.4.7 Space

The <space> key causes the emulator state display to be output. The generalized emulation system prints a message for the state display. The specialized emulation systems display labels and values for the registers of the emulated machine, the emulation control variables, and the simulated clock. The state display command interrupts the emulation execution, outputs the display, and causes the emulation to resume execution.

7.4.5 Execution Control Commands

7.4.5.1 BREAK

The command BREAK/address specifies a breakpoint in the target program. The target program is stepped until the emulated program counter equals the BREAK address. The parameter specifies the address and the conversion mode (octal=8, hexadecimal=16). The nominal mode is octal. The processing requires that the SASS conventions for the location of the emulated program counter and for the program step code must be followed.

7.4.5.2 GO

The command GO initiates or restarts emulator execution. The emulator continues execution until it relinquishes control to SASS for external processing or EASY interrupts for a higher priority task (e.g., processing another SASS command.)

7.4.5.3 HALT

The HALT command stops the execution of the emulator and sets the task inactive. Unlike the @ command, the SMITE subsystem commands are still available. The LOADCS command must be input to restart the emulator.

7.4.5.4 MBP

The CONTROLQ command MBP is used to set microbreakpoints within the SMITE emulation. When the microbreakpoint is executed, the CONTRLQ debug system is activated. The user has access to the debug facility until returning with the @ command.

7.4.5.5 MBPPRO

The CONTROLQ command MBPPRO restarts the emulator execution from a microbreakpoint. This command restores the original instruction at the breakpoint and then performs the same action as the GO command.

7.4.5.6 NEXT

The command NEXT inserts a microbreakpoint and restarts the emulation. An input parameter specifies the condition which determines the microbreakpoint address. The nominal mode causes the microbreakpoint to be

inserted at the next address to be executed. The JUMP mode inserts the microbreakpoint at the next conditional transfer microinstruction to be executed. The NOJUMP mode inserts a microbreakpoint at the next instruction executed, except when the instruction is reached through a BALN (call). In this case, the microbreakpoint is placed at the return address from the processor call.

7.4.5.7 TIMESLICE

The TIMESLICE input specifies the time slice allocate to concurrent tasks. The parameter is expressed in the number of milliseconds. The minimum value is the basic clock interval time, which is 20 milliseconds at most installations.

7.4.5.8 <Carriage Return>

The carriage return character causes a SMITE program step to occur. The SMITE step flag is set and processing identical to the GO command is performed. If the SASS step conventions have been followed and the emulator is not halted at a breakpoint, one target computer instruction is executed. If the emulator is breakpointed, an additional MBPPRO command is required.

7.4.6 Performance Measurement Control

7.4.6.1 PM

The command PM selects the performance measurement subsystem and specifies the input device along with the echo device. The input device may be *KBD, *CDR, or the name of the disk file containing the performance measurement inputs. The echo device may select CRT, LPT, or BOTH. The nominal condition accepts input from the keyboard and echoes the input and error messages to both the CRT and line printer. If the card reader or disk file is the input source and an end-of-file condition occurs, the remaining input is accepted from the keyboard. Only one set of performance measurement inputs should be used during a given invocation of the SASS subsystem. Section 7.5 describes the capabilities and syntax of the performance measurement subsystem.

7.4.6.2 RECORD

The command RECORD outputs the results of the performance measurement data collection. The first

parameter specifies the type of data to output (ALL, COUNT, SAMPLE, or TIMING). The second parameter designates the output device (CRT, LPT, or BOTH). The default case outputs all performance measurement data to both the CRT and line printer.

7.4.7 Miscellaneous Control

7.4.7.1 <ESCAPE>

The ESCAPE character transfers control of the QM-1 keyboard/CRT from EASY to the emulator. All subsequent characters are stored by SASS for use as input to the target software. User supplied external functions are used to transmit the data and status of the keyboard to the emulator in response to CPU instructions. The next ESCAPE character returns control of the keyboard and CRT to EASY for normal SASS input processing.

7.4.7.2 <at sign>

The at sign command terminates the emulation subsystem and returns control to the primary EASY command system. The SMITE command and LOADCS command are necessary to restart the emulator.

7.4.7.3 <LINEFEED>

The LINEFEED character causes the CRT screen to be output on the line printer. Due to the presence of non-printing control characters in the display output, the printout will not exactly match the screen.

7.4.7.4 >CONTROL-Q>

The CONTROL-Q character initiates the DEBUG subsystem. The available commands are listed in Table 7.4.1.2. Upon issuance of the @ Debug command, control returns to the SMITE subsystem for further SASS commands.

7.5 Performance Measurement

The performance measurement subsystem is activated by the PM command. A series of free field inputs from the keyboard, card reader, or disk file specify the performance measurement probes. The basic format of the performance measurement is a

COUNT, SAMPLE, or TIMING input, along with optional IF clause conditions or qualifiers. Comments, delimited by /* and */ or by /* and the end of line, may be present anywhere in the line. The backspace (character erase) and control-U (line erase) characters provide line editing capability from the keyboard. The @ input terminates the subsystem and returns control to the SASS subsystem.

The performance measurement data is collected at control points in the emulation. A control point occurs when a specified SMITE statement is executed and any qualifying condition is true. At a control point, the following data may be collected:

- a. The count of the number of times the control point was reached (COUNT)
- b. The accumulated simulated clock time between the occurrence of one control point and the occurrence of a second control point (TIMING). The first condition starts the timer and the second time halts it. No units are presumed for the clock.
- c. A particular data item may be sampled at a control point (SAMPLE). Each time the control point is reached, an array entry corresponding to data item value is incremented.

7.5.1 System Performance Measurement Application

The performance measurement system is designed to measure the emulator under specific program conditions or indirectly to measure the performance of the target software. Some data, such as instruction mixes, is obtained readily without much forethought on the part of the SMITE programmer. Other data, such as memory usage, is quite cumbersome to collect unless the program was designed in a modular fashion and the data allocated with the performances measurement goals in mind. Some complex measurement data, such as instruction sequence occurrences, can be obtained only if the description saves extraneous information needed only for performance measurement. The following paragraphs listed some system measurements which can be gathered (although some are non-trivial to specify).

The COUNT measurement can measure the number of page faults which occur on an entire run, within different memory ranges, or in different time intervals. It can measure the number of times any resource (memory, I/O channel,

interrupt line) is unavailable and causes a delay. COUNT can measure the number of times certain unfavorable conditions occur (e.g. addition of floating point numbers with exponent differences exceeding the mantissa width) which may justify an alternative implementation in the target machine. COUNT can also be used to detect sequences of instructions which may warrant new instructions on the target machine. It may also be used to detect conditions within the target software, such as the number of times a routine was called.

The SAMPLE input may measure the number of times each instruction is executed or the relative use of each addressing. It can produce a table of shift count values for shift instructions or floating point operations to determine if enough long shifts are present to warrant a faster shift technique on the target machine. SAMPLE can be used to derive the use of system resources such as memory pages or I/O channels. Or it may be used to produce a histogram of program counter usage according to memory region for the target software.

The TIMING input may be used to measure the timing delay due to unavailable resources, memory, or I/O busy states. The input may also be used to measure the amount of time spent in various program states (supervisory, interrupt, or various subroutines).

The performance measurement system measures the emulator performance while executing specific target software. It is especially advantageous when the computer description can be modified to measure the effect of additional hardware or faster devices. A somewhat esoteric use of the performance measurement capability is to analyze the emulator performance to identify bottlenecks or frequently used areas which can be recoded to improve the performance of the emulator itself.

7.5.2 Performance Measurement Syntax

The syntax definition of the performance measurement input is the following (where { } encloses an optional field and / signifies an alternative choice):

```
PM input = function-expression/@
```

```
function-expression = {IF test-expression  
{THEN}}function
```

function = COUNT control-probe/
 TIMING control-probe {TO} control-probe/
 SAMPLE data-item {AT} control-probe

 data-item =
 SMITE-symbol{.SMITE-processor}[[index]]{<width>}
 index = integer constant/octal constant/hexadecimal

 width = integer constant/octal constant/hexadecimal

 integer constant = positive integer

 octal constant = 0' octal number

 hexadecimal constant = X' hexadecimal number

 control probe = control point {AND/AND NOT status-word}

 control point =
 {\$}SMITE-label{.SMITE-processor}{+integer
 constant/-integer constant}

 status word = data item defined as 1 bit wide/data item
 restricted to 1 bit width

 test-expresssion = test-relation{AND/OR test-relation}

 test-relation = data-item op data-item/data-item op
 constant/TIME op constant

 op >or>=or<or<=or=or/=

The performance measurement inputs are subject to the following restrictions:

- 16 COUNT inputs
- 16 TIMING inputs
- 8 SAMPLE inputs
- 256 SAMPLE data collection bins
- 100 unique data item reference
- 11 or more unique IF clauses (depending on type and complexity of the expression)

The performance measurement input may be entered in free field format with minor restrictions if the input requires multiple lines. Data items and control probes must be contained on the same line. Likewise, a logical operator must be contained on the same line. All logical operations are restricted to 36 bit data comparison.

7.5.3 Performance Measurement Syntax Discussion

The performance measurement syntax is relatively simple. The three important elements to understand are what is meant by a data item, what is meant by a control probe, and what conditional capability exists.

A data item is equivalent to a reference to a SMITE variable with four minor exceptions:

- 1) An index specification must be a constant
- 2) An octal or hexadecimal constant does not have a terminating single quote
- 3) The data definition qualification is not permitted. To restrict the width of a data item, an explicit width restriction (e.g., A<>>) must be used instead of a qualification (e.g., A.SIGN)
- 4) When a SMITE variable is defined in more than one processor, the name of the defining processor is used to qualify the variable name (e.g., TEMP.MPY).

The following are valid data item specifications for the INTEL-8080 description:

REGISTER[5]<3:0>

L<0>

IR

VALUE.STORE<1:0>

MEM[X'F3]

Examples of invalid data item specifications are the following:

MEM[SP]

MEM[X'F']

A.SIGN-BIT

A[3]

When the data item has no width definition, then the width declaration found in the SMITE description is used. For the processor name qualification, the name of the processor containing the definition must be used. Only one level of processor qualification is permitted and shared variables may not be qualified according to the task invocation.

A control probe is a reference to a SMITE statement, along with an optional status test. The SMITE statement is indicated by referencing a SMITE program label with an optional constant offset. A negative offset should be preceded by a blank if a possible ambiguity exists (e.g., a label ADD-1). The compiler output listing should be examined to determine the statement number, since it may differ from the user's intuitive notion. When the SMITE label is preceded by "\$", the SMITE statement at the end of the context block defined for the label is used. The SMITE label may be qualified by the name of the processor in which it is defined. Examples of valid control point references for the INTEL-8080 are:

ADD.PEFORM-OP+1,

\$ADD-1

OPERAND

OPERAND+2

The status test examines the value of a one bit data item (FLAG) or a data item restricted to one bit. It is not possible to use a relational test as a control point qualifier. That is, ADD and A=0 is not a valid control probe reference for the INTEL-8080 description, but both ADD AND A<>> and ADD AND CARRY are legal.

Each performance measurement input may be preceded by a conditional IF clause. The clause consists of a binary comparison or the logical result of two binary comparisons. The terms in the comparison must be data items, constants, or TIME. Examples of valid conditional clauses are:

IF PC<X'44C

IF A=0 OR B/=0

IF A=B AND TIME>1000

Expressions may not be used in the comparisons. Only AND or OR may separate the binary relations. No unary conditions are allowed. Data comparisons must be less than 36 bits wide. No parenthesis are allowed in the conditional clause. The following violate the syntax rules:

IF A-1 =0

IF A=0 AND NOT B=0

IF CARRY

IF (A=1) OR B=0

For the performance measurement data to be gathered, the specified SMITE location must be reached, the optional control point qualifier true, and the optional IF condition true. However, for the TIMING input, the IF condition pertains only to the timing start control probe and not to the second control probe.

The syntax permits noise words THEN, TO, and AT to improve readability of the input.

7.5 4 Performance Measurement Error Conditions

The performance measurement system tests the validity of the inputs. Appendix J lists the performance measurement error conditions, their causes, and the recommended corrective action.

7.5.5 Performance Measurement Input Procedure

Successful application of the performance measurement capability starts with the design and coding of the emulation. Three basic guidelines should be followed during the emulation development:

1. All SMITE variables which may be used as sampled items, conditional tests, or status qualifiers should be allocated to unique permanent storage locations. Particular attention should be given to procedure arguments, syntax macro data items, re-entrant

2. Control points and data items should be given unique labels to facilitate references in the performance measurement input.

3. Organization of the description according to hardware functional units, such as the memory unit, permits the performance to be measured with centrally-located probes.

The efficiency of equivalent performance measurement inputs can differ drastically according to the location of the probe. Each execution of the control point location causes control to be transferred to SASS. Then the control probe qualifier and IF clause condition are tested. It is far more efficient to qualify a control probe by moving its location, if possible, in the description.

The performance measurement capability is designed to measure the emulated computer; but it may be extended to produce information about the application software executing on the emulator. For example, a target software data item may be referred to as MEM[n] and a target program location may be tested by the statement IF PC = m.

As an example, assume that the performance measurement capability is used to obtain (1) the total amount of time interrupts are disabled, (2) the number of reference to each memory bank; and (3) the number of unsuccessful I/O status tests.

Also, assume that the following is a skeleton description of the emulation:

```

~ EMULATOR:  PROCESSOR;

  DECLARE    EASY-REG<35:0>REGISTER;

             PC<12:0>DEFINED EASY-REG<12:0>,

             IO-STATUS FLAG,

             IR<15:0>

             MEM[0:8192]<15:0>MEMORY;

MEMREF:      PROCESSOR<15:0>(IADD1)

  DECLARE    IADD1<12:0>REGISTER,

```

32584-6015-RU-00

```

                                IADD<12:0>REGISTER;
                                IADD<-IADD1,
                                MEMREF<-MEM[IADD];
MEMREF:      END;
DECODE IR<15:12>;
    ENABLE-INT:      BEGIN;
                                .
                                .
                                .
    ENABLE-INT:      END;
    DISABLE-INT:     BEGIN;
                                .
                                .
                                .
    DISABLE-INT:     END;
    IO-STATUS-TEST:  BEGIN
                                IF IO-STATUS
                                THEN PC<- PC+2;
                                ELSE BEGIN;
                                        PC<-PC+1;
                                        END;
    IO-STATUS-TEST:  END
END DECODE;
EMULATOR:      END;
```

The total time that the interrupts are disabled is expressed as

TIMING ENABLE-INT TO DISABLE-INT.

Each time that the enable interrupt instruction is executed, SASS records the time; when the disable interrupt instruction is executed, the time differential is added to the accumulated time. If the interrupt system is more complicated, with an enable signal for each interrupt, the total time a particular interrupt is disabled may be obtained by using appropriate qualifiers. For instance, if bit 2 in the instruction register IR contains the particular interrupt is affected by the instruction, the input could be changed to

TIMING ENABLE-INTERRUPT AND IR <2> TO DISABLE-INTERRUPT AND IR<2>.

The number of references to each memory bank may be obtained by sampling the memory address at the memory reference processor. Assuming that the upper four bits of the address specify the memory bank, this performance measurement input could be specified as SAMPLE IADD<12:9> AT \$MEMREF. IADD, a permanent variable allocated in mainstore or control store, is sampled rather than the processor argument (assigned to a QM-1 register for a limited lifetime) to assure valid data. The data is sampled at the end of the processor to assure that IADD is updated before being sampled.

The I/O status test is a simple example of efficient use of the performance measurement capability. The input could be specified as:

1. COUNT IO-STATUS-TEST AND NOT IO-STATUS
2. IF IO-STATUS = 0 THEN COUNT IO-STATUS-TEST
3. SAMPLE IO-STATUS AT IO-STATUS-TEST
4. COUNT IO-STATUS-TEST+3

The first three methods cause an emulator trap to occur whenever the IO status instruction is executed. The last input is preferable since the trap occurs only when the emulator executes the path for an unsuccessful branch.

7.6 SMITE Concurrency Implementation

Statements within a PARALLEL-BEGIN/PARALLEL-END block are executed as concurrent subtasks of the emulator. The occurrence of a PARALLEL-BEGIN statement causes the compiler to generate a system recall instruction, distinct from the SMITE externals, which suspends execution of the emulator subtask and creates new subtasks for the statements inside the block. The next available task executes until it is terminated by the end of the statement, it is suspended, or the interval timer ends its time slice. When all its descendant tasks complete, a parent task is placed on the list of active tasks.

A significant amount of the concurrency implementation is transparent to the user. The concurrent subtasks are not tasks recognized by the TASK system. SASS creates a subtask storage area in control store and exchanges subtask registers at each swap. From the viewpoint of TASK, the emulator task was continuously executing. SASS also initializes the subtask registers, including the establishment of the stack frame register for the tasks. SASS protects shared-data accesses from interrupts by checking the protection flag at the time of the planned task swap.

Three implementation issues affect the user in a minimal way. First the number of parallel tasks is limited by the amount of available control store. Each active subtask requires a 70 word control block. Second, the user also may specify the length of the time slice. The nominal value of 10 seconds is designed for loosely-coupled systems. When more frequent interaction is required, the time slice can be reduced. The user must weigh the synchronization delays. Lastly, the user should be aware the SMITE clock is updated at the end of the parallel block to the maximum value of any subtask.

The SMITE programmer may be directly affected by the synchronization issue. Neither SASS nor the language provide any inherent clock or data synchronization. The time slice approach yields an equal amount of QM-1 processor time to each subtask; however, the simulated clocks of two tasks could diverge with user intervention. This would occur when a slow but simple processor is coupled with a fast and complex processor description. Likewise, SASS is unaware of any data dependencies necessary to resume. The user is faced with two alternatives:

- 1) Place the necessary synchronization test as a DO condition and continue to execute the fail case until the condition becomes true after one or more time slices.

- 2) Place a direct code statement to suspend the task (SYSTEM 0'4013') as the fail case of the DO statement.

7.7 SASS Modification

SASS is a generalized support system for SMITE emulators. Although the general purpose version of SASS is quite useful for the initial testing of an emulator, additions to or modifications of SASS are often desirable at later stages for:

1. Specialized input/output support,
2. Emulation-unique state display,
3. Special commands to alter emulation variables (registers, switches, controls) directly rather than through the CHANGEMS command, and
4. New external functions.

This section will discuss the SMITE-unique requirements for any system modification. The programmer should refer to the SIMPLQ Reference Manual and the EASY System reference manuals [14,15,16].

7.7.1 SASS Files

The SASS files delivered should not be modified by the SMITE programmer in order to maintain configuration control over the system. The proper method for modifying SASS is to create equivalent routines which perform the modified function. Table 7.7.1 describes the files used to create SASS. The following sections discuss the files which are generally updated for special versions of SASS and also their restrictions.

7.7.1.1 EZ:GEND

EZ GEN is an NCS EXEC file which controls the generation of the SASS disk boot file. New user supplied routines should be added to this file as additional inputs to the NCS program PREPD, which is executed by the EZ:GEN file. The EASY system is regenerated when new nanocode is added for DECODE on OPDEF statements. Appendix D describes the restrictions imposed on user nanocode.

TABLE 7.7.1 SASS PROGRAM FILES

EZGEN	Prep File for SASS
TCPGEN	Prep file for TASK
SMITE	SMITE Library File
CMDSMITE	SMITE Command File
ADVNS:SOU	Nanocode Files
SHFTX	
TRWNS:SOU	
AUXFCN	SASS Control Command
CSCHED	Concurrency Task Scheduler
CSTART	Concurrency Task Start
CTERM	Concurrency Task Termination
ITEM	PM Data Item Input
LINEGRD	PM Command Line Read
LOADMS	Load MainStore File
LSMITE	Load Compiler Object File
MEMDSP	Memory Display/Modify
NEXT	Micro-step
PARSE	PM Input Control
PMENTER	PM Input Table Construction
PMREL	PM IF Clause Control
PMTRAP	PM Trap Processor
PROBE	PM Control Probe Input
RECORD	PM Data Output
SDISP	State Display
SMITECOPY	Emulator File Copy
SMITELIB	SASS Common Decks
?CMDSMITE	SMITE Command File Source
?COPY	File Copy
?IHSMITE	SMITE Recall Handler
?SMITE	SMITE Driver

7.7.1.2 Command File

A specific version of SASS is initiated by the command SMITE,EMULATION=X. The processing for the SMITE command assumes that the SASS commands are on the file CMDX. This file should be constructed by adding commands to the standard SASS command file ?CMDSMITE and executing the EASY utility MAKECF (Make Command File).

7.6.1.3 Library File

Any SASS code developed or modified for a specific version of SASS must be contained on the EASY library file X, where X is the emulation name. The system searches for external names from libraries in the order X, SMITE, and SYSTEM. In particular, the recall and display modules have fixed names and are usually present on more than one library. All SASS code is compiled by the SIMPLQ command and inserted into the library by EDITLIB.

7.6.1.4 Recall Handler

The recall handler for SASS must be named IHSMITE. This routine processes system recall instructions issued by the emulator (externals) and declares all global data unique to the version of SASS.

The interrupt handler must observe the TASK and SMITE recall conventions. The recall number is contained in the upper 9 bits of the Event Control Block for the emulator. For SASS, the system recall instructions are defined as follows:

Recall Number	Function
0-5	SMITE Externals 1-6
6	Additional SMITE Externals R.ADR=External Ordinal
7-10	Unused
11	Task Suspension
12	Task Completion

13	Task Start
14	Performance Measurement Trap
15	Micro Breakpoint

The code in the standard SASS recall handler ?IHSMITE should be used as a model for new handlers.

The SASS functions are implemented as a series of overlays kept on disk. The only resident code are the server routine for the SMITE command and the recall handler. The SMITE command processor contains the global data base for SASS and several common routines. The interrupt handler for a particular emulation contains the global data base unique to that version of SASS and its common processing routines.

The recall handler also redefines the processing array which controls the display and alteration of emulated memory. The array is organized in the following manner:

WORD	Function
0 store	Location of memory relative to SMITE main store
1	Width of memory
2	Length of memory
3	Display format (8=octal,16=hexadecimal)
4	QM-1 words per parent item
5	Data items per parent (packing density)
6	Number item printed per line
7	Number digits in data item value
8	1 if memory addressing is [a:b] where a>b
9	Number digists in address field

This array is used by the CHANGEMS, MEMDISPLAY, MEMPRINT, and MEMVAL commands.

7.7.1.5 State Display

The state display for SASS must be named SDISP.

