# SERIES <br> 32000 <br> INSTRUCTION <br> SET <br> MANUAL 

6. September 2017

## CONTENTS

Chapter Page
1 INTRODUCTION ..... 1-1
2 PROGRAMMING MODEL ..... 2-1
2.1 GENERAL REGISTERS ..... 2-2
2.2 DEDICATED REGISTERS ..... 2-2
2.3 CONFIGURATION REGISTER (CFG) ..... 2-10
2.4 FLOATING-POINT REGISTERS ..... 2-11
2.4.1 Floating-Point Data Registers ..... 2-11
2.4.2 Floating-Point Status Register (FSR) ..... 2-12
2.5 MEMORY MANAGEMENT REGISTERS ..... 2-15
2.6 DEBUG REGISTERS ..... 2-16
2.7 MEMORY ORGANIZATION ..... 2-17
2.7.1 Addressing ..... 2-17
2.7.2 Memory Operand Formats ..... 2-17
2.7.3 Data Alignment ..... 2-19
2.8 DEDICATED MEMORY AREAS ..... 2-19
2.8.1 User and Interrupt Stacks ..... 2-20
2.8.2 Module Table ..... 2-21
2.8.3 Link Tables ..... 2-23
2.8.4 Interrupt Dispatch Table and Cascade Table ..... 2-24
2.8.5 Input and Output ..... 2-24
2.9 PRIVILEGE STATES AND PROTECTION ..... 2-25
3 INSTRUCTIONS AND DATA TYPES ..... 3-1
3.1 INTEGER INSTRUCTIONS ..... 3-2
3.2 PACKED DECIMAL INSTRUCTIONS ..... 3-7
3.3 FLOATING-POINT INSTRUCTIONS ..... 3-9
3.3.1 Floating-Point Operand Formats ..... 3-10
3.3.2 Normalized Numbers ..... 3-11
3.3.3 Zero ..... 3-12
3.3.4 Reserved Operands ..... 3-13
3.3.5 Integers ..... 3-13
3.3.6 Memory Representations ..... 3-13
3.3.7 Floating-Point Traps ..... 3-14
3.4 LOGICAL INSTRUCTIONS ..... 3-16
3.5 BIT INSTRUCTIONS ..... 3-18
3.6 BIT FIELD INSTRUCTIONS ..... 3-21
3.7 STRING INSTRUCTIONS ..... 3-24
3.8 BLOCK INSTRUCTIONS ..... 3-29
3.9 ARRAY INSTRUCTIONS ..... 3-31
3.10 PROCESSOR CONTROL INSTRUCTIONS ..... 3-33
3.11 PROCESSOR SERVICE INSTRUCTIONS ..... 3-35
3.12 MEMORY MANAGEMENT INSTRUCTIONS ..... 3-37
3.13 CUSTOM INSTRUCTIONS ..... 3-38
4 INSTRUCTION OPTIONS AND CONSTRUCTION ..... 4-1
4.1 SYNTAX PRESENTATION ..... 4-2
4.2 OPERAND ATTRIBUTES ..... 4-3
Chapter Page
4.2.1 Access Classes ..... 4-4
4.2.2 Length Attributes ..... 4-6
4.2.2.1 Integer Length Attributes ..... 4-7
4.2.2.2 Floating-Point Length Attributes ..... 4-8
4.2.3 Implied Operand Attributes ..... 4-9
4.3 BINARY INSTRUCTION FORMAT ..... 4-10
4.3.1 Basic Instruction ..... 4-12
4.3.1.1 Operation Code Fields ..... 4-12
4.3.1.2 Operation Length Fields: $i$ and $f$ ..... 4-12
4.3.1.3 General Addressing Mode Fields: gen ..... 4-12
4.3.1.4 Implied Operand Fields: reg,quick,short ..... 4-13
4.3.2 Extension Fields ..... 4-13
4.3.2.1 Index Bytes ..... 4-13
4.3.2.2 Addressing Extensions ..... 4-13
4.3.2.3 Implied Operand Extensions: imm, disp ..... 4-14
4.4 NS32000 ADDRESSING MODES ..... 4-15
4.4.1 Register Modes ..... 4-17
4.4.2 Register Relative Modes ..... 4-19
4.4.3 Memory Relative Modes ..... 4-20
4.4.4 Immediate Mode ..... 4-21
4.4.5 Absolute Mode ..... 4-22
4.4.6 External Mode ..... 4-23
4.4.7 Top of Stack Mode ..... 4-24
4.4.8 Memory Space Modes ..... 4-25
4.4.9 Scaled Indexing ..... 4-26
4.5 CONSTRUCTING COMPLETE BINARY INSTRUCTIONS: SOME EXAMPLES ..... 4-28
5 SERIES 32000 INSTRUCTION SET ..... 5-1
5.1 INSTRUCTION EXAMPLES ..... 5-3
5.1.1 Coding Examples ..... 5-3
5.1.2 Action Examples ..... 5-3
5.1.3 Operand Presentation Format ..... 5-5
5.2 INSTRUCTION DEFINITIONS ..... 5-7
Appendix
A INSTRUCTION SET LISTED BY FUNCTIONAL GROUPS ..... A-1
Figure Page
2-1 Series 32000 Register Set ..... 2-3
2-2 Processor Status Register ..... 2-6
2-3 Floating-Point Status Register ..... 2-12
2-4 Module Descripter Format ..... 2-21
2-5 Sample Link Table ..... 2-23
3-1 Floating-Point Operand Formats ..... 3-10
4-1 General Format ..... 4-11
5-1 Typical Instruction Definition ..... 5-2
5-2 Typical Instruction Example ..... 5-4
TABLES
Table ..... Page
2-1 PRIVILEGED INSTRUCTIONS ..... 2-26
3-1 SAMPLE F FIELDS ..... 3-10
3-2 SAMPLE E FIELDS ..... 3-11
3-3 NORMALIZED FLOATING-POINT RANGES ..... 3-12
3-4 EXECUTION SEQUENCES ..... 3-28
3-5 ROW MAJOR ORDERING ..... 3-32
3-6 COLUMN MAJOR ORDERING ..... 3-32
4-1 ADDRESSING MODE ACTIONS VS. ACCESS CLASS ..... 4-5
4-2 SERIES 32000 ADDRESSING MODES ..... 4-16

## Chapter 1

## INTRODUCTION

This document is a revised definition of the Series 32000 instruction set. It provides more specific information on architectural details, and also incorporates further information on compatibility issues.

This is not a full architectural description, and is intended to supplement and update other documentation already in print. Specific areas not included here are:

* Material which is primarily tutorial in nature.
* Details of memory management and exception processing.

The term "undefined" is used frequently as the outcome of an illegal instruction form. An outcome which is architecturally undefined is not guaranteed to remain the same under all conditions, in all component revisions, or in future expanded implementations of this architecture. Many of these illegal options may "work" in the current implementation, but they are nevertheless considered undefined by NSC, and should always be avoided. Illegal instruction forms, when executed in User mode, are guaranteed not to bypass any of the protection mechanisms implemented in the Series 32000 family.

The manual is divided as follows:

1. INTRODUCTION
2. PROGRAMMING MODEL

Definitions of the Series 32000 register set and other resources visible to the programmer.
3. INSTRUCTIONS AND DATA TYPES

A discussion of the instruction set by functional groups, including definitions of associated data types and exceptional conditions.
4. INSTRUCTION OPTIONS AND CONSTRUCTION

Definitions of the Series 32000 addressing modes and the construction of instructions in assembly language and binary.
5. INSTRUCTION SET

Individual definitions of the Series 32000 instructions, organized alphabetically by mnemonic.

Appendices:
A. LIST OF INSTRUCTIONS BY FUNCTIONAL GROUP

## PROGRAMMING MODEL

This chapter defines the programming model (resources visible to the programmer) presented by the Series 32000 architecture. More specifically, this chapter presents the Series 32000 register set, memory organization, and the functions of dedicated memory areas used by Series 32000 hardware. Also presented here is the mechanism used to protect privileged portions of the programming model.

This chapter is organized as follows:

| Topic | Section |
| :--- | :---: |
| General Registers | 2.1 |
| Dedicated Registers | 2.2 |
| Configuration Register | 2.3 |
| Floating-Point Registers | 2.4 |
| Memory Management Registers | 2.5 |
| Debug Registers | 2.6 |
| Memory Organization | 2.7 |
| Dedicated Memory Areas | 2.8 |
| Privilege States and Protection | 2.9 |

### 2.1 General Registers

There are eight 32-bit General-Purpose registers, named R0 through R7 (see Figure 2-1). The contents of any General-Purpose register can be used as:

1. Data, using the Register addressing modes (Section 4.4.1).
2. A base pointer, using the Register Relative addressing modes (Section 4.4.2).
3. An index value, using the Scaled Indexing modifier in an addressing mode (Section 4.4.9).

Data held within a General-Purpose register may be treated as an 8-bit, 16-bit, or 32-bit value. When an instruction operates on data of less than 32 bits, the value used is the low-order portion of the register. The remaining portion of the register is neither used nor affected.

For extended arithmetic (the MEIi and DEIi instructions), the General-Purpose registers are combined to form even/odd register pairs: R0/R1, R2/R3, R4/R5, and R6/R7. See Section 4.4.1 for details of this use.

### 2.2 Dedicated Registers

The Dedicated registers store memory addresses and general status information (see Figure 2-1). The nine Dedicated registers are:

* Program Counter (PC)
* Static Base Register (SB)
* User Stack Pointer (SP1)
* Interrupt Stack Pointer (SPO)
* Frame Pointer (FP)
* Interrupt Base Register (INTBASE)
* Module Register (MOD)
* Processor Status Register (PSR)
* Configuration Register (CFG)

The PC, SB, SP1, SP0, FP, and INTBASE registers each hold 32-bit memory addresses. The MOD and PSR registers are each 16 bits long. The MOD register contains a memory address, and the PSR register contains status information. The addresses contained in these registers are interpreted as virtual in memory-managed systems. The CFG register is 32 bits long.

DEDICATED

| Program Counter | PC |
| :---: | :---: |
| Static Base | SB |
| Frame Pointer | FP |
| User Stack Pointer | SP1 |
| Interrupt Stack Pointer | SP0 |
| Interrupt Base | INTBASE |
| Configuration | CFG |
| Module | MOD |
| Status | PSR |

GENERAL


MEMORY MANAGEMENT

| Page Table Base 0 | PTB0 |
| :---: | :---: |
| Page Table Base 1 | PTB1 |
| Invalidate Address 0 | IVAR0 |
| Invalidate Address 1 | IVAR1 |
| Translation Exception | TEAR |
| Control | MCR |
| Status | MSR |

DEBUG

| Control |
| :---: |
| Status |
| DCR |
| Compare Address |
| Breakpoint PC |
| CAR |
| BPC |

FLOATING-POINT

| F0:L | F1: F | FO: F |
| :---: | :---: | :---: |
| F1: L |  |  |
| F2: L | F3: F | F2: F |
| F3: L |  |  |
| F4: L | F5: F | F4: F |
| F5:L |  |  |
| F6:L | F7: F | F6: F |
| F7: L |  |  |
|  |  |  |
| FSR | Floating-Point Status |  |

$$
\begin{gathered}
\text { Figure } 2-1 \text { Series } 32000 \text { Register Set } \\
2-3
\end{gathered}
$$

A description of each Dedicated register follows.
PC The Program Counter is available as a Base register (using the Program Memory addressing mode, Section 4.4.8). It contains the memory address of the first byte of the instruction currently being executed. The PC is incremented (to point to the next instruction) only when the current instruction is completed. On occurrence of a Reset, the PC is set to zero, and the first instruction is fetched from this address.

SP1 The User Stack Pointer points to the top of the User Stack (Section 2.8.1). The SP1 register is selected for all stack operations while the $S$ bit in the Processor Status Register is set to 1.

SPO The Interrupt Stack Pointer points to the top of the Interrupt Stack (Section 2.8.1). The Interrupt Stack is selected for all stack operations while the $S$ bit in the $P S R$ is set to 0 . It is also automatically selected whenever an interrupt or trap occurs. In memory-managed systems, SP0 must always contain a valid Supervisor-Mode virtual address (see Section 2.8.1).

NOTE: The SP1 and SPO registers are never referenced directly by a program. Instead, the symbol "SP" is used, meaning the Stack Pointer which is currently selected. This SP register is available as a base pointer using the Stack Memory and Stack Memory Relative addressing modes (Sections 4.4.8 and 4.4.3). The Top of Stack addressing mode uses the SP register in performing "push" and "pop" references to the top of the stack (Section 4.4.7).

FP The Frame Pointer points to a dynamically-allocated data area created at the beginning of a procedure (by the ENTER instruction). This area is generally called the "activation record" for the procedure, and contains its parameters, local variables, saved registers, and return address. The FP register is available as a base pointer using the Frame Memory and Frame Memory Relative addressing modes (Sections 4.4.8 and 4.4.3).

INTBASE The Interrupt Base register contains the base address of the Interrupt Dispatch Table. This is a vector table which contains the descriptors of the trap and interrupt service procedures. In memory-managed systems, INTBASE must always contain a valid Supervisor-Mode virtual address (see Section 2.8.4).

The Module register points to the current module's Module Table entry. The Module Table entry is a 16-byte block of memory containing three pointers for the current module:

* SB (Static Base, a pointer to its static data area)
* LB (Link Base, a pointer to its Link Table)
* PB (Program Base, a pointer to the beginning of its code)

See Section 2.8.2.
The Static Base register contains the base address of data which has been statically allocated (i.e. allocated once, before program execution) to the current module. This address is a copy of the SB pointer in the current Module Table entry. It is available for use in the Static Memory and Static Memory Relative addressing modes (Sections 4.4.8 and 4.4.3). The Static Base register is automatically updated whenever control is transferred from one module to another.

$\mathrm{x}=$ reserved

Figure 2-2 Processor Status Register

PSR, The Processor Status Register (Figure 2-2) contains 16 mode and UPSR status flag bits, of which 11 bits are currently implemented. All implemented PSR flags are readable and writable. The bit positions marked "x" in Figure 2-2 are reserved for future use. They are not currently implemented, and do not retain information written to them. For upward compatibility reasons, no program should attempt to change these bits, nor should any program assume that they are always zero (even though they appear to be permanently zero in the current implementation).

The least-significant byte of the PSR contains flags which are always accessible. This byte is also called the UPSR, for "User PSR".

The most-significant byte of the PSR contains the Supervisor flags. Supervisor flags are accessible only by a program running in Supervisor mode (see the discussion of the $U$ bit which follows). Any attempt by a User Mode program to load, store or modify this byte causes the Illegal Operation trap, Trap (ILL), instead. See Section 2.8 for further details of protection features.

Upon occurrence of an interrupt or trap, the PSR is pushed onto the Interrupt Stack. Certain PSR bits are then automatically cleared (as stated in their descriptions) to establish the proper modes of operation for interrupt service.

NOTE: The PSR P bit is sometimes cleared before the PSR is pushed onto the Interrupt Stack. For further details see the datasheet of the NS32532 CPU.

All implemented PSR flags are cleared to zero on occurrence of a Reset.

C is the Carry flag. The Carry flag signals a carry condition during execution of an addition instruction or a borrow condition during a subtraction instruction. If a carry or borrow has occurred, the $C$ bit is set to 1. If no carry or borrow has occurred, the $C$ bit is set to 0 . See Section 3.1 for definitions of carry and borrow conditions.
$T$ is the Trace flag. This flag places a program in Trace mode, allowing step-by-step inspection of the effects of each instruction. While the $T$ bit is set, the Trace trap, Trap (TRC), occurs at the completion of each instruction. The $T$ bit interacts with the $P$ bit to ensure correct operation of Trace Mode regardless of any interrupts or other traps which may also be occurring. It is cleared on occurrence of any trap or interrupt.

L is the Low flag. The Low flag signals the result of an unsigned comparison between two integers. (All integer comparison instructions perform both signed and unsigned comparisons.) If the second operand of a comparison instruction is less than the first, the L bit is set to 1. If the second operand is greater than or equal to the first, the L bit is set to 0. The L flag is always cleared by the floating-point comparison instruction (CMPf).

V is the Overflow Trap Enable flag. It allows generation of a trap (OVF) when an integer arithmetic operation overflows.

F is the F Flag. The $F$ flag is a general condition flag, used by various instructions to signal exceptional conditions (e.g. integer overflow from addition or subtraction), or to distinguish among outcomes (e.g. what condition has caused a String instruction to terminate).

Z is the Zero flag. The Zero flag indicates the result of comparing two integers or two floating-point values. If they are equal, the $Z$ bit is set to 1 . If they are not equal, the $Z$ bit is set to 0 .
$N$ is the Negative flag. The Negative flag indicates the result of a signed comparison between two integers or two floating-point values.

NOTE: The integer comparison instructions, CMPi and CMPQi, perform both signed and unsigned comparisons.

If the second operand is less than the first, the $N$ bit is set to 1. If the second operand is greater than or equal to the first, the N bit is set to 0 .

The $N$, $Z, F, L$ and $C$ bits constitute $a$ "condition" which may be used by the Conditional Branch (Bcond) and Save Condition Code (Scondi) instructions. In addition, the $F$ bit may be used to cause a trap (by the FLAG instruction).

## Supervisor PSR Flags

U is the User Mode flag. If the $U$ bit is 1 , the current program is running in User mode, and may not use privileged instructions or reference protected registers. If the $U$ bit is 0 , the current program is running in Supervisor mode, and is not restricted. In memory-managed systems, address translation and memory protection features may also be affected by the state of this bit. The $U$ bit is automatically cleared on occurrence of any interrupt or trap. See Section 2.8 for further details of protection features.

S is the Stack flag. The $S$ bit selects which of the two stack pointers is to be used for stack operations. If the $S$ bit is 1, the User Stack Pointer (SP1) is selected. If the $S$ bit is 0 , the Interrupt Stack Pointer (SPO) is selected. The S bit is automatically cleared on occurrence of a trap or interrupt.

P is the Trace Trap Pending flag. The P bit interacts with the $T$ bit to ensure correct trace results in programs which are being interrupted or trapped. It is automatically cleared on occurrence of any trap or interrupt. The P bit in the PSR image which is pushed on occurrence of an interrupt or trap may also be cleared, depending on the trap or interrupt.

I is the Interrupt Enable flag. If the I bit is 1, both Maskable and Non-Maskable interrupts are accepted. If the $I$ bit is 0 , only Non-Maskable interrupts are accepted. The I bit is automatically cleared on occurrence of an interrupt or the Abort trap, Trap (ABT). No other traps affect this bit, and this bit does not disable traps when clear.

### 2.3 Configuration Register (CFG)

The Configuration register is used to enable or disable certain Series 32000 features which are currently optional. With the SETCFG instruction only the four LSBs are loaded. The remaining bits are set to zero except the bits 4 to 7 which are always set to 1. With the LPR instruction all implemented bits can be set.

The CFG register is 32 bits wide, of which ten bits are implemented.

```
+-----...-----+---+----+---+---++---+----+---+---+---+---+----+---+--------------
! (reserved) !PF !LIC!IC !LDC!DC !DE ! 1 ! 1 ! 1 ! 1 ! C ! M ! F ! I !
lol
```

The bits correspond to features as given below.
I Interrupt vectoring. This bit declares whether hardware support is available for direct vectoring of maskable interrupts. If the $I$ bit is set, service of a maskable interrupt includes reading an 8-bit value which selects an Interrupt Dispatch Table entry to use in locating the interrupt service procedure (see Section 2.8.4). This 8-bit value is supplied by an NS32202 Interrupt Control Unit. If the I bit is not set, maskable interrupts are not vectored, and use by default the first entry (NVI) of the Interrupt Dispatch Table, requiring no hardware support.

F Floating-Point instruction set. If this bit is set, the Floating-Point instruction set (Section 3.3) is enabled, and an attached NS32381 Floating-Point Unit will be used to execute these instructions. If the F bit is not set, all Floating-Point instructions generate Trap (UND) instead. (The trap mechanism employed by the Series 32000 architecture allows software to intercept this trap and fully emulate the functions of the NS32381.)

M Memory Management instruction set. If this bit is set, the LMR, SMR, RDVAL and WRVAL instructions (Section 3.12) are enabled, and the integrated NS32532 Memory Management Unit will be used to execute them. If the $M$ bit is not set, these instructions generate Trap (UND) instead. (Note: the Memory Management instructions MOVSUi and MOVUSi are not affected by this bit, and are always available.)

C Custom instruction set. If this bit is set, the Custom instruction set (Section 3.13) is enabled, and will use attached custom hardware (unique to a given system). If it is not set, all Custom instructions generate Trap (UND) instead.

DE Direct-Exception mode enable. This bit enables the Direct-Exception mode for processing exceptions. When this mode is selected, the CPU response time to interrupts and other exceptions is significantly improved.

DC Data Cache enable. This bit enables the on-chip Data Cache to be accessed for data reads and writes.

LDC Lock Data Cache. This bit controls whether the contents of the on-chip Data Cache are locked to fixed memory locations (LDC = 1), or updated when a data read is missing from the cache (LDC = 0).

$$
2-10
$$

IC Instruction Cache enable. This bit enables the on-chip Instruction Cache to be accessed for instruction fetches.

IDC Lock Instruction Cache. This bit controls whether the contents of the on-chip Instruction Cache are locked to fixed memory locations (LIC=1), or updated when an instruction fetch is missing from the cache (LIC=0).

PF Pipelined Floating-Point execution. This bit indicates whether the Floating-Point unit uses the pipelined slave protocol. When PF is 1 the pipelined protocol is selected. $P F$ is ignored if the $F$ bit is 0. The NS32381 FPU does not support the pipelined slave protocol.

### 2.4 Floating-Point Registers

Floating-Point registers are present in systems supporting the Floating-Point instruction set (either by using the NS32381 Floating-Point Unit or by software emulation). See Figure 2-1. There are eight Floating-Point Data registers (F0-F7) and one Floating-Point Status register (FSR).

### 2.4.1 Floating-Point Data Registers

The Floating-Point Data registers provide a high-speed workspace for floating-point operations. These registers are named F0 through F7, and are 64 bits in length. They are referenced whenever the Register addressing mode (Section 4.4.1) is used in a floating-point instruction to specify the location of a floating-point operand. Floating-Point operands are located in memory or in Floating-Point Data registers, and integer operands are located in memory or in General-Purpose registers.

The even double-precision floating-point register contains the respective even and the next following odd single-precision floating-point registers. The odd single-precision floating-point register is held in the high-order half of the double-precision floating-point register, the even register is held in the loworder half. See Figure 2-1.

### 2.4.2 Floating-Point Status Register (FSR)

The Floating-Point Status register (FSR) selects operating modes and records any exceptional conditions encountered during execution of a floating-point instruction. Figure 2-3 shows the format of the FSR.


Figure 2-3 Floating-Point Status Register

Bits 17 through 31 of the $F S R$ are reserved. The SWF field (bits 9 through 15) is currently reserved for NSC software use (floating-point extension software). Information written to this field is retained, but does not affect any hardware operations. The remaining bits (17 through 31) are not implemented, and do not retain information written to them. For upward compatibility reasons, no program should attempt to change either reserved field, nor should any program assume that their contents are always zero (even though bits 17-31 appear to be permanently zero in the current implementation). To change the contents of the FSR, the following procedure should always be followed:

1. Use the SFSR instruction to store the FSR in a temporary location.
2. Change the desired fields in this temporary copy.
3. Use the LFSR instruction to load the temporary copy into the FSR.

The FSR mode fields are set by the programmer to establish modes of operation for floating-point instructions. The mode fields are encoded as follows.

RM Rounding Mode: bits 7 and 8. This field selects the rounding method to be used whenever a floating-point result cannot be exactly represented in the format of the destination operand. The rounding modes are:

00 Round to nearest value. The value which is nearest to the exact result is selected. If the result is exactly halfway between the two nearest values, the even value (LSB $=0$ ) is delivered to the destination.

01 Round toward zero. The nearest value whose absolute value is less than, or equal to, the exact result is delivered to the destination.

10 Round toward positive infinity. The nearest value which is greater than, or equal to, the exact result is delivered to the destination.

11 Round toward negative infinity. The nearest value which is less than, or equal to, the exact result is delivered to the destination.

UEN Underflow Trap Enable: bit 3. If this bit is set, Trap (FPU) occurs whenever an underflow condition is encountered. See Section 3.3.7 for the definition of floating-point underflow. If it is not set, any underflow condition returns a result of positive zero (Section 3.3.3), and no trap occurs.

IEN Inexact Result Trap Enable: bit 5. If this bit is set, Trap (FPU) occurs whenever the result of a floating-point instruction is not exact. If it is not set, the result is rounded according to the selected rounding mode, and no trap occurs.

The FSR status fields record exceptional conditions encountered during the execution of a floating-point instruction. The meanings of the FSR status bits are as follows:

TT Trap Type: bits 0-2. This 3-bit field records any exceptional condition detected by a floating-point instruction. These conditions are defined in Section 3.3.7. They are reported as:

000 No exceptional condition occurred.
001 Underflow
010 Overflow
011 Division by Zero
100 Illegal Instruction
101 Invalid Operation
110 Inexact Result
111 (Reserved for future use.)

The TT field is loaded with zero whenever any floating-point instruction except LFSR or SFSR completes without encountering an exceptional condition. It is also set to zero by a Reset or by writing zero into it with the Load FSR (LFSR) instruction. Underflow and Inexact Result are always reported in the TT field, regardless of the settings of the UEN and IEN bits.

UF Underflow Flag: bit 4. This bit is set whenever an underflow condition is detected. See Section 3.3.7 for the definition of floating-point underflow. The function of the UF bit is not affected by the state of the UEN bit. The UF bit is cleared only by writing a zero into it with the LFSR instruction or by a Reset.

IF Inexact Result Flag: Bit 6. This bit is set whenever an Inexact Result condition is detected, and no other errors have occurred. See Section 3.3.7 for the definition of this condition. It is cleared only by writing a zero into it with the LFSR instruction or by a Reset.

RMB Register Modify Bit: Bit 16. This bit is set whenever writing to a floatingpoint data register. It is cleared only by writing a zero into it with the LFSR instruction or by a Reset. This bit can be used in context switching to determine whether the FPU registers should be saved.

### 2.5 Memory Management Registers

Memory Management registers are present in systems incorporating the Series 32000 memory management option. These registers are currently implemented in the NS32532 CPU and are made available by setting the $M$ bit in the CFG register (Section 2.3). There are seven 32-bit Memory Management registers (Figure 2-1):

| PTB0, PTB1 | Page Table Base Registers |
| :--- | :--- |
| IVAR0, IVAR1 | Invalidate Virtual Address Registers |
| TEAR | Translation Exception Address Register |
| MCR | Memory Management Control Register |
| MSR | Memory Management Status Register |

The Memory Management registers are each 32 bits in length. The following describes briefly the function of each register. Further informations can be found in the data sheet of the NS32532 CPU.

PTBO and PTB1 support virtual memory and address translation. These registers contain the base addresses of the Level 1 Page Tables.

IVARO and IVAR1 the Invalidate Virtual Address registers are write-only registers. When a virtual address is written to IVARO or IVAR1 using the LMR instruction, the translation for that virtual address is purged, if present, from the TLB.

TEAR the TEAR register is loaded by the MMU when a translation exception occurs. It contains the 32 -bit virtual address that caused the translation exception.

MCR contains the memory management control flags.
MSR contains the memory management status flags.

The Debug registers are each 32 bits in lenth. They can be accessed by the LPR and SPR instructions. The following describes briefly the function of each register. Further informations can be found in the data sheet of the NS32532 CPU.

DCR contains the debug control flags.

DSR contains the debug status flags.
CAR the CAR Register contains the address that is compared to operand reference addresses to detect an address-compare condition.

BPC the BPC Register contains the address that is compared with the PC contents to detect a PC-match condition.
2.7 Memory Organization

The Series 32000 architecture supports a memory addressing space of four gigabytes (corresponding to a 32-bit address).

### 2.7.1 Addressing

A memory address is a 32-bit unsigned integer. It uniquely identifies an 8-bit location (a byte) within the memory space. In decimal, the addressing range is 0 through 4,294,967,295.

NOTE: Except where otherwise indicated, all addresses and memory spaces given in this manual are virtual in memory-managed systems, and can be mapped to any "physical" (or "real") memory page.

### 2.7.2 Memory Operand Formats

The basic storage unit is the byte. A byte holds eight bits of data and has the following form:


Byte at Address A

Bit positions are numbered from 0 to 7. Bit 0 is the least-significant bit; bit 7 is the most-significant bit.

A 16-bit value is called a word. It is held in memory as a pair of contiguous bytes.


The byte at the lower address is the least-significant byte; the byte at the higher address is the most-significant byte. A word has the same address as its least-significant byte and may start at any address.

A 32-bit value is called a double-word. It is held in memory as four contiguous bytes. A double-word can hold either a 32-bit integer or a single-precision floating-point value.


The least-significant byte of a double-word is stored at the lowest address. A double-word has the same address as its least-significant byte and may start at any address.

A 64-bit value is called a quad-word. It is held in memory as eight contiguous bytes. A quad-word can hold a 64-bit integer or a double-precision floatingpoint value.


Quad-word at Address A

The least-significant byte of a quad-word is stored at the lowest address. A quad-word has the same memory address as its least-significant byte and may start at any address.

### 2.7.3 Data Alignment

With the sole exception of the Page Tables used for memory management, there are no alignment restrictions in the Series 32000 architecture. Operands of any length may start at any byte address.