7.7.1.6 Memory Allocation Considerations

The QM-1 main store allocated to the EASY system contains the emulator memory (SMITE) followed by the EASY stack at the low end of memory and the resident EASY program space at the top end. For the emulator to execute, it must be allocated enough memory space for all its mainstore variables and its stack.

There are two possible methods of causing this allocation:

- 1) Some EASY system routines may be reassembled with the new mainstore allocation
- 2) The EASY command DEADSTART may be input before invoking SASS.

The particular choice is a function of the relative system usage.

7.7.2 Specialized SASS: Case Study

The generalized SASS system may be used for the initial checkout of an emulator. This is usually a cumbersome process, requiring the user to map user variables to the QM-1 memory location and bit position. A specialized SASS system can be constructed to provide input control over the emulation and to supply the external environment.

As a hypothetical example, assume that an emulation has been developed for a Zilog Z80 microcomputer used in a telecommunication application. Further, assume that the emulation system is required to provide the following capabilities:

1. Simulate the external telecommunication device
2. Supply input stimuli to the device (e.g., dialing information) under used control
3. Support the Z80 input/output instructions which interface with the device

4. Provide the capability of modifying emulated registers and switches
5. Generate a Z80 state display
6. Contain the nanocode to implement the DECODE statement and one new microinstruction.

An analysis of these requirements may indicate that the file space must be created with the NOVA based NCS system for the following:

1. Command File (named CMDZ80)
2. Library File (name Z80)
3. Command File Source
4. Recall Handler
5. Input Command Processor
6. State Display
7. Device Simulator
8. Time Keyed Event Processor
9. Nanocode Source
10. Nanocode Object

The nanocode source may be developed and maintained with the NCS utility EDIT. Alternatively, the source may be constructed with the EASY utility EDIT and copied to a NOVA-format file through the COPYSN utility. The source file is assembled with the nano-assembler and the object file is created by MAP. The system generation is modified by including the user nanocode file and EASY is regenerated.

The command source file is a copy of the standard command source file (?CMDSMITE) or that used for the 8080 system (?CMD8080). Each new command for the Z80 system requires a 4 line entry as described in reference [15]. The command file is created with the MAKECF utility.

The library file Z80 is initialized with the BUILD

option of the EDITLIB utility. The ADD and REPLACE options are used to add or modify binary files on the library. The binaries are created by the SIMPLQ compiler on source files maintained by the EDIT utility. The emulation system is then referenced by the command SMITE,EMULATION=Z80.

The new program modules should be patterned after the existing SASS routines. The recall handler should declare all global data (e.g. device information) unique to the emulation system.

The recall handler may also declare the location of interface variable between SASS and the emulation (e.g. device buffers and interrupt flags). It must also process the external calls for I/O instructions and interface with the device simulations. The server routines for the new commands must observe the EASY protocol for parameter passing and control transfer.

8. Advanced Capabilities

8.1 Syntax Macros

Syntax Macros allow the user to extend the SMITE language to better suit the problem at hand via the definition and use of new statements and primitives. For a statement macro the user defines the syntax and semantics of a statement level operation within a SMITE description. A primitive macro provides for new elements of a SMITE expression. For example, one might use a primitive macro to define and implement a ones complement add. Ultimately, the macro capability provides another level of abstraction, similar in purpose to a processor, which allows for clearer and more properly structured computer descriptions.

A powerful tool in defining a new language feature is the provision to use direct code in a syntax macro. Direct code is a block, or blocks, of micro code assembly language which may be interspersed with SMITE statements. The ability to use this lower level language provides the user with added flexibility in developing solutions to his problem.

The macro capability consists of macro definitions which include syntax and semantic requirements used to recognize subsequent references to the macros. At the point of reference the macro is optionally expanded inline or simply referenced as a call to a closed processor expansion of the macros. All macro definitions must occur at the beginning of the program, before the processor descriptions. The format of the syntax macro definition is:

```
header clause (includes expansion type, macro use type)
template clause
WHERE clause (optional)
RETURNS clause (if any)
MEANS clause
    including: data declarations
               means body
END macro statement
```

For clarity the header and template are usually on the same line, as shown in the following example of a simple macro definition:

```
INLINE PRIMITIVE 'NEGATE PARM1:ID'
WHERE PARM1:REGISTER
```

```

RETURNS MINUS <35>
MEANS
''NO DATA DECLARATION NEEDED IN THIS EXAMPLE ''
MINUS <- -PARM1;
END PRIMITIVE;

```

8.1.1 The Header Clause

A syntax macro begins with a header clause which contains two elements describing the general nature of the macro expansion.

The first element is optional; it specifies whether the macro is to be expanded as an `INLINE` or a `CLOSED` function. If not specified the expansion will be a closed function. `CLOSED` functions will be expanded as separate closed processors, one for each macro call. Only if the parameters for two `CLOSED` macro calls are the same will their expansions be combined into a common processor.

The second element in the header clause describes the use of the macro as either a `STATEMENT` or `PRIMITIVE`. A `STATEMENT` macro defines a new `SMITE` statement type. The use of a `STATEMENT` macro is similar to that of a `SMITE` subroutine processor described in chapter 1. The macro call is a complete `SMITE` statement to perform whatever actions are prescribed by the macro definition. For example:

```
NEGATE PARM;
```

is a `SMITE` statement which would execute the `STATEMENT` macro satisfying the syntax: keyword `NEGATE` followed by a parameter. (This would not execute the `PRIMITIVE` macro defined in the example above.) The results of processing the statement macro is reflected by the change in the parameter used in the calling statement, e.g., change the sign of `PARM`.

A `PRIMITIVE` macro is similar in use to a function processor; it defines a set of actions which provide a specific result. A call to a `PRIMITIVE` macro is an element of a `SMITE` statement. For example:

```
X <- NEGATE PARM;
```

would take the output from the `PRIMITIVE` macro (as defined in the example above) and place it in a data item called "X". The parameter would not necessarily be changed, but

it could. The attributes of the PRIMITIVE macro output must be defined by the RETURNS statement described below.

8.1.2 The Template Clause

Following the header is the macro template. The template is a string of elements, enclosed in single quotes, which define the syntax or form of the new statement (the macro call or reference). The elements are either tokens or formal parameters. The elements must appear in the calling statement, in the same order and number, for it to be recognized as a macro call. A token can be a string of legal characters that is not a form of one of the following:

END, DEBUG, DECLARE, EXPRESSION, REFERENCE, ";", ":"

For STATEMENT macros the first token must be a keyword which will uniquely identify the macro call statement. This symbol cannot be one of the SMITE reserved words (Appendix I). While PRIMITIVE macros do not have this first token requirement it is necessary to avoid ambiguity in the definition of the macro.

A formal parameter is of the form:

id:parameter type

where:

id is any allowable SMITE identifier.

parameter type specifies the form of the actual parameters in the macro call. The allowed values are ID or REFERENCE, and EXPRESSION for statement macros.

The following are examples of the header and the template of a syntax macro.

1. INLINE STATEMENT 'DO-NOTHING'
2. CLOSE PRIMITIVE 'DO-SOMETHING PARM1:ID PARM2:REFERENCE
; PARM3:ID'
3. STATEMENT 'DOANYTHING PARM1:ID TIMES PARM2:EXPRESSION'

Macros may be referenced anywhere in a SMITE description. However, to avoid confusion in the recognition of new tokens defined in a macro's syntax (in the template clause), a macro must be defined before it is referenced within the defined expansion of another macro (means

clause). Additionally, recursive macro expansions should be carefully avoided.

8.1.3 The WHERE Clause

One of the powerful features of a syntax macro is that the same macro can be referenced with parameters defined by different attributes. If the parameter attributes are critical to the macro process, the WHERE clause can be used to filter out those cases where the expansion should not occur. The primary use for this is to have several macro definitions with the same template but which apply to parameters with different attributes. For example the macro processing to increment an 18 bit register would be different from that to increment one 32 bits wide. The WHERE clauses for the two functions could be:

1. `INLINE PRIMITIVE 'INCREMENT (PARM:ID)'
WHERE PARM:REGISTER,WIDTH(PARM)<=18`
2. `INLINE PRIMITIVE 'INCREMENT (PARM:ID)'
WHERE PARM:REGISTER,(WIDTH(PARM)>=19) AND
(WIDTH(PARM)<33)`

The first macro would expand calls where the actual parameters are registers less than or equal to 18 bits in width. The second macro would expand for calling parameters which are registers between 19 and 32 bits in width. If the calling parameter were not a register or its width were greater than 32 bits an error would be indicated if another similar macro definition could not be found to satisfy the attributes of the actual parameters.

The WHERE clause can contain any number of attribute requirements against the parameters contained in the template. Each such specification is separated by commas. There are two formats for the specifications: those for type requirements and those for the attribute requirements. The type requirements are of the form:

parameter:type

where type is one of the SMITE data types such as REGISTER or MEMORY.

The attribute requirements are stated as a compile time expression.

Both forms are shown in the examples above.

8.1.4 The RETURNS Clause

The RETURNS clause is always used with (and only with) a PRIMITIVE macro. It defines the attributes of the output returned by the macro processing to the "main line" code. The RETURNS clause describes a data item, within the macro definition, which will be used as the output variable. The format of this data description is a normal SMITE data description. It has the following form:

```
RETURNS data-id width
```

where:

data-id is a local data item. The data item named in the RETURNS clause may not be declared elsewhere in the macro definition. The use of the item on the left side of an assignment statement within the macro definition will determine the value of the macro output.

width is the normal SMITE width declaration for the output.

An example of a macro definition with the returns clause is:

```
INLINE PRIMITIVE 'INCREMENT P1'
RETURNS INC <35,0>
MEANS
INC <- P1 + 1,
END PRIMITIVE;
```

This macro could be used by a statement such as:

```
LC <- INCREMENT (LC),
```

8.1.5 The MEANS Clause

8.1.5.1 Data Declarations

All data items used in the macro must be formal parameters, defined in the RETURNS clause, or be defined locally in this section of the macro definition. No other data may be referenced. The data declarations use the normal SMITE data declaration format described in Chapter 1. In addition to the data types described in Chapter 1, a new data declaration, TEMPORARY, is permitted in the syntax macro. This data declaration is of the form:

id TEMPORARY <width>

where:

id is the SMITE identifier for this data item which is to be maintained in a QM-1 register.

width is a compile time expression for the data item width in terms of QM-1 registers. It must evaluate to 1, 2 or 4.

TEMPORARY data items may be used in a SMITE statement like any other data item. They may also be used as register designations in the micro code statements, thus allowing the communication between the direct code and the SMITE statements.

8.1.5.2 MEANS Body

The body of the MEANS clause may contain any number of SMITE statements and blocks of direct code which describe the processing to be performed by the macro expansion. Any statement allowed in a SMITE processor is permitted in this clause, with the exception of embedded processors.

The statements may use only locally defined data items and the formal parameters specified in the template. The formal parameters can appear in any part of a statement, consistent with their type, as specified in the template. That is: a parameter typed as EXPRESSION can appear in any place a expression is legal; an expression can only appear in a STATEMENT macro. A parameter typed as REFERENCE can appear in any part of a statement where a reference can be made. A reference, in this case, is an identifier reference, a parenthesized expression, or a constant. A parameter typed as ID can appear anywhere an identifier is legal. The parameter communication for macros uses the same CBVR and CBV methods described for processors in Chapter 1.

In addition to the allowable SMITE statements, blocks of direct code are permitted in the means clause. They may be interspersed between the SMITE statements and communicate with them using TEMPORARY data items.

The MEANS clause describes the processing to be performed by the macro expansion. This clause consists of the

keyword MEANS, data declarations, and the body of the MEANS clause.

8.1.6 The END Statement

The END statement terminates the macro definition and has one of the following forms, depending on the type of macro:

END STATEMENT, END PRIMITIVE, ENDSTATEMENT, or
ENDPRIMITIVE

8.2 Direct Code

Direct code provides the SMITE programmer with the capability of incorporating a sequence of microinstructions into the code to improve the efficiency of the emulator. Since the typical emulation spends a large percentage of its execution time in a small segment of code, a substantial throughput achievement may be attained by direct code. The primary reasons for direct code are:

- 1) New microinstructions developed especially for this emulation, such as an operand fetch or one's complement arithmetic, may be used.
- 2) Existing microinstructions can be used to access hardware capabilities not used by the SMITE compiler. These include the RMI, E store, the 16 bit arithmetic mode, and the 32 bit shifter.
- 3) The SMITE programmer may use the QM-1 status register FIST for the status of the emulated computer arithmetic operations rather than calculating it from SMITE level statements.
- 4) Some existing microinstructions (e.g., EXTR, ALUX, MPY) which are not generally used by the compiler may be optional for certain special cases.
- 5) The SMITE programmer may fine tune small segments of code by recognizing common subexpressions and the data flow.

Direct code is a powerful capability, fraught with potential problems. The user must understand the compiler interface to the code as well as the microinstruction set.

8.2.1 Micro Code Format

The micro code statements are similar to the normal MULTI format, but they must be specially formatted so that the SMITE compiler can recognize and pass them to the assembler. The primary difference is that each statement ends with a semicolon (;), and labels are followed by ":". In addition the following items must be observed:

- a. Numbers are nominally decimal, not octal as in the MULTI assembler.
- b. Comments are enclosed by ' ' or are at the end of the statement, prior to the ";" which terminates the statement.
- c. There are no predetermined assembler constants such as SIGN, RESULT, DOUBLE, LEFT, etc.
- d. The same level of assembly error messages are not produced.
- e. Label must be local to the direct code segment.
- f. There is no data definition facility for memory addresses. It is therefore difficult to perform load and stores except with absolute addressing.
- g. SMITE identifiers used in a macro definition should not be the same as the MULTI operation mnemonics.

8.2.2 Direct SMITE Statements

A direct code block may also contain SMITE statements interspersed among the assembly language statements. These statements are referred to as direct SMITE statements and are identified by the dollar sign, "\$", preceding the statement. One purpose for the direct SMITE statement is to perform load and store operations for the direct code.

The direct SMITE statements do not affect the context block of the direct code; the direct code labels are known across the direct SMITE statements. However, only QM-1 registers associated with TEMPORARY data items, macro parameters and returned values will be maintained across the direct SMITE statements.

Direct code consists of a set (block) of QM-1 micro assembly language (MULTI) statements, delimited by the following statements:

DIRECT;
SMITE;

Direct code blocks are allowed only in syntax macros; more than one block may be present within a single macro, but each block is a separate entity and communicates with any other block or with the remainder of the SMITE program only through the TEMPORARY data items described above.

8.3 The OPDEF Statement

Augmenting the MULTI statements allowed in a direct code block are user defined micro statements. These statements are defined by the OPDEF statement. The OPDEF statement or statements must be the first statements of a SMITE program. It has the form:

OPDEF operation-list;

The operation list contains one or more operations, each separated by a comma. Two different operations are allowed. The first operation, of the form:

COPY system-file

references a system file to retrieve previously defined OPDEF's.

The second operation actually defines the format of the new instruction which is to be generated. This format must be compatible with the QM-1 instruction format and the nanocode capabilities for operating on the new instruction. The form of this clause is:

mnemonic opcode <operand-lengths>

where:

the mnemonic is any unique identifier for the instruction.

the opcode is the decimal value of the operation code for

the instruction. The operation code will occupy 7 bits of the generated instruction; it may not conflict with existing MULTI operation codes.

the operand lengths describe the number of bits for each operand address in the generated instruction. The number and order of the operands is implicit from the number of lengths described.

For example:

```
OPDEF NEWOP 66 <5,6>
```

describes a new instruction called "newop" with an opcode of 66 and two operands. The operand addresses will occupy the next 5 and 6 bits, respectively, after the 7 bit opcode in the generated micro instruction. (As a practical matter the first two operands will usually be as shown in the example.)

Although the example would generate a single word instruction, multiple word instructions can be generated, e.g., LONG-OP 66 <5,6,18>. Appropriate nanocode must be implemented on the QM-1 to execute these instructions.

The use of the new instruction shown above is shown in the following example of a syntax macro:

```
INLINE PRIMITIVE 'USEOPDEF A:ID '
RETURNS B<35:0>
MEANS DECLARE C TEMPORARY <2>,
      D TEMPORARY <1>;
C <- A;
DIRECT;
NEWOP C,D;
SMITE;
B <- D;
END PRIMITIVE;
```

8.4 Examples

The following are the complete examples of the INCREMENT macro described in the paragraph on the WHERE clause. The first example is for the short register increment.

```
INLINE PRIMITIVE 'INCREMENT (A:ID)
WHERE A:REGISTER, WIDTH(A)<=18
RETURNS B<17:0>
```

```

MEANS
B<-A;
  DIRECT;
    ADI B,1
  SMITE;
END PRIMITIVE;

```

The second example is for a longer register increment (between 18 and 36 bit wide).

```

INLINE PRIMITIVE 'INCREMENT ( A:ID )'
WHERE A:REGISTER, (WIDTH(A)>=19) AND (WIDTH(A)<37)
RETURNS B<35:0>
MEANS
  B<-A,
  DIRECT;
    ADI B+1,1,
    BNZ B+1,J1,
    ADI B,1,
  J1:
  SMITE;
END PRIMITIVE;

```

Both macros would be referenced with the same statement format. For example, the statement:

```
LC <- INCREMENT (LC);
```

would use the correct macro expansion, depending on the declared width of LC. After having once declared the width of the data item the user no longer need be concerned about which macro to use for which width.

References


- [1] Bell, C. G. and A. Newell, Computer Structures: Readings and Examples, McGraw-Hill, Inc., 1971.
- [2] Boehm, B. W., D. W. Kosy, and N. R. Nielson, "Simulation Aids for Designing Integrated Information Systems: The ECSS Language," Astronautics and Aeronautics, November 1972, pp 68-74.
- [3] Knudsen, M. H., PMSL, An Interactive Language For System-Level Description and Analysis of Computer Structures, Carnegie-Mellon University (NTIS AD 762 513), 1973.
- [4] Hill, F. J., "Introducing AHPL," Computer, December 1972 (Vol. 7, No. 12), pp 28-30.
- [5] Chu, Y. "Introducing CDL," Computer, December 1972 (Vol. 7, No. 12), pp 31-33.
- [6] Siewiorek, D. "Introducing ISP," Computer, December 1977 (Vol. 7, No. 12), pp 39-41.
- [7] Press, B. and R. K. McClean, "The Flexible Analysis, Simulation, and Test Facility: A Practical Software First Capability," IEEE NAECON '76 Record, pp 264-268.
- [8] McClean, R. K., and B. Press, "The Flexible Analysis, Simulation, and Test Facility: Diagnostic Emulation," TRW Software Series TRW-SS-75-03, October, 1975.
- [9] Hilbing, Col. F. J., "The RADC Computer Architecture Program,"
- [10] MULTI Micromachine Description, Nanodata Corporation, Williamsville, New York.
- [11] QM-1 Hardware Level User's Manual, Nanodata Corporation, Williamsville, New York.
- [12] Programmable Run-time Operators Display (PROD) Manual, Nanodata Corporation, Williamsville, New York.
- [13] NCS User's Manual, Nanodata Corporation, Williamsville, New York.
- [14] SIMPL-Q Reference Manual, Naval Surface Weapons Center, NSWC/DL TR-3778, April 1978.

32584-6015-RU-00









[15] Emulation Aid System (EASY) System Programmer's Guide, Naval Surface Weapons Center, NSWC/DL TR-3774, December 1977.

[16] EASY - The Design and Implementation of an Intermediate Language Machine, Naval Surface Weapons Center, NSWC/DL TR-3765, October 1977.

[17] UNIVAC 642B Emulator Programmer's Reference Manual, TRW DSSG, December 1978.




SMITE SYNTAX DIAGRAMS (CONTINUED)



TEMP WIDTH CTE RESULTING IN 1
2 OR 4

SMITE SYNTAX DIAGRAMS (CONTINUED)



SMITE SYNTAX DIAGRAMS (CONTINUED)

Appendix B. SMITE Compiler Operation Procedures

The SMITE-2 compiler has been installed on the Honeywell 6180 computer at Rome Air Development Center, and is accessible over the ARPANET (RADCMULTICS).

The smite command invokes the SMITE compiler to translate a segment containing the text of a SMITE computer description into an object segment suitable for downline loading and executing on the QM-1. A listing file and an optional symbolic data dump file are also produced. The result segments are placed in the user's working directory.

ACCESS TO SMITE

The SMITE compiler is currently located in a subdirectory of the advanced SMITE project. The quickest way to access the compiler is to establish the following link:

```
link >udd>2529c0103>Prentice>asc>smite smite
```

COMMAND SYNTAX

The syntax of the smite command is

```
smite $ control_args †
```

where the control_args are 0 or more of the following:

-i path

specifies that path is the path name for the input to the SMITE compiler. If path does not have a suffix of smite, then one is assumed; however, the suffix smite must be the last component of the name of the input segment. The default path name is input.smite.

The input segment name with the smite suffix removed

becomes part of the default names for the other files manipulated by the compiler.

-l segname

specifies that segname is the segment name for the list output from the SMITE compiler. The default segment name is input_file_name.list

-b segname

specifies that segname is the segment name for the object output from the SMITE compiler. The default segment name is input_file_name.lgo.

-sda opt

where opt is one of

trees
symbols
both
none

specifies the symbolic data dump option for the compiler. The default value is none. Upper/lower case distinctions are ignored for the sda parameter.

-sdafile segname

specifies the segment name for the symbolic data dump output from the compiler. The default value for the segment name is input_file_name.sdaout.

-map

-nomap

control the allocation map print for the compiler. The default parameter is -map.

-rech

-norech

controls the recursion checking in the compiler. The default parameter is -rech.

-maxcs n

where n is an decimal number between 1 and 40000,

inclusive, controls the maximum amount of control store allocated by the compiler. The default value is 8192 (20000 octal).

-csall opt

where opt is one of

all
part
words
none

controls the allocation of REGISTER class variables to control store. The default value is all. Upper/lower case distinctions are ignored.

-csseg

-nocsseg

specifies whether the object code will be targetted to a QM-1 with control store segmentation hardware installed and operational. If -csseg is specified, the object code will be partitioned into a read only and a read/write segments, with read-only segment accesses based at address 200000 (octal). If -nocsseg is specified, the read-only section will begin at the first page boundary (256 words) after the end of the read/write section. The default value is -nocsseg.

-asmlist

-noasmlist

controls the assembly list of the compiler. The default parameter is -asmlist.

-pack

-nopack

controls allocation of multiple array elements of data structures to QM-1 double words. This trades reduced data space for increased code space. The default value is -nopack.