For optimal throughput, however, it is usually desirable to align data. A method for alignment which applies well to all memory bus size implementations (8, 16 or 32 bits) is to align operands on "integral" boundaries. By this method, words are stored at even addresses, double-words at multiples of four, and quad-words at multiples of eight.

### 2.8 Dedicated Memory Areas

A Series 32000-based system will make use of certain designated memory areas for the following purposes:

* User and Interrupt Stacks
* Module Table
* Link Tables
* Interrupt Dispatch Table and Cascade Table
* Input and Output


### 2.8.1 User and Interrupt Stacks

A stack is a block of memory used as a last-in/first-out (LIFO) buffer. The contents of a Stack Pointer register specify an address within the block, and the value at this address is considered to be at the top of the stack.

There are two stacks: a User Stack and an Interrupt Stack. The User Stack Pointer (SP1) specifies the address of the top of the User Stack, and the Interrupt Stack Pointer (SPO) specifies the address of the top of the Interrupt Stack. At any time, one of these stacks is selected for stack operations (by the PSR S bit, Section 2.2). The User stack is generally assigned to User-Mode programs, although programs running in Supervisor Mode may also select it. The Interrupt stack is identical in function to the User stack, except that it is always selected on a trap or interrupt to receive the return information (return address, MOD and PSR). An interrupt or trap service routine may continue to use the Interrupt Stack, or it may re-select the User stack.

Stacks grow downward in memory: i.e., toward lower addresses. To pop a value, the current Stack Pointer is incremented by the value's length in bytes after reading it ("post-increment"). To push a value, the current Stack Pointer is decremented by the value's length in bytes before writing it ("pre-decrement"). In either case, the Stack Pointer indicates the new top of the stack.

Data may be read from, or written to, the currently-selected stack at any time, using the Top of Stack addressing mode (Section 4.4.7), which performs an automatic push or pop, as appropriate. In addition, the current Stack Pointer may be used as a base pointer in the Stack Memory and Stack Memory Relative addressing modes (Sections 4.4.8 and 4.4.3).

The current stack also receives return addresses and other context information saved in the process of invoking a procedure. Examples of this use are the BSR (Branch to Subroutine) instruction and the ENTER (Enter Procedure Context) instruction. Instructions of this type always modify the Stack Pointer in multiples of four, so that the stack may always be kept aligned on 32-bit boundaries if desired for optimal throughput.

NOTES: 1. Information popped from a stack should never be considered still available in its original memory location after the popping instruction terminates, nor should any program ever store information in a memory area which is available for stack expansion but is not within the stack. These requirements are made for reasons of upward compatibility and compatibility between systems.
2. In memory-managed systems, the Interrupt stack must always be available in physical memory. On occurrence of an interrupt or trap, the contents of the Interrupt Stack pointer are treated as a SupervisorMode virtual address.

### 2.8.2 Module Table

The Series 32000 architecture supports software modules and modular programs through a Module Table. This table contains one 16-byte entry (a module descriptor) for each module in the program. The MOD Register (Section 2.2) holds the address of the Module Table entry for the currently-running module.

All Module Table entries need not be held in a single contiguous memory space, but they must all be contained within the first 64 K bytes of memory, due to the fact that the MOD register holds only a 16-bit address. A Series 32000-based system, therefore, can hold up to 4096 modules at a time (4096 modules per user, in memory-managed systems).

A module descriptor contains four 32 -bit pointers, of which the first three are used in the current implementation. These pointers are found relative to the contents of the MOD register as shown in Figure 2-4.

Address

|  | ! 31 |  | 0 ! |
| :---: | :---: | :---: | :---: |
| MOD : | ! | Static Base | ! |
| MOD + 4: | ! | Link Base | ! |
| MOD + 8: | ! | Program Base | ! |
| MOD + 12: | ! | (Reserved for future use) | ! |

Figure 2-4 Module Descriptor Format

The Static Base pointer contains the address of a memory area allocated to this module for static data; i.e., data which is allocated only once, before execution. This pointer is loaded into the Static Base register whenever control is transferred from one module to another.

The Link Base pointer contains the address of the Link Table assigned to this module. See Section 2.8.3.

The Program Base pointer contains the address of the first byte of the code section of this module. It is used by other modules (through their Link Tables) to transfer control to specific procedures within this module.

NOTES: 1. All Module Table entries must be entirely contained within the first 64 K bytes of memory. This means that MOD register values of FFF1 through FFFF (Hex) are reserved.
2. In memory-managed systems, all module descriptors for interrupt or trap service routines must always be in physical memory. The contents of the three pointers are interpreted as Supervisor-Mode virtual addresses.

### 2.8.3 Link Tables

One Link Table is allocated to each module of a program. The Link Base pointer of the current Module Table entry (Section 2.8.2) points to the Link Table for the currently running module.

Each Link Table provides information which is used for:

1. Sharing variables between modules. Such variables are available to other modules via the External addressing mode (Section 4.4.6).
2. Transferring control from one module to another. This is done directly from the current Link Table via the CXP instruction.

A module's Link Table is constructed by a linker program based on requests made by the module for external items. After allocating all of the modules comprising a program, the linker then fills each Link Table with the information necessary for communication between modules.

The format of a Link Table is given in Figure 2-5. A Link Table entry for an external variable contains the 32 -bit address of that variable. An entry for an external procedure contains a $32-$ bit procedure descriptor consisting of two 16-bit fields: Module and Offset. The Module field holds the new MOD register contents for the module containing the external procedure. The Offset field is an unsigned value giving the position of the external procedure's entry point relative to its module's Program Base pointer (Section 2.8.2).


Figure 2-5 Sample Link Table

### 2.8.4 Interrupt Dispatch Table and Cascade Table

The Series 32000 architecture supports handling of exceptions (traps and interrupts) through the Interrupt Dispatch Table. This table contains procedure descriptors (Section 2.8.3) for locating the service procedures assigned to each exception. The Interrupt Dispatch Table location is given by the INTBASE register.

The Interrupt Dispatch Table contains one 32-bit descriptor for each exception. A Series 32000 -based system can process up to 256 exceptions, depending on the system configuration. A Cascade Table may also exist, appended before the Dispatch Table.

For further details of interrupt and trap service, see the NS32532 data sheet.

NOTE: In memory-managed systems, the Interrupt Dispatch Table (and Cascade Table, if present) must always reside in physical memory. The INTBASE register contents are interpreted as a Supervisor-Mode virtual address. The Module portion of each procedure descriptor is also interpreted as a Supervisor-Mode virtual address.

### 2.8.5 Input and Output

Input and output ports are memory-mapped in Series 32000 -based systems. That is, all I/O devices are addressed as memory locations, and I/O operations are performed by reading from, or writing to, an I/O device as if it were a byte, word, or double-word of memory. There are no specific input and output instructions.

The hardware design of each individual system defines the number and type of $I / O$ devices as well as the addresses at which they are located. This is not defined by the Series 32000 architecture. However, the current implementation encourages two I/O assignments for interrupt handling, described below.

When a maskable interrupt occurs, an 8 -bit vector number is read from address FFFFFE00 (Hex). In memory-managed systems, this is a Supervisor-Mode virtual address, and must always have a valid mapping. Depending on the interrupt configuration mode (Vectored or Non-Vectored, Section 2.3), the vector value may not actually be used, but the read operation always occurs.

When a Non-Maskable Interrupt (NMI) occurs, the processor reads one byte from address FFFFFFOO (Hex). In memory-managed systems, this again is a SupervisorMode virtual address, and must always have a valid mapping. The processor does not use the data which was read.

Care should be taken in the system design to ensure that these read operations do not trigger side-effects.

For further details of interrupt service, see the NS32532 data sheet.

### 2.9 Privilege States and Protection

The Series 32000 family implements two privilege states: User Mode and Supervisor Mode.

The $U$ flag in the PSR determines the privilege state. When the $U$ flag is 1, the system is in User Mode, otherwise it is in Supervisor Mode.

A program running in User Mode is prevented from accessing privileged registers. These registers are:

* The most-significant byte of the Processor Status Register (PSR).
* The INTBASE register.
* The CFG register.
* The USP register.
* All Debug registers.
* All Memory Management registers.

The Interrupt Stack Pointer (SP0) is also implicitly protected by the fact that a User-Mode program cannot access the PSR S bit to select it for use.

User-Mode restrictions are enforced by the Illegal Operation trap, Trap (ILL), which occurs whenever a User-Mode program attempts to access a privileged register. Instructions which cause, or may cause, Trap (ILL) are listed in Table 2-1.

Programs running in Supervisor Mode have none of the above restrictions, as they are assumed to be trusted portions of an operating system.

In addition to the above restrictions, memory-managed systems can restrict access to memory pages based on the privilege state. Violations of such access restrictions cause the Abort trap, Trap (ABT). Since I/O devices are mapped as memory, they may also be protected by this mechanism as required.

| Instruction | Mnemonic |
| :--- | :--- |
| Load Processor Register (if INTBASE,PSR, USP, CFG, Debug) | LPRi |
| Store Processor Register (if INTBASE,PSR, USP, CFG, Debug) | SPRi |
| Bit Clear in PSR | (if Word length) |
| Bit Set in PSR | (if Word length) |
| Set Configuration | BICPSRW |
| Return from Trap | BISPSRW |
| Return from Interrupt | SETCFG |
| Load Memory Management Register | RETT |
| Store Memory Management Register | RETI |
| Move Value from Supervisor to User Space | LMR |
| Move Value from User to Supervisor Space | SMR |
| Validate Address for Reading | MOVSUi |
| Validate Address for Writing | MOVUSi |
| Cache Invalidate | RDVAL |

```
    Chapter 3
INSTRUCTIONS AND DATA TYPES
```

This chapter presents an overview of the Series 32000 instruction set by functional groups and describes the data types and structures on which they act.

The groups by which this chapter is organized are:

| Group | Section |
| :--- | :---: |
| Integer Instructions | 3.1 |
| Packed Decimal (BCD) Instructions | 3.2 |
| Floating-Point Instructions | 3.3 |
| Logical Instructions | 3.4 |
| Bit Instructions | 3.5 |
| Bit Field Instructions | 3.6 |
| String Instructions | 3.7 |
| Block Instructions | 3.9 |
| Array Instructions | 3.10 |
| Processor Control Instructions | 3.11 |
| Processor Service Instructions | 3.12 |
| Memory Management Instructions | 3.13 |

```
Instructions in each group are listed in three columns.
    Instruction: A brief instruction name.
    Mnemonic Forms: A list of all forms that the instruction mnemonic may
    take in assembly language.
Index: The general mnemonic form of the instruction. Chapter 5
    (Instruction Set) is organized alphabetically by this
    index.
```


### 3.1 Integer Instructions

Integer instructions operate on byte, word, and double-word integer operands. The following is a list of the Integer instructions:

## Instruction

Arithmetic

```
Add
Add Quick
Add with Carry
Subtract
Subtract with Carry [Borrow]
Negate
Absolute Value
Multiply
Multiply Extended Integer
Divide
Modulus
Quotient
Remainder
Divide Extended Integer
```

Movement and Conversion
Move
Move Quick
Move with Sign-Extension
Move with Zero-Extension

Comparison
Compare
Compare Quick

## Mnemonic Forms

## Index

| ADDB, ADDW, ADDD | ADDi |
| :--- | :--- |
| ADDQB, ADDQW, ADDQD | ADDQi |
| ADDCB, ADDCW, ADDCD | ADDCi |
|  |  |
| SUBB, SUBW, SUBD | SUBi |
| SUBCB, SUBCW, SUBCD | SUBCi |
|  |  |
| NEGB, NEGW, NEGD | NEGi |
| ABSB, ABSW, ABSD | ABSi |
| MULB, MULW, MULD |  |
| MEIB, MEIW, MEID | MULi |
|  |  |
| DIVB, DIVW, DIVD |  |
| MODB, MODW, MODD | DIVi |
| QUOB, QUOW, QUOD | MODi |
| REMB, REMW, REMD | QUOi |
| DEIB, DEIW, DEID | REMi |

MOVB, MOVW, MOVD MOVi MOVQB, MOVQW, MOVQD MOVQi MOVXBD, MOVXWD, MOVXBW MOVXii MOVZBD, MOVZWD, MOVZBW MOVZii

CMPB, CMPW, CMPD
CMPi
CMPQB, CMPQW, CMPQD

CMPQi

Integer operands are binary numbers. An integer operand may be a byte (8 bits), word (16 bits), or double-word ( 32 bits) in length. Its contents are interpreted as either signed or unsigned.

Unsigned integers range from 0 to 255 (byte), 0 to 65535 (word), and 0 to 4,294,967,295 (double-word). Each bit in an unsigned integer is a value bit, i.e., contributes to the integer's magnitude.

Signed integers are represented in two's-complement form. They range in value from -128 to 127 (byte), -32768 to 32767 (word), and $-2,147,483,648$ to 2,147,483,647 (double-word) and have the following form:

| +-+-------------+ |  |
| :--- | ---: |
| $!s!$ |  |
| $!-+-+-+-+-+-+-+-!$ |  |
| 7 | 0 |



The most significant bit in a signed integer indicates the sign of the number. A sign bit of zero specifies a positive value in which the remaining bits of the operand are in true binary form. A sign bit of one specifies a negative value, in which the remaining bits hold the two's complement of the absolute value of the operand. The sign bit does not contribute to the integer's magnitude.

The following illustrates a byte, word, and double-word integer and gives the signed and unsigned decimal interpretations for each.

| Binary | Signed <br> $\left(\begin{array}{l}\text { Decimal })\end{array}\right.$ | Unsigned <br> (Decimal) |
| ---: | ---: | :---: |
| 00000000000000000001001000110100 | -100 | -278 |
| 111111011101010 | 4660 | 65258 |
| 0660 |  |  |

Addition and subtraction operations yield the correct result regardless of whether the operands are interpreted as signed or unsigned. In the Quick instructions, however, one should note that the Quick immediate operand is sign-extended internally before use, and should therefore only be considered signed.

The other integer instructions treat integers as either signed or unsigned, as stated in their individual descriptions in Chapter 5.

## Integer Arithmetic

Integer arithmetic is performed to the length specified by the operation length appended to the instruction mnemonic by the programmer. This length may be byte, word or double-word (Section 4.1). Except where noted, the operands of these instructions are both general, meaning that general addressing mode expressions may be used independently to specify the location of each operand.

Addition instructions consist of ADDi, which adds two general operands, and ADDQi (Add Quick), which adds a small value (range -8 to +7) to a single general operand. Extended addition to any length can be performed using the ADDCi instruction, which adds also the contents of the PSR C bit (indicating a carry from a previous addition).

Subtraction (the SUBi instruction) may be modelled as adding together the second operand (the minuend), the one's complement of the first operand (the subtrahend), and the value 1. This definition, using the one's complement, is required to correctly define the overflow and borrow conditions (see "Exceptional Conditions" below). The result is placed in the location of the second operand. Extended subtraction to any length can be performed using the SUBCi instruction, which also subtracts the contents of the PSR C bit (indicating a borrow from a previous subtraction).

Negation (NEGi) and Absolute Value (ABSi) functions are provided. These instructions read a general (source) operand, convert it, and store the result in a second general operand location. Negation is performed by subtracting the source value from zero.

Multiplication is performed according to the standard rules of algebra. The length of the result may be selected as either the same length as the original operands (using the MULi instruction) or double that length (using the MEIi instruction). The MEIi instruction interprets its operands as unsigned integers, making it usable for multiplication to arbitrary length. The distinction between signed and unsigned operands is not relevant to the MULi instruction.

Division is performed according to three separate algorithms. The DIVi instruction divides the second operand by the first, producing as its result the nearest integer which is less than, or equal to, the exact quotient. The QUOi instruction produces the nearest integer whose absolute value is less than, or equal to, the exact quotient. These both interpret their operands as signed values. Note that they differ when the quotient is negative. The DEIi instruction divides a double-length integer (64, 32 or 16 bits) by a single-length divisor, and produces both a quotient and a remainder. It interprets its operands as unsigned for performing extended division; the distinction between the DIVi and QUOi algorithms is therefore irrelevant to this instruction. Remainder instructions are provided for both the DIVi and QUOi algorithms. The MODi (Modulus) instruction performs division according to the DIVi algorithm and produces the remainder as its result. The REMi (Remainder) instruction performs division as per the QUOi instruction and produces the corresponding remainder.

The MOVi instruction moves the first general operand to the second. A variation of this is the MOVQi instruction, which moves a small immediate value (range -8 to +7) into a general operand location.

An integer value can be converted to any greater length while being moved. The conversion for signed integers is provided by the MOVXii instructions, which perform sign-extension, and the conversion for unsigned integers is provided by the MOVZii instructions, which perform zero-extension.

## Comparison

Integer comparison instructions compare two operands and set the PSR Z, N and L bits to form a condition code. This condition code can be tested by subsequent instructions for program control or saved to generate operands for Boolean computations.

The CMPi instruction compares two general operands. The CMPQi instruction compares a general operand to a small immediate value (range -8 to +7).

The contents of the $P S R Z$ and $N$ bits indicate the result of comparing the operands as signed integers. The Z bit indicates equality when set. The N bit, when set, indicates that the first operand is greater than the second.

The contents of the $P S R Z$ and $L$ bits indicate the result of comparing the operands as unsigned integers. The $Z$ bit indicates equality when set. The L bit, when set, indicates that the first operand is greater than the second.

## Exceptional Conditions

Three exceptional conditions may occur in integer operations. These are a carry (or borrow), an overflow, or attempted division by zero.

Carry and borrow events are signaled in the Processor Status register C bit (Section 2.2). When an addition instruction is executed, the occurrence of $a$ carry out of the most significant bit position (bit 7, 15, or 31, depending on the selected operation length, Section 4.1) constitutes a "Carry" condition, and is indicated by setting the PSR C bit. If no carry occurs, the PSR C bit is cleared. When a subtraction instruction is executed, the lack of a carry out of the most significant bit position constitutes a "Borrow" condition, and the PSR C bit is set to indicate this exceptional condition. If a carry does occur, the PSR C bit is cleared. The result delivered follows the standard rules of binary two's-complement arithmetic, regardless of the occurrence of a carry or borrow condition.

Overflow events from addition and subtraction are signaled in the Processor Status Register F bit (Section 2.2). If the carry into the sign bit position and the carry out of the sign bit position do not agree, this constitutes an "overflow" condition, indicating that the correct result would be too great in magnitude to represent as a signed integer in the number of bits selected as the operation length (Section 4.1). If an overflow occurs in executing an addition or subtraction instruction, the PSR F bit is set, otherwise it is cleared. The result delivered follows the standard rules of binary two's-complement arithmetic (including alteration of the sign bit), regardless of the occurrence of an overflow.

Attempted division by zero always causes a trap, Trap(DVZ). This trap can occur in the DIVi, MODi, QUOi, REMi and DEIi instructions. A trapped instruction delivers no result, neither to the destination operand location nor to the PSR.

### 3.2 Packed Decimal Instructions

Packed Decimal instructions add and subtract packed decimal operands. There are two Packed Decimal instructions:

## Instruction

Add Packed Decimal

Subtract Packed Decimal

## Mnemonic Forms

ADDPB, ADDPW, ADDPD ADDPi<br>SUBPB, SUBPW, SUBPD

## Index

SUBPi

A packed decimal operand consists of two, four, or eight binary-coded decimal (BCD) digits stored in a byte, word, or double-word, respectively. A BCD digit is a 4-bit field whose value is within the range 0 to 9, encoded as binary 0000 to 1001, respectively. Each byte contains two BCD digits as illustrated below. Digit dO is the least-significant digit.

```
+-------+---------+
! d1 ! d0 1
!-+-++-+-+-++-+-+-!
7
0
```




Packed Decimal instructions operate on two general operands. Both operands are interpreted as unsigned numbers. The ADDPi instruction places the sum of the two operands, plus the contents of the PSR C bit, into the second operand location. The SUBPi instruction subtracts the first operand from the second, subtracting also the contents of the PSR C bit, and places the result into the second operand location. Incorporation of the $P S R$ C bit into the result facilitates use of these instructions in performing packed decimal calculations to arbitrary length.

Decimal subtraction can be modeled as adding the ten's complement of the subtrahend to the minuend.

Both operands must contain only legal BCD digits. If either operand contains digits which are not legal, the result value is undefined, and the setting of the PSR C bit is undefined.

## Exceptional Conditions

A decimal carry or borrow condition can occur from Packed Decimal instructions. Decimal carry and borrow events are signaled in the Processor Status register C bit (Section 2.2).

When the ADDPi instruction is executed, the occurrence of a carry out of the most-significant digit position constitutes a "carry" condition, and is indicated by the CPU by setting the PSR C bit. This indicates that the sum is too large to be held as a Packed Decimal number in the length of the original operands. The result produced is the least-significant portion of the entire result.

If no carry occurs, the $P S R C$ bit is cleared.

When the SUBPi instruction is executed, the lack of a carry out of the most significant digit position constitutes a "borrow" condition, and the PSR C bit is set to indicate this. A borrow condition indicates that a high-order "1" digit has been assumed to the left of the most-significant minuend digit in order to produce a positive result.

If a carry does occur from subtraction, the PSR C bit is cleared.

### 3.3 Floating-Point Instructions

Floating-Point instructions operate on floating-point numbers. Included also in this group are the instructions which load and store the Floating-Point Status register (FSR). The following is a list of the Floating-Point instructions:

## Instruction

```
Add Floating
Subtract Floating
Multiply Floating Divide Floating
Dot Product Floating
Polynomial Floating
```

Negate Floating
Absolute Value Floating
Compare Floating
Move Floating
Logartithm Binary Floating
Scale Binary Floating
Move Long Floating to Floating Move Floating to Long Floating

Move Integer to Floating

Round Floating to Integer
Truncate Floating to Integer
Floor Floating to Integer

Load FSR
Store FSR

## Mnemonic Forms

| ADDF, ADDL | ADDf |
| :--- | :--- |
| SUBF, SUBL | SUBf |
| MULF, MULL | MULf |
| DIVF, DIVL | DIVf |
| DOTF, DOTL | DOTf |
| POLYF, POLYL | POLYf |
| NEGF, NEGL |  |
| ABSF, ABSL | NEGf |
| CMPF, CMPL | ABSf |
| MOVF, MOVL | CMPf |

LOGBF, LOGBL LOGBf
SCALBF, SCALBL SCALBf
MOVLF MOVLF
MOVFL MOVFL

MOVBF, MOVWF, MOVDF, MOVif MOVBL, MOVWL, MOVDL ROUNDFB, ROUNDFW, ROUNDFD, ROUNDLB, ROUNDLW, ROUNDLD TRUNCFB, TRUNCFW, TRUNCFD, TRUNCLB, TRUNCLW, TRUNCLD FLOORFB, FLOORFW, FLOORFD, FLOORLB, FLOORIW, FLOORLD

| LFSR | LFSR |
| :--- | :--- |
| SFSR | SFSR |

Floating-Point arithmetic operations are performed by the ADDf, SUBf, MULf, DIVf, DOTf and POLYf instrucions. The NEGf and ABSf instructions move the negative or the absolute value of their first operand to the second operand location. The CMPF instruction compares two floating-point values, setting the PSR condition codes as per the CMPi (integer compare) instruction. The MOVf instruction moves a floating-point value. LOGBf and SCALBf instructions perform operation on the exponent of a floating-point number.

The full range of conversions are provided; between floating-point types, and between any integer and floating-point types. Conversion from floating-point to integers can be performed by rounding to nearest (ROUNDfi), toward zero (TRUNCfi) or toward negative infinity (FLOORfi).

The LFSR and SFSR instructions load and store the FSR, which holds mode and status information pertaining to floating-point operations (Section 2.4.2).

### 3.3.1 Floating-Point Operand Formats

The Series 32000 Floating-Point instruction set operates on two floating-point data types: single precision (32 bits) and double precision (64 bits). Floating-Point instruction mnemonics use the operation length suffix $F$ (Floating) to specify the single precision data type and the suffix $L$ (Long Floating) to specify the double precision data type.

A floating-point number is divided into three fields as shown in Figure 3-1.



Figure 3-1 Floating-Point Operand Formats

The $F$ field is the fractional portion of the represented number. The binary point is assumed to be immediately to the left of the most-significant bit of the F field, with an implied 1 bit to the left of the binary point. Thus, the $F$ field represents values from 1.0 (inclusive) to 2.0 (exclusive) as shown in Table 3-1.

Table 3-1 SAMPLE F FIELDS

| F Field | Binary Value | Decimal Value |
| :---: | :---: | :---: |
| $000 \ldots 0$ | $1.000 \ldots 0$ | $1.000 \ldots 0$ |
| $010 \ldots 0$ | $1.010 \ldots 0$ | $1.250 \ldots 0$ |
| $100 \ldots 0$ | $1.100 \ldots 0$ | $1.500 \ldots 0$ |
| $110 \ldots 0$ | 1.110............0.0 | $1.750 \ldots 0$ |
|  | Implied |  |

The E field holds an unsigned number which gives the binary exponent of the represented number. The value in the E field is biased; that is, a constant bias value must be subtracted from the value in the E field in order to obtain the true exponent. This bias value is 011...11 (binary), which is either the value 127 (in single precision) or 1023 (in double precision). Thus, the true binary exponent can be either positive or negative, as shown in Table 3-2.

Table 3-2 Sample E Fields

| E Field | F Field | Represented Value |
| :---: | :---: | :---: |
| $011 \ldots 110$ | $100 \ldots 0$ | $1.5 * 2^{-1}=0.75$ |
| $011 \ldots 111$ | $100 \ldots 0$ | $1.5 * 2^{0}=1.50$ |
| $100 \ldots 000$ | $100 \ldots 0$ | $1.5 * 2^{1}=3.00$ |

NOTE: Two forms of the E field represent special values, and are not interpreted as binary exponent values. 11...11 represents a value which is a Reserved operand (Section 3.3.4). 00...00 represents the value Zero (Section 3.3.3) if the $F$ field is also all zeroes, otherwise the represented value is a Reserved operand.

The S bit indicates the sign of the operand: 0 for positive and 1 for negative. Floating-Point numbers are represented in sign-magnitude form, such that only the $S$ bit is complemented in order to change the sign of the represented number.

### 3.3.2 Normalized Numbers

Normalized numbers are numbers in floating-point format, where the E field is neither all zeroes nor all ones.

The value represented by a normalized number is determined by the formula:

$$
(-1)^{S} * 2^{(\mathrm{E}-\mathrm{Bias})} * 1 . \mathrm{F}
$$

The ranges of normalized numbers are given in Table 3-3.

Table 3-3 NORMALIZED FLOATING-POINT RANGES

|  | Single Precision | Double Precision |
| :---: | :---: | :---: |
| Most Positive | 127 -23 | 1023 -52 |
|  | $2 *(2-2)$ | $2 *(2-2)$ |
|  | $=3.40282346 * 10$ | $=1.7976931348623157 * 10$ |
| Least Positive | -126 | -1022 |
|  | 2 | 2 |
|  | $=1.17549436 * 10^{-38}$ | $=2.2250738585072014 * 10^{-308}$ |
| Least Negative | -126 | -1022 |
|  | - (2) | -(2) |
|  | $=-1.17549436 * 10^{-38}$ | $=-2.2250738585072014 * 10^{-308}$ |
| Most Negative | 127 -23 | 1023 -52 |
|  | -2 * (2-2) | -2 * (2-2 ) |
|  | $=-3.40282346 \times 10$ | $=-1.7976931348623157 \times 10$ |

NOTE: The values given are extended one full digit beyond their represented accuracy to help in generating rounding and conversion algorithms.

### 3.3.3 Zero

There are two representations for zero -- a positive form and a negative form. Positive zero has all-zero $F$ and E fields, and its S bit is zero. Negative zero also has all-zero $F$ and E fields, but its sign bit is one. In spite of these differences, the two zeroes are considered equal to each other when compared using the CMPf instruction.

### 3.3.4 Reserved Operands

The proposed IEEE Standard for Binary Floating Point Arithmetic (IEEE Task P754) provides for certain exceptional forms of floating-point operands. The Series 32000 hardware currently treats these forms as reserved operands. The reserved operands are:

* Positive and Negative Infinity
* Not-a-Number (NaN) values
* Denormalized numbers

Both Infinity and NaN values have all one's in their E fields. Denormalized numbers have all zeroes in their E fields and non-zero values in their $F$ fields.

The Series 32000 hardware causes an Invalid Operation trap (Section 3.3.7) if it receives a reserved operand, unless the instruction being executed is a simple MOVf instruction (move without conversion). The Series 32000 hardware does not generate reserved operands as results of floating-point calculations. The trapping mechanism used in the Series 32000 family allows handing of these operand forms transparently in software.

### 3.3.5 Integers

Some floating-point instructions perform conversions between integer and floating-point data types. Integers are accepted and generated as two's complement values of byte, word or double-word length, as specified in the conversion instruction.

### 3.3.6 Memory Representations

Floating-Point operands are stored in memory with the least-significant byte at the lowest address, except in the Immediate addressing mode. In this mode, the operand is held within the instruction format with the most-significant byte at the lowest address.

### 3.3.7 Floating-Point Traps

## Trap (UND)

The Floating-Point instruction set is made available to a Series 32000-based system with an NS32381 Floating-Point Unit by setting the $F$ bit in the CFG register (Section 2.3). If the CFG $F$ bit is not set, any floating-point instruction causes the Undefined Instruction trap, Trap (UND). In systems without floating-point hardware, Trap (UND) can be used to transfer control to float-ing-point emulation software.