Examples

smite -i intel80

This causes the segment intel80.smite in the user's current working directory to be input to the SMITE compiler. The listing will go to segment intel80.list and the object code to intel80.lgo. The listing file will contain the source listing, an allocation map, and an assembly code listing. No symbolic data file will be produced. REGISTER class variables will be allocated to control store, arrays will not be automatically packed, and the object code will not be segmented. This is expected to be the nominal usage of the SMITE compiler.

smite -i <acctest>u642b.smite -csall words

The segment u642b.smite in parallel subdirectory acctest will be used as input, with the segments u642b.list and u642b.lgo in the working directory receiving the list and object code output, respectively. In addition, only non-dimensioned REGISTER class variables will be allocated to control store.

NOTES:

1. The SMITE compiler creates segments in the user's working directory; therefore he/she must have the proper permissions there.
2. The SMITE compiler can be executed only once per process. Failure to establish a new process between calls to the compiler will produce unpredictable (BAD!) results. A sure sign that this has not been done is that the first page of list output is not numbered 1.

Appendix C: Intel 8080 Microprocessor Definition

```

INTEL-8080 MICROPROCESSOR:

DECLARE DEFAULT<7:0> REGISTER,
PRUD-FLAGS<35:0>,
  STFP-FLAG FLAG DEFINED PRUD-FLAGS<35>,
  INTERRUPT-IR<7:0> DEFINED PRUD-FLAGS<25:18>,
  INTE FLAG DEFINED PRUD-FLAGS<17>,
  INTERRUPT FLAG DEFINED PRUD-FLAGS<16>,
  PC<15:0> DEFINED PRUD-FLAGS<15:0>,
REGS<131><15:0>,
  REGISTERS<17><7:0> DEFINED REGS,
  R DEFINED REGISTERS<0>,
  C DEFINED REGISTERS<1>,
  D DEFINED REGISTERS<2>,
  E DEFINED REGISTERS<3>,
  H DEFINED REGISTERS<4>,
  L DEFINED REGISTERS<5>,
  STATUS DEFINED REGISTERS<6>,
    CARRY FLAG DEFINED STATUS<6>,
    PARITY FLAG DEFINED STATUS<2>,
    AUX-CARRY FLAG DEFINED STATUS<4>,
    ZFRO FLAG DEFINED STATUS<6>,
    SIGN FLAG DEFINED STATUS<7>,
    A DEFINED REGISTERS<7>,
    CARRY-A<8:0> DEFINED REGS<31><6:0>,
SP<15:0>,
B-PAIR<15:0> DEFINED REGS<0>,
D-PAIR<15:0> DEFINED REGS<1>,
H-PAIR<15:0> DEFINED REGS<2>,
PSW<15:0> DEFINED REGS<3>,
IP,
  IF-OPFAND-SELECT<2:0> DEFINED IR<2:0>,
  IF-DEFT-SELECT<2:0> DEFINED IR<5:3>,
  IF-SUB-FUNCTION<2:0> DEFINED IR<2:0>,
  IF-TEST-SELECT<2:0> DEFINED IR<5:3>,
OP-PFC,
HOLD-A,
LP-PAIR<15:0>,
TEST FLAG,
HOLD-PFC,
PIM<14:0> MEMORY:

```

```

*****
UP-STEP EXTERNAL,
UP-HALT EXTERNAL,
UP-FRAME EXTERNAL,
UP-MOTISM EXTERNAL,
IN-PRFT PORT,
OUT-PRFT PORT;

** ADD-REG ADDS AN OPERAND IN OP-REG TO THE REGISTER SELECTED BY**
** IR-DEST-SELECT. **

ADD-REG: PROCESSOR:

    IF IR-DEST-SELECT = 6
    THEN IN 10 HOLD-REG <- MEM[H-PAIR] <-
        MEM[H-PAIR] + OP-REG;
    ELSE IN 5 HOLD-REG <- REGISTER[IR-DEST-SELECT] <-
        REGISTER[IR-DEST-SELECT] + OP-REG;
    END IF;

    SIGN <- HOLD-REG<7>;
    ZERO <- HOLD-REG = 0;
    PARITY <- HOLD-REG<7> XOR HOLD-REG<6> XOR HOLD-REG<5> XOR
        HOLD-REG<4> XOR HOLD-REG<3> XOR HOLD-REG<2> XOR
        HOLD-REG<1> XOR HOLD-REG<0> XOR 1;

    ADD-REG: END;

** LUIALU FETCHES A MEMORY ADDRESS WHICH IS LOCATED IN THE TWO **
** BYTES IMMEDIATELY FOLLOWING THE CURRENT INSTRUCTION. THE **
** INSTRUCTION COUNTER #PC# IS UPDATED. THE CLOCK IS NOT CHANGED**

LUIALU: PROCESSOR:

    OP-PAIR <- MEM[PC+1] // MEM[PC];
    PC <- PC + 2;

    LUIALU: END;

** LUIALU FETCHES AN 8-BIT OPERAND BASED ON THE OPERAND SELECT **
** BITS OF THE CURRENT INSTRUCTION. **

```

```

OPERAND: PROCESSOR<710>;

IF IP-OPERAND-SELECT /> 6
  THEN IN 1 OPERAND <- REGISTER[IN-OPERAND-SELECT];
  ELSE IN 4 OPERAND <- MEM[H-PAIR];
  END IF;

OPERAND: END;

** PERFORM-ADD ADDS THE CONTENTS OF OP-REG TO THE ACCUMULATOR **
** AND SETS THE STATUS BITS. THE CARRY IS NOT ALTERED. **

PERFORM-ADD: PROCESSOR(CARRY-IN);

UNCLARE CARRY-IN FLAG;

(AUX-CARRY // A<310>) <- CARRY-IN +
  (O//A<310>) + OP-REG<310>;
CARRY-A<R14> <- AUX-CARRY +
  (O//A<714>) + OP-REG<714>;
SET-STATUS;

PERFORM-ADD: END;

** PERFORM-OP EXECUTES ARITHMETIC OPERATIONS BASED ON THE **
** FUNCTION CODE IN IP. THE CARRY IS NOT ALTERED. THE OPERANDS **
** ARE *A* AND *OP-REG*. **

PERFORM-OP: PROCESSOR;

CASE IP<313>;
  ADD: BEGIN;
    PERFORM-ADD(C);
    ADD: END;

  ADD: BEGIN;
    PERFORM-ADD(CARRY);
    ADD: END;

  SUB: BEGIN;

```

```

OP-REG ← - OP-REG;
PERFORM-ADD(C);
CARRY ← NOT CARRY;
SUB: END;

SPR: BEGIN;
OP-REG ← OP-REG;
PERFORM-ADD( NOT CARRY );
CARRY ← NOT CARRY;
SHR: END;

ANA: BEGIN;
A ← A AND OP-REG;
SET-STATUS;
CARRY ← 0;
ANA: END;

XFA: BEGIN;
A ← A XOR OP-REG;
SET-STATUS;
CARRY ← 0;
XFA: END;

OFA: BEGIN;
A ← A OR OP-REG;
SET-STATUS;
CARRY ← 0;
OFA: END;

CMP: BEGIN;
HOLD-A ← A;
OP-REG ← - OP-REG;
PERFORM-ADD(C);
CARRY ← NOT CARRY;
A ← HOLD-A;
CMP: END;

JPD CASE;

PERFORM-OP: END;

```

PERFORM-OP REMOVES THE TOP TWO WORDS FROM THE CURRENT STACK AND

```

**
** RETURNS THE RESULT AS THE PROCESSOR RESULT.
**
PUSH PROCESOR<1510>;

POP <- MEM[SP+1] // MEM[SP];
SP <- SP + 2;

POP# END;

** PUSH STORES ITS OPERAND ONTO THE CURRENT STACK. THE CLOCK IS **
** NOT ALTERED. **
PUSH# PROCESOR(VALUE);

DECLARE VALUE<1510>;

SP <- SP - 2;
(MEM[SP+1] // MEM[SP]) <- VALUE;

PUSH# END;

** SET-TEST-STATUS EVALUATES TEST CONDITIONS FOR CONDITIONAL **
** JUMPS, CALLS, AND RETURNS. IF THE CONDITION IS TRUE THEN THE **
** FLAG #TEST# WILL BE SET TRUE. THE CLOCK IS NOT ALTERED. **
**
SET-TEST-STATUS# PROCESOR;

CASE IP-TEST-SELECT;

  ** N7 ** TEST <- NOT ZERO;
  ** 7 ** TEST <- ZERO;

  ** NC ** TEST <- NOT CARRY;
  ** C ** TEST <- CARRY;

  ** PN ** TEST <- NOT PARITY;
  ** PE ** TEST <- PARITY;

  ** P ** TEST <- NOT SIGN;

```

```

** M ** TEST <- SIGN;
END CASE;

SET-TEST-STATUS: END;

** SET-STATUS INTERPRETS THE CURRENT ACCUMULATOR CONTENTS TO SET
** THE STOP, ZERO, AND PARITY BITS IN THE PSM.
** THE CIRC IS NOT ALTERED.
**

SET-STATUS: PROCESSOR;

SIGN <- A<7>;
ZERO <- A = 0;
PARITY <- A<7> XOR A<6> XOR A<5> XOR A<4> XOR
A<3> XOR A<2> XOR A<1> XOR A<0> XOR 1;

SET-STATUS: END;

** STORE STORES A VALUE BASED ON THE DESTINATION SELECT BITS.
** THE CIRC IS UPDATED.
**

STORE: PROCESSOR(VALUE);

DECLARE VALUE;

IF IF-TEST-SELECT /> 6
THEN IN 3 REGISTER-DIST-SELECT <- VALUE;
ELSE IN 6 MEM[H-PAIR] <- VALUE;
END IF;

STORE: END;

** ***** MAIN PROCESSOR *****
** *****
** *****

```

```

**
** INITIALIZATION.
PC <- 1;
STATUS <- 2;
INTE <- INTERRUPT <- 0;
LD UNLVEP; BEGIN;

** FIRST OFF, CHECK IF PROD HAS SET THE EMULATOR STEP FLAG, AND **
** DRUP TO PROD IF SO.
IF STEP-FLAG
  THEN BEGIN;
    GP-STEP;
  END;
  END IF;

** PATCH INSTRUCTION.
II INTERRUPT AND INTE
  THEN BEGIN;
    INTERRUPT <- INTE <- 0;
    IP <- INTERRUPT-IR;
  END;
  ELSE BEGIN;
    IF <- MEMPC;
    PC <- PC + 1;
  END;
  END IF;

** INCLUDE CN THE UPPER TWO BITS.
CASE IR<7:6>;

** CC -- MISCELLANEOUS INSTRUCTIONS.
CODECC CASE IR-SUB-FUNCTION;

** POP ** IN 4 1;

```

```

** IXI AND DAD** IF IR<3>
  THEN IN 10 DAD: BEGIN:
  IF IR<5:4> /= 3
    THEN UP-PAIR <- REGS[IR<5:4>]]:
    ELSE UP-PAIR <- SP:
    END IF:
    (CARRY // H-PAIR) <- OP-PAIR + 0//H-PAIR:
    DAD: END:
  ELSE IN 10 LXI: BEGIN:
    GETADD:
    IF IR<5:4> /= 3
      THEN REGS[IR<5:4>] <- OP-PAIR:
      ELSE SP <- OP-PAIR:
      END IF:
      LXI: END:
    END IF:
** STA, LDA, LDAX, STAX, SHLD, LHLD ** CASE IR<5:13>:
  STAX-R: IN 7 MEM[B-PAIR] <- A:
  LDAX-B: IN 7 A <- MEM[B-PAIR]:
  STAX-D: IN 7 MEM[D-PAIR] <- A:
  LDAX-D: IN 7 A <- MEM[D-PAIR]:
  SHLD: IN 16 BEGIN:
    GETADD:
    MEM[OP-PAIR] <- L:
    MEM[OP-PAIR + 1] <- H:
    END:
  LHLD: IN 16 BEGIN:
    GETADD:
    L <- MEM[OP-PAIR]:
    H <- MEM[OP-PAIR + 1]:
    END:
  STA: IN 13 BEGIN:
    GETADD:

```

```

MEM(OP-PAIR) <- A;
END;

LDAB IN 13 REGIN;
GETADD;
A <- MEM(OP-PAIR);
END;

END CASE;

** INX, DCX ** IN 5 BEGIN;
IF IP<3>
    THEN OP-PAIR <- -1;
    ELSE OP-PAIR <- 1;
    END IF;
IF IR<514> / = 3
    THEN REGS[IR<514>] <- REGS[IR<514>] + OP-PAIR;
    ELSE SP <- SP + OP-PAIR;
    END IF;
END;

INR: BEGIN;
OP-REG <- 1;
ADD-REG;
END;

DCP: BEGIN;
OP-REG <- -1;
ADD-REG;
END;

MVI: BEGIN;
OP-REG <- MEM(PT);
PC <- PC + 1;
IF IR-DEST-SELECT / = 6
    THEN IN 7 REGISTERS[IR-DEST-SELECT] <- OP-REG;
    ELSE IN 10 MEM(H-PAIR) <- OP-REG;
    END IF;
END;

** ROTATIS, CMA, CAPPS, DAA ** IN 4 CASE IR<513>;

```

```

RRC: BEGIN:
  CARRY <- A<7>;
  A <- SLC(A,1);
  RRC: END;

RRC: BEGIN:
  CARRY <- A<0>;
  A <- SRC(A,1);
  RRC: END;

** PAL** CARRY-A <- SLC(CARRY-A, 1);
** PAP ** CARRY-A <- SRC(CARRY-A, 1);

DAA: BEGIN:
  IF (A<310> > 9) OR AUX-CARRY
  THEN (AUX-CARRY//A<310>) <-
    (C//A<310>) + 6;
  END IF;
  IF (A<714> > 9) OR CARRY
  THEN CARRY-A <- (C//A) + X#60;
  END IF;
  SFT-STATUS;
  DAA: END;

** CMA ** A <- NOT A;
** STC ** CARRY <- 1;
** CMC ** CARRY <- NOT CARRY;

END CASE;

CODE00: END CASE;

** ( ) -- DATA TRANSFER INSTRUCTIONS.

CODE01: IF IP = X#76#
  THEN IN 7 DP-HALT;
  ELSE STOPF(OPERAND);
  CODE01: END IF;

```

10 -- ACCUMULATOR FUNCTIONS.

```
CODE10: BEGIN;
  OP-RFC <- OPERAND;
  IN 3 PERFORM-OP;
  CODE10: END;
```

11 -- TRANSFERS AND OTHER STUFF.

```
CODE11: CASE IP-SUB-FUNCTION;
```

000 - CONDITIONAL SUBROUTINE RETURNS.

```
BEGIN;
  IN 5 SET-TEST-STATUS;
  IF TEST
    THEN IN 6 PC <- POP;
    END IF;
  END;
```

001 - RETURN, POP, SPHL, PCHL.

```
CASE IR<013>;
```

```
  POP-R: IN 10 R-PAIR <- POP;
```

```
  RETURN: IN 10 PC <- POP;
```

```
  POP-D: IN 10 D-PAIR <- POP;
```

```
  ## ILLEGAL OP-CODE ## UP-NOT-SIM;
```

```
  POP-H: IN 10 H-PAIR <- POP;
```

```
  PCHL: IN 5 PC <- H-PAIR;
```

```
  POP-PSW: IN 10 PSW <- POP;
```

```

SPHL IN 5 SP <- H-PAIR;
END CASE;

** C1C - CONDITIONAL JUMPS.
IN 10 BEGIN;
GETA00;
SFI-TEST-STATUS;
IF TEST
    THEN PC <- OP-PAIR;
    END IF;
END;

** 011 - MISCELLANEOUS.
CASE IR<513>;

    JMP IN 10 BEGIN;
    GETADD;
    PC <- OP-PAIR;
    END;

    ** ILLEGAL OP CODE ** OP-NOTSIM;

    OUT IN 10 BEGIN;
    IR <- MEM[PC];
    OUT-PORT <- A;
    PC <- PC + 1;
    END;

    INPUT IN 10 BEGIN;
    IR <- MEM[PC];
    A <- IN-PORT;
    PC <- PC + 1;
    END;

    XTHL IN 10 BEGIN;
    OP-PAIR <- MEM[SP+1] // MEM[SP];
    (MEM[SP+1] // MEM[SP]) <- H-PAIR;
    H-PAIR <- OP-PAIR;

```

```

      END;
X(HG) IN 4 BEGIN;
  OP-PAIR <- H-PAIR;
  H-PAIR <- D-PAIR;
  D-PAIR <- OP-PAIR;
  END;

  D1 IN 4 INTE <- 0;
  F1 IN 4 INTE <- 1;
  END CASE;

** JCC - CONDITIONAL CALLS.
BEGIN;
  IN 11 BEGIN;
    SET-TEST-STATUS;
    GETADD;
    END;
  IF TEST
    THEN IN 6 BEGIN;
      PUSH(PC);
      PC <- OP-PAIR;
      END;
    END IF;
  END;

** 101 - CALL AND PUSH.
IF IP<3>
  THEN IF IR<514> = 0
    THEN CALL IN 17 BEGIN;
      GETADD;
      PUSH(PC);
      PC <- OP-PAIR;
      END;
    ELSE ** ILLEGAL OPCODE ** OP-NOTSIM;
      END IF;
  ELSE PUSH-INSTR IN 11 PUSH(REGS[IR<514>]);
  END IF;

```

** 110 - IMMEDIATE OPERAND ARITHMETIC.

IN 7 BEGIN;
 OP-REG ← MEM(PC);
 PC ← PC + 1;
 PERFORM-OP;
 END;

** 111 - PST

PST: IN 11 BEGIN;
 PUSH(PC);
 PC ← 0;
 PC<513> ← IR<513>;
 END;

CONF11: END CASE;

END CASE;

END;
 INTEL-BLOCK: END;

Appendix D: Augmented MULTI Micromachine Definition

Several microinstructions have been added to the standard MULTI and 36 bit microinstruction set for use by the SMITE compiler. These new microinstructions are as follows:

MPY A,B

The positive quantity in register A is multiplied by the positive quantity in register B to produce a positive quantity in registers (A-1)//A.

DIV A,B

The positive quantity in registers (A-1)/A is divided by the positive quantity in register B to produce a positive quotient in register A, and a remainder in register A-1.

SHFTX A,B,CDE

The SHFTX instruction shifts a single or double register by the amount in $R(C) + E$ as specified by the shift control D. /If a single shift is specified R(B) is shifted into R(A). For double shifts $R(B-1)//R(B)$ is shifted into $R(A-1)//R(A)$.

STMSK A,B

The value in register A is stored under the mask in register B into the main store location addressed by R.MX. R.MX is incremented by one after the store. The stored value is $(R(A) \text{ AND } R(B)) \text{ OR } (MS(R.MX) \text{ AND NOT } R(B))$.

TICK AB

The contents of the SMITE clock register is incremented by the 11 bit parameter of the instruction.

STCSK A,B

The value in register R(A) is stored under the mask in register R(B) into the control store location addressed by

R.IY. R.IY is incremented by 1 after the store. The stored value is (R(A) and R(B)) or (CS(R.IY) and NOT R(B)).

STCSK A,B,V

The value in register R(A) is stored under the mask in register R(B) into the control store location addressed by V. The parameter V is then moved into R.ADR. The stored value equals (R(A) and R(B)) or (CS(V) and NOT R(B)).

MVR2 A,B

The register pair (R(B)//R(B+1)) is moved onto the register pair R(A)//R(A+1).

SHIFT4 A,B,C,DE

The register quadruple R(A)//R(A+1)//R(A+2)//R(A+3) is shifted back onto itself. The type shift is specified by the parameter B. The double shift control must be set for this microinstruction. The shift count is the sum of R(C) plus the unsigned 11 bit value in the DE parameter. The status word FIST is destroyed by the 72 bit shift instruction.

PMRET V2,V3

The three word PMRET instruction restarts execution of the emulator after a performance measurement trap. The second word V2 specifies the original instruction in the trap location and the third word V3 designates the address of the trap. PMRET remotely executes the instruction V2 as if it were located at V3, leaving the trap intact. This allows execution to resume as if the trap had not occurred.

SETG10 A,B

The SETG10 instruction is used to provide software protection from an interval timer interrupt during shared data access. This instruction stores A into G10 under the mask in B. G10 equals (G10 and NOT B) or (A and B). The lower order bit of G10 is used by SASS to continue or swap an interrupted task.

DECODE A,B

The DECODE statement in SMITE generates a DECODE microinstruction and a table of pointer addresses to the statements within the DECODE loop. The first parameter

specifies the QM-1 register containing the decode expression. The second parameter specifies the register containing the address of the address table. In the current implementation, the address table occurs immediately after the DECODE microinstruction; and so R.MPC also points to the table. The standard DECODE microinstruction supplied in SASS causes the DECODE to function as a CASE. The DECODE microinstruction must have an opcode of 276 (764XXX).

User Nanocode

User nanocode must be included in the system generation file (EZGEND or another corresponding file) and the system must be regenerated. The compiler is informed of the new microinstructions through OPDEF statements. The instructions may be referenced through direct code within syntax macros.

The current allocation of nanostore for SASS is the following:

FUNCTION	Half-Page Add	Full-Page Add
MULTI Entries	0-216	0-116
36 Bit Entries	217-240	117-140
SASS 1 Entries	241-246	141-146
*EASY Entries	247-261	147-161
Spare Entries	262-266	162-166
SASS 2 Entries	267-275	167-175
Decode Entry	276	176
System Entry	277	177
Nano Interrupt	400-423	200-223
MULTI Exits	424-456	224-256
36 Bit Exits	457-576	257-276
SASS 1 Exits	477-607	277-307

*EASY Exits	610-616	310-316
Spare	617-662	317-362
SASS 2 Exits	663-676	363-376

The user nanocode should reside in the areas labelled spare. At this time, the EASY nanocode is not included in SASS. However, for further compatibility, the user should avoid these areas.

In addition, the operation of two of the Nanodata standard microinstructions has been modified as follows:

SHIFT A,B,CDE

The SHIFT instruction does not perform a 37-bit operation involving the carry bit for double right arithmetic shifts. Instead, a true 36-bit shift is performed. (Note: the SHIFT instruction has been redefined using another opcode. Both shifts are therefore available.)

ALUX A,B,CDE

The status returned in FIST by the ALUX instruction has been modified to change the definition of zero status. Zero is now returned if and only if the ALUX operation generated a zero result, and zero status was set on input to the ALUX instruction.