## Trap (FPU)

Any exceptional conditions encountered during the execution of a floating-point instruction will cause a floating-point trap. This trap is labeled Trap (FPU) and uses the fourth entry (entry \#3) of the Interrupt Dispatch Table.

The following are true for any floating-point instruction causing Trap (FPU):

1. The status fields of the $F$ SR are updated before trapping.
2. No other result is delivered, neither to the destination operand location nor to the Processor Status Register (PSR).
3. The return address pushed onto the Interrupt Stack is the address of the first byte of the trapped instruction. This allows software analysis or emulation of the trapped instruction, or re-execution after the exception has been logged.

The conditions which cause Trap (FPU) are:

1. Underflow. A non-zero floating-point result is too small in magnitude to be represented as a normalized floating-point number in the format of the destination operand. This condition is always reported in the FSR TT field and UF bit, but causes a Trap (FPU) only if the FSR UEN bit is set. If the UEN bit is not set, a result of Positive Zero is produced, and no trap occurs.
2. Overflow. A result (either floating-point or integer) of a floatingpoint instruction is too great in magnitude to be held in the format of the destination operand. Note that rounding, as well as calculations, can cause this condition.
3. Divide by Zero. An attempt has been made to divide a non-zero floating-point number by zero. Dividing zero by zero is considered an Invalid Operation instead (below). Note that the trap caused by this condition is still Trap (FPU) and not Trap (DVZ), which is caused only by integer instructions.
4. Illegal Instruction. Two undefined floating-point instruction forms cause Trap (FPU) rather than Trap (UND). The binary formats causing this trap are:
xxxxxxxxxx0011xx10111110
xxxxxxxxxx1001xx10111110
5. Invalid Operation. One of the floating-point operands of a floatingpoint instruction is a Reserved operand (Section 3.3.4), or an attempt has been made to divide zero by zero using the DIVf instruction.
6. Inexact Result. The result (either floating-point or integer) of a floating-point instruction cannot be represented exactly in the format of the destination operand, and a rounding step must alter it to fit. This condition is always reported in the FSR TT field and IF bit unless any other exceptional condition has occurred in the same instruction. In this case, the $T T$ field always contains the code for the other exception and the IF bit is not altered. A Trap (FPU) is caused by this condition only if the FSR IEN bit is set; otherwise the result is rounded and delivered, and no trap occurs.

### 3.4 Logical Instructions

Logical instructions perform masking, shifting and Boolean arithmetic operations. The following table lists the logical instructions:

## Instruction

Arithmetic

Logical AND
Logical OR
Bit Clear
Exclusive OR
Complement

## Mnemonic Forms

## Index

| ANDB, ANDW, ANDD | ANDi |
| :--- | :--- |
| ORB, ORW, ORD | ORi |
| BICB, BICW, BICD | BICi |
| XORB, XORW, XORD | XORi |
| COMB, COMW, COMD | COMi |


| ASHB, ASHW, ASHD | ASHi |
| :--- | :--- |
| LSHB, LSHW, LSHD | LSHi |
| ROTB, ROTW, ROTD | ROTi |

NOTB, NOTW, NOTD NOTi

ScondB, ScondW, ScondD Scondi

The arithmetic instructions perform bitwise Boolean arithmetic on byte, word or double-word general operands. The shift instructions perform shifting on byte, word or double-word general operands. The Boolean instructions generate and complement Boolean values.

The ANDi, ORi and XORi instructions perform the bitwise Boolean AND, OR and Exclusive OR functions between two general operands. The BICi instruction performs an AND NOT operation, clearing all bits in the second operand which are set in the first. The COMi instruction moves the bitwise complement of the first operand to the second.

The shift instructions shift their second general operand in the direction and by the magnitude given by the first operand (a positive shift is left, a negative shift is right). The logical shift fills the emptied bit positions with zeroes always. The arithmetic shift fills these locations with zeroes if the shift is to the left, and with the original contents of the sign bit (the most-significant bit) if the shift is to the right. The rotation shift consecutively replaces each bit emptied with the contents of the bit shifted out of the operand.

NOTE: The result generated by shifting an operand by a count which is greater than, or equal to, its length in bits is undefined.

The Boolean instructions generate and handle unpacked Boolean values, defined as integers whose values are interpreted as $0=$ False and $1=$ True. This definition follows conventions established by several high-level languages which require that True be greater than False when compared and that conversions between Boolean and integer variables generate the above correlation between values.

All of the logical arithmetic instructions perform correct Boolean arithmetic on Boolean values except the COMi instruction. To allow complementing Boolean values (from True to False and vice versa), the NOTi instruction is provided, which complements only the least-significant bit of its first operand, placing the result in the second.

Because Boolean arithmetic often deals with values derived from relational operations (e.g. whether one value is greater than another), the Save Condition (Scondi) instruction is provided, which generates a Boolean value based on a condition code test.

### 3.5 Bit Instructions

Bit instructions perform or support manipulation of individual bits in General Purpose Registers or memory. The following is a list of the Bit instructions:

## Instruction

Test Bit

Set Bit

Clear Bit

Invert Bit

Find First Set Bit

Convert to Bit Pointer

Mnemonic Forms

| TBITB, TBITW, TBITD | TBITi |
| :--- | :--- |
| SBITB, SBITW, SBITD, | SBITi, |
| SBITIB, SBITIW, SBITID | SBITIi |
| CBITB, CBITW, CBITD, | CBITi, |
| CBITIB, CBITIW, CBITID | CBITIi |
| IBITB, IBITW, IBITD | IBITi |
| FFSB, FFSW, FFSD | FFSi |
| CVTP | CVTP |

The TBIT instruction tests a bit by copying its contents to the PSR F bit. The SBIT, CBIT and IBIT instructions test the specified bit, and then either set, clear or invert it. The SBITI and CBITI instructions, in addition, allow testing and either setting or clearing of $a$ bit in an indivisible operation for handling multiprocessor semaphores.

The FFSi and CVTP instructions do not operate on bits, but provide related functions to aid in bit handling. The FFSi instruction scans a byte, word or double-word for a set bit, producing its position as a one-byte offset value. The CVTP instruction generates the bit address of a specified bit.

Bit positions are specified using two general operand specifications: a base and an offset, as in the instruction.

## TBITi offset,base

The base operand specification is used only to determine a base location (either a memory address or a register) relative to which the bit is to be located, and does not itself reference an operand at that location. The offset is a general operand of byte, word or double-word length, as specified by the operation length selected by the programmer (Section 4.1). It contains a signed integer which specifies the position of the desired bit relative to bit 0 of the location specified as the base.

If the base is specified as a General Purpose register, the offset must be within the range 0 to 31, inclusive. If the offset is outside this range, the location of the bit is undefined.

If the base is specified as a memory address, the offset specifies a bit in memory.

Both positive and negative offsets are allowed and meaningful. An offset of 0 specifies bit 0 of the byte at the base address. An offset of 8 specifies bit 0 of the byte at the next higher address. An offset of -1 specifies bit 7 of the byte at the next lower address, and an offset of -8 specifies bit 0 of the byte at the next lower address.

The maximum range of a double-word offset is $-2,147,483,648$ to $+2,147,483,647$ bits, corresponding to an addressing range of $-268,435,456$ to $+268,435,455$ bytes from the specified base.

The address of the byte containing the desired bit is formally defined as

```
EA(base) + (offset DIV 8)
```

where "EA(base)" is the effective address calculated from the base operand specification and "offset DIV 8" is the nearest integer less than or equal to offset/8 (as per the DIVi instruction). The bit number of the desired bit is computed as

```
offset MOD 8
```

where MOD is the modulus function (as per the MODi instruction).

The following examples illustrate the interpretations of various bit specifications:

Example 1:


Example 2:


Offset: +28 EA(Base): 1000
Interpreted as bit 4 of the byte at address 1003.
In this example, the address of the byte containing the desired bit is $1000+(28$ DIV 8), or 1003, since 28 DIV $8=3$. The bit number within this byte is 28 MOD 8, or 4.

Example 3:


In this example, the address of the byte containing the desired bit is $1003+(-13$ DIV 8), or 1001, since -13 DIV $8=-2$. The bit number within this byte is -13 MOD 8, or 3. If these results look confusing, consult again the definitions of the DIV and MOD operations given above.

### 3.6 Bit Field Instructions

Bit Field instructions copy information to and from unaligned fields in General Purpose Registers or memory. The following is a list of the Bit Field instructions:

## Instruction

Extract Field<br>Extract Field Short<br>Insert Field<br>Insert Field Short

## Mnemonic Forms

EXTB, EXTW, EXTD
EXTSB, EXTSW, EXTSD
INSB, INSW, INSD

INSSB, INSSW, INSSD

## Index

EXTi
EXTSi
INSi
INSSi

Extract instructions read a bit field and place it into a byte, word, or doubleword general operand, right-justified. Insert instructions replace a bit field from aligned information in a general operand. A bit field may be one to 32 bits in length.

A bit field is fully specified by the position of its least-significant bit and its length in bits. The position of the least-significant bit is specified as in the Bit instructions (Section 3.5), using a general operand specification for the base and an offset contained either in a General Purpose Register or (in the "Short" forms of these instructions) in an immediate constant. The length of the field is specified as an immediate constant, which must specify a length in the range of 1 to 32 bits, inclusive. The interpretation of any length specified outside this range is undefined.

The general bit field instructions (EXTi and INSi) allow a 32-bit offset value to be dynamically specified in a General Purpose Register, supporting the indexing necessary to access structures such as Pascal packed arrays. The "Short" bit field instructions (EXTSi and INSSi) eliminate the overhead of loading a register when the offset is fixed, as is commonly the case in accessing structures such as Pascal packed records.

If the base is specified as a General Purpose Register, the bit field is in that register. The offset must be within the range 0 to 31 , and the entire bit field must be contained within the specified register, otherwise the location of the bit field is undefined.

If the base is specified as a memory address, the offset specifies a bit in memory as the least-significant bit of the field. Both positive and negative offsets are allowed and meaningful, as in Bit instructions (Section 3.5).

As in the Bit instructions, the address of the byte containing the leastsignificant bit of the field is defined as

$$
\text { EA(Base) }+ \text { (offset DIV 8) }
$$

where "EA(Base)" is the effective address calculated from the base operand specification and "offset DIV 8" is the nearest integer less than, or equal to, offset/8 (as per the DIVi instruction). The bit number of the least-significant bit in the field is computed as
offset MOD 8
where MOD is the modulus function (as per the MODi instruction).

NOTES: 1. The current implementation of bit field instructions places an alignment restriction on bit fields greater than 25 bits in length. This restriction is imposed due to the fact that a field in memory is accessed in a double-word transfer starting at the byte containing the least-significant bit of the field. A bit field in memory must be composed of bits from no more than four contiguous bytes. For a field of 25 bits or less, this imposes no restriction on alignment, as it is impossible for such a field to span more than four bytes.
2. Regardless of the length of a bit field in memory, it is always accessed by Bit Field instructions as a double-word starting with the byte which contains the least-significant bit of the field. The Extract instructions read a full double-word, and the Insert instructions read, modify and rewrite a full double-word. These instructions can therefore cause a page fault in memory-managed systems if the field is close to the end of a page. In multiprocessor systems, care should be taken to ensure that the processors do not attempt to modify adjacent fields simultaneously.

The following examples illustrate how a bit field is located in a register and in memory:

Example 1:


Base: R0 Offset: 6 Length: 11
Interpreted as an l1-bit field in register $R 0$ starting with bit 6.

Example 2:


BA(Base): 1000 Offset: +9 Length: 20
Interpreted as a 20-bit field starting at bit 1 of address 1001.
In this example, the address of the byte containing the least-significant bit of the field is $1000+(9$ DIV 8), or 1001. The bit number of the first field bit within that byte is 9 MOD 8, or 1.

Example 3:


EA(Base): 1003 Offset: -13 Length: 3
Interpreted as a 3-bit field starting at bit 3 of address 1001.
In this example, the address of the byte containing the least-significant bit of the field is $1003+(-13$ DIV 8), or 1001 , since -13 DIV $8=-2$. The bit number of the first field bit in that byte is -13 MOD 8 , or 3 . If these results look confusing, consult again the definitions of the DIV and MOD operations given above.

### 3.7 String Instructions

String instructions operate on strings of integer elements. The following is a list of the String instructions:

## Instruction

```
Move String
Move String, Translating
Compare Strings
Skip String
Skip String, Translating
```


## Mnemonic Forms

MOVSB, MOVSW, MOVSD MOVSi
MOVST MOVST

CMPSB, CMPSW, CMPSD
CMPST CMPST

SKPSB, SKPSW, SKPSD SKPSi
SKPST SKPST

A string is a sequence of integer elements, all of the same length, stored in consecutive memory locations. Elements of a string may be bytes, words, or double-words as specified by the operation length (Section 4.1), except when the Translating form (above) is used, in which case the elements must be bytes.

String instructions operate on either one or two strings. These strings are designated String 1 and String 2. The MOVS instructions copy elements from String 1 to String 2. The CMPS instructions compare String 1 elements to the corresponding String 2 elements. The SKPS instructions scan elements of String 1, without using a String 2.

String locations and length are specified by the General Purpose registers R0, R1, and R2. Before instruction execution, the registers must be set to the following:

R0 -- the maximum number of elements to be processed
R1 -- the address of the first element of String 1
R2 -- the address of the first element of String 2 (except for SKPS, which does not use or modify R2)

NOTE: The number of elements processed is undefined if register RO contains a negative number.

String instructions process the elements of the string(s) one at a time until a specified termination condition is reached. After each element is processed, the instructions modify the 32 -bit contents of registers R0, R1, and R2 so that they contain the following values:

```
R0 -- the number of elements left to be processed (old contents minus one)
R1 -- the address of the next element of String 1
R2 -- the address of the next element of String 2 (except in SKPS)
```

If the resulting value in $R 0$ is zero, the instruction terminates. The contents of register $R 2$ always remain unchanged by the SKPS instruction.

## Options

String instructions have the following options:

* Translation (T)
* Backward (B)
* Until Match (U)
* While Match (W)

Additional information required by these options is specified in General Purpose registers R3 and R4, as follows:

R3 -- the address of a translation table, required if the Translation option is specified

R4 -- a termination value, required if the Until Match or While Match option is specified

Registers R 3 and R 4 remain unchanged by the instruction.

The Translation option causes a string instruction to translate each String 1 element before using it. String instructions with the Translation option operate on l-byte elements only, and because of this the Translation option is specified as a mnemonic suffix "T" replacing the operation length suffix.

Translation is performed by using the String 1 element value as an unsigned index into a translation table, whose base address is taken from register R3. A byte is read from this table location, and is used in place of the original String 1 element.

The Backward option causes a string instruction to reverse its direction, processing string elements from successively lower memory addresses instead of successively higher addresses. This means that registers R1 and R2 are decremented by the element length after each element is processed instead of being incremented. The Backward option is specified in assembly language by listing the letter $B$ in the instruction as an operand. When used in conjunction with the Until Match or While Match option, it must be separated with a comma.

The Until Match and While Match options specify a termination condition based on whether the contents of each String 1 element match the contents of register R4 (after translation, if that option is also specified). In order to distinguish this termination condition from any other, the PSR F bit is set to 1 before termination. The Until Match and While Match options are mutually exclusive.

If the Until Match option is specified, the instruction terminates as soon as the current value matches R4. This option is specified in assembly language by listing the letter $U$ in the instruction as an operand.

If the While Match option is specified, the instruction terminates as soon as the current value does not match R4. This option is specified in assembly language by listing the letter $W$ in the instruction as an operand.

## Option Encoding

Each string instruction contains a 4-bit field defining which options are specified. The field has the following form:


The 1-bit $T$ field defines the state of the Translation option. If the field is 1, the Translation option is in effect; otherwise, the option is not in effect. If the $T$ bit is set, the operation length field (i) must contain binary 00 (Byte).

The 1-bit B field defines the state of the Backward option. If the field is 1 , the option is in effect; otherwise, the option is not in effect.

The 2-bit UW field defines the state of the Until and While options, as given below:

| 00 | neither option |
| :--- | :--- |
| 01 | While Match |
| 10 | (reserved) |
| 11 | Until Match |

## Interrupts During String Instructions

String instructions are interruptible. If an interrupt is asserted during a String instruction, the CPU first finishes processing the current string element. It then saves the address of the String instruction as the return address and passes control to the interrupt service procedure. When the interrupt service procedure returns, the String instruction is re-executed, but because the registers have been updated this has the effect of continuing string processing from the point where the instruction was interrupted. Note that the interrupt service procedure must follow the standard practice of restoring all registers used before returning.

A string instruction terminates for one of the following reasons:

1. The limit count originally specified in register R0 has been decremented to zero, or was zero at the beginning of the instruction.
2. The CMPS instruction has found a pair of string elements which are unequal and has, therefore, determined which string has the greater value.
3. The Until Match or While Match option is in effect and the string instruction has found an element in String 1 which meets the specified termination condition.

When a string instruction terminates due to its limit count, the resulting state of the machine is as follows:

PSR - bit $F=0$. If a CMPS instruction terminates for this reason, then also PSR bits $Z=1, N=0, L=0$.
R0 -- contains 0.
R1 -- contains the address of the next unprocessed String 1 element.
R2 -- contains the address of the next unprocessed String 2 element (except in SKPS).

When a CMPS instruction finds an unequal pair of string elements, the resulting state of the machine is:

PSR - bits $F=0$ and $Z=0$. The $N$ and L bits indicate the relation between the two unequal string elements.
R0 -- contains the number of element pairs left to be processed (this includes the element pair which caused termination).
R1 -- contains the address of the String 1 element which caused termination.
R2 -- contains the address of the String 2 element which caused termination.

Whenever the Until Match or While Match option terminates execution of a string instruction, the resulting machine state is:

```
PSR - bit F = 1. If a CMPS instruction terminates for this reason, then
        also PSR bits Z = 1, N = 0, L = 0.
R0 -- contains the number of elements left to be processed (this includes
        the element which caused termination).
R1 -- contains the address of the element in String 1 which caused
    termination
R2 -- contains the address of the element in String 2 which corresponds to
    the String 1 element which caused termination (except in SKPS)
```

The contents of registers $R 3$ and $R 4$ always remain unchanged.

Table 3-4 below gives the detailed execution sequences followed by the string instruction. A temporary holding location within the processor is referenced by the name "TEMP".

Table 3-4 Execution Sequences

|  | CMPS | MOVS | SKPS |
| :---: | :---: | :---: | :---: |
| 1 | In the PSR, set bits $\mathrm{Z}=1, \mathrm{~N}=0, \mathrm{~L}=0$ <br> If R0=0, set the PSR F bit to 0 and terminate the instruction. | If $R 0=0$, set the $P S R$ bit to 0 and terminate the instruction. | If $R 0=0$, set the PSR $F$ bit to 0 and terminate the instruction. |
| 2 | Read the current String 1 element (address in R1) from memory into TEMP. | Read the current String 1, element (address in R1) from memory into TEMP. | Read the current String 1 element (address in R1) from memory into TEMP. |
| 3 | If the translation option is selected, then zero-extend TEMP from 8 bits to 32 bits and add it to the contents of R3, generating the address of a translation table entry. Read a byte from this memory location and place it into TEMP. | If the translation option is selected, then zero-extend TEMP from 8 bits to 32 bits and add it to the contents of R3, generating the address of a translation table entry. Read a byte from this memory location and place it into TEMP. | If the translation option is selected, then zero-extend TEMP from 8 bits to 32-bits and add it to the contents of R3, generating the address of a translation table entry. Read a byte from this memory location and place it into TEMP. |
| 4 | If the Until Match or While Match option is specified, then compare TEMP to R4, interpreting both as integers of the size specified by the operation length. <br> If the Until Match option is specified, and TEMP and R4 are equal, then set the PSR $F$ bit to 1 and terminate the instruction. <br> If the While Match option is specified, and TEMP and R4 are unequal, then set the PSR $F$ bit to 1 and terminate the instruction. | If the Until Match or While Match option is specified, then compare TEMP to R4, interpreting both as integers of the size specified by the operation length. <br> If the Until Match option is specified, and TEMP and R4 are equal, then set the PSR $F$ bit to 1 and terminate the instruction. <br> If the While Match option is specified, and TEMP and R4 are unequal, then set the PSR F bit to 1 and terminate the instruction. | If the Until Match or While Match option is specified, then compare TEMP to R4, interpreting both as integers of the size specified by the operation length. <br> If the Until Match option is specified, and TEMP and R4 are equal, then set the PSR $F$ bit to 1 and terminate the instruction. <br> If the While Match option is specified, and TEMP and R4 are unequal, then set the PSR F bit to 1 and terminate the instruction. |
| 5 | Compare TEMP to the contents of the current String 2 location (address in R2) and update PSR bits Z, N and L to reflect the result. If the resulting $Z$ bit is zero (meaning not equal), then set PSR F bit to 0 and terminate the instrucion. | Write TEMP to the String 2 location (address in R2). | Do nothing; continue to Step 6. |
| 6 | If the Backward option is specified, decrement R1 and R2 by the length in bytes specified by the operation length. Otherwise increment R1 and R2 by this amount. | If the Backward option is specified, decrement R1 and R2 by the length in bytes specified by the operation length. Otherwise increment R1 and R2 by this amount. | If the Backward option is specified, decrement R1 and R2 by the length in bytes specified by the operation length. Otherwise increment R1 and R2 by this amount. |
| 7 | Decrement R0 by 1. | Decrement R0 by 1. | Decrement R0 by 1. |
| 8 | ```If an interrupt is pending, service it here. Otherwise go to Step 1.``` | ```If an interrupt is pending, service it here. Otherwise go to Step 1.``` | ```If an interrupt is pending, service it here. Otherwise go to Step 1.``` |

### 3.8 Block Instructions

Block instructions move and compare byte, word, and double-word elements stored in contiguous blocks of memory. There are two block instructions:

## Instruction

Move Multiple

Compare Multiple

## Mnemonic Forms

MOVMB, MOVMW, MOVMD

CMPMB, CMPMW, CMPMD

## Index

MOVMi

CMPMi

A block is a small string (16 bytes or less) of integers.
Block instructions differ from their string counterparts in three major ways:

1. They require no overhead in setting up registers, as both block operands are general.
2. They are not interruptible.
3. They are limited to blocks of 16 bytes or less so that they do not adversely affect interrupt latency.

Block instructions have three operands: block1, block2, and length. The MOVMi instruction copies block1 to block2. The CMPMi instruction compares the elements of block1 to the corresponding block2 elements, indicating in PSR bits Z, $N$ and L which block contains the greater value, or whether they are equal.

Block1 and block2 are general operands which must be in memory (access class addr, Section 4.2.1).

The length operand is an immediate value which specifies the length of each block. In assembly language, length is specified as the number of elements (bytes, words or double-words) in the block. (This is not the value which is encoded in the binary form of the instruction.) Since a block must contain at least one byte and no more than 16 bytes, the range of values for length depends on the instruction's operation length suffix (B, W, or D: Section 4.1) as shown by the following:

## Operation Length Suffix

B
W
D

## length

1 to 16
1 to 8
1 to 4

In the binary form of the instruction, the block length is encoded in a displacement field and appended to the basic instruction. The displacement field contents are to be computed from the specified length value as
(length - 1) * i
where i is the element size in bytes: 1 (for B), 2 (for $W$ ), or 4 (for D).
NOTES: 1. The two block operands of the MOVMi instruction must not overlap. If they do overlap, the resulting values in the destination block are undefined.
2. If the binary contents of the length operand differ from those values which can be derived from the expression above, the length of the blocks is undefined.

### 3.9 Array Instructions

Array instructions operate in conjunction with the Scaled Indexing addressing mode option (Section 4.4.9) to support random accesses into single- and multi-dimensional arrays. The following ia a list of the array instructions:

## Instruction

Bounds Check
Calculate Index

## Mnemonic Forms

CHECKB, CHECKW, CHECKD
INDEXB, INDEXW, INDEXD

## Index

CHECKi
INDEXi

An array consists of a number of elements of the same length, stored in a contiguous block of memory. An array can be of a single dimension (i.e., a vector) or of multiple dimensions (i.e., a matrix). Individual elements in an array are accessed using one subscript or index expression per dimension.

The CHECKi instruction performs a bounds check on any general operand, checking whether its value is within the range specified by a pair of values in another general operand. If so, it zero-adjusts the value by subtracting the lower bound from it, and places the result in any specified General Purpose Register. If not, it indicates an error in the PSR F bit, which can be used either as a branch condition or to cause a trap (see the FLAG instruction). If the value being checked is an index into a single-dimensional array, the result placed in the register is directly usable with Scaled Indexing to access the indicated array element.

The INDEXi instruction is used for accesses into multidimensional arrays. Its purpose is to calculate a single l-dimensional index based on the values of the indexes (one per dimension) by which the desired element is specified. The order in which the indexes are incorporated into the result depends on the scheme used for ordering the array elements in memory.

Depending on the high-level language, array storage ordering generally follows one of two schemes. Row major ordering, the most popular, and typical of the Pascal and C languages, is shown in Table 3-5. Column major ordering, typical of FORTRAN, is shown in Table 3-6. Note that in row major ordering it is the rightmost index which is incremented with consecutive element addresses, and in column major ordering it is the leftmost.

Table 3-5 Row Major Ordering
Pascal array declaration:

VAR A: ARRAY[1..2,1..3,1..2]
OF INTEGER;

Element size: 4 bytes Base address: 1000 (Hex)

| Array Element | Address (Hex) |
| :---: | :---: |
| A $[1,1,1]$ | 1000 |
| A $[1,1,2]$ | 1004 |
| A $[1,2,1]$ | 1008 |
| A $[1,2,2]$ | $100 C$ |
| A $[1,3,1]$ | 1010 |
| A $[1,3,2]$ | 1014 |
| A $[2,1,1]$ | 1018 |
| A $[2,1,2]$ | $101 C$ |
| A $[2,2,1]$ | 1020 |
| A $[2,3,1]$ | 1024 |
| A $[2,3,2]$ | 1028 |
|  | $102 C$ |

Table 3-6 Column Major Ordering

FORTRAN array declaration:
INTEGER A $(2,3,2)$

Element size: 4 bytes Base address: 1000 (Hex)

| Array Element | Address $\quad$ (Hex) |
| :---: | :---: |
| A $(1,1,1)$ | 1000 |
| A $(2,1,1)$ | 1004 |
| A $(1,2,1)$ | 1008 |
| A $2,2,1)$ | 100 C |
| A $(1,3,1)$ | 1010 |
| A $(2,3,1)$ | 1014 |
| A $(1,1,2)$ | 1018 |
| A $(2,1,2)$ | 101 C |
| A $(1,2,2)$ | 1020 |
| A $(2,2,2)$ | 1024 |
| A $(1,3,2)$ | 1028 |
| A $(2,3,2)$ | $102 C$ |

Note that the same memory location is referenced by the Pascal index sequence [I,J,K] and the FORTRAN index sequence (K,J,I).

The general expression for the one-dimensional index generated to access either $A[I, J, K, \ldots, Z]$ in Pascal or $A(Z, \ldots, K, J, I)$ in FORTRAN is:

$$
(\ldots((I a * D j+J a) * D k+K a) * \ldots) * D z+Z a
$$

 dimensions, respectively, and the values Ia, Ja, Ka, ... , Za are the index values, zero-adjusted by the CHECKi instruction (by subtracting their lower bounds).

The INDEXi instruction implements one step of the evaluation of this expression from the inside out, by providing the function

```
accum = accum * (length+1) + index
```

where accum is any register (R0-R7), used in consecutive INDEXi instructions as an accumulator location,
index is the current index value being processed, and
length is a general operand containing the current dimension length minus 1 (so that it always matches the size of the index operand).

Processor control instructions control the sequence of program execution. These instructions provide conditional and unconditional branches, calls to and returns from local and external procedures, and generation and returns from traps and interrupts. The following is a list of the processor control instructions:

## Instructions

## Branches

## Jump

Conditional Branch Unconditional Branch
Case Branch (Multiway)
Add, Compare and Branch

## Local Procedure Calls/Returns

Jump to Subroutine
Branch to Subroutine

| JSR | JSR |
| :--- | :--- |
| BSR | BSR |

Return from Subroutine
BSR
BSR
Return from Subroutine

| JUMP | JUMP |
| :--- | :--- |
| BCond | Bcond |
| BR | BR |
| CASEB, CASEW, CASED | CASEi |
| ACBB, ACBW, ACBD | ACBi |

External Procedure Calls/Returns
Call External Procedure
Call External Procedure with Descriptor
Return from
External Procedure
Explicit Trap Instructions

```
Breakpoint Trap
Flag Trap (Conditional)
```

Supervisor Call Trap

| CXP | CXP |
| :--- | :--- |
| CXPD | CXPD |
| RXP | RXP |

Trap/Interrupt Returns

| Return from Trap* | RETT | RETT |
| :--- | :--- | :--- |
| Return from Interrupt* | RETI | RETI |

* Privileged instruction (see note).