Appendix E: Advanced SMITE Compiler Error Messages

Numerous error diagnostics are built into the SMITE compiler. These diagnostics generate error messages when violations of the SMITE language definition or of limits of the SMITE compiler are violated.

Errors detected in the assembly phase are printed either immediately before the assembly listing and identify the offending value, or else are printed on the same line as the offending line of microcode. The error messages are of several forms. Errors detected while scanning the input stream are printed interspersed with the listing of the input program, with an arrow provided to indicate the approximate location of the input scan when the error was detected. Errors detected during the storage allocation phase of the compiler are printed above the allocation data line for the offending item. Errors detected during parser semantic analysis or code generation are printed with the source listing line number of the statement being processed when the error was detected.

"AMBIGUOUS MACRO EXPANSION, PRIMITIVE IGNORED"

PYCRPRT

The template and WHERE clause of multiple primitive macros resulted in non-unique identification the macro to be expanded.

"AMBIGUOUS MACRO USE, STATEMENT IGNORED"

PYCRSTT

Selection of the proper statement macro cannot be determined.

"AMBIGUOUS PRIMITIVE IGNORED"

PYPRIM

32584-6015-RU-00

The parser could not determine which primitive macro to be expanded.

"AMBIGUOUS PROCESSOR REFERENCE"

AAEXTPR

A processor has been defined with the same id as an external.

"AMBIGUOUS STATEMENT IGNORED"

PYSTAT

The compiler could not determine which statement macro to be expanded.

"ATTEMPT TO ALLOCATE DEDICATED REGISTER"

GRGTREG

Probable compiler error.

"ATTEMPT TO FREE AVAILABLE ATTRIBUTE"

GEFRATT

Compiler error.

"ATTEMPT TO READ DATA ITEM OF CLASS LIGHT AT LINE"

PMEXP

The data item is on the right side of a transfer operation.

"ATTEMPT TO USE BUSY ATTRIBUTE"

GEGTATT

Compiler error.

"BAD CONSTANT VALUE"

PYDCPHR

The value specified for a constant cannot be evaluated.

"BAD CTE IN DIRECT CODE INSTRUCTION"

GPDIRECT

An operand in direct code could not be evaluated at compile time. The operation is printed following this message.

"BAD DIMENSION LIST"

PYLENGH

The dimension cannot be parsed as a compile time expression.

"BAD DIMENSION LIST, NO RIGHT BRACKET"

PYLENGH

Right bracket missing.

"BAD EXPRESSION AS TEMPORARY REGISTER WIDTH"

PYDCPHR

The width expression cannot be evaluated at compile time.

"BAD EXPRESSION FOR LOW ORDER WIDTH BIT"

PYWIDTH

The value is not a compile time expression.

"BAD PARAMETER LIST"

PYFORML

The right parenthesis is missing.

"BAD PARAMETER LIST, NO RIGHT PAREN"

PYPRCCL

32584-6015-RU-00

Right parenthesis missing for call statement.

"BAD PARAMETER"

PYPRCCL

A processor call parameter cannot be parsed as an expression.

"BAD REGISTER GROUP FOR DROP"

GRDROP

Probable compiler error.

"BAD REGISTER GROUP SELECTION FOR DROP"

GRDROP

Probable compiler error.

"BAD TEMPORARY REGISTER WIDTH TERMINATOR"

PYDCPHR

The ">" is missing.

"BAD TEMPORARY WIDTH"

PYDCPHR

The width definition for a temporary declaration cannot be evaluated at compile time.

"BAD VALUE FOR FINAL INDEX"

PYLENGH

Second index of a length declaration cannot be evaluated.

"BAD VALUE FOR HIGH ORDER WIDTH BIT"

PYWIDTH

The value cannot be evaluated at compile time.

"BAD VALUE FOR INITIAL INDEX"

PYLENGH

Initial index of length declaration cannot be evaluated.

"BAD VALUE FOR INSTRUCTION OPCODE"

PYRDOPC

The opcode for an OPDEF instruction is not a constant value.

"BAD VALUE FOR LOW ORDER WIDTH BIT"

PYWIDTH

The value cannot be evaluated at compile time.

"BAD VALUE FOR PARAMETER WIDTH"

PYRDOPC

A parameter width for an OPDEF instruction cannot be evaluated as a constant.

"CANNOT ALLOCATE SPECIFIED REGISTERS"

GRGTREG

Probable compiler error.

"CANNOT SWITCH TO ALTERNATE FILE FROM ALTERNATE FILE"

PYRDOPC

The copy option has been used in an OPDEF already on an alternate input file.

"CONTAINER TOO SMALL"

GFFREDY

32584-6015-RU-00

Probable compiler error.

"CONTROL STORE OVERFLOW"

AADRIVE

Data allocation is greater than the specified size of control store on the QM-1.

"COPY FILE NAME NOT PROVIDED"

PYRDOPC

An OPDEF clause specifies definitions to be copied from a file but the file name is not provided.

"DECLARE AFTER EXECUTABLE STATEMENT"

PYRDMME

Within a processor or macro means clause all data declarations must precede the executable statements.

"DEFAULT DECLARATION ALLOWED IN MAIN PROCESSOR ONLY"

PYDCPHR

Self explanatory.

"DEFAULT MAY NOT BE EXTERNAL"

PYDCPHR

DEFAULT is a special declaration of default attribute.

"DEFINED ITEM DIRECTION INVALID"

AACMPTE

Bit ordering is not the same as the parent.

"DEFINED LENGTH PHRASE INVALID"

AADIMPA

Subscript specification is not within the range of the parent.

"DEFINED WIDTH INVALID"

AACMPTB

The DEFINED width is inconsistent with the width of the subject data item.

"DEFINED WIDTH PHRASE DELETED"

AADIMPA

The width of the defined element is the same as the parent.

"DEFINITION PHRASE ILLEGAL"

PYDFPH

The parent for a defined phrase cannot be an external, constant, clock, or the default declaration.

"DIMENSIONED TEMPLATES INVALID"

AAPASAT

"DATA" declarations cannot be dimensioned.

"DIMENSIONED VARIABLE SUBSCRIPT NOT PRESENT AT LINE"

PMEXP

Reference to dimensioned data item is not subscripted.

"DIRECT CODE DATA ILLEGAL - USE SMITE DECLARES"

PYDIRCD

Data may not be defined within direct code. Declarations of temporaries or constants must be used.

"DUPLICATE DATA ITEM DECLARATION"

PYDCPHR

Item declared twice within the same scope.

"DUPLICATE LABEL"

PMOPEN

The label has been defined twice within the same scope.

"DUPLICATE PARAMETER IN ARG LIST"

PYGTARG

Each element of a processor formal argument list should be unique.

"DUPLICATE PROCESSOR NAME"

PYPRCHD

The same name has been defined twice within the same scope.

"END MISSING, BUT ASSUMED"

PYENDIT

Context blocks are not properly nested.

"ERROR IN SUPPLIED DECODE LENGTH AT STMT"

PMDECOD

The optionally provided decode length does not match the number of statements provided.

"ERROR MAXIMUM CONTROL STORE SIZE EXCEEDED"

MMPASS3

The assembler has exceeded the specified control store limit of the QM-1.

"ERRORS EXPANDING MACRO NO."

PYEXPCL

Error summary for a single macro expansion. The macro no. is assigned by the compiler in the order macros are defined.

"ERRORS IN COMPILATION"

SECLNUP

Summary of errors during compiler execution

"FORMAL PARAMETER MUST NOT BE DEFINED ON ANOTHER ITEM"

AAFORPA

Processor parameter is illegally declared.

"FORMAL PARAMETERS MUST NOT BE DIMENSIONED"

AAFORPA

Illegal processor parameter declaration.

"FUNCTION AT LINE***DOES NOT HAVE A VALUE"

PMPROCE

The function processor does not return a value.

"FUNCTION RESULT WIDTH TOO LARGE"

AAFWITH

Function width is greater than 72 bits.

"FUNCTION WIDTH UNDEFINED"

AAFWITH

A function must define the width of the returned value. The width may be the default.

"GCALEX:PARAMETERS GIVEN FOR EXTERNAL PROC"

GCALEX

Probable compiler error.

"GECALL:CALL TO UNKNOWN PROCESSOR"

PECALL

Probable compiler error.

"GECONCT:CONCATENATE EXPRESSION TOO WIDE"

GECONCT

The result of a concatenate is greater than 72 bits.

"GEEEXTRA:GESUFIT:EXTRACT DIRECTIONS REVERSED"

GEEEXTRA

Illegal extract. Bit ordering does not match source of extract.

"GEEEXTRA:GESUFIT:EXTRACT SUBFIELD TOO WIDE"

GEEEXTRA

Illegal extract.

"GEEEXTRA:GESUFIT:INVALID EXTRACT OPERAND (HIGH)"

GEEEXTRA

Illegal extract.

"GEEEXTRA:GESUFIT:INVALID EXTRACT OPERAND (LOW)"

GEEEXTRA

Illegal extract.

"GEEXTRA:PROCESSOR WITH NO WIDTH USED IN EXTRACT" GEEXTRA

Illegal use of extract on non-function processor.

"GEMKTMP:TEMP TABLE FULL" GEMKTMP

Compiler size limit exceeded.

"GEPARMS.ACTUAL PARAMETER TOO WIDE" GEPARMS

Parameter from a call is larger than that expected by the processor.

"GEPARMS.TOO FEW ACTUAL PARAMETERS" GEPARMS

A processor call has too few parameters.

"GEPARMS.TOO MANY ACTUAL PARAMETERS" PEPARMS

A processor call has too many parameters.

"GESCAT1:CONCATENATE EXPRESSION TOO WIDE" GESCAT1

Result of a concatenate greater than 72 bits

"GESEXT1:GESUFIT:EXTRACT DIRECTIONS REVERSED" GESEXT1

Illegal extract. Bit ordering does not match source for extract.

32584-6015-RU-00

"GESEXT1:GESUFIT:EXTRRACT SUBFIELD TOO WIDE" GESEXT1

Illegal extract.

"GESEXT1:GESUFIT:INVALID EXTRACT OPERAND (HIGH)" GESEXT1

Illegal extract.

"GESEXT1:GESUFIT:INVALID EXTRACT OPRAND(LOW)" GESEXT1

Illegal extract.

"GESEXT1:PROCESSOR WITH NO WIDTH USE IN EXTRACT" GESEXT1

Illegal use of non-function processor in an extract.

"GESPAS1:UNRECOGNIED STORE OPERAND" GESPAS1

Probable compiler error.

"GESREG1:REG TREE NOT REGISTER" GESREG1

Probably compiler error.

"GESVASUB:SUBSCRIPT OUT OF RANGE" GESVSUB

Subscript not within the range declared for the data item.

"GEVAR:GEVTEMP:TEMP USED BEFORE BEING SET"

GEVAR

Reference to a user/compiler temporary item which has not previously been given a value.

"GEVCNST:SYMBOLIC CONSTANT HAS NO VALUE"

GEVCNST

Declared constant does not have a value.

"GEVSUBS.SUBSCRIPT OUT OF RANGE"

GEVSUBS

Subscript not within the range declared for the data item.

"ID IN CONSTANT VALUE NOT A CONSTANT ITSELF"

PYDCPHR

Declaration of a constant value symbolically references another value which has not been defined as a constant.

"ILLEGAL ARRAY DEFINITION"

PYNODIM

The declared data class cannot be dimensional (e.g., externals, light, port, switch).

"ILLEGAL BINARY OP ON LEFT SIDE OF TRANSFER OP AT LINE"

PMEXP

An expressions not valid as the target of a store.

"ILLEGAL CLASS FOR FORMAL PARAMETERS"

AAFORPA

The formal parameter does not have a class or its class is of type processor or "DATA".

"ILLEGAL CONTAINER"

GRALLOC

Probable compiler error.

"ILLEGAL CTE IN REQUIREMENTS FOR MACRO"

PYCMPSR

The compiler time expression could not be evaluated at the point the macro was expanded.

"ILLEGAL DATA ITEM NAME"

PYDCPHR

The data item name does not begin with a letter.

"ILLEGAL DATA ITEM TARGET OF TRANSFER OP AT LINE"

PMEXP

The data item is on the right side of a transfer operation.

"ILLEGAL DECLARATION OF A RESERVED WORD"

PYDCPHR

SMITE reserved words cannot be declared as data items.

"ILLEGAL DIMENSION REFERENCE ON A CONSTANT AT LINE"

PMEXP

Constants are not dimensional.

"ILLEGAL KEYWORD FOR STATEMENT MACRO"

PYRDMTM

The statement macro keyword cannot be the same as a SMITE statement keyword.

"ILLEGAL NAME IN MEANS CLAUSE"

PYRDMME

Attempt to declare a SMITE reserved word in a macro means clause.

"ILLEGAL NAME IN MEANS CLAUSE"

PYRDMME

Attempt to declare a SMITE reserved word in a macro means clause.

"ILLEGAL PARAMETER"

PYGTARG

The parameter is a SMITE or macro defined token.

"ILLEGAL PARENT DATA ITEM NAME"

PYDFPH

The parent data item name is a token.

"ILLEGAL PRIMITIVE REFERENCE"

PYREFER

Right paren missing from a parenthesized expression.

"ILLEGAL PRIMITIVE"

PYPRIM

Expected primitive could not be parsed.

"ILLEGAL REGISTER BLOCK"

GRSETRG

Probable compiler error.

"ILLEGAL STATEMENT"

PYSTAT

The statement is not recognized by the parser.

"ILLEGAL STORE OF FUNCTION NAME AT STATEMENT"

PMEXP

The function is on the left side of a store as well as not having parameters.

"ILLEGAL TARGET FOR ESCAPE AT LINE"

PMSCAP

Cannot jump completely out of an IN or parallel context.

"ILLEGAL TEMPLATE TOKEN"

PYRDMTM

Macro templates cannot contain any form of end, declare, debug, ":" or ";".

"ILLEGAL TOKEN"

PLNXTOK

Parser lexical analyzer cannot recognize the token.

"ILLEGAL USE OF PRIMITIVE NONTERMINAL 7"

PYPRNTM

Compiler error.

"ILLEGAL USE OF RESERVED WORD AS PARENT"

PYDFPH

Illegal name for a defined phrase of a declaration.

"ILLEGAL USE OF RESERVED WORD, IN ARG LIST"

PYGTARG

A reserved word may not be an argument.

"ILLEGAL USE OF STATEMENT NONTERMINAL 7"

PYSTNTM

Compiler error.

"ILLEGAL USE OF STATEMENT NONTERMINAL 6"

PYSTNTM

Compiler error.

"IMPROPER PARALLEL FORK TEMPORARY STACK NESTING"

GRCFORK

Probable compiler error.

"IMPROPERLY NESTED DO STATEMENT"

PYSTNTM

Self explanatory.

"IMPROPERLY NESTED IN STATEMENT"

Self explanatory.


F/G 9/2

NL

3 OF 3

255

08774



MICROCOPY RESOLUTION TEST CHART
NBS 1963-A

"INCORRECT END LABEL IGNORED"

PYENDIT

The label is wrong or the context blocks are not properly nested.

"INCORRECT END STATEMENT ACCEPTED"

PYENDIT

Incorrect form of END used or the context blocks are not properly nested.

"INCORRECT NUMBER OF OPERANDS FOR -"

PYDIRCD

Direct code instruction has wrong number of operands.

"INCORRECT NUMBER OF SUBFIELD ELEMENTS"

AADIMPA

Subfield elements of a DEFINED phrase will not properly fit in the declared item. Boundaries of the declared data item may not be spanned by a subfield.

"INVALID ATTRIBUTE REQUIREMENT"

PYRDMWH

A macro attribute requirement is not a compile time expression.

"INVALID CALL ON LEFT SIDE OF TRANSFER OP AT LINE"

PMEXP

Attempt to store into a function processor call. Undeclared data items are assumed to be function calls which may be the source of the error.

"INVALID DECLARATION"

PYDCPHR

The declaration follows executable statements.

"INVALID DEFINED LENGTH PHRASE DELETED"

AAUDIMP

An invalid length specification is ignored and the definition of the parent used.

"INVALID DEFINED SUBSCRIPT"

AADIMPA

The width of the defined element is the same as the parent.

"INVALID EXPRESSION NO SECOND OPERAND FOR BINARY OPERATION" PYEXP

Self explanatory.

"INVALID EXPRESSION"

PYEXP

The first term of an expected expression cannot be parsed.

"INVALID FACTOR"

PYFACT

The leading primitive of an expression element cannot be parsed.

"INVALID FACTOR, NO SECOND OPERAND FOR CONCAT"

PYFACT

The concatenation operation does not have a second operand.

"INVALID MACRO PARAMETER TYPE"

PYRDMTM

The macro parameter type is not ID, REFERENCE, or EXPRESSION.

"INVALID OPDEF OP CODE"

PYRDOPC

The opcode is not a constant.

"INVALID PARAM FOR MIN/MAX"

PYCTEPL

The parameter cannot be parsed as a compile time expression.

"INVALID PARAM FOR WIDTH/LENGTH, NO LEFT PARAM"

PYCTEPR

The parameter list cannot be found.

"INVALID PARAMETER LIST FOR MIN/MAX"

PYCTEPL

The parameter list cannot be found.

"INVALID SUBFIELD REFERENCE"

PYSUBFL

A subfield qualifier must be an id, not a token.

"INVALID TERM, NO FACTOR FOLLOWING UNARY OPERATOR"

PYTERM

Self explanatory.

"INVALID TERM, NO FACTOR"

PYTERM

Syntax for expected factor not found.

"INVALID TYPE ON SEMANTIC REQUIREMENT"

PYRDMWH

A macro type requirement is not one of the SMITE data types.

"INVALID UNARY OP ON LEFT SIDE OF TRANSFER OP"

PMEXP

Unary operators are not valid as the target of a store.

"INVALID USE OF DATA TEMPLATE AT LINE"

PMEXP

Attempt to use as a normal data item.

"INVALID VALUE FOR FINAL INDEX"

PYLENGH

The index is not a valid compile time expression.

"INVALID WIDTH DECLARATION"

PYDCPHR

The width is not a valid compile time expression or does not terminate with a ">".

"INVALID WIDTH/LENGTH, PARAM ELEMENT"

PYCTEPR

The parameter is a token, not a data item.

"LABEL FOUND IN EXPRESSION AT LINE"

PMEXP

A label is meaningless in an expression.

"LABEL MISSING ON END, BUT ASSUMED"

PYENDIT

Label at the beginning of a context (IN, DO, IF, CASE, DECODE, BEGIN) has not been used on the END statement.

"LABEL NOT ALLOWED ON THIS STATEMENT"

PYSTTRN

Labels are only allowed on context generating statements.

"LENGTH AND WIDTH INCONSISTENT"

AADIMPA

For DEFINED elements of declaration.

"LOCKED REGISTERS MAY NOT BE MOVED TO TEMPORARY STORAGE"

GRSAVE

Probable compiler error.

"MACRO *** CANNOT BE EXPANDED"

PYEXPMC

There were errors in definition of the nth macro. n is ***.

"MEANS CLAUSE NOT PROVIDED"

PYRDMME

Macro definitioin requires a MEANS clause.

"MIN/MAX PARAMETER LIST RIGHT PAREN MISSING"

PYCTEPL

Self explanatory.

"MISPLACED KEY WORD AT STATEMENT"

PMEXP

A SMITE keyword exists within an expression.

"MISSING OR INVALID MACRO TYPE"

PYRDMHD

STATEMENT or PRIMITIVE required in macro header.

"MISSING RIGHT WEDGE"

PYRDOPC

Opdef parameter width definition is missing a right wedge terminator.

"MULTIPLE CLOCK DEFINITION"

PYDCPHR

There is only one clock.

"NO ARGUMENTS FOR FUNCTION PROCESSOR"

PYFORML

A function processor must have formal parameters.

"NO MACRO TEMPLATE CLAUSE"

PYRDMTM

A template clause is required.

"NO REGISTER AVAILABLE FOR FTEMP"

GRSAVE

Probable compiler error.

"NO REGISTERS TO ALLOCATE"

GRGTREG

Probable compiler error.

"NO RIGHT BRACKET IN ARRAY REFERENCE"

PYSUBSC

Self explanatory.

"NO STORAGE CLASS"

AAPARNT

Class has not been provided for a declared data item.

"NO TEMPORARY REGISTER WIDTH"

PYDCPHR

The numer of registers for a temporary have not been declared.

"NO VALID DESCRIPTION"

PYDESCR

There is no main processor description.

"NON-FUNCTION PROCESSOR REFERENCE IN AN EXPRESSION AT LINE" PMEXP

The processor can only be CALLED.

"NUMBER ASSOCIATED WITH LENGTH TOO LARGE, ZERO USED" PYLENGH

The length is larger than that which can be held in 22 bits.

"NUMBER ASSOCIATED WITH WIDTH TOO LARGE, ZERO USED" PYWIDTH

The number will not fit in 22 bits.

"OPERAND TRUNCATION" GRMVREG

Probable compiler error.

"OUT OF ATTRIBUTE TABLE SPACE" GEGTATT

Compiler error.

"PARALLEL CONTEXT***HAS NO PROCESSOR REFERENCES" LPARLEL

Probable compiler error.

"PARENT CANNOT BE A PARAMETER" PYDEFNS

Another data item may not be defined on a processor parameter.

"PARENT CANNOT BE A TEMPORARY" PYDEFNS

Another data item may not be defined on a temporary.

"PARENT IS NOT DECLARED"

PYDEFNS

Parent for a defined data item is undeclared.

"PARENT/SUBFIELD CLASSES INCOMPATIBLE"

AAPASAT

Parent and subfield must have the same class or one register and the other memory.

"PRIMITIVE MACRO(S) DID NOT MEET SEMANTIC REQUIREMENTS"

PYCRPRT

The template for a primitive macro(s) has been recognized but requirements of the WHERE clause are not satisfied resulting in an undefined syntax.

"PROCESSOR REQUIRES AT LEAST ONE EXECUTABLE STATEMENT"

PYBODY

Self explanatory.

"REGISTER BLOCK ILLEGALLY LOCKED"

GRCLEAR

Probable compiler error.

"REGISTER TEMP ALREADY IN LOCAL STORE"

GRRETRV

Probable compiler error.

"REQUESTED CONTAINER SIZE GREATER THAN 4"

GRALLOC

Probable compiler error.

"REQUESTED CONTAINER TOO SMALL, COMPUTED SIZED USED" GRALLOC

Probable compiler error.

"REQUESTED RESIDENCE NOT SPECIFIED, ANY REGISTER USED" GRALLOC

Probable compiler error.