Branches transfer control to an instruction nonsequentially. The JUMP instruction allows the destination address to be specified using a general choice of addressing modes. The BR instruction also transfers control, but provides a more code-compact form for PC-relative references. The Bcond instruction performs a branch as per the BR instruction if a specified condition code is true. The CASEi instruction branches by adding the contents of any general operand to the

Program Counter. In conjunction with Scaled Indexing (Section 4.4.9), this implements a multiway branch which corresponds directly to the Pascal CASE statement and the C SWITCH statement. The ACBi (Add, Compare and Branch) instruction supports looping by adding a small increment (range -8 to +7) to any general operand and branching if the result is non-zero.

Local procedure calls (JSR and BSR) transfer control as per the JUMP and BR instructions, respectively, except that they first save the address of the next sequential instruction onto the current stack as a 32-bit return address. The called procedure returns control after such a call with the RET instruction.

External procedure calls are implemented by the CXP and CXPD instructions. An external procedure is defined as a procedure which is in another module from the procedure currently executing. See Section 2.8.2 for further details of the module environment implemented by the Series 32000 architecture. An external procedure call saves the current contents of the MOD register as well as the return address onto the current stack, sets up the MOD and $S B$ registers to match the environment of the destination module, and transfers control. In the CXP instruction, the destination procedure is specified with an index into the Link Table belonging to the current module, from which a descriptor is read, locating the destination. In the CXPD instruction, this descriptor is given as a general operand, greatly facilitating references to procedures which have themselves been passed as parameters. (A procedure can be passed as a parameter by passing its descriptor, using the ADDR instruction.) The RXP instruction is used to return control after an external procedure call, restoring the MOD and SB registers as well as the Program Counter.

Three instructions have the function of causing deliberate traps. The BPT, FLAG and SVC instructions each have unique vectors in the Interrupt Dispatch Table (Section 2.8.4). The BPT instruction is intended to support debug breakpointing of programs. The FLAG instruction causes a trap if the PSR F bit is set (e.g. if the previous ADD instruction overflowed), and the SVC instruction provides the mechanism to make requests of a protected operating system.

The RETT instruction returns control from a trap or the Non-Maskable or NonVectored interrupt, restoring the PSR, MOD and SB registers. Since traps are often caused deliberately to request service of an operating system, the RETT instruction also allows parameters on the top of the original stack to be discarded in the process of returning. The RETI instruction is used for returning from any vectored maskable interrupt, providing the function of the RETT instruction and also communicating with one or more NS32202 Interrupt Control Units to implement transparent interrupt control.

NOTE: The instructions RETT and RETI are privileged, because they may change the contents of the high-order byte of the PSR, which is protected. The Illegal Operation trap, Trap(ILL), will ocour if either of these instructions is attempted by a program in User Mode (i.e., while the PSR U bit is set).

Processor service instructions provide general housekeeping functions and services. The following is a list of the processor service instructions:

## Instructions

Effective Address
Calculate Effective Address
Context Instructions

| Save General Purpose Registers | SAVE | SAVE |
| :--- | :--- | :--- |
| Restore General Purpose Registers | RESTORE | RESTORE |
| Enter New Procedure Context | ENTER | ENTER |
| Exit Procedure Context | EXIT | EXIT |

Register/Stack Manipulation

| Adjust Stack Pointer | ADJSPB, ADJSPW, ADJSPD | ADJSPi |
| :--- | :--- | :--- |
| Bit Clear in PSR* | BICPSRB, BICPSRW | BICPSRB |
|  |  | BICPSRW |
| Bit Set in PSR* | BISPSRB, BISPSRW | BISPSRB |
|  |  | BISPSRW |
| Load Processor Register* | LPRB, LPRW, LPRD | LPRi |
| Store Processor Register* | SPRB, SPRW, SPRD | SPRi |
| Set Configuration Register* | SETCFG | SETCFG |

Miscellaneous

| No Operation | NOP | NOP |
| :--- | :--- | :--- |
| Wait for Interrupt | WAIT | WAIT |
| Diagnose | DIA | DIA |
| Cache Invalidate* | CINV | CINV |

* Privileged, or having privileged forms (see note).

There is one effective address instruction, ADDR, which calculates the effective address of its first operand and places that 32-bit address into its second operand location.

Context instructions allow the saving and restoring of portions of the processor context to and from the current stack. The SAVE instruction pushes the contents of any set of General-Purpose registers specified by the programmer. The RESTORE instruction undoes this by popping information from the top of the stack into any set of these registers. The ENTER and EXIT instructions deal with a larger context which is used by both local and external procedures. The ENTER instruction is generally the first instruction executed in a procedure, and has the function of completing the "activation record" or "stack frame". It saves the Frame Pointer (FP) register onto the current stack, allocates a specified number of bytes on the stack to be used for dynamic local variables, and sets up the Frame Pointer as a base pointer for this area. It also pushes the contents of any specified General-Purpose registers, as per the SAVE instruction. After executing this instruction, the Frame Pointer can be used in the Frame Memory and Frame Memory Relative addressing modes (Sections 4.4.8 and 4.4.3) to access both these local variables and any parameters passed to this procedure. The EXIT instruction is placed at the end of the procedure, undoing the action of the matching ENTER instruction. It restores the contents of the specified GeneralPurpose registers from the stack, discards the local variable space, and restores the Frame Pointer, leaving the return address at the top of the stack for the appropriate Return instruction.

Register/Stack Manipulation instructions provide the means to load, store and adjust the contents of CPU dedicated registers. (Corresponding instructions for manipulating dedicated Floating-Point and Memory Management registers are listed in Sections 3.3 and 3.12.) The ADJSPi instruction provides the means to directly adjust the current Stack Pointer register by the contents of any general operand in order to allocate or purge space on the stack or for alignment purposes. The BICPSR and BISPSR instructions allow specified bits in the PSR register to be cleared or set without affecting the rest of the PSR. The LPRi and SPRi instructions load or store a specified dedicated register. The SETCFG instruction sets up the CFG register (Section 2.3) to declare the presence of external interrupt control and slave processors.

Four instructions provide miscellaneous functions. The NOP (No Operation) instruction is a 1-byte instruction which does nothing except transfer control to the next sequential instruction. The WAIT instruction causes instruction processing to be suspended until an interrupt occurs. The DIA instruction provides a function similar to WAIT for hardware breakpointing purposes, but is not intended for use in programming. The CINV instruction invalidates entries in the on-chip caches.

NOTE: The instructions flagged with an asterisk ("*") have forms which are privileged. The Illegal Operation trap, Trap(ILL), will occur if they are attempted in User Mode (i.e., while the PSR U bit is set). The BICPSRW and BISPSRW instruction forms are privileged, as they may change the high-order byte of the PSR, which is protected. The LPRi and SPRi instructions are privileged when they reference the INTBASE, USP, CFG, Debug register or the enitre PSR. The SETCFG and the CINV instructions are privileged always.
3.12 Memory Management Instructions

The following is a list of the Memory Management instructions:

## Instruction

Load Memory Management Register
Store Memory Management
Register
Validate Address for Reading Validate Address for Writing

Move Value from Supervisor
to User Space
Move Value from User
to Supervisor Space

## Mnemonic Forms

LMR

SMR

RDVAL
WRVAL

MOVSUB, MOVSUW, MOVSUD
MOVUSB, MOVUSW, MOVUSD

Index

LMR

SMR

RDVAL
WRVAL

MOVSUi

MOVUSi

The LMR and SMR instructions load and store the contents of Memory Management registers (Section 2.5) as 32-bit values. The RDVAL instruction tests the protection level of a specified user memory location to determine whether the current user-mode program is allowed to read it. The WRVAL instruction tests whether the current user is allowed to write into a specified memory location. The MOVSUi instruction moves a byte, word, or double-word value from a specified location in the Supervisor addressing space to a location in the User space, and the MOVUSi instruction moves a value from User space to Supervisor space.

NOTES: 1. If the $M$ bit in the CFG register has not been set (by the SETCFG instruction), the LMR, SMR, RDVAL and WRVAL instructions will generate the Undefined Instruction trap, Trap(UND).
2. All Memory Management instructions are privileged. If attempted by a program running in User Mode (i.e., while the PSR U bit is set), the Illegal Operation trap, Trap(ILL), will occur instead.

### 3.13 Custom Instructions

A set of instructions has been set aside for custom use. These instructions are reserved for such use, and will not be defined otherwise by NSC.

A custom instruction starts with one of the following binary encodings as its least-significant byte.

1. 00010110
2. 00110110
3. 10110110

Note that each of these corresponds to the first byte of a Floating-Point or Memory Management instruction, the difference being that bit 3 is "0" instead of "1".

If the C bit in the CFG register is cleared (by the SETCFG instruction), these instructions cause the Undefined Instruction trap, Trap(UND). Since a trap pushes the address of this first byte as the return address, the format and length of the remainder of the instruction may be defined in any manner, as required by the custom application.

If the C bit in the CFG register is set, these instructions are executed by an external "Custom" Slave Processor. The remainder of each instruction must follow the format of its corresponding Floating-Point or Memory Management instruction. The custom instructions corresponding to Memory Management instructions are privileged. In executing a custom instruction, the operand definitions and the protocol followed in communicating with the Custom Slave are identical to those for the corresponding Floating-Point or Memory Management instruction.

See the applicable CPU data sheet for details of the instruction formats and the Slave Processor protocols used.

## Chapter 4

## INSTRUCTION OPTIONS AND CONSTRUCTION

This chapter defines the options available in Series 32000 instructions, how these options are denoted in Chapter 5 (Instruction Set), and how the binary form of an instruction is constructed based on the selections made.

The structure of an instruction is given in Chapter 5 by its format definition. A typical format definition follows:

Syntax:

## CMPf src1,

$\begin{array}{ll}\text { yrntax: CMPf } & \text { sre } \\ & \text { gen }\end{array}$


CMPf
Basic
Instruction
Format

The notations used are defined in the following sections.

| Syntax Line | 4.1 |
| :--- | :--- |
| Operand Attributes | 4.2 |
| Instruction Format | 4.3 |

Other information presented in this chapter:
Addressing Modes 4.4
Construction Examples 4.5

### 4.1 Syntax Presentation

The Syntax line presents the instruction mnemonic, followed by a list of operands, as shown. Lower-case items indicate options to be specified by the programmer.

Syntax.


Within the mnemonic, the following lower-case items may appear:
i An integer operation length suffix. It is specified by the programmer as

```
B = Byte (8-bit integer operation)
W = Word (16-bit integer operation)
D = Double-Word (32-bit integer operation)
```

and defines the length of the operation to be performed. In arithmetic operations, the carry and overflow tests use this specification to determine which bit positions are to be checked. When an implied operand of attribute "quick" appears (Section 4.2.3), it is internally sign-extended to this length before use. The lengths of integer general operands are usually taken from this length, but this depends on their individual length attributes, Section 4.2.2.
f A floating-point operation length suffix. It is specified by the programmer as

```
F = Single-precision Floating (32-bit floating-point operation)
L = Double-precision Long Floating (64-bit floating-point
                operation)
```

and defines the length of the operation performed. The lengths of floating-point general operands are usually taken from this length specification, but this depends on their individual length attributes, Section 4.2.2. In certain conversion instructions (e.g. ROUNDfi) both integer and floating-point operation lengths may appear.
cond A condition code, as in the Conditional Branch instruction:

## Syntax: Bcond dest

The specifications allowed and their interpretations are listed in the instruction description.

Operands are always given in lower case, and are to be specified by the programmer according to the attributes appearing below them (Section 4.2). The name given to an operand on the Syntax line serves to identify it in the instruction description.

### 4.2 Operand Attributes

Operands are defined in Chapter 5 by a set of attributes. These define what may be specified for each operand, and exactly how any valid operand specification will be interpreted when the instruction is executed.

A typical set of attributes is shown below:


Some operands listed as part of the instruction syntax are implied, meaning that their locations are not determined from a general choice of addressing modes. An implied operand is identified by the attribute "reg", "quick", "short", "imm" or "disp"; i.e., anything except "gen". For the explanations of implied operand attributes, see Section 4.2.3.

Most Series 32000 operands, however, are general, meaning that a general choice of addressing modes (Section 4.4) may be used to specify their locations. General operands are identified by the attribute "gen". A general operand has the additional attributes of an access class and also a length where relevant.

The access class attribute serves to define all cases of addressing mode usage including exceptional cases whose effects (or even legality) might not otherwise be obvious. The possible access classes for a general operand are read, write, rmw, addr and regaddr. Three addressing modes are affected by the access class: Register, Immediate and Top of Stack, as shown in Table 4-1 and described in Section 4.2.1.

The length attribute defines a general operand's data type and its size in bytes (see Section 4.2.2).

An operand with attribute i is an integer of the size given as the integer operation length by the programmer. An operand with attribute 2i is twice this size. An operand with attribute $B, W$ or $D$ is a byte, word or double-word integer, respectively, regardless of the operation length.

An operand with length attribute $f$ is a floating-point value of the size given as the floating-point operation length by the programmer. An operand with length attribute $F$ or $L$ is a single-precision or double-precision floating-point value, respectively, regardless of the operation length.

### 4.2.1 Access Classes

Computer architectures usually have exceptional cases of operand reference based on the context of the instruction making the reference. For example, if an architecture allows references to registers as general operands, and provides a Jump instruction specifying a general destination, an obvious question becomes whether in this context (Jump) it is still legal to specify a register.

Rather than defining the interpretations of operand references on an instruction-by-instruction basis, the Series 32000 architecture defines five standard contexts (access classes) within which an Series 32000 family CPU will interpret a reference to a general operand. Each instruction assigns one access class to each of its general operands, which in turn fully defines the action of any addressing mode in referencing that operand.

Only three addressing modes have interpretations which are affected by the access class of an operand. These are Register, Immediate and Top of Stack. The five access classes, defined below, are read, write, rmw, addr and regaddr. See also Table 4-1.
read: The addressing modes are interpreted in the context of an operand being read but not rewritten. If Register mode is used, the specified register contains the operand. Immediate mode is legal only for operands of this access class. If Top of Stack mode is specified, the Stack Pointer is post-incremented by the number of bytes corresponding to the length of the operand (as determined from its length attribute, Section 4.2.2), thus "popping" it from the stack.
write: The addressing modes are interpreted in the context of an operand being written without having been read. If Register mode is used, the specified register receives the operand. Immediate mode is undefined for this access class. If Top of Stack mode is specified, the Stack Pointer is pre-decremented by the number of bytes corresponding to the length of the operand (as determined from its length attribute, Section 4.2.2), thus "pushing" it onto the stack.
rmw: Read-Modify-Write. The addressing modes are interpreted in the context of an operand being read, modified and rewritten to the same location. If Register mode is used, the specified register contains the operand. Immediate mode is undefined for this access class. If Top of Stack mode is specified, the Stack Pointer provides the address of the operand, but is not altered.

Table 4-1 Addressing Mode Actions vs. Access Class

| Addressing <br> Mode | Access Class |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | read | write | rmw | addr | regaddr |
| Register | Rn, Fn | Rn, Fn | Rn, Fn | (Rn) | Rn, Fn |
| Immediate | legal | undefined | undefined | undefined | undefined |
| Top of Stack | Pop | Push | (SP) | (SP) | (SP) |

NOTES: 1. The notations (Rn) and (SP) signify use of the enclosed register as a pointer. The register is not altered.
2. Using Scaled Indexing in an addressing mode overrides the access class and forces it to "addr".

Address. The addressing modes are interpreted in the context of an operand which cannot be held in a register, or of an effective address calculation which does not correspond to an operand being fetched as data. Examples of this context are ADDR A,B (place the effective address of A into B), JUMP X (place the effective address of $X$ into the Program Counter) or any addressing mode using Scaled Indexing (since arrays cannot be held in registers; see Table 4-1). If Register mode is used, the operand is in memory, and the specified register contains its address. Immediate mode is undefined for this access class. If Top of Stack mode is specified, the Stack Pointer provides the address of the operand, but is not altered.

Note: The addr access class does not define the use to which an operand is put, but only the context in which the addressing modes are interpreted. An addr operand may be read, written, or neither read nor written, depending on the instruction being executed.
regaddr: Register/Address. The addressing modes are interpreted in the context of designating a base for locating a data item of nonstandard size and/or alignment. An example of this context is the operand $B$ in the instruction TBITW A, B (test the bit which is A bits from the beginning of base location B). If Register mode is used, the data item is held within the specified register. Immediate mode is undefined for this access class. If Top of Stack mode is specified, the Stack Pointer provides the address of the base, but is not altered.

Note: The regaddr access class does not define the use to which an operand is put, but only defines the context in which the addressing modes are interpreted. Information at the location given in a regaddr context may be read, written, or neither read nor written, depending on the instruction being executed.

### 4.2.2 Length Attributes

The length attribute of a general operand defines its data type and its length (in bytes). Operands with length attribute $B, W, D, i$ or $2 i$ are integers. Operands with length attribute $F$, $L$ or $f$ are floating-point values.

The length in bytes of an operand affects the following three addressing modes:
Register: If the length of an operand is smaller than the designated General-Purpose register, it is only the low-order portion of the register which is referenced or modified. The rest of the register is unchanged. Operands with length attribute $2 i$ are a special case; see Section 4.2.2.1 below.

Immediate: The length of the value held within the binary instruction format matches the length in bytes of the operand.

Top of Stack: If the access class attribute (Section 4.2.1) indicates that the Stack Pointer is to be modified, it is modified by the operand length in bytes.

### 4.2.2.1 Integer Length Attributes

The length attributes which identify an integer are $B, W, D, i$ and $2 i$. For integers, the Register addressing mode assumes that the General-Purpose registers (R0-R7) are to be used. Floating-Point registers cannot be specified for integer operands. The integer length attributes are defined as follows:

B The operand is a 1-byte integer.
W The operand is a 2-byte (word) integer.
D The operand is a 4-byte (double-word) integer.
i The operand is either one, two, or four bytes in length, depending on the operation length suffix (B, W or D: Section 4.1) appended to the instruction mnemonic by the programmer.

2i The operand is twice the length given as the operation length suffix (Section 4.1) appended to the instruction mnemonic by the programmer.

The MEI and DEI instructions (Multiply/Divide Extended Integer) present special cases in which operands with length attribute $2 i$ can be held in registers. If an operand with length attribute $2 i$ is specified as being within a register, it occupies a pair of General-Purpose registers (R0 and R1, R2 and R3, R4 and R5, or R6 and R7), and the even-numbered register of the pair must be specified as the operand location. The operand is held with its least-significant half in the even-numbered register (right-justified) and its most-significant half in the odd-numbered register (also right-justified). Any portions of the two registers not used to hold the operand are neither referenced nor modified.

### 4.2.2.2 Floating-Point Length Attributes

The length attributes which identify a floating-point operand are F , L and f . For floating-point operands the Register addressing mode assumes that the Floating-Point registers (F0-F7) are to be used. General-Purpose registers cannot be specified for floating-point operands. The floating-point length attributes are defined as follows:

F The operand is a 4-byte single-precision floating-point value.
L The operand is an 8-byte double-precision ("Long") floating-point value.
f The operand is either a single-precision or double-precision floatingpoint value, depending on the operation length suffix (F or L, Section 4.1) appended to the instruction mnemonic by the programmer. See the description of "L" above for the format of a double-precision operand within registers.

### 4.2.3 Implied Operand Attributes

Implied operands are specified without using addressing modes. Their attributes define how they may be specified.
reg: The operand location is a General-Purpose register (R0-R7). Any General-Purpose register may be specified. The entire register is always used and/or modified by the instruction. The register number is encoded in the binary instruction format within a 3-bit field marked "reg".
quick: The operand is a signed, 4-bit immediate value. Its range is -8 to +7 . Before use, it is internally sign-extended to the length given by the operation length suffix appended to the instruction mnemonic. A quick operand is encoded in the binary instruction format within a 4-bit field marked "quick".
short: The operand occupies a 4-bit field within the binary instruction format. The interpretation of the field depends on the instruction.
imm: The operand is a l-byte immediate value, appended to the instruction following any addressing extensions. Its interpretation is determined by the instruction.
disp: The operand is an immediate signed integer value, encoded as a displacement field and appended to the instruction following any addressing extensions. Its use is determined by the instruction.

A displacement field is stored with the most-significant byte at the lowest address. Its format is determined by its most-significant bits as shown below.


The binary format of an Series 32000 instruction is shown in Figure 4-1. It is divided into two sections.

1. The Basic Instruction portion defines the operation performed and the number and kinds of operands. It is presented in Chapter 5 individually for each instruction, using field nomenclature as defined in Section 4.3.1 below.
2. Extension fields are optionally appended as defined by the instruction and the addressing modes chosen by the programmer. These extensions fall into a general instruction format, defined in Section 4.3.2.

Because the Series 32000 family implements a full two-address architecture, most instructions have two general operands (with attribute "gen", Section 4.2). To distinguish between them, the first general operand appearing in the Syntax line of an instruction description will be designated operand $A$ and the second operand B.


Figure 4-1 General Format

### 4.3.1 Basic Instruction

The Basic Instruction portion defines the operation performed and the addressing modes used for referencing general operands, and provides fields within it for holding all implied operands with attribute reg, quick or short (Section 4.2.3). It is one, two or three bytes in length.

The format of the Basic Instruction is diagrammed for each instruction under the Syntax line of the instruction description. The format used for storing the Basic Instruction in memory is the same as for data elements; that is, the least-significant byte appears first, at the lowest address. Fields within the Basic Instruction are presented as defined below.

### 4.3.1.1 Operation Code Fields

Operation code fields are presented explicitly in binary. All fields presented in this manner are derived from the instruction mnemonic and define the basic operation to be performed.

### 4.3.1.2 Operation Length Fields: i and f

Operation Length fields define the length to which calculations are performed within a basic data type (integer or floating point). They also define the lengths of most general operands (indirectly, through each operand's own length attribute, Section 4.2.2). They are derived from the Operation Length mnemonic suffix (Section 4.1) chosen by the programmer, as shown below.

| Field | Mnemonic <br> Suffix | E |
| :---: | :---: | :---: |
| i | W | 00 |
| f | F | 11 |
|  | L | 1 |
|  |  | 01 |

### 4.3.1.3 General Addressing Mode Fields: gen

These are 5-bit fields which define the addressing mode used to access each operand. The name of the operand from the Syntax line appears above the field. The encodings of these fields are given in the definitions of the addressing modes, Section 4.4.

### 4.3.1.4 Implied Operand Fields: reg, quick, short

These fields hold the necessary information for implied operands which are defined with the corresponding attribute (reg, quick or short: Section 4.2.3). The name of the operand from the Syntax line appears above the field.

A reg field is a 3-bit field holding a register number (0-7).
A quick field is a 4-bit field holding a signed value (range -8 to +7).
A short field is a 4-bit field holding information which is required by the individual instruction. Its contents are defined in the instruction description.

### 4.3.2 Extension Fields

The following fields extend the length of the instruction beyond the Basic Instruction field. They appear as required by the individual instruction or by the addressing modes chosen for specifying its general operands.

### 4.3.2.1 Index Bytes

The first form of extension is in the form of Index Bytes. The instruction is extended in this manner whenever Scaled Indexing (Section 4.4.9) is used in specifying a general operand. Either or both of the general operands may be specified using Scaled Indexing. If both operands are specified in this form, then the Index Byte for Operand A appears before the Index Byte for Operand B. See Figure 4-1. The format of an Index Byte is given in the definition of Scaled Indexing, Section 4.4.9.

### 4.3.2.2 Addressing Extensions

An addressing extension is appended for each general operand as required. Its contents depend on the addressing mode chosen for each. See Section 4.4 for the usages of addressing extensions in addressing modes. The addressing extension for operand A appears before the one for operand B (Figure 4-1).

Addressing extensions are constructed from two basic elements: displacement fields and immediate values.

NOTE: Unlike other values in memory, addressing extensions are ordered with the most-significant byte at the lowest address.

An addressing extension contains either:

1. One immediate value, or
2. One displacement field, labelled "disp" in the addressing mode definitions (Section 4.4), or
3. Two displacement fields, labelled "disp1" and "disp2". In this form, disp1 is appended first, followed by disp2.

If a Register or Top of Stack addressing mode is used to specify a general operand, no addressing extension appears for that operand.

A displacement field holds a signed two's-complement addressing constant. It is stored with the most-significant byte at the lowest address. Its length is determined by its most-significant bits as shown below.


An immediate value appears as an addressing extension only when the Immediate addressing mode is specified (Section 4.4.4). The length of the value is determined from the operand's length attribute (Section 4.2.2). The value is ordered with its most-significant byte at the lowest address.

### 4.3.2.3 Implied Operand Extensions: imm, disp

Implied operands, of attribute "imm" or "disp" (Section 4.2.3), appear last, after all addressing extensions. If there is more than one imm or disp operand appearing in the instruction, then the operands are appended in the order in which they are listed on the Syntax line.

Any general operand (Section 4.2) may be specified by the programmer using a general choice of addressing modes. This section defines addressing mode syntax, functions and encodings.

Table 4-2 lists the addressing modes provided for specifying a general operand. It also serves as an index to this section. The Encoding column gives the binary encoding used in a gen field (Section 4.3.1.3) to select each mode. The Name column gives the name of the addressing mode as used in this manual, and the Syntax column shows the syntax used in assembly language to express it. (Note: What is given is only the lowest level of expression, which most directly relates to the action of the addressing mode. See the applicable assembler manual for full details of expression syntax and symbolic features.)

Scaled Indexing is an option available as part of any addressing mode except Immediate. It does not stand alone as an addressing mode, but is listed with the addressing modes because of the binary encodings used to select the option.

Table 4-2 Series 32000 Addressing Modes

|  | Encoding | Name | Syntax |
| :---: | :---: | :---: | :---: |
| Register | 00000 | Register 0 | R0 or F0 |
|  | 00001 | Register 1 | R1 or F1 |
|  | 00010 | Register 2 | R2 or F2 |
|  | 00011 | Register 3 | R3 or F3 |
|  | 00100 | Register 4 | R4 or F4 |
|  | 00101 | Register 5 | R5 or F5 |
|  | 00110 | Register 6 | R6 or F6 |
|  | 00111 | Register 7 | R7 or F7 |
| Register Relative | 01000 | Register 0 Relative | disp(R0) |
|  | 01001 | Register 1 Relative | disp(R1) |
|  | 01010 | Register 2 Relative | disp(R2) |
|  | 01011 | Register 3 Relative | disp(R3) |
|  | 01100 | Register 4 Relative | disp(R4) |
|  | 01101 | Register 5 Relative | disp(R5) |
|  | 01110 | Register 6 Relative | disp(R6) |
|  | 01111 | Register 7 Relative | disp(R7) |
| Memory Relative | 10000 | Frame Memory Relative | disp2(disp1 (FP)) |
|  | 10001 | Stack Memory Relative | disp2(disp1(SP)) |
|  | 10010 | Static Memory Relative | disp2(disp1(SB)) |
| (reserved) | 10011 | (Reserved for future use.) |  |
| Immediate | 10100 | Immediate | value |
| Absolute | 10101 | Absolute | @disp |
| External | 10110 | External | EXT (disp1) +disp2 |
| Top of Stack | 10111 | Top of Stack | TOS |
| Memory Space | 11000 | Frame Memory | disp(FP) |
|  | 11001 | Stack Memory | disp(SP) |
|  | 11010 | Static Memory | disp(SB) |
|  | 11011 | Program Memory | * + disp |
| Scaled Indexing | 11100 | Byte Indexed | basemode [ Rn : B$]$ |
|  | 11101 | Word Indexed | basemode[Rn:W] |
|  | 11110 | Double-Word Indexed | basemode[Rn: D] |
|  | 11111 | Quad-Word Indexed | basemode[Rn:Q] |

### 4.4.1 Register Modes

| Mode | Syntax | Encoding |
| :---: | :--- | :---: |
| Register 0 | R0 or F0 | 00000 |
| Register 1 | R1 or F1 | 00001 |
| Register 2 | R2 or F2 | 00010 |
| Register 3 | R3 or F3 | 00011 |
| Register 4 | R4 or F4 | 00100 |
| Register 5 | R5 or F5 | 00101 |
| Register 6 | R6 or F6 | 00110 |
| Register 7 | R7 or F7 | 00111 |

None.

The interpretation of these modes is formally defined below. However, rule 6 defines the general case, which is that the specified General-Purpose register (R0-R7) holds the operand.

The following rules are listed in order of decreasing precedence. Lower-numbered rules take precedence over higher-numbered rules.

1. If the access class of the operand (Section 4.2.1) is "addr", then the operand is in memory. The effective address of the operand is held in the specified General-Purpose register.
2. If Scaled Indexing is used, the access class of the operand is redefined as "addr", and rule 1 above applies.
3. If the operand length attribute (Section 4.2.2) is "2i", then a pair of General-Purpose registers (R0 and R1, R2 and R3, R4 and R5, or R6 and R7) holds the operand. The even-numbered register of the pair must be specified, and if the odd-numbered register is specified the location of the operand is undefined. The least-significant half of the operand is held in the low-order portion of the even-numbered register, and the remaining portion of the register is neither used nor affected. The most-significant half of the operand is held in the low-order portion of the odd-numbered register, and any remaining portion of the register is neither used nor affected.
4. If the operand length derived from its length attribute (Section 4.2.2) is single-precision floating-point, then the operand is held in the specified Floating-Point register (F0-F7).
5. If the operand length derived from its length attribute (Section 4.2.2) is double-precision floating-point, then the operand is held in the specified Floating-Point register (F0-F7).
6. When none of the above exceptions apply, the operand is an integer held within the specified General-Purpose register (R0-R7). If the operand length derived from its length attribute is shorter than the full 32-bit length of the register, then the operand occupies the low-order portion of the register, and the remaining portion of the register is neither used nor affected.
4.4.2 Register Relative Modes

| Mode | Syntax | Encoding |
| :---: | :---: | :---: |
| Register 0 Relative | disp(R0) | 01000 |
| Register 1 Relative | disp(R1) | 01001 |
| Register 2 Relative | disp(R2) | 01010 |
| Register 3 Relative | disp(R3) | 01011 |
| Register 4 Relative | disp(R4) | 01100 |
| Register 5 Relative | disp(R5) | 01101 |
| Register 6 Relative | disp(R6) | 01110 |
| Register 7 Relative | disp(R7) | 01111 |

## Extensions

One displacement field: disp.

The operand is in memory. Its effective address is the sum of the 32 -bit contents of the specified General-Purpose register (R0-R7) and the displacement value sign-extended to 32 bits.

### 4.4.3 Memory Relative Modes

| Mode | Syntax | Encoding |
| :---: | :---: | :---: |
| Frame Memory Relative | disp2(disp1(FP)) | 10000 |
| Stack Memory Relative | disp2(disp1(SP)) | 10001 |
| Static Memory Relative | disp2(disp1(SB)) | 10010 |

## Extensions

```
Two displacement
fields: disp1
followed by disp2.
```

The operand is in memory, at the address given by the sum of disp2 (sign-extended to 32 bits) and a 32-bit pointer in memory. The address of this pointer is generated by adding displ (sign-extended to 32 bits) and the contents of the specified register (FP, SP or SB). The symbol "SP" means the stack pointer which is currently selected by the $S$ bit in the PSR (Section 2.2).

NOTE: The Stack Memory Relative mode uses the contents of the selected stack pointer as it was at the beginning of the instruction. The effective address is therefore independent of any changes made to the stack pointer by any Top of Stack mode appearing in the same instruction.

| Mode | Syntax | Encoding |
| :---: | :---: | :---: |
| Immediate | value | 10100 |

## Extensions

The value, placed
most-significant
byte first.

The operand value is input from the addressing extension portion of the instruction. The value appears most-significant byte first. Its length in bytes is determined from the operand length attribute (Section 4.2.2). Floating-Point as well as integer instructions may use Immediate mode.

NOTES: 1. Immediate mode is legal only for operands of access class "read". Any other use is undefined.
2. Immediate mode may not be used as the base mode for Scaled Indexing.
4.4.5 Absolute Mode

| Mode | Syntax | Encoding |
| :---: | :---: | :---: |
| Absolute | @address | 10101 |

Extensions
One displacement
field: address.

The absolute address is specified. This address is encoded in the binary instruction as a displacement field of any length required to hold the address.

| Mode | Syntax | Encoding |
| :---: | :---: | :---: |
| External | EXT (disp1) +disp2 | 10110 |
|  | Or |  |
|  | EXT (disp1) |  |

## Extensions

> Two displacement fields: disp1 followed by disp2. If disp2 is omitted in assembly language, it must still be included as a disp2 field containing zero.

The External addressing mode provides the means for a software module to access data within a data space outside of that module. The operand is referenced through the Link Table of the current module (Section 2.8.3). The value displ is a Link Table entry number, and disp2 is a final displacement added to the address provided from that Link Table entry.

The operand is in memory, at the address given by the sum of disp2 (sign-extended to 32 bits) and a 32 -bit pointer in the current Link Table. The address of this pointer is generated by adding displ, multiplied by four, and the contents of the 32-bit value at memory address MOD + 4. "MOD" is the contents of the MOD register, interpreted as a 16-bit unsigned number.

| Mode | Syntax | Encoding |
| :---: | :---: | :---: |
| Top of Stack | TOS | 10111 |

## Extensions

None.
The operand is in memory, at the top of the current stack. It is pushed, popped, or neither pushed nor popped, as appropriate to the usage of the operand.

The stack pointer used is the stack pointer that is currently selected by the $S$ bit in the PSR (Section 2.2).

The stack pointer is used by Top of Stack mode according to the access class of the operand. The rules below are listed in order of decreasing precedence. Lower-numbered rules take precedence over higher-numbered rules.

1. If the operand is of access class "rmw", "addr" or "regaddr", then the effective address of the operand is given by the contents of the stack pointer, and no increment or decrement is performed.
2. If Scaled Indexing is used, the access class of the operand is redefined as "addr", and rule 1 above applies.
3. If the operand is of access class "read", the operand is read from the address given by the contents of the stack pointer. The stack pointer is then incremented by the length in bytes of the operand, as determined from its length attribute (Section 4.2.2).
4. If the operand is of access class "write", the stack pointer is decremented by the length in bytes of the operand, as determined from its length attribute (Section 4.2.2). The operand is then written to the address given by the new contents of the stack pointer.

NOTES: 1. If Top of Stack mode is used for both general operands of an instruction, the operands are accessed and the stack pointer modified in left-to-right operand order. The rightmost addressing mode uses as its initial stack pointer value the contents of the stack pointer after any increment or decrement has been performed by the leftmost addressing mode.
2. The Stack Memory and Stack Memory Relative modes use as their stack pointer value the contents of the selected stack pointer as they were at the beginning of the instruction. The actions of these modes are therefore independent of any modifications made to the stack pointer by any Top of Stack mode appearing within the same instruction.

### 4.4.8. Memory Space Modes

| Mode | Syntax | Encoding |
| :---: | :---: | :---: |
| Frame Memory | disp(FP) | 11000 |
| Stack Memory | disp(SP) | 11001 |
| Static Memory | disp(SB) | 11010 |
| Program Memory | $*+$ disp | 11011 |

## Extensions

One displacement field: disp.

The operand is in memory, at the address given by the sum of the contents of the specified register and the displacement value sign-extended to 32 bits.

The symbol "SP" means the stack pointer (SPO or SP1) which is currently selected by the $S$ bit in the $P S R$ (Section 2.2). The symbol "*" means the contents of the Program Counter.

NOTES: 1. The Stack Memory mode uses the contents of the selected stack pointer as it was at the beginning of the instruction. The effective address is therefore independent of any changes to the stack pointer contents made by any Top of Stack mode occurring in the same instruction.
2. The Program Counter always contains the address of the first byte of the instruction being executed.
4.4.9 Scaled Indexing

| Mode | Syntax | Encoding |
| :---: | :---: | :---: |
| Byte Indexed | basemode[Rn:B] | 11100 |
| Word Indexed | basemode[Rn:W] | 11101 |
| Double-Word Indexed | basemode[Rn:D] | 11110 |
| Quad-Word Indexed | basemode[Rn:Q] | 11111 |

## Extensions

```
basemode = base addressing mode
                        (see below)
        Rn = any General-Purpose Register,
        used as the index register.
```

1. Index Byte.
2. Any extensions required by basemode.

Any addressing mode except Immediate is allowed to include indexing by the contents of any General-Purpose register (R0-R7), interpreted as a signed 32-bit integer. The index value is scaled (multiplied) by a factor of 1,2 , 4 or 8 before use, so that it can be used as an element number for an array of 1-, 2-, 4- or 8-byte elements. An indexed addressing expression has the form
basemode[Rn:l]
where basemode is an addressing mode expression,
Rn is any General-Purpose register, and
1 is an element length qualifier, chosen from:
$B=$ Byte, scale factor $=1$
W = Word, scale factor $=2$
D = Double-word, scale factor $=4$
$\mathrm{Q}=$ Quad-word, scale factor $=8$.
In the binary instruction format, addressing modes with Scaled Indexing are encoded within the Basic Instruction gen field as one of four special codes which specify only the length qualifier (see table above). The basemode and Rn components are specified in an Index Byte appended to the Basic Instruction. See Section 4.3 for the position of an Index Byte in the general instruction format. The Index Byte has the following format:


Any further addressing extensions required by basemode are appended as given in Section 4.3.2.2, in exactly the same manner as if basemode were not indexed.

NOTES: 1. Any operand specified using Scaled Indexing is redefined as being of access class "addr" regardless of the operand's access class in the instruction definition. This affects the interpretation of basemodes Register and Top of Stack, and makes the use of an Immediate basemode illegal. See Section 4.2.1.
2. Scaled Indexing may be applied only once in an addressing expression. Basemode is therefore not allowed to include Scaled Indexing within itself.

### 4.5 Constructing Complete Binary Instructions: Some Examples

The following examples illustrate the process of assembling the binary form of an Series 32000 instruction from its assembly-language form.

Example 1:
The simple example below as generated from the Move instruction (MOVi).

```
MOVB R0, R1
```

This instruction copies the low-order byte of register R0 to the low-order byte of register R1. The format definition of the MOVB instruction is taken from Chapter 5 as shown below.

## Syntax:




#### Abstract

the programmer as follows:


```
i = B (Byte operation length, Section 4.1)
src = R0 (Register 0 addressing mode, Section 4.4.1)
dest = R1 (Register 1 addressing mode, Section 4.4.1)
```

To complete the Basic Instruction, the gen fields for the two general operands src and dest and the i field for the operation length must be provided. The encoding for the src operand (R0 Register addressing mode) is 00000. The encoding for the dest operand (R1 Register addressing mode) is 00001 . The encoding for the operation length (B) is 00 . Thus, the Basic Instruction is:

```
! R0 ! R1 ! MOVB !
!-+-+-+-+-+-++-+-!--+-+-++-+-+-++-+-!
    15 8 7 0
```

$100000!00001!0101!00$ ! consecutive bytes: 5400 (Hex).
and appears in memory as the two

The Register addressing modes R0 and R1 require no addressing extensions. Therefore, the Basic Instruction above is the complete binary form of the example instruction.

Example 2:
The next example is generated from the JUMP instruction.

$$
\text { JUMP } 0(4(\mathrm{SB}))
$$

This instruction performs an indirect jump through a 32 -bit pointer in memory. The pointer's address is calculated by adding 4 to the contents of the $S B$ register.

The format definition of the JUMP instruction is:

## Syntax: JUMP dest

gen
addr


This instruction has only one operand, the general operand dest, which is specified by the programmer with the addressing expression $0(4(\mathrm{SB})$ ). This form of addressing expression specifies that the Static Memory Relative addressing mode (Section 4.4.3) is to be used to calculate the address to which the instruction will jump. The code for this addressing mode is placed in the gen field as binary 10010. Thus, the Basic Instruction is:

```
! Static !
!Mem. Rel.! JUMP !
+---------+---------------------+
!1 0 0 1 0!!0 1 0 0 1 1 1 1 1 1 1 1 1!
!-+-+-++-+-+-++-+-!-++-+-+-++-+-+-+--!
    15 8 7 0
```

The Memory Relative addressing modes require that two displacements be appended to the Basic Instruction. These are designated disp1 and disp2. From the expression provided in the assembly-language example, the displacement values are to be:

$$
\begin{aligned}
& \operatorname{disp1}=4, \text { and } \\
& \operatorname{disp} 2=0 .
\end{aligned}
$$

(continued)

From the format given for a displacement field (Section 4.3.2.2), we see that a small value can be represented in either one, two or four bytes. Obviously, we wish to choose the smallest field which works, so we will use the 1-byte format for each displacement field.

Appending the two displacements to the Basic Instruction, we get the complete binary instruction as shown below.

```
! Static !
!Mem. Rel.! JUMP !
+---------+---------------------+
!1 0}0
!-+-+-++-+-+-++-+-!--+-+-++-+-+-++-+-!
    15 8 7 0
disp1:
    +-+---------------+
    !0!0 0 0 0 1 0 0!
    !-++-+-+-++-+-+-+-!
    7
                            0
disp2:
    +-+--------------+
    !0!0 0 0 0 0 0 0!
    !-+-++-+-+-++-+-+-!
    7 0
```

The complete binary instruction is represented in consecutive memory bytes as 7F 920400 (Hex).

Example 3:
The following example is generated from the ADDi instruction.

```
ADDD EXT(8)+80, -4(FP)
```

This instruction adds a 32-bit value from the memory location specified as EXT (8) +80 to a 32-bit value at the memory location specified as -4(FP).

The format definition of the ADDi instruction is:

```
Syntax: ADDi src, dest ADDB
gen gen ADDW
read.i rmw.i ADDD
```



This instruction has two general operands. For purposes of constructing its binary form, the src operand is labeled operand A and the dest operand is labeled operand B, as discussed in Section 4.3.

The operation length suffix is $D$, encoded as 11 in the i field. The src operand is specified using the External addressing mode (Section 4.4.6), which is encoded in the binary instruction as 10110 in the corresponding gen field. The dest operand is specified using the Frame Memory addressing mode (Section 4.4.8), which is encoded in the corresponding gen field as 11000. The Basic Instruction appears then as shown below.

```
                    Frame
! External! Memory ! ADDD !
+---------+---------+-------+---+
```



```
!-+-+-++-+-+-++-+-!-++-+-+-++-+-++-+-!
    15 8 7 0
```

(continued)

Since neither operand uses Scaled Indexing, the first extensions appended to the Basic Instruction are the addressing extension fields required by the External addressing mode used to specify the src operand (Operand A). The External addressing mode requires two displacement fields: disp1 (containing 8) followed by disp2 (containing 80). The displ displacement value can be held in a single-byte displacement field. The disp2 displacement value cannot, as it is outside the range (-64 to +63 ) which can be represented in a signed 7 -bit number. It can, however, be held in a two-byte displacement field. Appending the displacement fields for Operand A yields the result shown below.

Frame


After the addressing extensions required for Operand A, the addressing extensions required for Operand B are appended. Since Operand B (the dest operand) is specified using the Frame Memory addressing mode, there is one displacement field required, containing the value -4. This value is within the range -64 to +63, and so it can be held in the single-byte displacement format. It is appended as shown:

## Frame



Disp1 (A): $\quad!0!0001000$ !
!-+-+-+-+-+-+-+-!

$\begin{array}{ll}! & 80 \\ +---+----------+\end{array}$
Disp2 (A): $\quad!10!000000!$
+-+-+-+-+-+-+-+-+
! $01101010000!$
! -+-++-+-+-++-+-+-!
7
0
$\begin{array}{llll}! & -4 & ! \\ +-+---------------+~\end{array}$
Disp (B): ! 0! 1 1 1 1 1 0 0!
$!-+-+-+-+-+-+-+-!$
7

The complete instruction appears in consecutive memory bytes as:
03 B6 $0880507 C$ (Hex).

Example 4:
A final example of how an instruction is assembled uses the Extract Field
(EXTi) instruction.

$$
\text { EXTB R0, } 10(S B), 0(S B)[R 1: B], 5
$$

This instruction copies a 5-bit field from a point in memory determined by a bit offset (contained in R0) from the address $10(\mathrm{SB})$ to the address specified by $0(S B)[R 1: B]$. The format definition of the Basic Instruction is:


In this more complex instruction, there are several items which must be placed in the Basic Instruction. These are the addressing modes specified by the expressions $10(S B)$ and $0(S B)[R 1: B]$, the $i$ field corresponding to the $B$ operation length suffix, and the reg field corresponding to the reg operand specified as R0. The code for the expression $10(S B)$, specifying the Static Memory addressing mode, is 11010. The code for the expression $0(S B)[R 1: B]$, specifying the Static Memory addressing mode with Scaled Indexing (scale factor = 1), is 11100. (Note that when Scaled Indexing is used, it is the code for Scaled Indexing which is placed in the Basic Instruction. See Section 4.4.9.) The i field is 00, for the B operation length suffix. The reg field is 000 , for $R 0$. Thus, the Basic Instruction is:

(continued)

The expression $10(S B)$ specifies Operand A and $0(S B)[R 1: B]$ specifies Operand B. Because it is indexed, Operand B requires an Index Byte. The Index Byte is the first extension to be appended to the Basic Instruction. It contains the code for the basemode $0(S B)$ and the register number for R1. The basemode (Static Memory) is encoded as 11010 and the register number is encoded for R1 as 001.

Static Byte


Static
! Memory ! R1 !
+---------+-----+
Index Byte (B) :
$\begin{array}{lllllll}1 & 1 & 0 & 1 & 0! & 0 & 0\end{array} 1!$
!-+-++-+-+-+-+-+-!
(continued)

The next extensions to be appended are the addressing extensions required by the addressing modes for the general operands. Since Operand $A$ is specified using the Static Memory addressing mode, it requires one displacement field, containing 10. This displacement is placed in single-byte format after the Index Byte.

The Static Memory basemode $0(S B)$ for Operand B requires one displacement field containing 0. This displacement is placed in single-byte format after the displacement field for Operand A.

(continued)

Finally, the length operand (specified as 5) is an implied displacement which is appended after all addressing extensions. It also can be encoded in single-byte format due to its small contents. Thus, the complete machine instruction is:


The complete binary form of this instruction therefore appears in consecutive memory bytes as

## Chapter 5

SERIES 32000 INSTRUCTION SET

This chapter contains the detailed definitions of each of the instructions in the Series 32000 instruction set.

Instructions are presented in the format shown in Figure 5-1. The items indicated there are described below.

1. Mnemonic index. Instructions are alphabetized according to this index, which gives a general form of the mnemonic(s) for each instruction. For a listing of instructions by functional groups, see instead Appendix A or Chapter 3.
2. Enumerated mnemonics. This area holds a list of all valid mnemonic forms for the instruction, if there are alternative forms.
3. Format definition. This area defines the assembly-language and binary formats of the instruction, and the number and kinds of operands. The information contained here is explained in Chapter 4.
4. Instruction description. The operation performed by the instruction is defined here.
5. Flags Affected. All flags in the Processor Status Register which are affected by the instruction are listed. See Section 2.2 for the general definitions of these flags.
6. Traps. Any trap that may be caused by the instruction is listed.

NOTE: Since the Abort trap, Trap (ABT), may occur on any instruction for memory management purposes, it is not listed unless there is a cause which is unique to that instruction.
7. Examples. One or more examples are given, where required, in order to clarify the operation performed by the instruction. Conventions used in presenting example instructions and operands are given in Section 5.1.

Add Quick Integer
Syntax:


The ADDQi instruction adds the src and dest operands and places the result in the dest operand location. Before the addition is performed, src is sign-extended to the length of dest.

Flags Affected: $C$ is set on a carry from addition, cleared if no carry.
$F$ is set on an overflow from addition, cleared if no overflow. Integer carry and overflow conditions are defined in Section 3.1.

Traps:


6 ) $\left\{\begin{array}{l}\text { Integer Overflow Trap (OVF) is activated if the V } \\ \text { and the result cannot be represented exactly in dest. }\end{array}\right.$

## Example:

$$
\text { ADDQB }-8, \mathrm{RO}
$$

OC 04

The above example adds the quick integer -8 to the low-order byte of register R0. The remaining bytes of $R 0$ are unaffected.

The action of the above instruction is illustrated below.

| Operands | Operand Values: Before | $\begin{aligned} & \text { Hex } \text { (Dec) } \\ & \text { After } \end{aligned}$ |
| :---: | :---: | :---: |
| $\begin{gathered} -8 \\ (q u i c k) \end{gathered}$ | $\begin{gathered} \text { F8 } \\ (-8) \end{gathered}$ | -- |
| R0 | $\begin{array}{r} \text { AAAAAA } 78 \\ (+120) \end{array}$ | $\begin{aligned} & \text { AAAAAA } 70 \\ & (+112) \end{aligned}$ |
| UPSR | nzfvxltc | $\mathrm{nz} \underline{0} \mathrm{Vxlt} \underline{1}$ |

* This shows the internal format of the quick operand after sign-extension to byte length. The operand is encoded within the instruction as binary 1000.

Figure 5-1 Typical Instruction Definition

### 5.1 Instruction Examples

Figure 5-2 shows an instruction example from Section 5.2. Each example shows the encodings and the actions of one or more typical forms of the instruction being described.

### 5.1.1 Coding Examples

Example instructions are shown coded both in assembly-language source form and in machine-language form.

The machine-language form is presented in hexadecimal as would be expected in a "dump" format. The leftmost byte displayed occupies the lowest memory address. The entire instruction is presented, including all extensions.

### 5.1.2 Action Examples

The actions of an example instruction are shown in three columns.
The "Operands" column identifies all operands of the instruction: both those explicitly stated in assembly language and those which are implicitly affected by "side-effects" (e.g. the PSR and SP registers where relevant). When a number is presented it generally refers to an operand at that memory address, and is a hexadecimal value. However, if the comment "(immediate)" or "(disp)" appears below it, it is a literal value provided from within the instruction itself, and is presented symbolically as in the assembly-language form of the instruction. Its value appears in the "Before" column.

The "Before" and "After" columns present the values of operands before and after execution of the example instruction. The radixes used in presenting these values are listed in the column heading, as

```
"Hex" = Hexadecimal,
"Binary" = Binary,
"Boolean" = Boolean interpretation of the value (True or False), or
"Dec" = Decimal interpretation of the value. Where a value can be
                                    interpreted as either signed or unsigned, and the distinction is
                                    relevant to the action of the instruction, the terms "Signed"
                                    and "Unsigned" are used.
NOTE: An immediate or displacement value is not considered to have an "After" value, even though it never changes, because it is not available as an immediate or displacement value to any subsequent instructions.
```



* The instruction has not itself changed the contents of these memory locations. However, information that is outside the stack should be considered unpredictable for other reasons. See Section 2.8.1.


### 5.1.3 Operand Presentation Format

The memory format convention used by the Series 32000 family places the leastsignificant byte of a memory operand at the first (i.e. lowest) address. The correct interpretation of a multiple-byte value in memory, therefore, is produced by assembling consecutive bytes of the value from right to left. The address of an operand in memory is also the address of its least-significant byte.

Operand values in examples are presented in units of bytes, words, double-words or quad-words. Each unit is shown in the form corresponding to the interpretation of its contents, so that the least-significant digit of its leastsignificant byte always appears as the rightmost digit.

Units appearing consecutively in memory are separated from each other either horizontally (by a space) or vertically. Memory addresses of consecutive units increase to the right and downward. The value given in the Operand column is the address of the first unit (i.e. the address of its least-significant byte). For example,

$$
\begin{array}{llll}
5000 & 12345678 \text { and } 9000 & 1234 \\
& & 5678 \\
& 9 \mathrm{ABC}
\end{array}
$$

both show three consecutive 16-bit words in memory starting with the value 1234 at address 5000. If the same memory information were presented as consecutive bytes, it would appear as

## 500034127856 BC 9A.

Because an immediate or displacement value is encoded within the instruction format with its most-significant byte at the lowest address (i.e. backward from the ordering used elsewhere in memory), any such value is presented in the form of consecutive bytes.

Hexadecimal and binary operand representations are always presented fully, including any leading zeroes, in order to define the length of each unit unambiguously.

The character "x" means "don't care". Within a value in the Before column, any field made up of these characters is ignored. Within a result in the After column, these represent a field which may be changed unpredictably. In a binary value, each "x" represents one don't care bit. In a hexadecimal value, each "x" represents four bits, all of which are don't care bits.

Filler values of hexadecimal A...A, B...B or C...C are used in examples instead of x...x whenever there is information which is ignored but also not changed. Any decimal interpretation given with the operand ignores these fields. The values 0...0 and F...F are never used as filler, as they occur very often within the significant portion of an operand.

The Processor Status register (PSR) is presented in binary, in the form xxxxIPSU/NZFVxLTC. In the Before column of an example, lower-case letters (e.g. xxxxipsu/nzfvxltc) represent initially unknown values of the corresponding bits. Any bits appearing in the After column which still contain these lowercase symbols have not been changed by the instruction being illustrated, with the exception of all bits shown as "x", which are don't care bits as defined above. Any bits which are changed by the instruction are shown in the After column with their new values underlined. In situations where the most-significant half of the PSR is never used or affected by an instruction, only the least-significant half of the PSR is shown, labeled UPSR for "User PSR".

### 5.2 Instruction Definitions

This section defines the individual Series 32000 instructions. The instructions are ordered alphabetically by their general mnemonic form. For listings of instructions by functional groups, see Appendix A. For help in interpreting the information presented here, see the beginning of this chapter.

ABSf
Absolute Value Floating

```
Syntax: ABSf src, dest ABSF
    gen gen ABSL
    read.f write.f
```



```
The ABSf instruction computes the absolute value of the src operand and places the result in the dest operand location.
Flags Affected: No PSR flags.
The FSR TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes. See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.
Traps: Undefined Instruction Trap(UND) is activated if the \(F\) bit in the CFG register is clear.
Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3.
```


## Example:

```
ABSF F0, F2
BE B5 00
```

This example computes the absolute value of the single-precision number in register FO and places the result in register F 2.

The instruction is illustrated below:

| Operands | Operand Values: Hex <br> Before | (Dec) <br> After |
| :---: | :---: | :---: |
|  |  |  |
| F0 | C2250000 | C2250000 |
|  | $(-41.25)$ | $(-41.25)$ |
| F2 | AAAAAAAA | 42250000 |
|  |  | $(+41.25)$ |

## Absolute Value

## Syntax:



The ABSi instruction computes the absolute value of the src operand and places the result in the dest operand location.

The absolute value of a positive number is the number itself. The absolute value of a negative number is taken by subtracting it (as two's complement) from zero.

Flags Affected: $F$ is set if an overflow from subtraction occurs, cleared otherwise. An overflow condition will occur if the src operand is the most negative number that can be represented in the operand length specified by the programmer. For bytes, this value is -128 (Hex 80): for words it is -32768 (Hex 8000) and for double-words it is $-2,147,483,648$ (Hex 80000000). These values have no corresponding positive values in the same operand length. The result produced on an overflow is the original src operand value.

C is not affected.
Traps: Integer Overflow Trap (OVF) is activated if the V flag is set and the result cannot be represented exactly in dest.

ABSi
Absolute Value (continued)

## Examples:

1. ABSB R5, R6
4E B0 29
2. ABSD $8(\mathrm{SP}), \mathrm{R} 7$
4E F3 C9 08

Example 1 computes the absolute value of the low-order byte of register R5 and places the result in the low-order byte of register $R 6$. The remaining bytes of R6 are not affected.

Example 2 computes the absolute value of the double-word at the memory address specified by $8(\mathrm{SP})$ and places the result in register R7.

These instructions are illustrated below:

|  | Operands | Operand Values: Hex Before | $\begin{array}{r} \text { (Dec) } \\ \text { After } \\ \hline \end{array}$ |
| :---: | :---: | :---: | :---: |
| Ex. 1: | R5 | $\begin{gathered} \text { AAAAAA13 } \\ (+19) \end{gathered}$ | $\begin{gathered} \text { AAAAAA13 } \\ (+19) \end{gathered}$ |
|  | R6 | BBBBBBBB | $\begin{gathered} \text { BBBBBB13 } \\ (+19) \end{gathered}$ |
|  | UPSR | $n z f v x l t c$ | $n z \underline{0} \mathrm{vxltc}$ |
| Ex. 2: | 8 (SP) | $\begin{gathered} \text { FFFFFFFF } \\ (-1) \end{gathered}$ | $\begin{gathered} \text { FFFFFFFF } \\ (-1) \end{gathered}$ |
|  | R7 | AAAAAAAA | $\begin{gathered} 00000001 \\ (+1) \end{gathered}$ |
|  | UPSR | nzfvxltc | $n z \underline{0} \mathrm{vxltc}$ |

## Syntax:



The ACBi instruction adds the inc value to the index operand (after sign- extending the 4 -bit inc value to the length of index) and places the sum in the index operand location. If the sum is not zero, the instruction branches to the location specified as dest. If the sum is zero, the instruction ignores dest and passes control to the next sequential instruction.

In the machine instruction, dest is specified as a displacement from the current contents of the Program Counter; i.e., from the address of the first byte of this instruction. Using the NSC Series 32000 assembler, this displacement may be given explicitly in the form *+disp or *-disp, or dest may be specified as a statement label or as any addressing expression that evaluates to an address accessible via Program Counter Relative addressing. See the applicable assembler manual for further information.

Flags Affected: None.

Traps: Integer Overflow Trap (OVF) is activated if the V flag is set and the result cannot be represented exactly in index.

## Example:

$$
\begin{array}{llll}
\text { LOOP: } & \text { MULD R2, R1 } & \text { CE } 6310 \\
& \text { ACBB -1, R0, LOOP } & \text { CC } 07 & 7 D
\end{array}
$$

In this example, the $A C B B$ instruction adds -1 to the low-order byte of register RO and passes execution control to the MULD statement labeled LOOP as long as the result is not zero. The combined instructions form an iterative loop.

ACBi
Add, Compare and Branch (continued)

The action of each execution of the above ACBB instruction is illustrated below. Initial values for registers R0, R1, and $R 2$ are assumed to be 3, 2, and 2, respectively. Note that at the first execution of the ACBB instruction the first MULD instruction has already been executed. The MULD instruction, labeled LOOP, is assumed to be at address 9000 Hex, and the ACBB instruction is assumed to be at address 9003 Hex.


```
Syntax: ADDf src, dest ADDF
gen gen ADDL
read.f rmw.f
```



The ADDf instruction adds the src and dest operands and places the result in the dest operand location. Results for normalized and zero operands are given in the table below. The symbols "m" and "n" represent any non-zero normalized numbers. The symbols "+z" and "-z" represent positive zero and negative zero, respectively.


* These cases, when the result is zero, select the result based on the current rounding mode selected in the FSR. If the "Round toward Negative Infinity" mode is selected, then the result returned is negative zero. Otherwise, the result returned is positive zero.