"REQUESTED WIDTH INADEQUATE, ACTUAL WIDTH USED" GFDWDTH

A data item used as a subscript may have an improper width.

"REQUIREMENT IS NOT FOR A PARAMETER" PYRDMWH

A macro requirement for type can only be for one of the parameters.

"RETURNS CLAUSE NOT ALLOWED WITH STATEMENT MACRO" PYRDMRE

Returns clause is: only applicable to primitive macro.

"RETURNS CLAUSE REQUIRED FOR PRIMITIVE MACRO" PYRDMRE

Self explanatory.

"RIGHT SIDE OF STORE TOO WIDE" GESTORE

Larger than target of the store.

"RTA BLOCK LENGTH 3"

GRSSETRG

Probable compiler error.

"RTA TABLE FULL NO MORE ENTRIES AVAILABLE"

GRGTRTA

Compiler size limit exceeded.

"SEMICOLON MISSING BUT ASSUMED"

PYRDOPS

The compiler procedes as if the semicolon were input.

"SIZE OF CLUMP IS 3"

GRCLUMP

Probable compiler error.

"STATEMENT DID NOT MEET SEMANTIC REQUIREMENTS"

PYCRSTT

The template for a statement macro(s) has been recognized but requirements of the WHERE clause are not satisfied resulting in an undefined syntax.

"STATEMENT MACRO DOES NOT BEGIN WITH A KEYWORD"

PYRDMTM

A keyword is required.

"STORE INTO FUNCTION WITH PARAMETERS AT STATEMENT"

PMEXP

A function reference is not valid as the target of a store.

"SUBFIELD TOO WIDE"

AACMPTB

The subfield definition for the DEFINED phrase of a declaration is wider than the parent.

"SUBFIELD/PARENT WIDTH DISCREPANCY"

AACMPTB

Discrepancy in the DEFINED phrase of a declaration.

"SUBSCRIPT PRESENT FOR NON-DIMENSIONED DATA ITEM AT LINE"

PMEXP

Self explanatory.

"TEMPORARIES ALLOWED WITHIN MACROS ONLY"

PYDCPHR

Temporaries along with direct code are legal within macros only.

"TEMPORARY ALREADY RELEASED"

GEDRTMP

Compiler error.

"TEMPORARY WIDTH IS NOT 1, 2 OR 4"

PYDCPHR

The number of QM-1 registers to be allocated to the temporary.

"TOO FEW SUBSTATEMENTS IN CASE"

GPCASE

Based upon the width of the CASE expression there are too few statements following.

"TOO FEW SUBSTATEMENTS IN DECODE"

GPDECOD

There are less statements in the decode than optionally specified by the user.

"TOO MANY PARALLEL FORKS FOR REGISTER ALLOCATION STACK" GROFORK

Probable compiler error.

"TOO MANY SUBSTATEMENTS IN CASE"

GPCASE

Based upon the width of the case expression there are too many statements following.

"TOO MANY SUBSTATEMENTS IN DECODE"

GPDECOD

There are more statements in the decode than optionally specified by the user.

"TOO MANY TOKENS IN MEANS CLAUSES - INCREASE PY_MNSZ"

PYRDMME

A compiler limit has been exceeded.

"UNDECLARED DATA ITEM NOT ALLOCATED"

AAITEM

The allocator ignores undeclared data items.

"UNDECLARED ITEMS--NOT ALLOCATED"

AAPWRAP

The allocator ignores allocation of undeclared items.

"UNDECLARED PROCESSORS"

PYCLOSE

Summary of processors referenced but not defined. In some instances these are declared data items because at the point of reference they are assumed to be processors which may be subsequently defined.

"UNDEFINED DATA ITEM IN STATEMENT"

PMEXP

The data item has not been declared.

"UNDEFINED TARGET FOR ESCAPE AT LINE **"**

PMESCAP

Label not provided an ESCAPE.

"UNKNOWN OPCODE IN DIRECT CODE"

PYDIRCD

The operation mnemonic cannot be recognized.

"UNRECOGNIZED ATTRIBUTE, DECLARATION IGNORED"

PYATTR

The declaration is not one of the SMITE data types.

"UNRECOGNIZED CONSTANT EXPRESSION"

GECONST

Probable compiler error.

"UNRECOGNIZED EXPRESSION"

GECONT

Probable compiler error.

"USE OF UNDECLARED FORMAL PARAMETER AT STATEMENT"

PMEXP

A processor parameter has not been declared.

"WIDTH AND FLAG CONFLICT"

PYDCPHR

A flag can only be one bit wide.

"WIDTH LIMIT EXCEEDED"

AAWIDTH

Data item width not declared.

"WIDTH NOT ALLOWED ON A CLOCK, DEFAULT TO <1:36>"

PYDCPHR

Clock width and bit numbering is fixed in the compiler.

"WIDTH NOT ALLOWED ON TEMPORARY DECLARATIONS"

PYDCPHR

Temporaries cannot have a defined bit width, only the number of registers.

"WIDTH NOT DEFINED"

AAWIDTH

Data item width not declared.

"WIDTH/LENGTH PARAM RIGHT PAREN MISSING"

PYCTEPR

Self explanatory.

"WRONG END STATEMENT"

PYRDMME

Macro definition is terminated with the wrong form of end statement.

Appendix F: FTSC Emulator Requirements Specification

F-1.0 SCOPE

This specification defines the requirements for the development of the Fault-Tolerant Spaceborne Computer (FTSC) emulator. A basic emulation of the Fault-Tolerant Spaceborne Computer will be developed at the simplex processor instruction level to provide bit-identical results except where code execution results are time dependent or where fault-tolerant features affect execution. The basic FTSC emulation will be implemented using a higher-order programming language or a macro-structured assembly language technique to provide easy separation of critical functions for later enhancement and for compatibility with both 18-bit and 36-bit versions of Nanodata equipment used to support the emulation. The FTSC emulation will be tested to demonstrate that the requirements have been met and to measure the specific performance achieved.

F-2.0 APPLICABLE DOCUMENTS

The following documents of the issue in effect on the date of May 3, 1976 form a part of this specification to the extent specified herein. In the event of conflict with the documents referenced herein and the contents of this specification, the contents of this specification shall be considered a superseding requirement.

SPECIFICATIONS

Military

CDRL AO11

For BFTSC Program, FO4701-75-C-0149, Computer Program Development Specification BFTSC Executive/Recovery Software Design.

CDRL A007

32584-6015-RU-00

Brassboard Fault Tolerant Spaceborne Computer (BFTSC)
(FO4701-75-C-0149) Configuration Item Development
Specification BFTSC Architecture Design.

SAMSO SS-SDSF-01.0

Space Data Systems Facility Project Plan

Reports

ER75-4400

Users Programming Information for the Brassboard Fault
Tolerant Computer.

System Specification for Interim Operational Configuration
Digital Logic Emulation System

Nanodata, Programmable Run-Time Operator Display Users
Guide

Nanodata, Task Control Program Overview

F-3.0 REQUIREMENTS

The FTSC instructions will be emulated to provide bit for bit agreement with the operation of the FTSC computer. Two memory modules and the active CPU will be simulated. The FTSC I/O operations will emulate the transfer of data and commands between the CPU and I/O ports. The simulation of the actual external I/O devices may be achieved through user-coded routines called by the emulator. The FTSC emulator will provide the necessary linkage to add user-coded device simulations. The interrupt logic of the FTSC will be emulated, although not all the interrupt conditions may be generated on the emulation. The FTSC execution time clock will be approximated by a count of the number of FTSC instructions executed. This simulated clock will be scaled at 120,000 instructions per second, which is the average instruction rate of the Brassboard Fault Tolerant Spaceborne Computer.

F-3.1 Computer Program Definition

F-3.1.1 CPU Emulation

The FTSC emulation will provide a functional simulation of all the FTSC instructions except those which are fault related. Table F-1 defines the FTSC instruction set to be implemented. The STH, LMO, SBAO, SBAZ, SBDC, SBDZ instructions are considered to be fault related and will be treated as no-op instructions. These instructions either induce fault effects, are used solely to recover from fault conditions, or are used only in the monitor mode. The remaining instructions noted in Table F-1 will be emulated to perform the operations described by References ER75-4400 and TBD.

The emulation will maintain all quantities which are apparent to the FTSC software, providing bit-for-bit agreement between emulated instructions with the operation of FTSC hardware (excluding time dependent or fault related sequences). Internal FTSC quantities which are not directly available to the FTSC software will not be maintained. These include the internal busses, working registers, data encoding, and timing/response lines. The FTSC emulation will maintain the following items:

- 2 modules of Main Memory
- 8 General Purpose Registers
- Program Counter
- Extension Register
- Overflow Status
- Carry out Status
- Interrupt Enable/Disable Status

The normal operation of the active FTSC processor will be emulated. No hardware faults will be simulated.

The required operation of the emulation in response to certain anomalies which may be classified as software faults is described in Table F-2. This table defines the emulator processing when conditions occur which do not have a well defined result.

F-3.1.2 Input/Output (I/O) Emulation.

The FTSC emulator shall provide a functional emulation of the same number of serial interface ports and direct memory access channels as the Brassboard Fault Tolerant Spaceborne Processor. Data shall be passed to or accepted from external, user supplied, peripheral device simulation routines through shared blocks of QM-system memory (I/O ports). Calls to these external user routines, which are referred to as User I/O simulation routines or external interface events, may be initiated by the TASK/PROD system controlling the emulator to facilitate user initiated external I/O interrupts and, periodically, by the emulator.

F-3.1.2.1 Serial Interface Subsystem.

The serial interface subsystem emulation shall transfer data directly between the simulated FTSC memory and the simulated serial I/O ports, without attempting to mimic the timing and sequencing of signals on the FTSC serial interface buss. Serial I/O ports will simulate the data transfer between the FTSC Device Interface Units and the peripheral devices.

F-3.1.2.2 Device Initiated Interrupts.

The emulator shall examine the interrupt request signal at each serial I/O port following each external interface event. The higher priority serial I/O interrupt or other processing, as applicable, will be initiated at this time.

F-3.1.2.3 Block Transfers.

Upon initiation by the emulator, block transfers of data between the simulated serial I/O ports and the simulated FTSC memory will proceed until an acknowledge signal is required from one of the ports. The emulator will then resume instruction emulation until the next external interface event. Following each external interface event the transfer of I/O data will proceed until another acknowledge signal is required. Simulated serial block transfers shall proceed in this intermittent manner until the end of the block, at which time the serial I/O end of block interrupt shall be initiated.

F-3.1.2.4 Direct Memory Access Module.

The FTSC emulator shall emulate a DMA port similar to those provided for serial I/O, with request - acknowledge communications carried out through a set of status flags similar in function to the control lines to the FTSC. The simulated DMA port will provide for the transfer of up to (TBD) words of information between the emulator and the external routines at each external interface event. Once a simulated DMA transfer has been initiated, up to (TBD) words will be transferred to each external interface event until the requested number of words have been processed, at which time a DMA end of block interrupt will be initiated.

F-3.1.2.5 I/O Status Words.

The I/O status words, soft addresses F440 through F445, shall be maintained as applicable to permit FTSC software to monitor the status of an I/O operation in progress.

F-3.1.3 Interrupts.

All ten FTSC interrupt levels shall be supported. The following table defines the initiating conditions for each interrupt.

Fault

TBD

Power Down/Illegal OP Code

Illegal OP Code detected during instruction decode.

Arithmetic Fault

As defined for disallowed FTSC arithmetic operations.

Real Time interrupt

Initiated each (TBD)th FTSC instruction.

SIU (General)

Initiated upon user request.

DMA1 (General)

Initiated upon user request.

DMA2 (General)

Not emulated.

SIU (End of block)

Serial I/O block transfer word count satisfied.

DMA1 (End of block)

DMA word count satisfied.

DMA2 (End of block)

Not emulated.

F-3.1.4 Deferred Emulation Items.

Any features described in the performance and interface description sections of reference 1 are beyond the scope of the initial FTSC emulation. The emulation shall be designed to add these capabilities at a later time. The following deferred items are discussed in the order they appear in Reference CDRL A007.

F-3.1.4.1 BFTSC System Limitations.

The FTSC emulation will emulate the active CPU and the two active memory modules. The Configuration Control Unit shall not be simulated because it is used only to control the system during fault conditions. The Reconfiguration Read Only Memory is used only for hardware recovery and so is beyond the scope of this effort. The Bus Arbiter is an internal computer mechanism which shall not be considered in the functional level emulation. The Power Module, Circumvention Unit, and Timing Module shall not be simulated. The internal FTSC buses shall be emulated only to the extent that they are needed to perform the functional emulation.

F-3.1.4.2 Internal Bus Configuration.

The error code appended to the data and address words shall not be simulated. The byte rippler and cyclic

code generation logic are outside of the scope the functional emulation. The control lines between the CPU and the memory modules are to be emulated to the extent that an invalid address is recognized.

F-3.1.4.3 Fault Recovery and Reconfiguration.

Section 3.2.3 of CDRL A007, Fault Recovery and Reconfiguration, is beyond the scope of the initial FTSC emulation.

F-3.1.4.4 Functional Characteristics.

The FTSC emulation shall emulate the functional characteristics of the FTSC computer with certain limitations. The FTSC emulation by itself will generate the conditions to trigger only Illegal OP Code, Arithmetic, and Real Time interrupts. The user I/O simulation routines may cause the emulator interface routine to generate an I/O interrupt request. The Real Time interrupt shall be generated after a specified number of FTSC instructions have been executed as an approximation to the actual timing. The FTSC working registers and control registers shall be simulated only to the extent that they are required for the functional emulation. Their values need not be kept unless an FTSC instruction uses that variable explicitly.

F-3.1.4.5 System Modules

Most of the items described in section 3.2.5 of CDRL A007, System Modules, are beyond the scope of the initial effort. The internal busses, special function decoders, working registers, and microprogram operation are beyond a functional level emulation. The CCU, CPU/Bus state sequencer, Hardware Status Word generator, and watchdog timer will not be simulated. The fault interrupt, fault category, reconfiguration control, CCU input, and CCU output shall not be simulated since they are fault related items.

F-3.1.4.6 System Software

The FTSC emulation is not required to execute fault recovery or fault testing FTSC software. The emulation is not required to execute FTSC software requiring an interface to an I/O device for which no

appropriate user device simulation is linked to the emulator.

F-3.1.5 FTSC Emulation Interface

The initial FTSC emulation will interface directly or indirectly with several other software elements. As the emulation task on the QM-1, it must interface with the TASK/PROD operating system. The object code generated by the SMITE compiler must be loaded into the QM-1. The FTSC program to run on the emulator must be read from a file into the simulated memory region on the QM-1. The FTSC emulator must also communicate with the Interim Control System and user programs written to supply input/output device simulations. This interface is illustrated in Figure F-1.

F-3.1.5.1 SMITE Interface

A CDC CYBER program will process the output of the SMITE compiler to generate a magnetic tape containing the control store image of the emulator code and flags needed to load and run the emulator. A function will be added to the QM-1 operating system to read the SMITE tape and initiate the emulator.

F-3.1.5.2 FTSC Assembler/Loader Interface

A CDC CYBER program will be developed to process the output of the FTSC Assembler/Loader (ER75-4400) and create a magnetic tape which can be read into the QM-1. The format of this tape shall be made acceptable to the PROD function LOADMT. No modifications to the QM-1 operating system are required for the FTSC memory load interface.

F-3.1.5.3 QM-1 Operating System Interface

The QM-1 emulation system has three inter-related elements. The TASK system controls the scheduling, physical I/O operations, interrupt processing, and PROD/emulator communication. The PROD system interfaces the emulation with the operator, contains debugging features, display routines, simulation control commands, and the means of communicating with the I/O device simulation routines. The emulator must interface with the TASK and PROD systems, particularly regarding the location of interface and

display variables and the protocol for transferring control.

The changes to TASK and the basic PROD system shall be minimized to avoid problems at the time of TASK/PROD system updates and the later transition to the SDSF site. Required operating system changes should be implemented by extensions to the PROD system. The following operator commands shall be added for the FTSC emulation:

PC

Set FTSC program counter

LSMITE

Load SMITE object code into QM-1 control store

DSMITE

Dump control store onto magnetic tape

REG

Modify an FTSC register

PORT

Modify an FTSC simulated port

STAT

Modify an I/O status line

WP

Set Write Protect Boundary in a memory module

An FTSC state display routine will be developed to output the following quantities to the extent that they are needed to debug the FTSC emulator:

FTSC Program Counter

FTSC Instruction Register

8 FTSC General Purpose Registers

FTSC Status as applicable

Instruction Count

Operand Address

Operand Value

The layout of the state display is shown in Figure F-2.

A TASK/PROD routine shall be developed to extend the basic PROD capabilities for storing and displaying simulated memory, stepping the target program through a complete emulated instruction, and breakpointing an emulated instruction word for the FTSC emulation. The emulator and the QM-1 system shall be consistent about the location of interface variables and external calling sequences so that the basic PROD capabilities work. A control procedure shall be developed to assure that the SMITE code satisfies these interface requirements.

F-3.1.5.4 User Routine Interface

The user I/O simulation routine shall act as a PROD subroutine called by the emulator. The FTSC effort shall document PROD register conventions required for the user routine, PROD utility routines available to the user, conventions for using the PROD utility routines, and the emulator linkage conventions.

Appendix G: Recommended SMITE Coding Conventions and Standards

This appendix documents standards and conventions TRW has found useful in the development of SMITE computer descriptions and emulations. Two separate issues are presented: the development and maintenance of the emulator data base, and standards of coding style.

G.1 Data Base Development and Maintenance

Names chosen for identifiers in the computer description should conform to the naming established for the computer being described whenever possible. Where additional items are required, such as temporary registers not available to the target machine assembly language programmer, they should be named to be descriptive of their function.

Commentary should also be used to define the exact function of a variable when appropriate. Examples of the suggested commentary may be found in the FTSC description, Appendix H.

The size of the microcode may be reduced by declaring the most used (i.e. referred to by the most number of statements) variables first in the description so that the compiler allocates them to the first 64 words of QM-1 memory. This allows the compiler to use the short LDI instruction when generating address references, rather than the longer LDN instruction.

When a DEFAULT declaration is used, it should be the first data item declared in the main processor.

G.2 Coding Style

Each processor should be headed or prefaced by a sequence of descriptive and explanatory comments. These comments should address the topics found in MIL-STD-490-B5 (Paragraph 3) documentation; indeed, for the FTSC description, the relevant

portions of the MIL-STD-490 specification was coded directly into the description. As a minimum, a well commented processor description should include the following:

1. The inputs to and outputs from the processor, taking particular notice of the type and format of the variables.
2. A description of the use made of any global data base variables referenced or set.
3. A functional description of the processing performed within the processor, including error conditions returned or fielded.

The number of comments and the degree of detail in the comments is a function of individual style and the complexity of the computer being described. For example, the Intel 8080 description (Appendix C) contains relatively few comments due to the simplicity of the machine and the resulting clarity of the description. The FTSC, in comparison, is a much more complex computer, and therefore the description benefited from the heavy use of comments.

A primary goal of SMITE coding style should be the clarity and readability of the computer description. Writability should be a secondary, if not irrelevant, consideration. Therefore, the description should utilize indenting to delineate levels of declarations and control flow, and should incorporate blank lines between logically distinct units of the description of suggest the distinction present.

Indentation in declaration statements should be used to indicate relationships created by the use of the DEFINED phrase. For example,

```

DECLARE
  REGS[0:3]<15:0>,
    REGISTERS[0:7]<7:0> DEFINED REGS,
      B DEFINED REGISTERS[0],
      C DEFINED REGISTERS[1],
      D DEFINED REGISTERS[2],
      E DEFINED REGISTERS[3],
      H DEFINED REGISTERS[4],
      L DEFINED REGISTERS[5],
      STATUS DEFINED REGISTERS[6],
        CARRY FLAG DEFINED STATUS<0>,
        PARITY FLAG DEFINED STATUS<2>,
        AUX-CARRY DEFINED STATUS<4>,
        ZERO FLAG DEFINED STATUS<6>,

```

```

SIGN DEFINED STATUS<7>,
A DEFINED REGISTERS[7],
CARRY-A DEFINED REGS[3]<8:0>;

```

describes the main register set of the Intel 8080. The use of indentation clearly shows the use of subfields and the dependency of subfield definitions on their parents. The job of tracing a series of subfield definitions to the ultimate top-level parent is also simplified by this visual technique.

Subprocessor structure is also usefully shown by the selective use of indentation. For example, the (partial) structure of the Intel 8080 description is

```

INTEL-8080: PROCESSOR;
  ADD-REG: PROCESSOR;
    ADD-REG: END;
  GETADD: PROCESSOR;
    GETADD: END;
  .
  .
  .
INTEL-8080: END;

```

The above example also shows the suggested way for using indentation for PROCESSOR/END, BEGIN/END, CASE/END CASE, etc. Thus,

```

GETADD: PROCESSOR;

GETADD: END;

```

for processors,

```

BEGIN;

END;

```

for BEGIN/END,

```

CASE OPCODE;

END CASE;

```

for case statements, and similarly for the others. IF statements can be handled in this way as well:

```

IF INTERRUPT AND INTE
  THEN BEGIN;

```

```

    INTERRUPT <- INTE <- 0;
    IR <- INTERRUPT-IR;
    END;
ELSE BEGIN;
    IR <- MEM[PC];
    PC <- PC + 1;
    END;
END IF;

```

Context labels can be integrated easily into this scheme:

```

CODE10: BEGIN;
    OP-REG <- OPERAND;
    IN 3 PERFORM-OP;
CODE10: END;

```

or

```

JMP: IN 10 BEGIN;
    GETADD;
    PC <- OP-PAIR;
    END;

```

Finally, IN statements may be treated as prefixes to other statements, with no effect on indenting, as illustrated in the last example.

G.3 Pitfalls

1. Nestled Macro Usage

The syntax and semantics of a macro MEANS clause is not analyzed until the macro is expanded. The macro is simply read and stored as stream of tokens until expanded. If the macro, however, references another macro which has a new token in its template, the new token will not be recognized unless it has been defined first. For example:

```

STATEMENT 'MACRO1 PARM:ID'
MEANS
    PARM <- PARM ++ PARM:
END STATEMENT;

```

```
PRIMITIVE 'PARM:ID ++ PARM2:ID'
```

```
RETURNS PARM3 < 35Z
```

```
MEANS
```

```
PAM3 <- PARM1 + PARM1 + PARM2;
```

```
END PRIMITIVE;
```

Since "++" has not been defined at the point it is used, the tokens read and saved will be two "+"s. When expanded this naturally is an illegal statement.

To avoid possible token ambiguities as well as recursive macro expansion, a macro definition should always occur before any case of the macro.

2. Compiler Register Allocation

Certain data elements in Advanced SMITE are permanently allocated to QM-1 registers, although not always resident in the same register. These are declared temporaries, processor parameters, and function return values. Use of the QM-1 registers in this manner improves emulation performance and enables proper incorporation of direct code into a description. It does create, however, two possible pitfalls in using performance measurement and concurrency.

a. Performance Measurement

Elements which are not allocated fixed storage (i.e., are in registers) cannot be symbolically referenced in SASS performance measurement commands. Therefore, if one desired to sample the interface between two processors, the SMITE description should not model the interface using processor parameters. The interface could alternately be modeled as a global data element representing a bus (the SMITE data type register would be used). This data element will have a fixed storage location in the QM-1 which can be sampled.

b. Concurrency

Concurrency has been implemented in SASS as time sliced execution of the subtasks defined in the PARALLEL-BEGIN construct of the SMITE description. In this process, each subtask is supplied its own

32584-6015-RU-00

copy of the full set of QM-1 registers. When all subtasks are complete, the register state reverts to that prior to entrance into parallel execution. Therefore, any attempt to store into a parameter or function result from within a subtask will not provide expected results.

Appendix H: FTSC Description

```

##          IFFFFF  ITTTTT  SSSSSS  CCCCCC  ##
##          TT      SS      CC      CC      ##
##          TT      SS      CC      CC      ##
##          TTTT    SSSSSS  CC      CC      ##
##          TT      SS      CC      CC      ##
##          TT      SS      CC      CC      ##
##          TT      SS      CC      CC      ##
##          TT      SSSSSS  CCCCCC  ##
##
##          FAULT-TOLERANT SPACERORNE COMPUTER  ##
##
FTSCPROCESSOR:
DECLARE
  L1FAULT<310>  REGISTER,      ## NOMINAL ATTRIBUTES
  PAR1<3510>,      ## MERGED FLAG WORD 1
  STEP-ILAC  FLAG DEFINED PAR1<35>,  ## INTERFACE WITH STEP COMMAND
  APY-SIGN  FLAG DEFINED PAR1<32>,  ## SIGN OF MULTIPLY RESULT
  PLAM-SIGN  FLAG DEFINED PAR1<31>,  ## SIGN OF VALUE TO NORMALIZE
  AMZLNG  FLAG DEFINED PAR1<30>,  ## ADDRESS MODE ZERO FLAG
  SHIT-DIFFECTION FLAG DEFINED PAR1<29>,
  L2-SIGN FLAG DEFINED PAR1<28>,  ## SIGN OF EXTENSION REGISTER
  EXECUT  FLAG DEFINED PAR1<1C>,  ## SET TO 1 FOR LEFT SHIFTS
  INTERRUPT STUFF  ## EXECUTE INSTRUCTION FLAG
  ##
  JAPIL<31>  DEFINED PAR1<910>,  ## INTERRUPT REQUEST FLAGS
  FAULT-RFC FLAG DEFINED  PAR1<4>,## FAULT INTERRUPT REQUEST
  POWER-RFC FLAG DEFINED  PAR1<6>,## POWER DOWN INTERRUPT REQUEST
  ARITH-RFC FLAG DEFINED  PAR1<7>,## ARITHMETIC ERROR INT REQUEST
  RTIME-RFC FLAG DEFINED  PAR1<6>,## REAL TIME INTERRUPT REQUEST
  SIUG-RFC FLAG DEFINED  PAR1<5>,## SIU GENERAL INTERRUPT REQ
  DMA1C-RFC FLAG DEFINED  PAR1<4>,## DMA1 GENERAL INT REQUEST
  DMA2C-RFC FLAG DEFINED  PAR1<3>,## DMA2 GENERAL INT REQUEST
  SIUL-RFC FLAG DEFINED  PAR1<2>,## SIU END OF BLOCK INTERRUPT
  DMA1-RFC FLAG DEFINED  PAR1<1>,## DMA1 END OF BLOCK INT REQ
  DMA2-RFC FLAG DEFINED  PAR1<0>,## DMA2 END OF BLOCK INT REQ
  ##
  PAR2,
  STATUS<710>  DEFINED PAR2<25118>,  ## SECOND MERGED DATA WORD
  EN1-STAT FLAG DEFINED STATUS<7>,  ## STATUS REGISTER
  DIVCK  FLAG DEFINED STATUS<6>,  ## ENABLE INTERRUPT NETWORK
  OVERFLOW FLAG DEFINED STATUS<5>,  ## OVERFLOW STATUS
  ##

```

```

FLAG DEFINED STATUS<4>, ## ILLEGAL DPCODE STATUS
CARRY FLAG DEFINED STATUS<3>, ## CARRYOUT STATUS
INTLEV<210> INTLEV STATUS<210>, ## INTERRUPT LEVEL
PC<1210> PC DEFINED PAR2<151 0>, ## PROGRAM COUNTER

```

```

**                                     **
FISC REGISTERS                                     **
**                                     **
CPR1 1C177,                                     ** GENERAL PURPOSE REGISTER S
1X3C<3>1C>,                                     ** EXTENSION REGISTER 3C BITS
FX DEFINED FX3C<3>1C>,                         ** FISC EXTENSION REGISTER
UW1-1X<1>C> OFFIN0 FX36<3>3132>,              ** 2 OVERFLOW BITS - FISC MPY
IPK,                                             ** INSTRUCTION REGISTER
UFCUNE<6>C>   DEFINED INR<3112>>,** OPERATION CODE
FC<210>      DEFINED INR<24122>>-** OPERATON R0C0'PAKZAI WURU
JPKUC <610>  DEFINED PAP3<271118>, ** INTERRUPT IN PROCESS FLAGS
IN1MASK<01C> OFFIN0 PAP3< 91 0>, ** INTERRUPT MASK
**                                     **

```

```

**      WORKING STORAGE
**
**      LUP1RANDB3<3210>,      ** 33 BIT INSTRUCTION OPERAND
**      UPRAND<3110>  DEFINED OPERAND33<3110>,
**      ** INSTRUCTION OPERAND
**      IF1AD32<,      ** 32 BIT ADDRESS (INDIRECTS)
**      IF1AD<1710>  DEFINED EFFAD32<1710>,** OPERAND ADDRESS
**
**      M10-UP33<3210>,      ** GPXK(K8) OPERAND,RESULT
**      M10-UP<3110>  DEFINED MFG-OP33<3110>,
**
**      M10K11,      ** TEMPORARY 32 BIT CELL
**      M10K12,      ** TEMPORARY 32 BIT CELL
**      M10K13,      ** TEMPORARY 32 BIT CELL
**      ICOUNT,      ** FTSC INSTRUCTION COUNT
**      SHIF1-MASK33<3210>,      ** N BIT MASK, N=SHIFT COUNT
**      SHIF1-MASK<3110>  DEFINED SHIF1-MASK33<3110>,

```

```

P4<3510>,
SHIFT-COUNT<1710> DEFINED P4<351116>,, COUNTER FOR DBL SHIFT PASS **
PASS-COUNT<1710> DEFINED P4<1710>,, DBL SHIFT COUNTER (LE 32) **

P5<3510>,
OPERAND-EXP<1710> DEFINED P5<351116>,, EXPONENT OF REG-OP **
REG-OP-EXP<1710> DEFINED P5<1710>,, EXPONENT OF OPERAND **

P6<3510>,
NORMALIZE-COUNT<1710> DEFINED P6<351118>,
** NUMBER OF SHIFTS TO NORMALIZE **

P7<3510>,
TYPE<1710> DEFINED P7<351118>, ** WRITE PROTECT THIS MODULE **
PRTYPE<1710> DEFINED P7<1710>, ** BIT CODED WRITE PROTECT WORD **

** I/O STUFF **

IO-STATUS<0151>,
DMA1-STAT1 DEFINED IO-STATUS<01>,
DMA1-STAT2 DEFINED IO-STATUS<11>,
DMA2-STAT1 DEFINED IO-STATUS<21>,
DMA2-STAT2 DEFINED IO-STATUS<31>,
SIO-STAT1 DEFINED IO-STATUS<41>,
SIO-STAT2 DEFINED IO-STATUS<51>,

** MEMORY STUFF **

MEM<01122871> MEMORY: ** 3 MEMORY MODULES **

** SHIFT DATA STRUCTURES **

MEM<01122871> DATA,
MANTISSA<291 0> DATA DEFINED WORD<3116>,, I.P. MANTISSA **
EXPONENT< 71 0> DATA DEFINED WORD< 71 0>,, I.P. EXPONENT **
INDEX<1710> DATA DEFINED WORD<1710>, ** REGISTER INDEX VAL **

```

```

EXP-SIGN DATA DEFINED WORD<7>, ## EXPONENT SIGN BIT ##
EXP-CARRY DATA DEFINED WORD<8>, ## EXPONENT CARRY ##

SIGN-BIT DATA DEFINED WORD<31>, ## FTSC SIGN BIT ##
BIT-30 DATA DEFINED WORD<30>, ## FTSC BIT 30 ##
BITS-29-30<31:0> DATA DEFINED WORD<31:30>, ## FTSC BITS 30,29 ##
BITS-28-29<31:0> DATA DEFINED WORD<30:29>, ## FTSC BITS 30,29 ##
LOW-BIT DATA DEFINED WORD<0>, ## FTSC LOW ORDER BIT ##
BIT DATA DEFINED WORD<1>

```

```

##
DECLARE
  EXTERNAL PROCESSORS
  UP-STEP EXTERNAL, ## FTSC INSTRUCTION STEP ##
  UP-HALT EXTERNAL, ## HALT INSTRUCTION ##
  UP-ERROR EXTERNAL, ## INSTRUCTION ERROR ##
  UP-MOTSTM EXTERNAL, ## UNSIMULATED INSTRUCTION ##
  IP-PORT PORT, ## CRT INPUT ##
  OUT-PORT PORT, ## CRT OUTPUT ##
  PUTOUT EXTERNAL, ##
DECLARE ID-CONTROL EXTERNAL, ## TEMPORARY FOR NOW ##

```

3.2.2 OPERAND-FETCH

OPERAND-FETCH FORMS THE INSTRUCTION OPERAND AND THE OPERAND ADDRESS.

```

3.2.2.1 INPUTS -- ADDRESS CONTAINS THE 16 BIT ADDRESS FIELD
                  OF THE INSTRUCTION.
                  AM CONTAINS THE ADDRESSING MODE.
                  RA POINTS TO THE INDEX REGISTER.

```

```

3.2.2.2 PROCESS-- THE EFFECTIVE ADDRESS IS SET TO THE ADDRESS
                  FIELD OF THE INSTRUCTION.

```

A BRANCH BASED ON THE ADDRESSING MODE IS TAKEN.

C - REGISTER TO REGISTER
 OPERAND = CONTENTS OF RA REGISTER
 SET AMZPRO TRUE.

```

***
***      1 - IMMEDIATE
***          OPERAND = SIGN EXTENDED ADDRESS FIELD.
***
***      2 - DIRECT
***          OPERAND = CONTENT OF EFFECTIVE ADDRESS
***
***      3 - INDIRECT
***          EFFECTIVE ADDRESS = MEM(ADDRESS)
***
***      4 - INDEXED,POST INCREMENT
***          EFFECTIVE ADDRESS = RA REGISTER + ADDRESS
***          RA REGISTER IS INCREMENTED BY 1.
***
***      5 - INDEXED,PRI DECREMENT
***          RA REGISTER IS DECREMENTED BY 1.
***          EFFECTIVE ADDRESS = RA REGISTER + ADDRESS
***
***      6 - INDEXED
***          EFFECTIVE ADDRESS = RA REGISTER + ADDRESS
***
***      7 - INDEXED INDIRECT
***          EFFECTIVE ADDRESS = MEM(RA + ADDRESS)
***
***      FOR MODES 2-7, OPERAND = CONTENTS OF EFFECTIVE ADDRESS
***
***      3.2.2.3    OUTPUTS-- FFFAD IS THE EFFECTIVE ADDRESS
***                      OPERAND IS THE INSTRUCTION OPERAND.
***                      AP7FPO IS SET TRUE IF THE REGISTER-REGISTER MODE
***                      REGISTER MODIFIED IN MODE 4 OR 5 ##
***
***      COMMAND-FETCH PROCESSOR:
***          EFFECTIVE ADDRESS COMPUTATION
***
***          LITAU <- ADDRESS;
***          CASE AM:
***              BEGIN;
***                  CIPRAN <- CPXN(FAT);
***                  AP7FPO <- 1;
***              END;
***
***          LITD:
***              UPTRAPD <- SE(ADDRESS);
***              ## MODE 0 REGISTER-REGISTER ##
***              ## MODE 1 IMMEDIATE ##

```

```

OPERAND <- FETCH(EFFAD);      ** MODE 2 DIRECT ADDRESSING **
BEGIN:
    EFFAD32 <- FETCH(EFFAD);
    OPERAND <- FETCH(EFFAD);
END:

```

```

BEGIN:      ** MODE 4 INDEXED, POST INCR **
    EFFAD <- EFFAD + GPXR[KA].INDEX;
    GPXR[RA] <- GPXR[RA] + 1;
    OPERAND <- FETCH(EFFAD);
END:

```

```

BEGIN:      ** MODE 5 INDEXED, PRE DEC **
    GPXR[RA] <- GPXR[RA] - 1;
    EFFAD <- EFFAD + GPXR[KA].INDEX;
    OPERAND <- FETCH(EFFAD);
END:

```

```

BEGIN:      ** MODE 6 INDEXED **
    EFFAD <- EFFAD + GPXR[KA].INDEX;
    OPERAND <- FETCH(EFFAD);
END:

```

```

BEGIN:      ** MODE 7 INDEXED INDIRECT **
    EFFAD <- EFFAD + GPXR[KA].INDEX;
    EFFAD32 <- FETCH(EFFAD);
    OPERAND <- FETCH(EFFAD);
END:

```

END CASE:

OPERAND-FETCH: END:

PERFORM MEMORY FETCH REFERENCES

5.4.2.1 FETCH PROCESSOR(ADDRESS)

FETCH READS DATA FROM EMULATED MEMORY TO AN ASSIGNED VARIABLE

```

***
****
*****
*****
*****
3.2.2.1.1 INPUTS -- THE CALLING PARAMETER TO FETCH IS THE
FTSC MEMORY ADDRESS.
*****
3.2.3.1.2 PROCESS-- IF THE ADDRESS IS BETWEEN 0 AND #19],
THE DATA IS FETCHED FROM MEMORY.
IF THE ADDRESS IS NON-EXISTENT A VALUE OF
ZERO IS RETURNED.
IF THE ADDRESS IS IN MODULE F, THE SIMULATED
VALUE IS RETURNED (RESIDES AFTER
SIMULATED MODULE I).
IF THE ADDRESS IS AN I/O STATUS LOCATION,
STATUS WORD IS RETURNED
*****
3.2.3.1.3 OUTPUTS-- THE SIMULATED VALUE AT THE FTSC ADDRESS IS
LOADED INTO THE ASSIGNMENT VARIABLE.
I.E. THE USAGE OF FETCH IS AS FOLWS
VARIABLE <- FETCH(ADDRESS);
THIS RESEMBLES A FORTRAN FUNCTION **
*****
FETCH=PROCESSOR<310>(ADDR);
DECLARE ADDR<170>;
IF ADDR<150> < X#2000#
THEN FETCH <- MEM[ ADDR<150> ] ;
ELSE BEGIN;
IF ADDR<150> < X#F000#
THEN BEGIN;
FETCH <- 0;
END;
ELSE BEGIN;
FETCH <- MEM[ ADDR<110> + X#2000# ] ;
IF (ADDR<150> >= X#F440#) AND
(ADDR<150> <= X#F445#)
THEN FETCH <- IO-STATUS[ ADDR<20> ] ;
END IF;
END;
END IF;
END;
END IF;
END;
END IF;
RETCEND;

```

```

***
***      3.2.3.2      FETCH-MULTIPLE (NUMBER)
***
***      FETCH-MULTIPLE PERFORMS A VECTOR LOAD INTO THE
***      FTSC REGISTER SET
***
***      3.2.3.2.1    INPUTS -- THE CALLING PARAMETER INDICATES THE
***      LENGTH OF THE LOAD VECTOR
***
***      EFFAD IS THE EFFECTIVE OPERAND ADDRESS
***      AMZERO IS NON-ZERO FOR REGISTER-REGISTER MODE
***
***      3.2.3.2.2    PROCESS-- THE RESULT OF THE OPERAND FETCH PROCESSOR
***      IS MOVED TO GPXR(RB).
***      A LOOP IS PERFORMED TO MOVE THE REMAINING
***      ELEMENTS (NUMBER-1).
***      IN THIS LOOP THE PROCESSOR LOAD-NEXT IS CALLED
***      TO FETCH THE NEXT MEMORY/REGISTER VALUE
***      INTO REG-OP. THIS VALUE IS THEN MOVED
***      TO GPXR(RB).
***
***      3.2.3.2.3    OUTPUTS-- GPXR CONTAINS THE RESULT OF THE VECTOR LOAD
***      REG-OP CONTAINS THE LAST VALUE LOADED.
***      RA IS INCREMENT BY NUMBER-1
***      RB IS INCREMENT BY NUMBER-1
***      EFFAD IS INCREMENTED BY NUMBER-1
***
***      FETCH-MULTIPLE PROCESSOR(N)
***      UCLANE N<178> ;
***      GPXR(RB) <- OPERAND;
***      DO FOR 2 TO N ;
***      GPXR(RB) <- REG-OP <- LOAD-NEXT;
***      END;
***
***      3.2.3.3      LOAD-NEXT
***
***      LOAD-NEXT LOADS THE NEXT OPERAND VECTOR FROM
***      THE REGISTER FILE OR MEMORY
***
***      3.2.3.3.1    INPUTS -- EFFAD IS THE CURRENT EFFECTIVE ADDRESS

```

```

***      RA      IS THE CURRENT OPERAND REGISTER
***      AMZKND IS THE REGISTER-REGISTER MODE FLAG
***
3.2.3.3.2 PROCESS-- BUMP RA, RB, AND EFFAD BY 1
***      IF AMZKND IS SET, LOAD GPXR(RA)
***      OTHERWISE LOAD MEM(EFFAD)
***
3.2.3.3.3 OUTPUTS-- RA, RB, EFFAD ARE INCREMENTED BY 1
***      THE REGISTER/MEMORY OPERAND IS LOADED
***      INTO THE ASSIGNED VARIABLE.
***      I.e. VARIABLE C~ LOAD-NEXT;
***      REG-OP IS LOADED WITH NEXT REGISTER VALUE**
LOAD-NEXT;
      REG-OP <- GPXR(RA);
      IF AMZKND
      THEN REGIN;
      RA <- RA+1;
      LOAD-NEXT <- GPXR(RA);
      END;
      ELSE REGIN;
      EFFAD <- EFFAD + 1;
      LOAD-NEXT <- FETCH(EFFAD);
      END;
      END IF;
LOAD-NEXT;
      END;

3.2.3.4      LDN
***
***      LDN LOADS THE TWO'S COMPLEMENT OF THE INSTRUCTION
***      OPERAND INTO THE RESULT REGISTER AND TEST FOR OVERFLOW.
3.2.3.4.1 INPUTS -- OPERAND CONTAINS THE INSTRUCTION OPERAND.
3.2.3.4.2 PROCESS-- THE NEGATIVE OF OPERAND IS STORED IN REG-UP.
***      IF THE VALUE HAS ONLY THE SIGN BIT SET,
***      THEN SET-OVERFLOW IS CALLED.
3.2.3.4.3 OUTPUTS-- REG-OP CONTAINS THE RESULT.
***      OVERFLOW IS SET IF THE OPERAND = F0000000**
LDN;
      PROCESSOR;

```



```

***
*** STORE (VALUE,ADDRESS)
***
*** STORE PLACES THE PARAMETER VALUE IN THE SIMULATED
*** MEMORY ADDRESS INDICATED BY THE PARAMETER ADDRESS
***
3.2.3.6.1 INPUTS -- THE CALLING PARAMETERS INDICATE THE
*** VALUE AND THE MEMORY ADDRESS.
***
3.2.3.6.2 PROCESS-- IF THE ADDRESS IS IN MODULE F :
*** TYPE (WRITE-PROTECT) IS FORMED FROM
*** ONE OF THE 4 UPPER BITS OF MEMTYPE
*** IF THE ADDRESS FALLS BETWEEN F440 AND
*** F445, THE I/O CONTROL PROCESSOR IS CALLED
*** IF THE ADDRESS FALLS BETWEEN F404 AND F407,
*** INTMASK OR ENI-STAT ARE ALTERED.
*** THE ADDRESS IS MAPPED TO A OM-1 ADDRESS
***
*** IF THE ADDRESS DOES NOT EXIST, TYPE = 0
*** NO ERROR MESSAGE AT THIS TIME.
***
*** OTHERWISE, TYPE IS FORMED FROM MEMTYPE TO
*** DESIGNATE THE WRITE PROTECT STATUS OF
*** THAT QUARTER OF A MEMORY MODULE
***
*** THEN ALL PATHS REJOIN
*** IF TYPE IS NON-ZERO, THE INPUT VALUE
*** IS STORED AT THE OM-1 ADDRESS CORRESPONDING
*** TO THE INPUT ADDRESS.
***
3.2.3.6.3 OUTPUTS-- THE SIMULATED MEMORY ADDRESS IS CHANGED
*** UNLESS THE WRITE PROTECT STATUS IS SET.
*** INTMASK AND ENI-STAT ARE MODIFIED BY STORES
*** IN F404-F407
***
STORE PROC(SOF(VALUE),ADDR);
DECLARE
VALUE<31:0>,
ADDR<31:0>;
*****JNDIFX <- ADDR;
IF ADDR<31:0> = XXX#
THEN BEGIN ;
    TYPE <- SLL(1,ADDR<31:0>) AND MEMTYPE ;