Flags Affected: No PSR flags. FSR flags are affected as follows:
UF is set if an underflow occurs; unaffected otherwise.
IF is set on an inexact result; unaffected otherwise.
TT field is set to reflect any exceptional conditions
encountered in executing the instruction. If none is
encountered, TT is set to all zeroes.
See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap (UND) is activated if the $F$ bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3.

## ADDf

Add Floating (continued)

## Examples:

1. ADDF F0, F7
BE C1 01
2. ADDL F2, 16 (SB)
BE 801610

Example 1 adds the single-precision numbers in registers $F 0$ and $F 7$ and places the result in register F 7 .

Example 2 adds the double-precision numbers in register $F 2$ and at the address $16(\mathrm{SB})$ and places the double-precision result at address $16(\mathrm{SB})$.

| Operands | Operand Values: <br> Before | Hex (Dec) <br> After |
| :---: | :---: | :---: |
| F0 | 40840000 | 40840000 |
| $(+4.125)$ | $(+4.125)$ |  |
| F7 |  |  |
|  | 41 D40000 | 41 F50000 |
|  | $(+26.5)$ | $(+30.625)$ |

Ex. 2:
F2
$41 C 0200888300000$
$(+541069584.375$
41C0200888300000
$(+541069584.375) \quad(+541069584.375)$
$16(\mathrm{SB}) \quad 4114 \mathrm{C} 86300000000 \quad 41 \mathrm{C0} 22 \mathrm{~A} 194900000$
$(+340504.75)(+541410089.125)$

Syntax: ADDi src, dest ADBB

read.i rmw.i
ADDD


The ADDi instruction adds the src and dest operands and places the sum in the dest operand location.

Flags Affected: $C$ is set on a carry from addition, cleared if no carry. $F$ is set on an overflow from addition, cleared if no overflow. Integer carry and overflow conditions are defined in Section 3.1

Traps: Integer Overflow Trap (OVF) is activated if the V flag is set and the result cannot be represented exactly in dest.

ADDi
Add (continued)

## Examples:

1. ADDB R0, R1
4000
2. $\operatorname{ADDD} 4(\mathrm{SB}),-4(\mathrm{FP})$
03 D6 04 7C

Example 1 adds the low-order byte of register $R 0$ to the low-order byte of register R1 and places the result in the low-order byte of register R1. The remaining bytes of R1 are not affected.

Example 2 adds double-words. 4(SB) and $-4(F P)$ specify the operand addresses. The instruction places the double-word sum in memory at the address specified by -4 (FP).

|  | Operands | Operand Values: Before | Hex (Dec) After |
| :---: | :---: | :---: | :---: |
| Ex. 1: | R0 | AAAAAA 9 F (-97) | AAAAAA9F (-97) |
|  | R1 | BBBBBB62 $(+98)$ | BBBBBB01 $(+1)$ |
|  | UPSR | $n z f v x l t c$ | $n z \underline{0} \mathrm{vxlt} \underline{1}$ |
| Ex. 2. | 4 (SB) | $\begin{gathered} 20401110 \\ (+541069584) \end{gathered}$ | $\begin{gathered} 20401110 \\ (+541069584) \end{gathered}$ |
|  | -4 (FP) | $\begin{gathered} 0334 A 001 \\ (+53780481) \end{gathered}$ | $\begin{gathered} 2374 \mathrm{~B} 111 \\ (+594850065) \end{gathered}$ |
|  | UPSR | $n z f v x l t c$ | $n z \underline{0} v x l t \underline{0}$ |


| gen | gen | ADDCW |
| :--- | :--- | :--- |
| read.i | rmw.i | ADDCD |



The ADDCi instruction adds the src operand, dest operand, and the $C$ flag and places the sum in the dest operand location.

Flags Affected: $C$ is set on a carry from addition, cleared if no carry. $F$ is set on an overflow from addition, cleared if no overflow. Integer carry and overflow conditions are defined in Section 3.1.

Traps: Integer Overflow Trap (OVF) is activated if the $V$ flag is set and the result cannot be represented exactly in dest.

ADDCi
Add with Carry (continued)

## Examples:

1. $\operatorname{ADDCB} 32$, R0
10 A0 20
2. ADDCD $8(\mathrm{SB})$, R0
13 D0 08

Example 1 adds 32, the low-order byte of register R0, and the $C$ flag contents and places the result in the low-order byte of register $R 0$. The remaining bytes of register $R 0$ are unaffected.

Example 2 adds the double-word at the address specified by $8(\mathrm{SB})$, the contents of the register $R 0$, and the contents of the $C$ flag and places the result in register R0.

In the following illustration, the $C$ flag is assumed to be 1.

Operands \begin{tabular}{c}

| Operand Values: |
| :---: |
| Before |


 

Hex (Dec) <br>
After
\end{tabular}

Ex. 1:
32
(immediate)

20
$(+32)$

| R0 | AAAAAA0F | AAAAAA30 |
| :---: | ---: | ---: |
|  | (+15) | $(+48)$ |
| UPSR | nzfvxlt1 | nźvxltㅇ |

Ex. 2:
8 (SB)
FFFFFFFF
FFFFFFFF
(-1)
(-1)
R0

UPSR
00000030
00000030
(+48)
(+48)

R
nzfvxlt1
$n z \underline{0} v x l t \underline{1}$
gen gen


The ADDPi instruction adds the src operand, dest operand, and the $C$ flag and places the result in the dest operand location as a packed decimal (BCD) integer.

The src and dest operands are interpreted as unsigned packed decimal (BCD) integers. If either operand contains invalid digits, the result is undefined. See Section 3.2 for details of packed decimal arithmetic.

Flags Affected: C is set on a carry, cleared if no carry.
F is cleared.
The packed decimal carry condition is defined in Section 3.2.

Traps:
None.

## ADDPi

Add Packed Decimal (continued)

## Examples:

1. ADDPD R0, R1
4E 7F 00
2. ADDPB $5(\mathrm{SB})$, TOS
4E FC D5 05

Example 1 adds the packed decimal double-word integers contained in registers R0 and R1 and the $C$ flag and places the result in register R1.

Example 2 adds two byte-long packed decimal integers. The integers are at the addresses specified by $5(\mathrm{SB})$ and $T O S$. The instruction places the one-byte result on the top of the stack.

In the following illustrations, the $C$ flag value is assumed to be 0.

|  | Operands | Operand Before | After |
| :---: | :---: | :---: | :---: |
| Ex. 1: | R0 | 75308643 | 75308643 |
|  | R1 | 12345678 | 87654321 |
|  | UPSR | $n z f v x l t 0$ | $n z \underline{0} v x l t \underline{0}$ |
| Ex. 2: | 5 (SB) | 99 | 99 |
|  | SP | 0000FFDE | O000FFDE |
|  | Stack: |  |  |
|  | 0000FFDE | 01 | 00 ** |
|  | 0000FFDF | AA | AA |
|  | UPSR | $n z f v x l t 0$ | $n z \underline{0} \mathrm{vxlt} \underline{1}$ |

* The hexadecimal representation also expresses the decimal interpretation of the value.
** In Example 2, a carry occurs.


The ADDQi instruction adds the src and dest operands and places the result in the dest operand location. Before the addition is performed, src is sign-extended to the length of dest.

Flags Affected: C is set on a carry from addition, cleared if no carry. $F$ is set on an overflow from addition, cleared if no overflow. Integer carry and overflow conditions are defined in Section 3.1.

Traps: Integer Overflow Trap (OVF) is activated if the $V$ flag is set and the result cannot be represented exactly in dest.

## Example:

ADDQB $-8, \mathrm{R} 0$
OC 04

The above example adds the quick integer -8 to the low-order byte of register R0. The remaining bytes of $R 0$ are unaffected.

The action of the above instruction is illustrated below.

| Operands | Operand <br> Before | $\begin{aligned} & \text { Iex } \text { (Dec) } \\ & \text { After } \\ & \hline \end{aligned}$ |
| :---: | :---: | :---: |
| $\begin{gathered} -8 \\ \text { (quick) } \end{gathered}$ | $\begin{gathered} \text { F8 * } \\ (-8) \end{gathered}$ | -- |
| R0 | AAAAAA 78 $(+120)$ | $\begin{array}{r} \text { AAAAAA70 } \\ (+112) \end{array}$ |
| UPSR | nzfvxltc | $n z \underline{0} \mathrm{vxlt} \underline{1}$ |

[^0]
## ADDR

Compute Effective Address

## Syntax: ADDR src, dest

gen gen
addr write.D


The ADDR instruction places the effective address of the src operand into the dest operand location. The src operand itself is not referenced.

Flags Affected: None.

Traps: None.

## Example:

ADDR 4(FP), R0
27 C0 04

This example places the effective address specified as 4(FP) into register R0.

| Operands | Operand Values: <br> Before | Hex <br> After |
| :---: | :--- | :---: |
| FP | 00001000 | 00001000 |
| R0 | AAAAAAAA | 00001004 * |

* The effective address of $4(F P)$ is the sum of the contents of the FP register (H'1000) and the displacement 4, as defined for the Frame Memory addressing mode.
gen
read.i
ADJSPD

```
! src ! ADJSPi
+---------+-----------------+---+
! gen !1 0 1 0 1 1 1 1 1 1!! i !
!-++-+-+-++-+-+-++-!-+-++-+-++-+-+-++-!
    15 8 7 0
```

The ADJSPi instruction adjusts the value of the current stack pointer by subtracting the src operand from it. This has the effect of lengthening the stack by the number of bytes given in the src operand if positive, and shortening it if src is negative. The $S$ flag in the PSR determines whether the current stack pointer register is SPO or SP1. Regardless of the length of the src operand, the entire stack pointer is modified. The src operand is interpreted as a signed integer, and is sign-extended to 32 bits before the subtraction is performed.

Flags Affected: None.
Traps: None.

## Example:

ADJSPD -4(FP) 7F C5 7C

This instruction subtracts the double-word at address -4(FP) from the contents of the current stack pointer, lengthening the stack by that number of bytes.

In the following illustration, the PSR $S$ flag is assumed to be set, selecting register SP 1 as the current stack pointer.

| Operands | Operand Values: <br> Before | Hex <br> After |
| :---: | :--- | :---: |
| $-4($ FP $)$ | 00000010 | 00000010 |
| SP1 | 00001010 | 00001000 |

ANDi
And
Syntax

| ANDi | src, | dest | ANDB |
| :--- | :--- | :--- | :--- |
|  | gen | gen | ANDW |
|  | read.i | rmw.i | ANDD |



The ANDi instruction performs a bit-wise logical AND on the src and dest operands and places the result in the dest operand location.

The instruction ANDs each bit in src with the corresponding dest bit. If two corresponding bits are both "1", the dest bit is set to "1"; otherwise, the dest bit is set to "0".

Flags Affected: None.
Traps: None.

## Example:

ANDB R0, R1
6800

This example ANDs the low-order bytes of registers R0 and R1 and places the result in the low-order byte of R1. The instruction affects only the low-order byte of R1.

The instruction is illustrated below:

| Operands | Operand Values: <br> Before | Binary <br> After |
| :---: | :---: | :---: |
| R0 | 10010010 | 10010010 |
| (low byte) | 01110111 | 00010010 |
| R1 |  |  |

## Arithmetic Shift

## Syntax:



The ASHi instruction performs an arithmetic shift on the dest operand in the manner specified by the count operand. The sign of count determines the direction of the shift. The absolute value of count gives the number of bit positions to shift the dest operand.

The count operand value must be within the range -7 to +7 for the ASHB form, -15 to +15 for the ASHW form, and -31 to +31 for the ASHD form. A positive count specifies a left shift; a negative count specifies a right shift. In an arithmetic left shift, high-order bits (including the sign bit) shifted out of dest are lost, and low-order bit positions emptied by the shift are zero-filled. In an arithmetic right shift, low-order bits shifted out of dest are lost, and all high-order bit positions emptied by the shift are filled from the original sign bit of dest.

The count and dest operands are interpreted as signed integers.

Flags Affected: None.
Traps: Integer Overflow Trap (OVF) is activated if the V flag is set and the result cannot be represented exactly in dest.

## ASHi

## Arithmetic Shift (continued)

## Examples:

1. ASHB 2, 16(SB)
4E 84 A6 0210
2. ASHB TOS, $16(\mathrm{SB})$
4E 84 BE 10

Example 1 shifts the byte specified by $16(S B)$ two bit positions to the left.
Example 2 pops a byte from the top of the currently-selected stack. Based on this value, it shifts the byte specified by $16(\mathrm{SB})$ accordingly.

|  | Operands | Operand Values: Before | $\begin{array}{r} \text { Binary (Dec) } \\ \text { After } \\ \hline \end{array}$ |
| :---: | :---: | :---: | :---: |
| Ex. 1. | $\begin{gathered} 2 \\ \text { (immediate) } \end{gathered}$ | $\begin{array}{r} 00000010 \\ (+2) \end{array}$ | -- |
|  | 16 (SB) | $\begin{array}{r} 00011111 \\ (+31) \end{array}$ | $\begin{array}{r} 01111100 \\ (+124) \end{array}$ |
| Ex. 2: | Stack: |  |  |
|  | (48000) | 11111110 (-2) | xxxxxxxx |
|  | (48001) | 10101010 | 10101010 |
|  | 16 (SB) | $\begin{array}{r} 11111000 \\ (-8) \end{array}$ | $\begin{array}{r} 11111110 \\ (-2) \end{array}$ |
|  | SP | (48000) | (48001) |

## Syntax: Bcond

## dest

disp


The Bcond instruction branches to the location specified as dest if the condition specified by cond is true. If the condition is false, execution continues with the next sequential instruction.

Cond is a two character condition name that specifies the state of a flag or flags in the PSR. If the flag(s) have the specified state, the condition is true! otherwise, the condition is false.

The Conditional Branch instruction may specify the following conditions:

| Condition | Condition <br> Name | True <br> State | Short <br> Field |
| :--- | :--- | :--- | :--- |
| Equal | EQ | Z flag set | 0000 |
| Not Equal | NE | Z flag clear | 0001 |
| Carry Set | CS | C flag set | 0010 |
| Carry Clear | CC | C flag clear | 0011 |
| Higher | HI | L flag set | 0100 |
| Lower or Same | LS | L flag clear | 0101 |
| Greater Than | GT | N flag set | 0110 |
| Less Than or Equal | LE | N flag clear | 0111 |
| Flag Set | FS | F flag set | 1000 |
| Flag Clear | FC | F flag clear | 1001 |
| Lower | LO | Z and L flags clear | 1010 |
| Higher or Same | HS | Z or L flag set | 1011 |
| Less Than | LT | Z and N flags clear | 1100 |
| Greater Than or Equal | GE | Z or N flag set | 1101 |

The condition name is appended to the instruction mnemonic as illustrated in the following examples. The name is translated at assembly time to the corresponding 4-bit Short field of the basic instruction.

The interpretation of condition codes is such that the instruction sequence

| CMPB | A, B |
| :--- | :--- |
| BGT | ERROR |

will cause a branch if operand $A$ is greater than operand $B$ in the CMPB instruction.

## Bcond

In the machine instruction, dest is specified as a displacement from the current contents of the Program Counter; i.e., from the address of the first byte of this instruction (see Section 4.2.3 for displacement formats). Using the ASM16 assembler, this displacement may be given explicitly in the form *+disp or *-disp, or dest may be specified as a statement label or any addressing expression which evaluates to an address accessible via Program Counter Relative addressing. See the applicable assembler manual for further information.

Flags Affected: None.

Traps: None.

## Examples:

1. BLO LOOP
AA BF 66
2. BNE *+10
1A 0A

Example 1 passes execution control to the instruction labeled LOOP if the $Z$ and $L$ flags in the PSR are 0 .

Example 2 passes execution control to a nonsequential instruction if the Z flag is 0. The instruction passes execution control by adding 10 to the PC register.

In the following illustrations, the $Z$ and $L$ flags are assumed to be zero. LOOP is assumed to be the label of a statement beginning at address 9000 Hex.

Operand Values: Hex (Dec)

Ex. 1:

| Operand | Before | After |
| :---: | :---: | :---: |
| PC | $\begin{aligned} & 0000909 \mathrm{~A} \\ & (37018) \end{aligned}$ | $\begin{aligned} & 00009000 \\ & (36864) \end{aligned}$ |
| $\begin{gathered} \text { LOOP } \\ (\text { disp) } \end{gathered}$ | $\begin{aligned} & \mathrm{BF} 66 \\ & (-154) \end{aligned}$ | -- |
| UPSR | n0fvx0tc | n0fvx0tc |
| PC | $\begin{aligned} & 00009 \mathrm{FFO} \\ & (40944) \end{aligned}$ | $\begin{aligned} & 00009 \mathrm{FFA} \\ & (40954) \end{aligned}$ |
| $\begin{gathered} \text { *+10 } \\ (\text { disp) } \end{gathered}$ | $\begin{gathered} 0 \mathrm{~A} \\ (+10) \end{gathered}$ | -- |

UPSR n0fvx0tc n0fvx0tc

## Bit Clear

## Syntax:



The BICi instruction clears (sets to 0) the bits in the dest operand that correspond to the "1" bits in the src operand.

Flags Affected: None.
Traps: None.

## Example:

```
BICB R0, 3(SB)

This example clears the bits, in the byte at address 3(SB), corresponding to the "1" bits in the low-order byte of register R0.
\begin{tabular}{llc} 
Operands & \begin{tabular}{c} 
Operand Values \\
Before
\end{tabular} & \begin{tabular}{c} 
(Binary) \\
After
\end{tabular} \\
\hline \begin{tabular}{c} 
R0
\end{tabular} & 10011001 & 10011001 \\
(low byte) & & \\
\(3(\) SB \()\) & 11110000 & 01100000
\end{tabular}

\section*{BICPSRB}

\section*{BICPSRW}

Bit Clear in PSR

Syntax: BICPSRB

\section*{src}
gen
read. B


Syntax: BICPSRW src
gen
read.w


The Bit Clear in PSR instructions clear (set to 0) the bits in the PSR corresponding to the "1" bits in the src operand. The BICPSRB instruction affects only the low-order byte of the PSR; the BICPSRW instruction affects the entire PSR.

Flags Affected: Flags specified by src "1" bits are cleared.
Traps: Illegal Operation Trap (IIL) is activated if a BICPSRW instruction is attempted while the PSR \(U\) flag is set.

\section*{Example:}

\section*{BICPSRB B'10100010}

7C A1 A2

This instruction clears bits 1, 5 and 7 in the low-order byte of the PSR. These are the \(T, F\) and \(N\) flags, respectively.

The instruction is illustrated below:
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values \\
Before
\end{tabular} & \begin{tabular}{c} 
(Binary) \\
After
\end{tabular} \\
\hline B'10100010 & 10100010 & -- \\
(immediate) & & \\
PSR & xxxxipsu/nzfvxltc & xxxxipsu/ \(\underline{0} \underline{0} \mathbf{v x l} \underline{0} c\)
\end{tabular}
```

    src
            gen
            read.B
    ```


\section*{Syntax: BISPSRW src}
gen
read.W


The BISPSRB and BISPSRW instructions set the bits in the PSR corresponding to the "1" bits in the src operand.

Flags Affected: Flags specified by src "1" bits are set.

Traps: Illegal Operation Trap (ILL) is activated if a BISPSRW instruction is attempted while the PSR U flag is 1.

\section*{Example:}
```

BISPSRB B'10100010
7C A3 A2

```

This instruction sets bits 1, 5 and 7 in the low-order byte of the PSR. These are the \(T, F\) and \(N\) flags, respectively.

Operand Values (Binary)
Operands Before After

> B'10100010
(immediate)

PSR
```

                                xxxxipsu/nzfvxltc
    ```
xxxxipsu/1z1vxĺc

BPT
Breakpoint Trap

\section*{Syntax: BPT}
\begin{tabular}{|c|c|}
\hline ! & BPT \\
\hline & 10010 ! \\
\hline & +-+-+-+-+-! \\
\hline 7 & 0 \\
\hline
\end{tabular}

The BPT instruction activates the Breakpoint Trap (BPT). The return address pushed on the Interrupt Stack is the address of the BPT instruction itself.

Flags Affected: None.
Traps: Breakpoint Trap (BPT) is activated.

\section*{Example:}
BPT
F2

\section*{Syntax: BR dest}
disp


The \(B R\) instruction branches to the location specified as dest.
In the machine instruction, dest is specified as a displacement from the current contents of the Program Counter; i.e., from the address of the first byte of this instruction. Using the NSC Series 32000 assembler, this displacement may be given explicitly in the form *+disp or *-disp, or dest may be specified as a statement label or any addressing expression which evaluates to an address accessible via Program Counter Relative addressing. See the applicable assembler manual for further information.

Flags Affected: None.

Traps: None.

\section*{Examples:}
\begin{tabular}{lll} 
1. & BR ERROR & EA BF 66 \\
2. & BR & \(\star+10\)
\end{tabular}

Example 1 passes execution control to the instruction labeled ERROR.
Example 2 passes execution control to a nonsequential instruction by adding 10 to the PC register.

In the following illustration, ERROR is assumed to be the label of a statement beginning at address 9000 Hex.
\begin{tabular}{|c|c|c|c|}
\hline & Operand & Operand Values: Before & \[
\begin{aligned}
\text { Hex } & \text { (Dec) } \\
& \text { After }
\end{aligned}
\] \\
\hline \multirow[t]{2}{*}{Ex. 1:} & PC & \[
\begin{aligned}
& 0000909 \mathrm{~A} \\
& (37018)
\end{aligned}
\] & \[
\begin{array}{r}
00009000 \\
(36864)
\end{array}
\] \\
\hline & \[
\begin{aligned}
& \text { ERROR } \\
& \text { (disp) }
\end{aligned}
\] & \[
\begin{gathered}
\mathrm{BF} 66 \\
(-154)
\end{gathered}
\] & -- \\
\hline \multirow[t]{2}{*}{Ex. 2:} & PC & \[
\begin{aligned}
& 00009 \mathrm{FFO} \\
& (40944)
\end{aligned}
\] & \[
\begin{aligned}
& 00009 \mathrm{FFA} \\
& (40954)
\end{aligned}
\] \\
\hline & \[
\begin{gathered}
*+10 \\
(\text { disp })
\end{gathered}
\] & \[
\begin{gathered}
0 \mathrm{~A} \\
(+10)
\end{gathered}
\] & -- \\
\hline
\end{tabular}

\section*{BSR}

Branch To Subroutine

\section*{Syntax: BSR dest}
disp
\begin{tabular}{|c|c|}
\hline ! & BSR \\
\hline \multicolumn{2}{|l|}{\(10000000 c c c l\)} \\
\hline & -+-+-+-+-! \\
\hline 7 & 0 \\
\hline
\end{tabular}

The BSR instruction calls the local procedure at the address specified as dest. It does so by pushing the address of the next sequential instruction onto the currently-selected stack and branching.

In the machine instruction, dest is specified as a displacement from the current contents of the Program Counter: i.e., from the address of the first byte of this instruction. Using the NSC Series 32000 assembler, this displacement may be given explicitly in the form *+disp or *-disp, or dest may be specified as a statement label or any addressing expression which evaluates to an address accessible via Program Counter Relative addressing. See the applicable assembler manual for further information.

Flags Affected: None.
Traps: None.

\section*{Example:}

BSR CALC 0210
This example causes a program to branch to the local procedure labeled CALC after saving the address of the next sequential instruction on the stack.

The action of the above instruction is illustrated below. The statement labeled CALC is assumed to be at address 9010 Hex, 16 bytes forward from the first byte of this instruction. Since the displacement field is one byte long, the total length of this instruction is two bytes. The return address is therefore the original PC contents plus two.
\begin{tabular}{ccc} 
Operand & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex \begin{tabular}{c} 
(Dec) \\
After
\end{tabular} \\
\hline PC
\end{tabular} 00009000
\end{tabular}

CASEB
gen
read.i


The CASEi instruction branches to a nonsequential instruction by adding the src operand to the PC register. The src operand is interpreted as a signed integer, and is sign-extended to 32 bits before the addition is performed.

A Case Branch instruction, using Scaled Indexing and a table of branch offsets, may be used to implement a multiway branch. See example below.

Flags Affected: None.
Traps: None.

\section*{Example:}
```

CASEB TABLE[R7:B] 7C E7 DF 04

```

TABLE: ; (starts here)
This example branches to a nonsequential instruction by adding the byte at the address specified by TABLE[R7:B] to the PC register. The entire contents of register R7 determine the location of the operand to be added.
\begin{tabular}{ccc} 
Operand & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex \\
After
\end{tabular} \\
PC & 00009000 & \(0000906 A\) \\
R7 & 00000005 & 00000005 \\
TABLE & \\
\(00009004 *\) & 0A 1A 3A 5A 7A 6A 4A & 0A 1A 3A 5A 7A 6A 4A \\
TABLE[R7:B] & 6A & 6A
\end{tabular}
* Address 9004 (Hex) marks the beginning of a table of branch offsets. In this example, the table is located directly after the CASEB instruction, and is accessed via the Program Memory addressing mode, with a displacement of 4. Since register R7 contains 5, the effective address of TABLE[R7:B] is 9009 (Hex). This means the sixth branch offset (6A ) is selected to be added to the PC register.

CBITi
CBITII
Clear Bit, Clear Bit Interlocked
\begin{tabular}{lll} 
Syntax: CBITi & offset, base & CBITB \\
& gen & gen \\
& read.i & regaddr
\end{tabular}


Syntax: CBITIi offset, base
gen gen
read.i regaddr


The CBITi and CBITIi instructions clear (set to 0) the register or memory bit specified by base and offset after copying the bit value to the \(F\) flag in the PSR.

The CBITIB, CBITIW, and CBITID instructions, in addition, activate the Interlocked Operation output pin on the CPU, which may be used in multi-processor systems to interlock accesses to semaphore bits. See the applicable CPU data sheet for further details.

The location of the bit is determined from offset and base. Offset is a general operand, whose length is given by the operation length suffix. Base is an addressing expression giving a byte address from which offset specifies a bit position. See Section 3.5 for details of specifying bit positions.

If base is a register, then the bit is within that register, at the bit position given by the offset operand. If base is a memory location, then the bit is at bit position
\[
\text { offset MOD } 8
\]
within the memory byte whose address is
EA(base) + (offset DIV 8),
where EA(base) is the effective address of base. See Section 3.5 for definitions of the operators MOD and DIV above, and for further details of bit instructions.

Offset is interpreted as a signed integer.

\section*{Clear Bit (continued)}

Flags Affected: \(F\) is set to the original value of the specified bit.
Traps: None.
Example:
CBITW R0, 0 (R1) 4E 490200
This example clears a bit in memory after copying the bit value to the f flag. For designating the location of the target bit, the low-order word of register R0 supplies the bit offset, and \(0(R 1)\) is specified as the base address.

In the following illustration, the target bit is assumed to be 1 prior to instruction execution.
\begin{tabular}{|c|c|c|}
\hline Operands & Operand Values: Before & Hex (Dec) [Binary]
\(\qquad\) \\
\hline \[
\begin{gathered}
\mathrm{R0} \\
(\mathrm{offset})
\end{gathered}
\] & \[
\begin{aligned}
& \text { AAAA004C } \\
& (+76)
\end{aligned}
\] & AAAA004C
\[
(+76)
\] \\
\hline R1 & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] \\
\hline base address 0 (R1) & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] & -- \\
\hline \[
\begin{gathered}
00001009 \text { * } \\
(+4105)
\end{gathered}
\] & \[
\begin{gathered}
10 \\
{[00010000]}
\end{gathered}
\] & \[
\begin{gathered}
00 \\
{[000 \underline{0} 0000]}
\end{gathered}
\] \\
\hline UPSR & \(n z f v x l t c\) & \(n z \underline{1}\) vxltc \\
\hline
\end{tabular}
* The address 1009 (Hex) is the effective address of the byte containing the desired bit. This address is computed from the offset and the base address as follows:
\[
\begin{array}{cc}
\text { base address } & + \text { (offset DIV 8) } \\
4096 & + \\
4105 \text {, or } 1009(\mathrm{Hex})
\end{array}
\]

The bit number within this byte is calculated as:
\[
\begin{array}{ccc}
\text { Offset } & \text { MOD } & 8 \\
76 & \text { MOD } & 8 \\
4 & . &
\end{array}
\]
\begin{tabular}{lllll} 
Syntax: & checki \\
& regt, bounds, & grc & gen & gen \\
& & addr & read. & CHECKB \\
& & CHECKW
\end{tabular}


The CHECKi instruction compares the src operand against an upper and lower bound from the bounds operand, determining whether it is within those bounds. The instruction then subtracts the lower bound from src, placing the result as a 32-bit value into the general-purpose register specified as dest. This "zero-adjusted" result is usable directly as either an index into a one-dimensional array (within an addressing mode using Scaled Indexing) or as an input value to the INDEX instruction for generating an index into a multidimensional array. See Section 3.9 for details of array access.

The bounds operand contains two values--an upper bound followed by a lower bound--as shown:

\section*{bounds:}


The upper and lower bounds each have the same length as the src operand. Thus, the entire bounds operand is twice the length of the src operand.

If src is greater than the upper bound or less than the lower bound, it is "out of bounds" and the \(F\) flag is set to 1 . If src is within the upper and lower bounds, the \(F\) flag is cleared.

The instruction places the zero-adjusted result into the dest register. The zero-adjusted value is computed as "src - lower bound". The result is zeroextended to 32 bits. If src is out of bounds, the result placed in dest is undefined.

The src and bounds operands are interpreted as signed integers. The result placed in the dest register is a 32-bit unsigned integer.

Flags Affected: \(F\) is set if src is out of bounds, cleared otherwise.
Traps: Integer Overflow Trap (OVF) is activated if the \(V\) flag is set and the src operand is out of bounds.

\section*{Example:}
```