```

```

**      WORK1<15:12> <- 2;
      ASSUMES F4XX IS NOT WRITE PROTECTED **
      IF ( ADDR<15:10> >= X#F440#)
      AND (ADDR<15:10> <= X#F445#)
      THEN BEGIN;
        ** SET UP VALUE AND ADDR AS PARAMETERS **
        IO-CONTROL;
      END;
      END IF;
      IF SRL( ADDR<15:10>-X#F440#, 2) = 0
      THEN BEGIN;
        CASE ADDR<15:10> ;
          INTMASK<7:0> <- VALUE<7:0> ;
          NULL;
          ENI-STAT <- 0 ;
          ENI-STAT <- 1 ;
        END CASE;
      END;
      END IF;
      END;
      ELSE BEGIN;
        IF ADDR<15:12> > 1
        THEN BEGIN ;
          TYPE <- 0;
        END;
        ELSE TYPE <- SLL(1,ADDR<13:11>) AND MEMTYPE ;
        END IF;
      END;
      END IF;
      IF TYPE /= 0
      THEN MEMT WORK1<15:10> ] <- VALUE ;
      END IF;
    SIGREPEND;

```

```

**      3.2.3.7  STORE-DOUBLE
***
***      STORE DOUBLE STORES A VALUE IN TWO ADJOINING MEMORY MODULES
***
***      3.2.3.7.1  INPUTS -- OPERAND CONTAINS THE DATA TO STORE.
***

```

```

***
***      EFFAD CONTAINS THE OPERAND (STORE) ADDRESS
***      ANZPRC IS 1 FOR THE REGISTER-REGISTER MODE
***
3.2.3.7.2 PROCESS-- THE PROCESSOR IS BYPASSED IN THE REGISTER TO
***      REGISTER MODE OF ADDRESSING.
***      THE STORE PROCESSOR IS CALLED TWICE
***      OPERAND IS THE VALUE PARAMETER
***      EFFAD, EFFAD+4096 ARE THE ADDRESS PARAMETERS
***
***      3.2.3.7.3 OUTPUTS-- MEM(EFFAD), MEM(EFFAD + 4096) ARE STORED
***      INTO BY THE STORE PROCESSOR
***
STORE-DOUBLE:  PROCESSOR
    IF NOT ANZPRC
        THEN REGIN:
            STORE(OPRAND, EFFAD)
            STORE(OPRAND, EFFAD + 4096)
        END:
    END:
STORE-DOUBLE:  END:

***
***      3.2.3.8 STORE-MULTIPLE (NUMBER)
***
STORE-MULTIPLE STORES A REGISTER VECTOR INTO TWO
***      ADDJNING MEMORY MODULES.
***
3.2.3.8.1 INPUTS -- THE INPUT PARAMETER NUMBER IS THE LENGTH
***      OF THE REGISTER VECTOR.
***      EFFAD IS THE STORE ADDRESS.
***      REG-OP INITIALLY EQUALS GPXK(RB).
***      3.2.3.8.2 PROCESS-- OPERAND AND REG-OP ARE LOADED WITH GPXK(RB).
***      STORE-DOUBLE IS CALLED.
***      RB AND EFFAD ARE INCREMENTED BY 1.
***      THIS LOOP IS PERFORMED THE NUMBER OF TIMES
***      INDICATED BY THE CALLING PARAMETER.
***
3.2.3.8.3 OUTPUTS-- RB IS INCREMENTED BY NUMBER-1.
***      EFFAD IS INCREMENT BY NUMBER-1.
***      REG-OP = GPXK LAST RB )

```

THE MEMORY VECTORS BEGINNING AT PIM(OLD EFFAD)
AND MEM(OLD EFFAD +4C96) ARE ALTERED

```

***
***
***
STORE-MULTIPLY: PROCESSOR(N);
  CLEAR( N*710 );
  OPERAND <- REG-OP;
  DO FOR 1 TO N;
    BIC( N );
    STORE-DOUBLE;
    EFFAD <- EFFAD + 1;
    IF <- N + 1;
    OPERAND <- REG-OP <- GPXK(KB);
  END;
STORE-MULTIPLY: END;

3.2.3.5. RAD-STORE
***
***
***
RAD-STORE EMULATES THE SBAD,SBZ,SBDO,SBDOZ INSTRUCTIONS
***
***
***
WHICH CAUSE MEMORY FAULT ERRORS TO OCCUR.

3.2.3.5.1 INPUTS -- AMZERO IS TRUE ONLY FOR THE REGISTER-REGISTER MODE.
***
***
***
REG-OP CONTAINS THE CONTENTS OF REGISTER KB

3.2.3.5.2 PROCESS-- THE DATA IS STORED IF THE REGISTER TO REGISTER
***
***
***
ADDRESSING MODE IS USED.
***
***
***
OTHERWISE, THE FAULT REQUEST IS SET.

3.2.3.5.3 OUTPUTS-- THE REGISTER FILE IS MODIFIED IN THE REGISTER
***
***
***
IF REGISTER MODE.
***
***
***
OTHERWISE, NO STORE OCCURS. **

BAD-STORE: PROCESSOR;
  IF AMZERO
    THEN GPXK(KB) <- REG-OP;
    ELSE FAULT-REQ <- 1;
  END IF;
BAD-STORE: END;

ADD: SUCCESS(ADDEND, ACN-FLAC);
  **REPLACE NEXT STATEMENT AFTER COMPILER BUG IS FIXED.**

```

```

**DECLARE ACC-FLAG FLAG**
ALG-UP3? <- (ACC-FLAG AND CARRY) +
  SE(ADDEND) + (REG-OP33 <- SE(REG-OP))
IF (CARRY <- REG-OP33<32>) /# REG-OP.SIGN-BIT THEN
  SET-OVERFLOW;
END IF;
ADD: END;

**
3.2.4.2 INTEGER MULTIPLY PROCESSOR **
MULTIPLY PERFORMS A 32 BIT BY 32 BIT MULTIPLICATION

3.2.4.2.1 INPUTS -- REG-OP CONTAINS ONE MULTIPLIER
  OPERAND CONTAINS THE OTHER.
  NORMALLY THESE REPRESENT THE INSTRUCTION
  OPERAND AND AN FISC REGISTER.

3.2.4.2.2 PROCESS-- THE ABSOLUTE VALUES OF THE MULTIPLIERS
  ARE FORMED. A 2 BIT AT A TIME
  MULTIPLICATION IS PERFORMED TO PRODUCE
  THE 64 BIT PRODUCT.
  IF NECESSARY, THE PRODUCT IS THEN NEGATED.

3.2.4.2.3 OUTPUTS-- THE MOST SIGNIFICANT 32 BITS OF THE RESULT
  IS STORED IN REG-OP. THE LEAST SIGNIFICANT
  32 BITS IS STORED IN EX.
  MPY-SIGN DENOTES THE SIGN OF THE RESULT.
  OPERAND IS LEFT UNCHANGED. **

MULTIPLY: PROCESSOR;
EX <- 0;
MPY-SIGN <- 0;
IF (OPERAND.SIGN-BIT
  THEN REG-OP.SIGN-BIT
  OPERAND <- - OPERAND;
  MPY-SIGN <- - 1;
  END;
END IF;
IF (FISC-UP.SIGN-BIT

```

```

      THEN BEGIN;
      REG-OP ← -REG-OP;
      MPY-SIGN ← NOT MPY-SIGN;
      IFD;
      END IF;
      DO 104 1 TO 16
      BEGIN;
        EX36 ← SLL (EX36,2);
        CASE REG-OP <30129> ;
        NULL;
        EX36 ← EX36 + OPERAND;
        EX36 ← EX36 + SLL ( OPERAND33,1 );
        EX36 ← EX36 + SLL ( OPERAND33,1 ) + OPERAND ;
      END CASE;
      REG-OP ← SLL (REG-OP,2 ) + OVF-EX ;
      END;
      IF MPY-SIGN
      THEN BEGIN;
        EX ← - EX;
        IF EX = 0
        THEN REG-OP ← - REG-OP;
        ELSE REG-OP ← NOT REG-OP;
      END IF;
      IFD;
      END IF;
      END;
      MULTPLY;
      END;

```

**

3.2.4.4 SHIFT PREPARATION PROCESSOR

SHIFT-PREP EXTRACTS THE SHIFT COUNT AND DIRECTION FROM THE INSTRUCTION OPERAND.

3.2.4.4.1 INPUTS -- OPERAND CONTAINS THE MAGNITUDE AND DIRECTION OF THE SHIFT

3.2.4.4.2 PROCESS-- THE SIGN OF THE LOWER 8 BITS OF OPERAND IS TESTED. A NEGATIVE VALUE INDICATES A LEFT SHIFT. THE MAGNITUDE OF THE LOWER 8 BITS FORMS THE SHIFT COUNT. THE SHIFT COUNT IS LIMITED TO 64 BITS

3.2.4.4.3 OUTPUTS-- SHIFT-DIRECTION IS 1 FOR LEFT SHIFTS

```

***
***
SHIFT-COUNT IS THE SHIFT LENGTH (LIMITED
TC 64)
**

SHIFT-PRGR: PROCESSOR:
SHIFT-COUNT ← SE(OPERAND.EXPONENT)
IF SHIFT-COUNT < 17
THEN BEGIN
  SHIFT-DIRECTION ← - 1
  SHIFT-COUNT ← - SHIFT-COUNT
  END
ELSE SHIFT-DIRECTION ← - 0
  END IF
  IF SHIFT-COUNT > 64
  THEN SHIFT-COUNT ← - 64
  END IF
SHIFT-PRGR: END :

**
**
3.2.4.5 DOUBLE LENGTH SHIFT PROCESSOR
DOUBLE-SHIFT PERFORMS THE FISC LONG SHIFTS

3.2.4.5.1 INPUTS -- SHIFT-DIRECTION IS 1 FOR LEFT SHIFTS.
SHIFT-COUNT IS THE SHIFT LENGTH.
REG-OP IS THE CONTENTS OF THE RB REGISTER
EX IS THE CONTENTS OF THE EXTENSION REGISTER
IN TWO PASSES, 32 BITS AND THE REMAINDER.
A MASK WORD IS FORMED OF LENGTH N, N =
THE SHIFT LENGTH FOR THIS PASS. LEFT
AND RIGHT BRANCHES ARE CONSIDERED
SEPARATELY.
THE BITS SHIFTED OUT ARE SAVED IN WORK1.
A LOGICAL SHIFT IS PERFORMED ON REG-UP,EX.
A BRANCHED BASED ON THE OP-CODE IS TAKEN
TO FIX UP THE RESULT FOR DIFFERENT SHIFTS.

3.2.4.5.2 PROCESS-- REG-OP,EX ARE THE RESULT.
SHIFT-COUNT, PASS-COUNT ARE USED AND ZEROED
OVERFLOW MAY BE SET FOR ARITHMETIC LEFT
SHIFTS
**

3.2.4.5.3 OUTPUTS--

```

```

UNPFL-SHIFT:      PROCESSOR;
SHIFT-PRIP:
  IF SHIFT-COUNT > 32
  THEN PASS-COUNT <- 32 ;
  ELSE PASS-COUNT <- SHIFT-COUNT ;
  IFD IF;
  DO UNTIL SHIFT-COUNT = 0 ;
  RECLP ;
    SHIFT-MASK33 <- SRAL D=4000000000000000, PASS-COUNT ) ;
    IF SHIFT-DIFFERENTIATION = 0
    THEN BEGIN ;
      WORK1 <- SLL( FX, 32 - PASS-COUNT) AND SHIFT-MASK;
      FX <- SRL(FX, PASS-COUNT) OR
        ( SLL (REG-OP, 32 - PASS-COUNT)
          AND SHIFT-MASK) ;
      REG-OP <- SRL(REG-OP, PASS-COUNT);
      CASE OPCODE <211> ;
        NULL ;
        NULL ;
        BEGIN ;
          IF SLL(REG-OP, PASS-COUNT) < 6
          THEN REG-OP <- REG-OP OR SHIFT-MASK ;
          END IF ;
        END ;
        REG-OP <- REG-OP OR WORK1;
      END CASE ;
    ELSE BEGIN ;
      WORK1 <- REG-OP AND SHIFT-MASK;
      REG-OP <- SLL(REG-OP, PASS-COUNT) OR
        SRL( EX, 32 - PASS-COUNT) ;
      FX <- SLL (FX, PASS-COUNT) ;
      CASE OPCODE <211> ;
        NULL ;
        NULL ;
        BEGIN ;
          IF REG-OP.SIGN-BIT
          THEN WORK1 <- WORK1 - SHIFT-MASK;
        END ;
      END CASE ;
    END ;
  END ;
  UNPFL-SHIFT:

```

```

      END IF ;
      IF WORK1 /= 0
      THEN SET-OVERFLOW ;
      END IF ;

```

```

      END ;
      EX <- EX OR SKL( WORK1, 32 - PASS-COUNT ) ;
      END CASE ;

```

```

      END ;
      END IF ;
      SHIFT-COUNT <- SHIFT-COUNT - PASS-COUNT ;
      PASS-COUNT <- SHIFT-COUNT ;

```

```

      END ;
      DOUBLE-SHIFT:      END;

```

```

3.2.5.1  FLOATING-PREP

```

```

      FLOATING-PREP SEPARATES REG-OP AND OPERAND
      INTO MANTISSA AND EXPONENT WORDS.

```

```

3.2.5.1.1 INPUTS -- REG-OP AND OPERAND

```

```

3.2.5.1.2 PROCESS-- THE SIGN-EXTENDED RIGHT-JUSTIFIED
      MANTISSA AND EXPONENT PARTS OF BOTH
      INPUTS ARE FORMED.

```

```

3.2.5.1.3 OUTPUTS-- OPERAND-EXP AND REG-OP-EXP CONTAIN THE EXPONENTS
      OPERAND AND REG-OP CONTAIN THE MANTISSAS

```

```

      FLOATING-PREP:      PROCESSOR:
      OPERAND-EXP <- X(COCCG) // (OPERAND.EXPONENT) ;
      REG-OP-EXP <- X(COCCG) // REG-OP.EXPONENT ;
      OPERAND <- SF (OPERAND.MANTISSA) ;
      REG-OP <- SF (REG-OP.MANTISSA) ;
      FLOATING-PREP:      END;

```

```

***
***      3.2.5.2  NORMALIZE-SHIFT (VALUE)
***
***      NORMALIZE-SHIFT PERFORMS A 32 BIT NORMALIZATION
***
***      3.2.5.2.1  INPUTS -- THE CALLING PARAMETER SPECIFIES THE
***                   VARIABLE TO BE NORMALIZED.
***
***      3.2.5.2.2  PROCESS-- THE VARIABLE IS SHIFTED ITERATIVELY
***                   UNTIL THE MOST SIGNIFICANT DATA BIT
***                   DIFFERS FROM THE SIGN BIT.
***                   THE CALLING ROUTINE MUST PROTECT AGAINST
***                   AN ALL-ZERO VALUE.
***
***      3.2.5.2.3  OUTPUTS-- THE NORMALIZED DATA WORD IS STORED BACK
***                   INTO THE CALLING PARAMETER WORD.
***                   NORMALIZE-COUNT IS THE NUMBER OF SHIFTS
***                   REQUIRED TO NORMALIZE THE VALUE
***                   NORM-SIGN IS THE SIGN OF THE VALUE
***
NORMALIZE-SHIFT:  PROCESSOR( VALUE )
DECLARE VALUE
NORMALIZE-COUNT <- 0
NORM-SIGN <- VALUE.SIGN-BIT
DO WHILE VALUE.PIT-30 = NORM-SIGN
BEGIN
    NORMALIZE-COUNT <- NORMALIZE-COUNT + 1
    VALUE <- SL(VALUE,1)
END
NORMALIZE-SHIFT:  END;

***
***      3.2.5.3  NORMALIZE
***
***      NORMALIZE PERFORMS A NORMALIZATION OF THE FLOATING
***                   POINT WORD IN REG-OP.
***
***      3.2.5.3.1  INPUTS -- REG-OP IS THE FLOAT POINT NUMBER TO
***                   NORMALIZE. THE FORMAT OF THIS WORD
***                   MUST BE A SIGN-EXTENDED RIGHT-JUSTIFIED
***                   WORD.

```



```

***
***
***
3.2.5.4.3 OUTPUTS-- REG-OP CONTAINS THE RESULT OF THE CONVERSION
***
***
***
PROCESSOR( IBASE )
  DECLARE IBASE<710>
  REG-UP <- OVERAND
  IF REG-UP = 0
    THEN REG-OP <- X*BC
  ELSE BEGIN
    NORMALIZE-SHIFT( REG-OP )
    REG-OP-EXP <- SET( IBASE<710> ) - NORMALIZE-COUNT
    TEST-FP-OVERFLOW
  END
  IFD IIB
  FND
  FLDAT

```

```

***
***
***
3.2.5.4 TEST-FP-OVERFLOW
  TEST-FP-OVERFLOW TESTS FOR FLOATING POINT OVERFLOW
  OR FLOATING POINT UNDERFLOW
  IT ALSO STORES THE EXPONENT.
  3.2.5.4.1 INPUTS -- REG-OP-EXP CONTAINS THE 18 BIT EXPONENT.
  3.2.5.4.2 PROCESS-- THE CARRY AND SIGN BITS ARE ANALYZED
  FOR AN OVERFLOW CONDITION.
  SET-OVERFLOW IS CALLED FOR OVERFLOWS.
  THE EXPONENT IS STORED AS THE EXPONENT OF
  REG-OP
  3.2.5.4.3 OUTPUTS-- REG-OP-EXP IS STORED IN THE EXPONENT PART
  OF REG-OP. OVERFLOW MAY BE SET.
  **
  TEST-FP-OVERFLOW PROCESSOR
  IF REG-UP-EXP. EXP-CARRY /= REG-UP-EXP. EXP-SIGN
    THEN SET-OVERFLOW
  END IF
  REG-UP-EXPONENT <- REG-OP-EXP.EXPONENT
  TEST-FP-OVERFLOW END

```

```

UNNORMALIZE: PROCESSOR;
  *ADJUST EXPONENTS OF FLOATING-POINT OPERANDS SO THAT
  THEY MATCH, AND SHIFT THE MANTISSAS ACCORDINGLY.*
  IF (WORK) <- OP-EXP - REG-OP-EXP > 0 THEN
    BEGIN
      REG-OP <- SPA(REG-OP, WORK<17:0>);
      REG-OP-EXP <- OP-EXP;
    END;
  ELSE
    OPERAND <- SPA(OPERAND, - WORK<17:0>);
    *OPERAND-EXP IS NOT NEEDED*
  END IF;
UNNORMALIZE: END;

```

```

FLOATING-ADD: PROCESSOR;
  IF (OPERAND-EXP > REG-OP-EXP + 23) OR
    (REG-OP = 0) THEN
    BEGIN
      REG-OP <- OPERAND;
      REG-OP-EXP <- OPERAND-EXP;
    END;
  ELSE
    IF (OPERAND-EXP > REG-OP-EXP - 23) AND
      (OPERAND /= 0) THEN
      REG-OP <- REG-OP-EXP + OPERAND;
      REG-OP-EXP <- REG-OP-EXP + OPERAND;
    ELSE
      REG-OP IS UNCHANGED.*
    END IF;
  END IF;
  NORMALIZE;
FLOATING-ADD: END;

```

```

***          1.2.6.2    FLOATING-MULTIPLY
***

```

```

***      FLOATING MULTIPLY MULTIPLIES TWO FTSC FLOATING
***      POINT OPERANDS.
***
3.2.6.2.1 INPUTS -- REG-OP CONTAINS THE REGISTER OPERAND.
                     OPERAND CONTAINS THE MEMORY OPERAND
***
3.2.6.2.2 PROCESS-- FLOATING-PRFP IS CALLED TO EXTRACT
                     THE MANTISSA AND EXPONENT OF EACH OPERAND
                     REG-OP IS LEFT SHIFTED 8 BITS (THIS REMOVES
                     NON-SIGNIFICANT BITS)
                     AN INTEGER MULTIPLY IS PERFORMED
                     THE PRODUCT OF REG-OP * OPERAND GOES TO REC-UP
                     IF THE RESULT IS ZERO, FP ZERO IS STORED
                     OTHERWISE, DBL-NORMALIZE (REG-OP,EX)*
***
FLOATING-MULTIPLY: PROCSSOR
    ;
    FLOATING-PRFP;
    REG-OP <- SII(NEG-OP,8);
    MULTIPLY;
    IF REG-OP = 0
        THEN FFC-OP <- 128
        ELSE BEGIN
            REG-OP-EXP <- OPERAND-EXP + REG-OP-EXP;
            DBL-NORMALIZE;
        END;
    END IF ;
    FLOATING-MULTIPLY: END
    ;

DBL-NORMALIZE:     PROCSSOR;
    IF SIII(FFC-OP + 1, 1) = 0
        THEN BEGIN
            REG-OP<30:0> <- FX<31:1>;
            FX <- 0;
            REG-OP-EXP <- REG-OP-EXP -31;
        END;
    END IF;
    NORMALIZE-SHIFT (REG-OP);
    REG-OP <- REG-OP OP SIII(FX, 32 - NORMALIZE-COUNT) ;
    FX <- SIII(FX, NORMALIZE-COUNT);
    REG-OP-EXP <- REG-OP-EXP + B - NORMALIZE-COUNT ;
    TEST-FP-OVERFLOW;

```

```

DNL-NORMALIZE:      END:

***
***      3.2.6.3      VECTOR-MULTIPLY
***      VECTOR-MULTIPLY PERFORMS A THREE WORD VECTOR MULTIPLICATION
***      BETWEEN THE REGISTER VECTOR AND THE OPERAND VECTOR.
***
***      3.2.6.3.1  INPUTS -- REG-OP AND OPERAND CONTAIN THE INITIAL VALUES
***      FOR THE REGISTER AND OPERAND VECTORS.
***      EFFAD IS THE EFFECTIVE ADDRESS
***      AMZERO IS TRUE FOR THE REGISTER-REGISTER MODE
***
***      3.2.6.3.2  PROCESS-- A FLOATING MULTIPLY IS PERFORMED
***      THE RESULT (REG-OP) IS STORED IN THE REGISTER FILE
***      LOAD-NEXT IS CALLED TO READ THE NEXT VALUES
***      OF THE REGISTER AND OPERAND VECTOR.
***      THE LOOP IS EXECUTED 3 TIMES.
***
***      3.2.6.3.3  OUTPUTS-- THE REGISTER FILE CONTAINS THE RESULT
***      OF THE FIRST TWO MULTIPLICATIONS.
***      THE THIRD RESULT IS IN REG-OP.
***      RB HAS BEEN INCREMENTED BY 2 **
***
VECTOR-MULTIPLY:    PROCESSOR
LC FOR 1 TO 2      ;
BEGIN:
    FLOATING-MULTIPLY:
    (PXRGRPH) <- REG-OP:
    LOAD-NEXT:
    END:
    FLOATING-MULTIPLY:
    VECTOR-MULTIPLY:    END
***
***      3.2.6.4      SQUARE-ROOT
***      SQUARE-ROOT PERFORMS A FLOATING POINT SQUARE ROOT.
***
***      3.2.6.4.1  INPUTS -- OPERAND CONTAINS THE DATA TO USE.
***

```

3.2.6.4.2 PROCESS-- THIS SQUARE ROOT ALGORITHM IS VERY SIMILAR TO THE ELEMENTARY SCHOOL METHOD OF FINDING SQUARE ROOTS. A RULE OF 20 WAS USED TO AID IN CALCULATING THE NEXT DIGIT (PARTIAL REMAINDER / 20*PARTIAL ANSWER = ESTIMATE OF NEXT DIGIT).