CHECKB R0, 4(SB), R2
EE 80 DO 04

```

This example compares the low-order byte of register \(R 2\) with each of the two one-byte bounds at the address specified by \(4(S B)\). The instruction sets or clears the \(F\) flag to indicate the comparison result and then subtracts the lower bound from the low-order byte of \(R 2\), placing the result in register \(R 0\).

The instruction is illustrated below. An array index in the low-order byte of R2 is being checked against its range of [1..10] and then "zero-adjusted" to its corresponding value for the range [0..9]. The result is being placed into R0 for use in an addressing mode using Scaled Indexing.
\(\left.\begin{array}{cc}\text { Operands } & \begin{array}{c}\text { Operand Values: } \\ \text { Before }\end{array}\end{array} \begin{array}{c}\text { Hex (Dec) } \\ \text { After }\end{array}\right]\)
* The result in \(R 0\) represents the result of adjusting the value 3 from a range of [1..10] to the zero-based range of [0..9]. The corresponding adjusted value is 2.

\section*{CINV}

Cache Invalidate

\section*{Syntax: CINV options, sr}


The CINV instruction invalidates the contents of locations in the on-chip Data Cache and Instruction Cache. The instruction can be used to invalidate either the entire contents of the on-chip caches or only a 16-byte block (one cache line). In the latter case, the 28 most-signifcant bits of the source operand specify the physical address of the aligned 16-byte block; the 4 least-significant bits of the source operand are ignored. If the specified block is not located in the onchip caches, then the instruction has no effect. If the entire cache contents is to be invalidated, then the source operand is read, but its value is ignored.

Options are specified by listing the letters A (invalidate All), I (Instruction Cache), and D (Data Cache). If neither the I nor D option is specified, the instruction has no effect.

In the instruction encoding, the options are represented in the A, I and D fields as follows:
```

    A: 0 - invalidate only a 16-byte block (one cache line)
    1 - invalidate the entire cache
    I: 0 - do not affect the Instruction Cache
    1 - invalidate the Instruction Cache
    D: 0 - do not affect the Data Cache
1 - invalidate the Data Cache

```

Flags Affected: None.
Traps: Illegal Operation Trap (ILL) is activated if this instruction is attempted while the PSR U bit is set.

\section*{Examples:}
1. CINV A,D,I,R3
1E A7 1B
2. CINV I,R3
1E 2719

Example 1 invalidates the entire Instruction Cache and Data Cache.

Example 2 invalidates the 16-byte block (one cache line) whose physical address in the Instruction Cache is contained in R3.

CMPf
Compare Floating
Syntax: CMPf src1, src2 CMPF
gen gen CMPL
read.f read.f


The CMPf instruction compares the src1 and src2 operands and sets the \(Z\) and \(N\) flags to indicate comparison results. Positive and negative zero are equal.

Flags Affected: Z is set if src1 equals src2, cleared otherwise.
N is set if src1 is greater than src2, cleared otherwise. L is cleared always. The FSR TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes.
See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3.

\section*{Example:}

CMPF F0, F2 BE 8900

This example compares the single-precision numbers in registers FO and F2.
\begin{tabular}{ccr} 
Operands & \begin{tabular}{c} 
Operand Values: Hex \\
Before
\end{tabular} & \begin{tabular}{c}
\((\) Dec \()\) \\
After
\end{tabular} \\
\hline F0 & 42250000 \\
\((+41.25)\)
\end{tabular}\(\quad\)\begin{tabular}{r}
42250000 \\
\((+41.25)\)
\end{tabular}

\section*{Compare}

\section*{Syntax:}


The CMPi instruction compares the src1 and src2 operands and sets the \(Z\), \(N\), and \(L\) flags in the PSR to indicate the comparison result. The \(N\) flag indicates the result of a signed integer comparison; the \(L\) flag indicates the result of an unsigned integer comparison. Both types of comparison are performed.

Flags Affected: Z is set if src1 is equal to src2, cleared otherwise.
N is set if src1 is greater than src2 (signed comparison), cleared otherwise.
L is set if src1 is greater than src2 (unsigned comparison), cleared otherwise.

\section*{Traps: \\ None.}

\section*{Example:}

CMPB 7(SB), 4(R0)
04 D2 0704

This example compares byte operands. \(7(S B)\) and \(4(R 0)\) specify the operand addresses.

In the following illustration, operand values before instruction execution are assumed: Z, \(N\) and \(L\) flag values are unknown.

gen gen disp


The CMPMi instruction compares the contents of block1 and block2 and sets the \(Z\), N , and L flags to indicate the comparison result. The blocks are comprised of integers of length i. The number of integers is specified by length.

The instruction compares two integers (one from each block) at a time. If the current integers are equal, the instruction continues with the next two integers; otherwise, the instruction sets the PSR flags and terminates.

The \(N\) flag indicates the result of signed integer comparison. The L flag indicates the result of unsigned integer comparison. Both types of comparison are performed.

In assembly language, the length operand is specified as the number of integers in each block. In the machine instruction, however, the length operand is encoded according to the formula
\[
(\text { num }-1) * i
\]
where num is the number of integers in each block, and i is the number of bytes per integer.

A block may not be greater than 16 bytes in length.
Flags Affected: Z is set if block1 and block2 are equal for their entire length; cleared otherwise.
\(N\) is set if, in the first unequal pair of integers, the block1 integer is greater than the block2 integer (signed comparison); cleared otherwise.

L is set if, in the first unequal pair of integers, the block1 integer is greater than the block2 integer (unsigned comparison); cleared otherwise.

Traps: None.

\section*{Compare Multiple (continued)}

\section*{Example:}

CMPMW 10(R0), 16(R1), 4
CE 45 42 OA 1006

This instruction compares four word-long integers from the block starting at the address specified by \(10(R 0)\) to the block starting at the address specified by 16 (R1).
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & Hex \\
R0 & (Signed) [Unsigned] \\
After
\end{tabular}

\footnotetext{
* The address of the first block as specified by 10 (R0).
** The address of the second block as specified by 16(R1).
}

\section*{CMPQi}

Compare Quick Integer
```

Syntax: CMPQi src1, src2 CMPQB
quick gen CMPQW
read.i CMPQD

```
```

! src2 ! src1 ! CMPQi !
+---------+-------+---------+---+
! gen ! quick !0 0 1 1 1! i !
!-+-+-+-++-+-+-++-!-+-+-+-++-+-+-++-!
15 8 7 0

```

The CMPQi instruction compares the src1 and src2 operands and sets the \(Z\), \(N\), and L flags to indicate the comparison result. Before the comparison, src1 is sign-extended to the length of src2. Leading "0"s are supplied for a positive src operand value; leading "1"s for a negative value.

The \(N\) flag indicates the result of signed integer comparison. The L flag indicates the result of unsigned integer comparison. Both types of comparison are performed.

Flags Affected: Z is set if src1 is equal to src2, cleared otherwise. N is set if src1 is greater than src2 (signed comparison), cleared otherwise.
L is set if src1 is greater than src2 (unsigned comparison), cleared otherwise.

Traps: None.

\section*{Example:}

CMPQB -8, R0
1C 04

This example compares the quick integer -8 with the low-order byte of register R0.

Operand Values: Hex (Signed) [Unsigned]
Operands
Before
After

* This shows the internal format of the quick operand after sign-extension to Byte length. The operand is encoded within the instruction as binary 1000.


\section*{Syntax: CMPST options}


Operands of the CMPSi and CMPST instructions are specified in General Purpose Registers:
```

R0 - Number of string elements to be processed.
R1 - Address of current String 1 element.
R2 - Address of current String 2 element.
R3 - Address of translation table (CMPST form only).
R4 - Match value (with Until Match or While Match option only).

```

The CMPSi instruction compares corresponding integer elements from String 1 (address in R1) and String 2 (address in R2) and sets the \(Z, N\) and L flags to indicate the comparison results (see "Flags Affected" below). If the current two elements are equal, the instruction compares the next two elements; otherwise, it terminates. After each comparison, the instruction sets register R0 to the number of elements remaining to be compared and sets registers \(R 1\) and \(R 2\) to the addresses of the next elements to be compared. See Section 3.7 for the exact sequences performed by String instructions.

The \(N\) flag indicates the result of signed integer comparison. The L flag indicates the result of unsigned integer comparison. Both types of comparison are performed.

The CMPST instruction compares one-byte elements in String 1, after translation, to one-byte elements in String 2. The translated value to be compared is found by adding the current element from the first string as an unsigned integer to the translation table address found in register R3. The instruction compares elements, sets flags, and sets registers as described above. See Section 3.7 for details of string translation.

Options may be specified by listing the letters B (Backward), U (Until Match) and W (While Match) as operands. The U and W options are mutually exclusive. See Section 3.7 for details of the options available in String instructions.

In the machine instruction, the options are encoded in the \(B\) and UW fields as follows:
\begin{tabular}{rl} 
B field \(=\)\begin{tabular}{l}
0 \\
1
\end{tabular} & \begin{tabular}{l} 
Forward direction. \\
Backward direction. \\
UW field \(=0\) \\
01
\end{tabular} \\
10 & Neither Until Match nor While Match. \\
11 & While Match. \\
& (reserved) \\
& Until Match.
\end{tabular}

String instructions are interruptible. See Section 3.7.
Flags Affected: Z, N and L are affected, as given below.
F is set if the \(U\) or \(W\) option is specified and the corresponding Until/While condition is met, otherwise it is cleared.

Because of the variety of termination conditions possible in a CMPS instruction, the following sequence is recommended to interpret the flag settings:
1. If the \(U\) or \(W\) option is specified, then check the \(F\) flag. If it is set, then the CMPS instruction has terminated because of the Until Match or While Match test, and the other flag settings are \(Z=1, N=0, L=0\). Register R1 holds the address of the String 1 element which caused termination, and register R2 holds the address of the corresponding element in String 2. Register R0 contains the number of elements left to be processed, including the element which terminated the instruction.
2. If \(F=0\), check the \(Z\) flag. If it is set, then the CMPS instruction has terminated because the limit count in \(R 0\) has been decremented to zero, and the strings are equal up to that point. Registers R1 and R2 hold the addresses of the next (unprocessed) string elements, and the remaining flag settings are \(\mathrm{N}=0\), \(\mathrm{L}=0\).
3. If neither the \(F\) or \(Z\) bit is set (above), then the CMPS instruction has terminated because the strings are unequal. Registers R1 and R2 hold the addresses of the first two string elements that are unequal, and the \(N\) and L flags show their relation. Register R0 holds the number of remaining elements, including the element at which the instruction has stopped. If the \(N\) bit is set, the element from String 1

\section*{Compare Strings (continued)}
(indicated by R1) is greater than the element from String 2 (indicated by R2), where both are interpreted as signed integers. If the L bit is set, the element from String 1 is greater than the element from String 2, where both are interpreted as unsigned integers.

Traps: None.

\section*{Example:}
CMPSB
OE 0400

This example compares l-byte elements from two strings until either an unequal pair is found or the limit count in RO decrements to zero.

The action of the above instruction on two unequal strings is illustrated below. The underlined string elements show the point at which the instruction terminates.
\begin{tabular}{|c|c|c|}
\hline Operands & Operands Values: Before & \[
\begin{aligned}
\text { Hex } \begin{array}{c}
\text { (Dec) } \\
\\
\\
\text { After }
\end{array} \\
\hline
\end{aligned}
\] \\
\hline R0 & \[
\begin{array}{r}
00000020 \\
(+32)
\end{array}
\] & \[
\begin{array}{r}
0000000 \mathrm{~F} \\
(+15)
\end{array}
\] \\
\hline R1 & 00002000 & 00002011 \\
\hline R2 & 0000F000 & 0000F011 \\
\hline UPSR & \(n z f v x l t c\) & \(\underline{100 v x 1 t c}\) \\
\hline
\end{tabular}

Starting String Contents
Addresses
\begin{tabular}{lllllllllllllllll}
2000 & \(1 E\) & 04 & 05 & \(1 C\) & \(0 A\) & 14 & \(0 C\) & \(0 B\) & 09 & 07 & \(1 F\) & \(0 F\) & 17 & 01 & 00 & 11 \\
& \(1 F\) & \(1 E\) & \(1 A\) & 09 & 01 & 12 & 14 & \(0 E\) & \(1 E\) & \(0 A\) & 00 & 03 & 09 & 06 & 16 & 18 \\
F000 & \(1 E\) & 04 & 05 & \(1 C\) & \(0 A\) & 14 & \(0 C\) & \(0 B\) & 09 & 07 & \(1 F\) & \(0 F\) & 17 & 01 & 00 & 11 \\
& \(1 F\) & \(1 D\) & \(1 A\) & 09 & 01 & 12 & 14 & \(0 E\) & \(1 E\) & \(0 A\) & 00 & 03 & 09 & 06 & 16 & 18
\end{tabular}

\section*{COMi}

Complement

Syntax:


The COMi instruction places the one's complement of the src operand in the dest operand location.

The one's complement is the logical NOT operation performed on each bit of src. If the src operand bit is "0", then the corresponding dest operand bit is set to "1". Otherwise, the dest bit is set to "0".

Flags Affected: None.
Traps: None.

\section*{Example:}
```

COMB R0, -4 (FP) 4E 3406 7C

```

This example places the one's complement of the low-order byte of register R0 in the byte at the address specified by -4(FP).

The instruction is illustrated below:

Operand Values: Binary
\begin{tabular}{ccc} 
Operands & Before & After \\
\hline R0 & 10101010 & 10101010 \\
(low-order) & & \\
\(-4(\) FP ) & 00000000 & 01010101
\end{tabular}

\section*{Convert to Bit Pointer}

\section*{Syntax:}

\section*{CVTP offset, base, dest \\ reg gen addr write.D}


The CVTP instruction places in the dest operand location the absolute bit address of the memory bit specified by base and offset. See Section 3.5 for the use of a base and offset in specifying a bit position.

The bit address specifies the number of bits from the first bit in the memory space (bit 0 of the byte at address 0) to the specified bit. The bit address is computed as
\[
8 \text { * EA (base) + offset }
\]
where EA(base) is the effective address calculated for base, and offset is a signed byte, word or double-word as given by the operation length.

Flags Affected: None.

Traps: None.

\section*{Example:}

CVTP R0, 32(SB), R2 6E 83 D0 20
This example computes the absolute bit address of the memory bit specified by register R 0 and the address \(32(\mathrm{SB})\). The instruction places the resulting bit address into register R2 as a double-word.

Operand Values: Hex
\begin{tabular}{ccc} 
Operands & Before & After \\
\hline R0 & 00001234 & 00001234 \\
SB \\
\begin{tabular}{c} 
base \\
address \\
\(32(\mathrm{SB})\)
\end{tabular} & 00000 FE 0 & 00000 FE 0 \\
R2 & 00001000 & ------- \\
& AAAAAAAA & 00009234
\end{tabular}

\section*{CXP}

Call External Procedure

\section*{Syntax: CXP index}
disp
\begin{tabular}{|c|c|}
\hline ! & CXP \\
\hline \multicolumn{2}{|l|}{\(\begin{array}{llllllll}0 & 0 & 1 & 0 & 0 & 0 & 1 & 0\end{array}\)} \\
\hline -+ & -+-+-+-+-! \\
\hline 7 & 0 \\
\hline
\end{tabular}

The CXP instruction calls a procedure which is outside the current module (an "external" procedure).

The entry point of the external procedure is specified by an external procedure descriptor, which is located in the Link Table (Section 2.8.3) of the current module. The index operand gives the Link Table entry number of the descriptor.

The descriptor is a 32-bit value in the following format:


The descriptor address is the sum of index, multiplied by four, and the contents of the double-word at memory address MOD+4 (the Link Base pointer, Section 2.8.2), where MOD is the contents of the MOD register.

Once the descriptor has been located, the instruction does the following:
1. Decrements the current stack pointer by two, then pushes the contents of the MOD register (16 bits) onto the currently-selected stack. The stack pointer is modified by a total of four in this step. The extra two bytes placed on the stack are reserved for future use.
2. Saves the address of the next sequential instruction (32 bits) onto the currently-selected stack. This double-word is the return address.
3. Copies the low-order word of the descriptor to the MOD register. The low-order word is the address of the new Module Table entry.
4. Copies the double-word at address MOD+0 to the \(S B\) register. This double-word is the Static Base pointer for the new module.

\section*{Call External Procedure (continued)}
5. Copies to the PC register the sum of the high-order word of the descriptor (interpreted as an unsigned value) and the double-word at address MOD+8. This sum is the address of the external procedure entry point in the new module.

Program execution continues at the address placed in the PC register. The procedure has been invoked, and is running in its own module environment.

In the machine instruction, index is encoded as a displacement field appended to the basic instruction. In assembly language, index is specified as the name of the external procedure or in the form of an External addressing mode expression.

Flags Affected: None.

\section*{Traps: None.}

\section*{Examples:}
1. CXP OUTSIDE
2200
2. CXP EXT (1)
2201

Example 1 calls the external procedure named OUTSIDE.
Example 2 calls the external procedure whose descriptor is located in the second entry of the current Link Table (entry number 1).

\section*{CXP}

\section*{Call External Procedure (continued)}

The action of the instruction in Example 2 is illustrated below:

* 90A4 is the descriptor address. It is computed as the sum of the index 1 in the expression EXT(1), scaled by 4, and the Link Table address. The Link Table address is from address 14 (Hex) in the Module Table.
** The 16-bit field shown as "xxxx" is reserved for future use, and should be treated as don't-care bits.

\section*{Call External Procedure with Descriptor}

\section*{Syntax: CXPD desc}
gen
addr


The CXPD instruction calls the external procedure specified by the desc (descriptor) operand. The descriptor is a 32-bit value in the following format:


The instruction does the following:
1. Decrements the current stack pointer by two, then pushes the contents of the MOD register (16 bits) onto the currently-selected stack. The stack pointer is modified by a total of four in this step. The extra two bytes placed on the stack are reserved for future use.
2. Saves the address of the next sequential instruction (32 bits) onto the currently-selected stack. This double-word is the return address.
3. Copies the low-order word of the descriptor to the MOD register. The low-order word is the address of the new Module Table entry.
4. Copies the double-word at address MOD+0 to the SB register. This double-word is the Static Base pointer for the new module.
5. Copies to the \(P C\) register the sum of the high-order word of the descriptor (interpreted as an unsigned value) and the double-word at address MOD+8. This sum is the address of the external procedure entry point in the new module.

Program execution continues at the address placed in the PC register. The procedure has been invoked, and is running in its own module environment.

Flags Affected: None.
Traps: None.

\section*{CXPD}

\section*{Call External Procedure with Descriptor (continued)}

\section*{Example:}

CXPD 8 (SB)
7F D0 08

This example calls an external procedure whose descriptor is contained at memory address \(8(\mathrm{SB})\).

The instruction is illustrated below:

Operand Values: Hex
\begin{tabular}{|c|c|c|}
\hline Operands & Before & After \\
\hline \multirow[t]{3}{*}{\[
\begin{aligned}
& 00009088 \text { * } \\
& \text { (descriptor) }
\end{aligned}
\]} & 00100020 & \multirow[t]{3}{*}{00100020} \\
\hline & (module 0020) & \\
\hline & (offset 0010) & \\
\hline MOD & 0010 & 0020 \\
\hline PC & 00009005 & 0000F010 \\
\hline SB & 00009080 & 0000F100 \\
\hline SP & 0000 FFF 8 & 0000FFFO \\
\hline Stack: & & \\
\hline 0000 FFF 0 & xxxxxxxx & 00009007 \\
\hline 0000 FFF 4 & xxxxxxxx & xxxx0010 \\
\hline 0000 FFF 8 & AAAAAAAA & AAAAAAAA \\
\hline
\end{tabular}

Module Table:
\begin{tabular}{rrr}
00000010 & 00009080 & (SB) \\
14 & 000090 O & (LB) \\
18 & 00009000 & \((\mathrm{~PB})\) \\
1 C & xxxxxxxx & \\
00000020 & 0000 F 100 & (SB) \\
24 & 0000 F 110 & (LB) \\
28 & 0000 F 000 & (PB) \\
2 C & xxxxxxxx &
\end{tabular}
* 9088 (Hex) is the descriptor's effective address, as specified by \(8(S B)\).
** The 16-bit field shown as "xxxx" is reserved for future use, and should be treated as don't-care bits.

Syntax: DEIi src, dest DEIB
gen gen DEIW
read.i rmw.2i DEID


The DEIi instruction divides the entire dest operand by the src operand and places the quotient and the remainder in the dest operand location.

The instruction places the quotient in the high-order half of dest and the remainder in the low-order half. The dest operand may be specified as an even-odd General Purpose register pair. In such cases, the instruction places the remainder in the even register and the quotient in the next consecutive (odd) register. The register pair must be specified in assembly language by the name of the even register of the pair.

The src and dest operands are interpreted as unsigned integers.
Flags Affected: None.
Traps: DVZ (Divide by Zero) activated if src equals zero.
Integer Overflow Trap (OVF) is activated if the \(V\) flag is set and the result cannot be represented exactly in dest.

\section*{Example:}
DEIW R2, R0

CE 2D 10
This example divides the double-word value contained in the low-order words of R0 and R1 by the low-order word of register R2. The result is a double-word containing a quotient and a remainder. The remainder is written to the low-order word of register R0! the quotient is written to the low-order word of register R1. The high-order words of registers R0 and R1 are not used or affected.

The instruction is illustrated below:
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex (Dec) \\
After
\end{tabular} \\
\hline R2 & & \\
& AAAA0001 \\
R0 & & AAAA0001 \\
R1 & BBBBFFFF & \((+1)\) \\
& CCCC0000 \\
\((+65535)\) & BBBB0000 \\
& & CCCCFFFF \\
\((+65535\), rem. 0)
\end{tabular}

The above case divides 65535 by 1 (H'0000FFFF by H'0001). The quotient is 65535 (in R1), and the remainder is 0 (in RO).

DIA
Diagnose

\section*{Syntax: DIA}
\begin{tabular}{|c|c|}
\hline ! & DIA \\
\hline & 00010 ! \\
\hline & -+-+-+-+-! \\
\hline 7 & 0 \\
\hline
\end{tabular}

The DIA instruction is intended to support breakpointing circuitry, and is not intended for use in a program. It is a l-byte instruction which performs a branch to itself, establishing an "infinite loop" which is interruptible. When the loop thus established is interrupted, the return address pushed onto the Interrupt Stack is the address of the DIA instruction itself.

Flag Affected: None.

Traps: None.

\section*{Example:}

DIA
C2
```

                gen gen
                                    DIVL
                                read.f rmw.f
    ```


The DIVf instruction divides the dest operand by the src operand and places the result in the dest operand location.

Results for normalized and zero operands are given in the table below. The symbols "+n" and "-n" represent non-zero normalized numbers, positive and negative, respectively. The symbols "+z" and "-z" represent positive and negative zero, respectively.
\begin{tabular}{|c|c|c|c|c|c|}
\hline dest: & & n & -n & + Z & - Z \\
\hline \multicolumn{6}{|l|}{SrC} \\
\hline & ! & & & & \\
\hline \multirow[t]{2}{*}{+n} & ! & * & * & + z & - Z \\
\hline & ! & & & & \\
\hline -n & ! & * & * & - Z & + Z \\
\hline
\end{tabular}
* The result in these cases is the quotient of the two operands.

Flags Affected: No PSR flags. FSR flags are affected as follows:
UF is set if an underflow occurs; unaffected otherwise.
IF is set on an inexact result; unaffected otherwise.
\(T T\) field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes.
See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps:
Undefined Instruction Trap (UND) is activated if the \(F\) bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3. Particularly relevant to floating-point division are the Divide by Zero exception, caused by attempting to divide a non-zero number by zero, and the Invalid Operation exception, caused by attempting to divide zero by zero.

\section*{DIVf}

Divide Floating (continued)

\section*{Examples:}
1. DIVF F0, F7
2. DIVL \(-8(F P), 16(S B)\)

BE E1 01
BE A0 C6 7810

Example 1 divides the single-precision number in register \(F 7\) by the number in register FO and places the result in F 7 .

Example 2 divides the double-precision number at address \(16(\mathrm{SB})\) by the number at address \(-8(\mathrm{FP})\) and places the result at address 16(SB).

The instructions are illustrated below:
\begin{tabular}{|c|c|c|c|}
\hline & Operands & Operand V Before & Hex (Dec) After \\
\hline Ex. 1: & F0 & \[
\begin{aligned}
& 42250000 \\
& (+41.25)
\end{aligned}
\] & \[
\begin{aligned}
& 42250000 \\
& (+41.25)
\end{aligned}
\] \\
\hline & F7 & \[
\begin{aligned}
& 434 \mathrm{E} 4000 \\
& (+206.25)
\end{aligned}
\] & \[
\begin{array}{r}
40 \mathrm{~A} 00000 \\
(+5.0)
\end{array}
\] \\
\hline Ex. 2: & -8(FP) & \[
\begin{gathered}
409 \mathrm{~F} 440000000000 \\
(+2001.0)
\end{gathered}
\] & \[
\begin{gathered}
409 \mathrm{~F} 440000000000 \\
(+2001.0)
\end{gathered}
\] \\
\hline & 16 (SB) & \[
\begin{aligned}
& \text { 41A2B128DDC00000 } \\
& (+156800110.875)
\end{aligned}
\] & \[
\begin{gathered}
40 \mathrm{~F} 3218 \mathrm{E} 00000000 \\
(+78360.875)
\end{gathered}
\] \\
\hline
\end{tabular}
```

Syntax: DIVi src, dest DIVB

```
            gen gen DIVW
                            read.i rmw.i DIVD
\begin{tabular}{|c|c|c|c|}
\hline ! src & ! dest & ! & DIVi \\
\hline ! gen & gen & & \\
\hline 23 & 161 & & \\
\hline
\end{tabular}

The DIVi instruction divides the dest operand by the src operand, rounds the quotient to the next lower (or more negative) integer, and places the result in the dest operand location. The src and dest operands are interpreted as signed integers.

Flags Affected: None.

Traps: DVZ (Divide by Zero) activated if src equals zero.
Integer Overflow Trap (OVF) is activated if the \(V\) flag is set. It occurs only if the largest negative integer in a data format is divided by -1.

\section*{Examples:}
\begin{tabular}{llllll} 
1. DIVW & \(10(\mathrm{SP}), 4(\mathrm{SP})\) & CE 7D CE OA 04 \\
2. & DIVD & \(-6(\mathrm{FP})\), & \(12(\mathrm{SB})\) & CE BF C6 7A & OC
\end{tabular}

Example 1 divides the word at the address specified by 4 (SP) by the word at the address specified by \(10(S P)\). The instruction rounds the quotient and places the result at 4(SP).

Example 2 divides the double-word at the address specified by \(12(S B)\) by the double-word at the address specified by \(-6(F P)\). The instruction rounds the quotient and places the result in the double-word at address 12 (SB).

These instructions are illustrated below:
\begin{tabular}{|c|c|c|c|}
\hline & Operands & Operand Values: Before & \[
\begin{array}{r}
\text { Hex (Dec) } \\
\text { After } \\
\hline
\end{array}
\] \\
\hline \multirow[t]{2}{*}{Ex. 1:} & 10 (SP) & \[
\begin{aligned}
& 000 \mathrm{~A} \\
& (+10)
\end{aligned}
\] & \[
\begin{aligned}
& 000 \mathrm{~A} \\
& (+10)
\end{aligned}
\] \\
\hline & 4 (SP) & \[
\begin{gathered}
006 \mathrm{~F} \\
(+111)
\end{gathered}
\] & \[
\begin{aligned}
& 000 \mathrm{~B} \\
& (+11)
\end{aligned}
\] \\
\hline \multirow[t]{2}{*}{Ex. 2:} & -6 (FP) & \[
\begin{array}{r}
00000014 \\
(+20)
\end{array}
\] & \[
\begin{array}{r}
00000014 \\
(+20)
\end{array}
\] \\
\hline & 12 (SB) & FFFFFF9F
\[
(-97)
\] & FFFFFFFB (-5) \\
\hline
\end{tabular}

In example 1, 111 divided by 10 is 11.1. The next lower integer is 11. In example 2, -97 divided by 20 is -4.85 . The next lower integer is -5.

\section*{DOTf}

Dot Product Floating

Syntax: DOTf src1, src2 DOTF
gen gen DOTL
read.f read.f


The DOTf instruction multiplies src1 and src2 and than adds the result to the floating-point register F0. ( FO := FO + (src1 * src2) )

Flags Affected: No PSR flags.
The FSR TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes. See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap(UND) is activated if the \(F\) bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3.

\section*{Example:}

\section*{DOTF F2, F3}

FE CC 10

This example multiplies the single-precision numbers in \(F 2\) and \(F 3\) and adds the result to register FO.