OPERATING ON BINARY DIGITS (2 BITS AT A TIME) A RULE OF 2 IS USED TO CALCULATE THE NEXT BIT.

THE ALGORITHM IS COMPLICATED SOMEWHAT BY THE FACT THAT THE PARTIAL REMAINDER IS CALCULATED IN A LOOK-AHEAD MANNER, ASSUMING AN ANSWER-1.

3.2.6.4.3 OUTPUTS-- REG-OP CONTAINS THE RESULT.

 DURING THE ALGORITHM IT IS THE REMAINDER WORK1 IS THE REMAINING VALUE OF OPERAND AFTER EACH STEP WHEN 2 MORE BITS ARE SHIFTED OUT. WORK2 IS THE PARTIAL ANSWER. WORK3 IS THE PARTIAL REMAINDER (WITH LOOK AHEAD). **

```

SQUARE-ROOT:      PROCEDURE
  IF OPERAND.LENGTH = 0
  THEN BEGIN
    REG-OP      <- OPERAND.BITS-30-29
    WORK1 <- SLIC WORK1 <- OPERAND.MANTISSA,11
    END
  ELSE BEGIN
    REG-OP      <- OPERAND.BITS-30
    WORK1 <- SLIC WORK1 <- OPERAND.MANTISSA,10
    END
  END IF
  WORK2 <- 0
  WORK3 <- 0
  **LOOP FOR THE ANSWER
  FIND THE PARTIAL REMAINDER
  SHIFT OVER 2 BITS AND OR IN THE NEXT TWO BITS OF
  THE ORIGINAL OPERAND
  SHIFT THOSE TWO BITS OUT OF THE OPERAND **
  DO FOR 1 TO 23
  BEGIN
    WORK3 <- REG-OP - WORK1
  
```

```

NEG-OP ← SLL(WORK3,2)
OR WORK1.BITS-31-30
WORK1 ← SLL(WORK1,2)
IF WORK3.SIGN-BIT
THEN BEGIN
  WORK2 ← SLL(WORK2,1)
  WORK3 ← -(SLL(WORK2,2) + 3)
END
ELSE BEGIN
  WORK2 ← SLL(WORK2,1) + 1
  WORK3 ← SLL(WORK2,2) + 1
END
END IF
END ;

*TEST FOR ZERO OR NEGATIVE HERE
IF THE OPERAND IS NEGATIVE, SET THE OVERFLOW AND
DO NOT STORE
OTHERWISE STORE THE RESULT IN THE FTSC REGISTER

THEN TEST IF THE OPERAND WAS ZERO
IF IT WAS, STORE FLOATING POINT ZERO
IN THE RESULT
IF OPERAND.SIGN-BIT
THEN SET-OVERFLOW
ELSE BEGIN
  REG-OP-EXP ← SRAL(OPERAND.EXPONENT + 1,1) - 63
  REG-OP ← SLL(WORK2,4)
  NORMALIZE
END
END IF
IF (OPERAND.PARTIALS = 0
  THEN REG-OP ← 128
  END IF
SRAFL-RESULT
END ;

```

```

***
***
***
***
***

```

3.2.7.1 ILLEGAL

ILLEGAL PROCESSING ILLEGAL FTSC INSTRUCTIONS.

3.2.7.1.1 TRUTHS -- H(0)

```

*** 3.2.7.1.2 PROCESS-- THE ILLEGAL OP CODE BIT IN STATUS IS SET
*** THE POWER DOWN INTERRUPT IS REQUESTED.
*** 3.2.7.1.3 OUTPUTS-- THE FISC STATUS WORD IS ALTERED **

```

```

ILLEGAL:
    PWRCSOR:
    ILDP <- 1:
    PWRP-RFO <- 1:
    END:

```

```

ILLEGAL:

```

```

*** 3.2.7.2 SET-OVERFLOW

```

```

*** SET-OVERFLOW SETS THE FISC OVERFLOW STATUS

```

```

*** 3.2.7.2.1 INPUTS -- NONE

```

```

*** 3.2.7.2.2 PROCESS-- THE OVERFLOW STATUS BIT IS SET.

```

```

*** 3.2.7.2.3 OUTPUTS-- THE STATUS WORD IS ALTERED **
SET-OVERFLOW:
    PWRCSOR:
    OVERFLOW <- 1:
    SET-OVERFLOW:
    END:

```

```

JUMP: PROCESSOR:

```

```

    IF ANZED THEN

```

```

        PC <- OPFAND<151C>:

```

```

    ELSE

```

```

        PC <- FFFAD<151C>:

```

```

    END IF:

```

```

JUMP: END:

```

```

** DEFINE INITIAL MACHINE STATE

```

```

INITIALIZEPROCESSOR:

```

```

    IF MFTYPE = C

```

```

        THEN MEMTYPE <- X#FOFF:

```

```

    END IF:

```

```

** IF NOT SET BY USER-ALL WRITE**

```

```

INITIALIZE:END;

** MACHINE MASTER CLEAR **

MASTER-CLEAR:PROCESSOR;
  INTRC <- 0;
  INTRUC <- 0;
  INTRASK <- 0;
  STATUS <- P000000111;
  EXECUTE <- 0;
MASTER-CLEAR:END;

** BEGIN EMULATION **

LPCULE-CX: PROCESSOR;
  CASE LPCODE<310>;
    REG-OP <- OPERAND;
    EX <- OPERAND;
    FETCH-MULTIPLE(2);
    FETCH-MULTIPLE(3);
    FETCH-MULTIPLE(7);
    LDN;
    LDNF;
    IF OPERAND.SIGN-BIT
      THEN LDN;
      ELSE REG-OP <- OPERAND;
      END IF;
    IF OPERAND.SIGN-BIT
      THEN LDNF;
      ELSE REG-OP <- OPERAND;
      END IF;
    REG-OP <- NOT OPERAND;
    REG-OP <- OPERAND;
    NULL;
  *OC - ANDF - ADD (FLOATING)*
  BEGIN;

** CLEAR INTERRUPT REQUESTS **
** CLEAR IN PROCESS FLAGS **
** CLEAR INTERRUPT MASK **
** CLEAR STATUS **
** CLEAR EXECUTE INST FLAG **

** LDR 00 LOAD REGISTER **
** LDE 01 LOAD EXTENSION REG **
** LDR2 02 LOAD 2 REGISTERS **
** LDR3 03 LOAD 3 REGISTERS **
** LDR7 04 LOAD 4 REGISTERS **
** LDN 05 LOAD NEGATIVE **
** LDNF 06 LOAD NEG FLOATING **
** LDA 07 LOAD ABSOLUTE **

** LDAF 08 LOAD ABS FLOATING **

** LDC 09 LOAD ONES COMPLEMENT **
** LAD 0A LOAD ACTIVE ONLY **
** LMD 0B LOAD MONITOR ONLY **

```

```

FLOATING-PREP:  FLOATING-ALU;
END;

```

```

**CD - SUBF - SUBTRACT (FLOATING)**
BEGIN;
FLOATING-PREP:  OPERAND <- - OPERAND;
FLOATING-ADD;
END;

```

```

FLOATING-MULTIPLY;      ** MPYF OE FLOATING MULTIPLY **

```

```

NULL;
END CASE;
LPCODE-UX;      END;

```

```

OPLCODE-1X;      PROCESSOR;
CASE OPCODE<310>;
  SQUARE-ROOT;      ** SRTF 10 SQUARE ROOT FLOAT **
  **11 - VADDF - VECTOR ADD (FLOATING)**
  BEGIN;
  FLOATING-PREP:  FLOATING-ADD;
  DO FOR 1 UP TO 2;
    BEGIN;
    CPXRTPT <- REG-OP;  LOAD-NEXT;
    FLOATING-PREP:  FLOATING-ADD;
  END;
END;

```

```

*** - VSURT - VECTOR SUBTRACT (FLOATING)**
DO FOR WORK2 <- C UP TO 2;
  BEGIN;
  IF WORK2 /= 0 THEN
    BEGIN;
    CPXR(RB) <- REG-OP;  LOAD-NEXT;
  END;
  END IF;
  FLOATING-PREP;
  OPERAND <- - OPERAND;  FLOATING-ADD;
END;

```

```

VECTOR-MULTIPLY;
BEGIN;
    ## VMPYF 13 VECTOR MULTIPLY ##
    ## VIPF 14 INNER PRODUCT ##
    ;
    VECTOR-MULTIPLY
    OPERAND <- GPXR[RB - 1];
    FLOATING-PREP;
    FLOATING-ADD
    OPERAND <- GPXR[RB - 2];
    FLOATING-PREP;
    FLOATING-ADD
    ;
END;
BEGIN;
    WORK3 <- OPERAND;
    DO FOR 1 TO 3;
        BEGIN;
            FLOATING-MULTIPLY;
            GPXR[RB] <- REG-OP;
            OPERAND <- WORK3;
            REG-OP <- GPXR[RB <- RB+1];
            END;
        ;
    ;
    ## VSMF 15 SCALAR MULTIPLY ##
    ## SAVE OPERAND ##
    ;
END;
BEGIN;
    IF SE ( OPERAND.EXPONENT) <= 0
    THEN REG-OP <- 0;
    ELSE BEGIN ;
        IF OPERAND.EXPONENT > 31
        THEN SET-OVERFLOW;
        ELSE REG-OP <- SKA(REG-OP,31- OPERAND.EXPONENT);
        END IF;
    END;
END IF;
END;
BEGIN;
    GPXR[RB] <- SE(OPERAND.EXPONENT);
    RB <- RB+1;
    REG-OP <- OPERAND;
    REG-OP.EXPONENT <- 0;
    END;
END;
BEGIN;
    REG-OP-EXP <- SE( OPERAND.EXPONENT);
    LOAD-NEXT;
    FINAT( REG-OP-EXP);
    END;

```

```

##?? - ADD - INTEGER ADD##
ADD(OPERAND, 0);

##?? - SUB - INTEGER SUBTRACT##
ADD(-OPERAND, 0);

BEGIN ;
  MULTIPLY;
  ## MULTIPLY AND ADJUST EX ##
  END;

  NULL;
  NULL;

  FLOAT(31);
  ## CFL 1E CONVERT TO FLOATING ##
  REG-OP <- REG-OP AND OPERAND; ## AND 1F ##
  END CASE;
  ENDCODE-1X;
  END;

OPCODE-2X;
PROCESSOR;
CASE OPCODE<310>;
  REG-OP <- REG-OP XOR OPERAND; ## XOR 20 ##
  REG-OP <- REG-OP OR OPERAND; ## OR 21 ##
  REG-OP <- REG-OP AND NOT OPERAND; ## AND 22 AND INVERTED ##
  BEGIN;
    ## ARS 23 ARITHMETIC SHIFT ##
    SHIFT-PREP;
    IF SHIFT-DIRECTION
      THEN REG-OP <- SRA (REG-OP,SHIFT-COUNT) ;
    ELSE BEGIN;
      REG-OP <- SLA(REG-OP, SHIFT-COUNT);
      IF GPXR[RB] /# SRA (REG-OP,SHIFT-COUNT)
        THEN SET-OVERFLOW;
      END IF;
    END;
  END;
  END IF;

  BEGIN;
    EX-SIGN <- EX-SIGN-BIT;
    DOUBLE-SHIFT;
    EX <- EX-SIGN // EX<31:1> ;
    END;
  ## RRS 25 CIRCULAR SHIFT ##

```

```

SHIFT-PPFP;
IF SHIFT-DIRECTION
  THEN REG-OP <- SRC (REG-OP, SHIFT-COUNT);
  ELSE REG-OP <- SLC (REG-OP, SHIFT-COUNT);
END IF;

FND;
DOUBLE-SHIFT;
BEGIN;

  ## RRL 26 DBL CIRCULAR SHIFT ##
  ## LSS 27 LOGICAL SHIFT ##

SHIFT-PPFP;
IF SHIFT-DIRECTION
  THEN REG-OP <- SRL (REG-OP, SHIFT-COUNT);
  ELSE REG-OP <- SLL (REG-OP, SHIFT-COUNT);
END IF;

FND;
DOUBLE-SHIFT;
NULL;
NULL;
##2R - ACN - INTEGER ADD WITH CARRYOUT##
ADD(OPERAND, 1);

##2C - AS7 - LOGICAL AND AND SKIP IF ZERO##
IF (OPERAND AND REG-OP) = 0 THEN
  PC <- PC + 1;
END IF;

##2D - DISO
  - LOGICAL OR INVERTED AND SKIP IF ONES##
IF (NOT OPERAND) OR (REG-OP) = X#FFFFFFF# THEN
  PC <- PC + 1;
END IF;

ILLEGAL;
ILLEGAL;
END CASE;
UPCODE-2X;
FND;

UPCODE-4X;
PROCESSOR;
CASE UPCODE<380>;

```

```

STORE (REG-OP, EFFAD);
STORE (FX, EFFAD);
BEGIN;
  OPERAND <- REG-OP;
  STORE-DOUBLE;
END;
BEGIN;
  OPERAND <- 0;
  STORE-DOUBLE;
END;
STORE-MULTIPLE(2);
STORE-MULTIPLE(3);
STORE-MULTIPLE(7);
STORE(C, EFFAD);
NULL;
BEGIN;
  OPERAND <- STATUS // PC;
  STORE (OPERAND, EFFAD);
END;
BEGIN;
  OPERAND <- STATUS // PC;
  STORE-DOUBLE;
END;
BAD-STORE;
BAD-STORE;
BAD-STORE;
BAD-STORE;
#4F - JP7 - JUMP IF PLUS OR ZERO;
IF REG-OP >= 0 THEN
  JUMP;
END IF;

END CASE;
CPCODE-4X;

UNCODE-5X;
PROCESSOR;
CASE CPCODE(3:0);
#50 - JMP - JUMP;
JUMP;

```

** STK 40 STORE REGISTER **
 ** STE 41 STORE EXTENSION REG **
 ** STD 42 STORE REGISTER DBL **

 ** SZD 43 STORE ZERO DOUBLE **

 ** STD2 44 STORE 2 REGISTER DBL **
 ** STD3 45 STORE 3 REGISTER DBL **
 ** STD7 46 STORE 7 REGISTER DBL **
 ** SZD 47 STORE ZERO **
 ** STH 48 STORE HARD ADDRESS **
 ** SPS 49 STORE PC, STAT SINGLE **

 ** SPC 4A STORE PC, STAT DOUBLE **

 ** SBAD 4B BADD ADDRESS PARITY **
 ** SBAB 4C BADD ADDRESS PARITY **
 ** SBDD 4D BADD ADDRESS PARITY **
 ** SB0Z 4E BADD ADDRESS PARITY **
 ** 4F - JP7 - JUMP IF PLUS OR ZERO **

```

#51 - JMI - JUMP IF NEGATIVE##
IF PFC-OP < 0 THEN
  JUMP;
END IF;

#52 - J7E - JUMP IF ZERO##
IF PFC-OP = 0 THEN
  JUMP;
END IF;

#53 - J7E1 - JUMP IF ZERO (FLOATING)##
IF PFC-OP.MANTISSA = 0 THEN
  JUMP;
END IF;

#54 - JN7 - JUMP IF NON-ZERO##
IF PFC-OP /= 0 THEN
  JUMP;
END IF;

#55 - JN7F - JUMP IF NON-ZERO (FLOATING)##
IF PFC-OP.MANTISSA /= 0 THEN
  JUMP;
END IF;

#56 - JPS - JUMP IF POSITIVE AND NON-ZERO##
IF PFC-OP > 0 THEN
  JUMP;
END IF;

#57 - JPSF
- JUMP IF POSITIVE AND NON-ZERO (FLOATING)##
IF PFC-OP.MANTISSA > 0 THEN
  JUMP;
END IF;

#58 - JN7 - JUMP IF NEGATIVE OR ZERO##
IF PFC-OP <= 0 THEN
  JUMP;
END IF;

```

```

#59 - JM7
- JUMP IF NEGATIVE OR ZERO (FLOATING)##
IF PIC-OP.MANTISSA <= 0 THEN
  JUMP#
END IF#

#5A - JDN - DECREMENT RB, JUMP IF NON-ZERO##
IF (REG-OP <- REG-OP - 1) /- 0 THEN
  JUMP#
END IF#

#5B - DSI - DISABLE INTERRUPT NETWORK##
ENI-STAT <- 1#

#5C - JDS
- JUMP IF OVERFLOW SET AND RESET OVERFLOW##
IF OVERFLOW THEN
  BEGIN#
  OVERFLOW <- 0# JUMP#
  END#
END IF#

#5D - JCS - JUMP IS CARRYOUT SET##
IF CARRY THEN
  BEGIN#
  CARRY <- 0# JUMP#
  END#
END IF#

#5E - JSR - JUMP TO SUBROUTINE##
BEGIN#
INTLEV <- 7# REG-OP <- STATUS // PC#
JUMP#
END#

#5F - ENI - ENABLE INTERRUPT NETWORK##
ENI-STAT <- 0#

END CASE#
END#
CPLD1-5X1

```

[illegible]

```

INITIALIZE:
PASTIN-CLEAR:

** DEFINE MEMORY PROPERTIES **
** DEFINE INITIAL STATE **

**
INSTRUCTION EMULATION LOOP
DO FOREVER:
    BEGIN
    IF (ICOUNT <- ICOUNT+1) > 0
    THEN BEGIN:
        ICOUNT <- -1000
        END:
    END IF:

    IF STEP-FLAG
    THEN OP-STEP:
        END IF:

    IF WORK <- INTRFO AND INTMASK) /= 0
    THEN #PINSIM# EX <- 1:
        END IF:

    IF CALCUTE
    THEN INR <- OPFRAND:
    ELSE BEGIN:
        INR <- FFCH(PC):
        END:
    END IF:
    PAM1<2810> <- C:
    PC <- PC + 1:
    OPLAND-FFCH:
    KIC-UP <- CPXW(PB):

**CLEAR EMULATOR FLAG**
** RET. EFFAD,OPFRAND**

```

```

(CASE UPCODE<F14> :
  UPCODE-0X:
  UPCODE-1X:
  UPCODE-2X:
  ILLEGAL:
  UPCODE-4X:
  UPCODE-5X:
  UPCODE-6X:
  ILLEGAL:
  END CASE:
  GPX(MB) <- REG-OP:
  END:
  FTSC:END:

```

RESULT TO REGISTER##

Appendix I SMITE Keywords

and	expression	parallel
begin	external	parallel-end
case	flag	port
clock	for	primitive
closed	forever	processor
constant	id	reference
copy	if	register
data	in	returns
debug	inline	s
declare	length	seconds
decode	light	smite
default	max	statement
defined	means	step
direct	memory	switch
do	microseconds	temporary
down	milliseconds	then
else	min	to
end	ms	until
endcase	nanoseconds	up
where		
enddecode	not	us
endir	ns	while
endprimitive	null	width

32584-6015-RU-00

endstatement	opdef	xor
escape	or	

Appendix J: Performance Measurement Error Conditions

ERROR	CAUSE AND ACTION
Address Range	A data item address exceeds memory bounds. Restart SASS.
Bad Index	A data item index is not in the proper format.
Data Item Overflow	The data item table has been exceeded. No more data items may be input.
Dimension Range	The specified index does not fall in the defined dimension range for the symbol.
End of Deck	EOF on input card deck. Complete inputs from keyboard.
IFSTACK Limit	The processing array for IF clause has overflowed. No more IF clauses may be input.
Illegal Function	The function following an IF clause is not allowable.

32584-6015-RU-00

Illegal Logical Op

The operator in a logical relation is not one of the allowed set.

Illegal Offset

The offset to a control point is not an integer

Illegal Qualifier

The SMITE label is not defined in the specified processor's scope.

Illegal Stack Operator

A SASS table is incorrect. Restart SASS.

Illegal Statement

A control point is not a valid SMITE statement number.

Illegal Trap Address

A PM trap instruction occurred without a corresponding table entry. Restart SASS.

Illegal Width

A data item width specification was not integer numbers.

_Not a Valid Qualifier

The data item is not defined in the scope of the processor used to qualify it.

Qualifier Not Matched

The label is not defined in the scope of the processor used to qualify it.

Qualifier Not 1 Bit

The data item used in the qualifier clause is not one bit wide.

Reentrant Processor

A control probe in a TIMING input is defined within a re-entrant processor. The results would be difficult to interpret since the processor is not restricted to one task invocation.

Reversed Width Specification A data item width specification is input in the opposite order from the SMITE width specification. SASS reverses the input specification to correspond.

Stack Variable

A data item is maintained on the SMITE main store stack. The item may be accessed correctly only when the control probe is located in its defining processor (i.e. the stack pointer is correct).

Temporary QM-1 Register

The data item specified is allocated as a SMITE temporary (QM-1 register). The data item may be accessed correctly only during the life span of the variable. An examination of the generated code is recommended to determine the range of SMITE statement which have valid data for the temporary.

Timing Overflow

The accumulated TIMING value overflows the 36 bit word. The length of the run should be shortened or the clock units changed.

32584-6015-RU-00

Too many Inputs

The maximum number of COUNT, TIMING, or SAMPLE inputs has occurred.

Too Many Sample Points

The maximum number of SAMPLE bins has been used. The width of the data item being sampled should be reduced.

Unknown Function

The PM function is not valid.

Unknown Qualifier


The processor name to qualify a data item or label was not found in the list of SMITE processors.

Unknown Symbol

The label or data item name was not found in the symbol table.

Width Range

The specified width parameter for a data item was not in the width range defined for the SMITE variable.



MISSION of Rome Air Development Center

RADC plans and executes research, development, test and selected acquisition programs in support of Command, Control Communications and Intelligence (C³I) activities. Technical and engineering support within areas of technical competence is provided to ESD Program Offices (POs) and other ESD elements. The principal technical mission areas are communications, electromagnetic guidance and control, surveillance of ground and aerospace objects, intelligence data collection and handling, information system technology, ionospheric propagation, solid state sciences, microwave physics and electronic reliability, maintainability and compatibility.

9-