The instruction is illustrated below:
\begin{tabular}{ccr} 
Operands & \begin{tabular}{c} 
Operand Values: Hex \\
Before
\end{tabular} & \begin{tabular}{c} 
(Dec) \\
After
\end{tabular} \\
\hline & & \\
F0 & C2250000 \\
& \((-41.25)\) & 42 E 88000 \\
F2 & 418 C 0000 & \((+116.25)\) \\
& \((+17.50)\) & 418 C 0000 \\
& & \((+17.50)\) \\
F3 & 41100000 & \((+9.00)\)
\end{tabular}

\section*{ENTER}

Enter New Procedure Context

\section*{Syntax: ENTER reglist, constant}
imm disp
\begin{tabular}{|c|c|}
\hline ! & ENTER \\
\hline & 00010 ! \\
\hline & -+-+-+-+-+ \\
\hline 7 & \\
\hline
\end{tabular}

The ENTER instruction creates a "Frame" on the current stack for use by a procedure. A Frame is a block of memory on the stack that provides local storage for the current procedure. The constant operand specifies the number of bytes to be reserved on the stack for local data storage. The Frame Pointer (FP) register is saved and then set up as a pointer from which frame information can be located.

The instruction does the following:
1. Pushes the contents of the \(F P\) register (32 bits) onto the stack.
2. Copies the contents of the current stack pointer to the \(F P\) register.
3. Subtracts the constant operand from the value of the current stack pointer, lengthening the stack by that number of bytes.
4. Pushes the General-Purpose registers specified by reglist onto the stack.

The reglist operand is specified in assembly language by a list of zero or more General-Purpose register names, enclosed in brackets "[]". The instruction pushes the contents of each register in the list as a double-word onto the currently-selected stack. Register names may appear in any order within reglist but must be separated by commas. Brackets are required even if no register names are given.

In the machine instruction, the reglist operand is encoded in an 8-bit field as shown. Each bit in the field corresponds to one general-purpose register. When the instruction is executed, the instruction reads the bits in the field from right to left beginning with bit 0. If a bit is "0", the instruction ignores the corresponding register. If a bit is "1", it saves the corresponding register.
\[
\begin{aligned}
& +--+--+--+--+--+--+--+--+ \\
& \text { !R7!R6!R5!R4!R3!R2!R1!R0!} \\
& \text { !--+--+--+--+--+--+--+--! } \\
& 7
\end{aligned}
\]

Flags Affected: None.

Traps:
None.

\section*{Enter New Procedure Context (continued)}

\section*{Example:}
```

ENTER [R0, R2, R7], 16
82 85 10

```

This instruction creates a frame on the stack consisting of 16 bytes for local data storage and the contents of register R0, R2, and R7.

In the following illustration, the PSR S flag is assumed to be 1, selecting SP1 as the current stack pointer.

Operand Values: Hex
\begin{tabular}{lrr} 
Operands & Before & After \\
\hline R0 & 00000010 & 00000010 \\
R2 & FFFFFFEF & FFFFFFEF \\
R7 & FFFFF9AB & FFFFF9AB \\
16 & 10 & \\
\((\) disp \()\) & \((+16)\) & -- \\
FP & \(000010 F 8\) & 000010 EC \\
SP1 & \(000010 F 0\) & 000010 D0
\end{tabular}

Stack:
\begin{tabular}{|c|c|c|}
\hline 000010D0 & mxxxxxxx & FFFFF9AB \\
\hline 000010D4 & xxxxxxxx & FFFFFFEF \\
\hline 000010D8 &  & 00000010 \\
\hline 000010DC & xxyxxxyx & xxxxxxxx \\
\hline 000010 E 0 & mxxexxyx & mxxexxex \\
\hline 000010E4 & xxxxxxxx & mxxxxxxx \\
\hline 000010E8 & mxxxxxix & xxxxxxxx \\
\hline 000010 EC & xxxxxxxx & 000010F8 \\
\hline 000010F0 & AAAAAAAA & AAAAAAAA \\
\hline
\end{tabular}
* 16 bytes of uninitialized local data storage.

\section*{EXIT}

\section*{Exit Procedure Context}

\section*{Syntax: EXIT reglist}
imm
\begin{tabular}{|c|c|}
\hline ! & EXIT \\
\hline & 10010 ! \\
\hline & -+-+-+-+ \\
\hline 7 & \\
\hline
\end{tabular}

The EXIT instruction removes the frame of the current procedure from the stack, restores the former contents of the specified General-Purpose registers (i.e., their contents prior to entering the current procedure), and restores the frame of the previous procedure as the current procedure context.

The instruction does the following:
1. Restores the General-Purpose registers specified by reglist by popping them from the current stack.
2. Copies the contents of the \(F P\) register to the current stack pointer.
3. Pops the old frame address (32 bits) from the stack to the FP register.

In assembly language, the reglist operand is specified as a list of zero or more General-Purpose register names, enclosed in brackets "[]". The instruction copies to each register in the list a double-word popped from the stack. Register names may appear in any order within reglist but must be separated by commas. Brackets are required even if no register names are given.

In the machine instruction, the reglist operand is encoded in an eight-bit field as shown below. Each bit in the field corresponds to one general-purpose register. When the instruction is executed, the instruction reads the bits in the field from right to left beginning with bit 0. If a bit is "0", the instruction ignores the corresponding register. If a bit is "1", it restores the corresponding register from the stack. Note that the format of the reglist operand is reversed from its format in the ENTER instruction; i.e. bit 0 corresponds to register \(R 7\) instead of \(R 0\).
```

+--+--+--+--+--+--+--+--+
!R0!R1!R2!R3!R4!R5!R6!R7!
!--+--+--+--+--+--+--+--!
70

```

Flags Affected: None.
Traps: None.

\section*{Exit Procedure Context (continued)}

\section*{Example:}
EXIT [R0, R2, R7]
92 A1

This instruction restores the contents of the listed General-Purpose registers, reclaims the frame of the current procedure, and restores the frame of the previous procedure as the current context.
\begin{tabular}{lll} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex \\
After
\end{tabular} \\
\hline R0 & CCCCCCCC & 00000010 \\
R2 & CCCCCCCC & FFFFFFEF \\
R7 & CCCCCCCC & FFFFF9AB \\
FP & 000010 EC & 00001000 \\
SP & 000010 D0 & 000010 F0
\end{tabular}

Stack:
\begin{tabular}{|c|c|c|}
\hline 000010D0 & FFFFF9AB & xxxxxxxx \\
\hline 000010D4 & FFFFFFEF & mxxxxxxx \\
\hline 000010D8 & 00000010 & xxxxxxxx \\
\hline 000010DC & BBBBBBBB & xxxxxxxx \\
\hline 000010E0 & BBBBBBBB & xxxxxxxx \\
\hline 000010E4 & BBBBBBBB & mxxxxxxx \\
\hline 000010 E 8 & BBBBBBBB & xxxxxxxx \\
\hline 000010EC & 00001000 & XXXXXXXX \\
\hline 000010F0 & AAAAAAAA & AAAAAAAA \\
\hline
\end{tabular}
* The EXIT instruction does not itself change the contents of these memory locations. However, information that is outside the stack should be considered unpredictable for other reasons. See Section 2.8.1.

EXTi
Extract Field
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline Syntax: & EXTi & offset, & bas & dest, & length & EXTB \\
\hline & & reg & gen & gen & disp & EXTh \\
\hline & & & regaddr & write & & EXTD \\
\hline
\end{tabular}


The EXTi instruction copies the bit field specified by base, offset and length to the dest operand location. The field is right-justified in dest. High-order bits are zero-filled if the field is shorter than dest or discarded if the field is longer than dest.

The location of the field is taken from the position of its least-significant bit, given by offset and base as follows:

If base is a register, then the field is within that register, starting at the bit position given by offset. If base is a memory location, then the field starts at bit position
offset MOD 8
within the memory byte whose address is
EA (base) + (offset DIV 8),
where EA(base) is the effective address of base. See Section 3.6 for definitions of the operators MOD and DIV above.

Offset is interpreted as a 32-bit signed integer.
Length specifies the number of bits in the field. It must be in the range 1 through 32.

See Section 3.6 for further details of specifying bit fields.
NOTE: Although a bit field may contain up to 32 bits, an alignment restriction appears for fields containing more than 25 bits: a field may not span more than four bytes. See Section 3.6.

Flags Affected: None.
Traps: None.

\section*{Extract Field (continued)}

\section*{Example:}
```

EXTW R0, 0(R1), R2, 7 2E 81 40 00 07

```

This example copies a 7-bit field from memory into the low-order word of register R2. Bits 7 through 15 of register \(R 2\) are set to zero and the remaining bits of R2 are unaffected. For designating the location of the field, register R0 supplies the bit offset, and 0 (R1) is specified as the base address.
\begin{tabular}{|c|c|c|}
\hline Operands & Operand Values: Hex Before & (Dec) After \\
\hline \[
\begin{gathered}
\mathrm{RO} \\
\text { (offset) }
\end{gathered}
\] & \[
\begin{array}{r}
0000004 \mathrm{C} \\
(+76)
\end{array}
\] & \[
\begin{array}{r}
0000004 \mathrm{C} \\
(+76)
\end{array}
\] \\
\hline R1 & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] \\
\hline base address 0 (R1) & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] & -- \\
\hline R2 & AAAAAAA & AAAA0071 \\
\hline \[
\begin{gathered}
00001009 \text { * } \\
(+4105)
\end{gathered}
\] & EF10 & EF10 \\
\hline
\end{tabular}
* The address 1009 (Hex) is the effective address of the byte containing the least-significant bit of the specified field. This address is computed as \(4096+(76\) DIV 8) \(=4105\), where 4096 is the base address specified by 0 (R1) and 76 is the bit offset given by the contents of register R0.
** The bit field starts at bit position 4 ( \(=76\) MOD 8) in the byte at address 1009 (Hex) and is seven bits long as illustrated.
! 7-bit field !
+---------+-------------+-------+


\begin{tabular}{llll}
\(!7\) & \(0!7\) & \(0!\)
\end{tabular}
    regaddr write.i EXTSD


The EXTSi instruction copies the bit field specified by base, offset, and length to the dest operand location. The field is right-justified in dest. High-order bits are zero-filled if the field is shorter than dest or discarded if the field is longer than dest.

The offset and length operands are encoded together as an immediate byte appended to the basic instruction. The offset is encoded as the high-order three bits of this byte! the length operand, minus one, is encoded as the low-order five bits. The byte has the following form:
```

+---------+---------------+
! offset ! length - 1 !
+--+--+--+--+---+--+---+--+
7

```

The offset value must be in the range 0 through 7. The length value specifies the number of bits in the field. It must be in the range 1 through 32 .

The location of the field is taken from the position of its least-significant bit. If base is a register, then the field is within that register, starting at the bit position given by offset. If base is a memory location, then the field starts at the bit position given by offset within the memory byte whose address is given as base.

See Section 3.6 for further details of specifying bit fields.

NOTE: Although a bit field may contain up to 32 bits, an alignment restriction appears for fields containing more than 25 bits: a field may not span more than four bytes. See Section 3.6.

Flags Affected: None.

Traps:
None.

\section*{Extract Field Short (continued)}

\section*{Example:}
```

EXTSW 16(SB), R2, 4, 7
CE 8D DO 10 86

```

This example copies a bit field to the low-order word of register R2. The field begins at bit position 4 of the byte at the address specified as \(16(\mathrm{SB})\) and is seven bits long.
\begin{tabular}{cccc} 
Operands & \multicolumn{2}{c}{ Operand Values: Hex } \\
Aefore & After \\
\cline { 1 - 3 } & AAAAAAAA & AAAA0071 ** \\
\(16(\mathrm{SB})\) & EF10 & EF10 *
\end{tabular}
* The bit field starts at bit number 4 in the byte at address 16(SB) and is seven bits long as illustrated:

** The bit field is right-justified in the low-order word of register R2. Nine leading zero bits are added to the bit field to fill the low-order word.

FFSi
Find First Set Bit

\section*{Syntax:}


The FFSi instruction searches for the first "1" bit in the base operand. The search starts at the bit specified by the offset operand and proceeds in ascending order to the first "1" bit or to the last bit in base.

If a "1" bit is found, the instruction sets the offset operand value to the bit number of the first "1" bit found in the base operand and clears the \(F\) flag in the PSR.

If a "1" bit is not found, the instruction sets the offset operand value to zero and sets the \(F\) flag in the PSR tc 1.

The offset is interpreted as an unsigned number. Its value must be within the range 0 to 7 (in FFSB instruction), 0 to 15 (in FFSW instruction), and 0 to 31 (in FFSD instruction), otherwise the result placed in the offset operand and in the \(F\) bit is undefined.

Note: If the FFSi instruction finds a "1" bit and it is desired to scan the remaining portion of the base operand, the offset operand must be incremented first or the "1" bit previously found must be cleared. Otherwise, the FFSi instruction will detect that bit again.

Flags Affected: \(F\) is set if a "1" bit is not found, cleared if found.

\section*{Traps: None.}

\section*{Examples:}
\begin{tabular}{llllll} 
1. & FFSW & \(8(\mathrm{SB})\), R0 & 6 E & 05 & D 0 \\
08 \\
2. FFSB & \(-4(\mathrm{FP})\), TOS & 6 E & C 4 & C 5 & 7 C
\end{tabular}

Example 1 searches the word at the address specified by 8(SB) for the first "1" bit. The search begins at the bit specified by the low-order byte of register R0, and the result is placed in the low-order byte of \(R 0\). The remaining portion of RO is not used or affected.

Example 2 searches the byte at the address specified by -4 (FP). The search begins at the bit specified by the byte on the top of the stack, which is replaced by the resulting bit number.

These instructions are illustrated below:
\begin{tabular}{|c|c|c|c|}
\hline & Operands & Operand Values: Before & Hex (Dec) [Binary]
\(\qquad\) \\
\hline \multirow[t]{3}{*}{Ex. 1:} & R0 & AAAAAA 05 (+5) & AAAAAA0 8 \((+8)\) \\
\hline & 8 (SB) & \[
\begin{gathered}
\text { EF10 } \\
{[1110111100010000]}
\end{gathered}
\] & \[
\begin{gathered}
\text { EF10 } \\
{[1110111100010000]}
\end{gathered}
\] \\
\hline & UPSR & \(n z f v x l t c\) & \(n z \underline{0} \mathrm{vxltc}\) \\
\hline \multirow[t]{4}{*}{Ex. 2:} & SP & 0000FFDE & 0000FFDE \\
\hline & \[
\begin{gathered}
\text { Stack: } \\
0000 \text { FFDE }
\end{gathered}
\] & \[
\begin{gathered}
05 \\
(+5)
\end{gathered}
\] & \[
\begin{aligned}
& 00 \\
& (0)
\end{aligned}
\] \\
\hline & -4 (FP) & \[
\begin{gathered}
10 \\
{[00010000]}
\end{gathered}
\] & \[
\begin{gathered}
10 \\
{[00010000]}
\end{gathered}
\] \\
\hline & UPSR & \(n z f v x l t c\) & \(n z \underline{1} \mathrm{vxltc}\) \\
\hline
\end{tabular}

In Example 1, the instruction finds the first "1" bit at bit position 8.
In Example 2, the instruction finds no "1" bits; that is, the bits at bit positions 5, 6, and 7 are all "0" bits.

\section*{FLAG}

Trap on Flag

\section*{Syntax: FLAG}
\begin{tabular}{|c|c|}
\hline ! & FLAG \\
\hline & 10010 ! \\
\hline ! - & -+-+-+-+-! \\
\hline 7 & \\
\hline
\end{tabular}

The FLAG instruction activates the Flag Trap (FLG) if the \(F\) flag in the PSR is set. The Flag Trap passes control to the Flag service procedure. The return address pushed on the Interrupt Stack is the address of the FLAG instruction instruction itself. If the \(F\) flag is not set, program execution continues with next sequential instruction.

Flags Affected: None.
Traps: The Flag Trap (FLG) is activated if the F flag is set.

\section*{Example:}

FLAG D2
\(\begin{array}{ll}\text { src, } & \text { dest } \\ \text { gen } & \text { gen } \\ \text { read.f } & \text { write.i }\end{array}\)
\begin{tabular}{ll} 
FLOORFB & FLOORLB \\
FLOORFW & FLOORLW \\
FLOORFD & FLOORLD
\end{tabular}
\(\begin{array}{ll}\text { FLOORFB } & \text { FLOORLB } \\ \text { FLOORFW } & \text { FLOORLW }\end{array}\)
FLOORFD FLOORLD


The FLOORfi instruction rounds the src operand to the nearest integer less than, or equal to, it (i.e., toward negative infinity) and places the result in the dest operand location as a signed integer.

Flags Affected: No PSR flags. FSR flags are affected as follows: IF is set on an inexact result; unaffected otherwise. TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes.
See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3. Particularly relevant to this instruction is the Overflow exception, which is caused by attempting to convert a floating-point number that is too great in absolute value to be held in a signed integer of the size specified for dest.

\section*{FLOORfi}

Floor Floating to Integer (continued)

\section*{Examples:}
1. FLOORFB F0, R0
2. FLOORLD F2, 16 (SB)

3B 3C 00
3E BB 1610
Example 1 rounds the single-precision number in register \(F 0\) to a byte-long integer and copies the integer to the low-order byte of register R0.

Example 2 rounds the double-precision number in register \(F 2\) to a double-word integer and copies the integer to address \(16(S B)\).
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex (Dec) \\
After
\end{tabular} \\
\hline F0 & C0280000 \\
\((-2.65)\) & C0280000 \\
R0 & & \((-2.65)\) \\
& AAAAAAAA & AAAAAAFD \\
\((-3)\)
\end{tabular}

Ex. 2: F2
\(41 C 0200888700000\)
\((+541069584.875\)
41C0200888700000 \((+541069584.875) \quad(+541069584.875)\)

16 (SB)
AAAAAAAA
20401110
(+541069584)

\section*{Invert Bit}

\section*{Syntax:}


The IBITi instruction inverts (complements) the register or memory bit specified by base and offset after copying the bit value to the \(F\) flag in the PSR.

The location of the bit is determined from offset and base. Offset is a general operand, whose length is given by the operation length suffix. Base is an addressing expression giving a byte address from which offset specifies a bit position. See Section 3.5 for details of specifying bit positions.

If base is a register, then the bit is within that register, at the bit position given by the offset operand. If base is a memory location, then the bit is at bit position
offset MOD 8
within the memory byte whose address is
EA(base) + (offset DIV 8),
where EA(base) is the effective address of base. See Section 3.5 for definitions of the operators MOD and DIV above, and for further details of bit instructions.

Offset is interpreted as a signed integer.
Flags Affected: \(F\) is set to the original value of the specified bit.

Traps:
None.

IBITi
Invert Bit (continued)

\section*{Example:}

IBITW R0, 1 (R1)
4E 790201

This example inverts a bit in memory after copying the bit value into the \(F\) flag. For designating the location of the target bit, the low-order word of register R0 supplies the bit offset, and \(1(R 1)\) is specified as the base address.

In the following illustration, the target bit is assumed to be 0 prior to instruction execution.
\begin{tabular}{|c|c|c|}
\hline Operands & Operand Values:
\(\qquad\) & Hex (Dec) [Binary]
After \\
\hline \[
\begin{gathered}
\mathrm{R0} \\
(\mathrm{offset})
\end{gathered}
\] & \[
\begin{aligned}
& \text { AAAAO } 04 \mathrm{C} \\
& (+76)
\end{aligned}
\] & \[
\begin{aligned}
& \text { AAAA004C } \\
& (+76)
\end{aligned}
\] \\
\hline R1 & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] & \[
\begin{aligned}
& 00001000 \\
& (+4096)
\end{aligned}
\] \\
\hline base address 1 (R1) & \[
\begin{aligned}
& 00001001 \\
& \quad(+4097)
\end{aligned}
\] & -- \\
\hline \[
\begin{gathered}
0000100 \mathrm{~A} \text { * } \\
(+4106)
\end{gathered}
\] & \[
\begin{gathered}
\text { EF } \\
{[11101111]}
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{FF} \\
{[1111111]}
\end{gathered}
\] \\
\hline UPSR & \(n z f v x l t c\) & \(n \mathrm{z}\) 0vxltc \\
\hline
\end{tabular}
* The address 100A (Hex) is the effective address of the byte containing the desired bit. This address is computed from the offset and the base address as follows:
```

base address + (offset DIV 8)
4097 + 9
4106, or 100A (Hex) .

```

The bit number within this byte is calculated as:
\[
\begin{array}{ccc}
\text { Offset } & \text { MOD } & 8 \\
76 & \text { MOD } & 8 \\
4 & . &
\end{array}
\]

\section*{Syntax:}


The INDEXi instruction assists the programmer in accessing multidimensional arrays by providing a 1-dimensional index which can subsequently be used directly in an addressing mode with Scaled Indexing. The 1-dimensional index is calculated from the values of the indices along each dimension of the array.

This instruction is intended to be executed iteratively, as discussed in Section 3.9, once for each dimension except the first. Each iteration accumulates its result into the general-purpose register specified as accum. The length operand defines the length of the current dimension, giving the difference between the upper and lower index bounds (this is the actual dimension length minus one). The index operand is the zero-adjusted value along the current dimension. The result placed in the accum register is.
```

accum * (length + 1) + index .

```

The length and index operands are interpreted as unsigned integers, and are zero-extended to 32 bits internally before use. The accum operand is interpreted as an unsigned 32-bit integer.

Flags Affected: None.

Traps: None.

\section*{Example:}
```

INDEXB R0, 20(SB), -4(FP) 2E 04 D6 14 7C

```

This example performs one step of an index calculation. R0 is the accum operand, memory location \(20(S B)\) holds a byte defining the length of the current array dimension, and memory location \(-4(F P)\) holds the index value along this dimension.

The case below shows the application of the above instruction to calculate the 1-dimensional index of array element A[I,J], where A has been declared (in the Pascal language) as being of dimensions [1..7, 0..16]. The array is assumed to be stored in row major order (Section 3.9). Since it is an array of only two dimensions, one INDEXi instruction serves to calculate the one-dimensional index.

The value of index I (assumed to be 4) has been zero-adjusted to 3 by a CHECK instruction (q.v.), and the result placed in register R0 as a double-word. The value of index J, held in one byte at address \(-4(F P)\), is assumed to be 3. The byte at location \(20(S B)\) holds the length operand for the second dimension of the array (16-0 - 16).

The result in \(R 0\), 54, is the final 1-dimensional index of element [4,3] of array A. This value can be used directly in any addressing mode with a Scaled Indexing modifier to access this array element.
\begin{tabular}{|c|c|c|}
\hline Operands & Operand Values: Before & \[
\begin{aligned}
& \text { Hex }(\text { Dec) } \\
& \text { After } \\
& \hline
\end{aligned}
\] \\
\hline R0 & \[
\begin{array}{r}
00000003 \\
(+3)
\end{array}
\] & \[
\begin{array}{r}
00000036 \\
(+54)
\end{array}
\] \\
\hline 20 (SB) & \[
\begin{gathered}
10 \\
(+16)
\end{gathered}
\] & \[
\begin{gathered}
10 \\
(+16)
\end{gathered}
\] \\
\hline -4 (FP) & \[
\begin{gathered}
03 \\
(+3)
\end{gathered}
\] & \[
\begin{gathered}
03 \\
(+3)
\end{gathered}
\] \\
\hline
\end{tabular}

\section*{Insert Field}
\begin{tabular}{lllll} 
Syntax: INSi \begin{tabular}{lll} 
offset, \\
reg
\end{tabular} & \begin{tabular}{l} 
src, \\
gen \\
read.i
\end{tabular} gen regaddr
\end{tabular}\(\quad\)\begin{tabular}{ll} 
length & INSB \\
&
\end{tabular}


The INSi instruction inserts the src operand into the bit field specified by base, offset, and length. The src operand is right-justified in the field. High-order bits are zero-filled if src is shorter than the field or discarded if src is longer than the field.

The location of the field is taken as the position of its least-significant bit, given by offset and base as follows:

If base is a register, then the field is within that register, starting at the bit position given by offset. If base is a memory location, then the field starts at bit position
offset MOD 8
within the memory byte whose address is
EA(base) + (offset DIV 8),
where EA(base) is the effective address of base. See Section 3.6 for definitions of the operators MOD and DIV above.

Offset is interpreted as a 32-bit signed integer.
Length specifies the number of bits in the field. It must be in the range 1 through 32.

See Section 3.6 for further details of specifying bit fields.
NOTE: Although a bit field may contain up to 32 bits, an alignment restriction appears for fields containing more than 25 bits: a field may not span more than four bytes. See Section 3.6.

Flags Affected: None.
Traps: None.

INSi
Insert Field (continued)

\section*{Example:}
INSW R0, R2, \(0(R 1), 7\)
AE \(41 \quad 1200 \quad 07\)

This example inserts seven bits from the low-order word of register R2 into a bit field in memory. For specifying the location of the field, register R0 supplies the bit offset, and \(0(R 1)\) is specified as the base address.

The instruction is illustrated below:
\begin{tabular}{|c|c|c|}
\hline Operands & \begin{tabular}{l}
Operand Values: \\
Before
\end{tabular} & Hex After \\
\hline \[
\begin{gathered}
\mathrm{R} 0 \\
(\mathrm{offset})
\end{gathered}
\] & \[
\begin{array}{r}
0000004 \mathrm{C} \\
(+76)
\end{array}
\] & \[
\begin{array}{r}
0000004 \mathrm{C} \\
(+76)
\end{array}
\] \\
\hline R1 & \[
\begin{aligned}
& 00001000 \\
& (+4096)
\end{aligned}
\] & \[
\begin{aligned}
& 00001000 \\
& (+4096)
\end{aligned}
\] \\
\hline base address 0 (R1) & \[
\begin{aligned}
& 00001000 \\
& (+4096)
\end{aligned}
\] & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] \\
\hline R2 & AAAAAA 67 & AAAAAA 67 \\
\hline \[
\begin{gathered}
00001009 \text { * } \\
(+4105)
\end{gathered}
\] & BBBB & BE7B \\
\hline
\end{tabular}
* The address 1009 (Hex) is the effective address of the byte containing the least-significant bit of the specified field. This address is computed as \(4096+(76\) DIV 8) \(=4105\), where 4096 is the address specified by \(0(R 1)\) and 76 is the bit offset given by the contents of register R0.
** The bit field starts at bit position 4 (= 76 MOD 8) in the byte at address 1009 (Hex) and is seven bits long as illustrated:


\section*{Insert Field Short}

\section*{Syntax:}


The INSSi instruction inserts the src operand into the bit field specified by base, offset, and length. The src operand is right-justified in the field. High-order bits are zero-filled if src is shorter than the field or discarded if src is longer than the field.

The offset and length operands are encoded together as an immediate byte appended to the basic instruction. The offset is encoded as the high-order three bits of this byte! the length operand, minus one, is encoded as the low-order five bits. The byte has the following form:
```

+---------+----------------
! offset ! length - 1 !
+--+--+--+--+--+--+--+--+
$\begin{array}{llllllll}7 & 6 & 5 & 4 & 3 & 2 & 1 & 0\end{array}$

```

The offset value must be in the range 0 through 7. The length value specifies the number of bits in the field. It must be in the range 1 through 32 .

The location of the field is taken from the position of its least-significant bit. If base is a register, then the field is within that register, starting at the bit position given by offset. If base is a memory location, then the field starts at the bit position given by offset within the memory byte whose address is given as base.

See Section 3.6 for further details of specifying bit fields.
NOTE: Although a bit field may contain up to 32 bits, an alignment restriction appears for fields containing more than 25 bits: a field may not span more than four bytes. See Section 3.6.

Flags Affected: None.
Traps: None.

\section*{INSSi}

Insert Field Short (continued)

\section*{Example:}
```

INSSW R2, 16(SB), 4, 7
CE 89 16 10 86

```

This example inserts seven bits from the low-order word of register R2 into a bit field in memory. The bit field begins at bit position 4 in the byte at the address specified by \(16(\mathrm{SB})\) and is seven bits long.

The instruction is illustrated below:
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex \\
After
\end{tabular} \\
\hline R2 & AAAAAA67 & AAAAAA67 \\
\(16(\mathrm{SB})\) & BBBB & BE7B *
\end{tabular}
* The bit field starts at bit number 4 in the byte at address \(16(S B)\) and is seven bits long as illustrated:


\section*{Jump to Subroutine}

\section*{Syntax: JSR dest}
gen
addr


The JSR instruction jumps to the procedure at the address specified by dest after saving the return address on the stack. The return address is the address of the next sequential instruction.

Flags Affected: None.

Traps: None.

\section*{Example:}

JSR 0(4(SB)) 7F 960400

This example causes the program to jump to a procedure at the address held within a double-word at address 4(SB). This is accomplished via the Static Memory Relative addressing mode. The instruction saves the address of the next sequential instruction on the stack.

The instruction is illustrated below:
\begin{tabular}{clc} 
Operand & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex \\
After
\end{tabular} \\
\hline PC & 00009000 & 00001 FFF \\
\(4(\mathrm{SB})\) & 00001 FFF & 00001 FFF \\
SP & 0000 FFD 4 & \(0000 \mathrm{FFD0}\) \\
Stack: & & \\
0000 FFD & & 00009004 \\
0000 FFD4 & AXXXXXXX & AAAAAAAAA
\end{tabular}

JUMP
Jump

\section*{Syntax: JUMP dest}
gen
addr


The JUMP instruction jumps to the address specified by dest by loading the effective address of dest into the PC register.

Flags Affected: None.
Traps: None.

\section*{Example:}
```

JUMP 0(-8(FP))
7F 82 78 00

```

This example loads the address held in the double-word at address -8(FP) into the PC register. This is accomplished via the Frame Memory Relative addressing mode. Program execution continues at that address.

The instruction is illustrated below:
\begin{tabular}{clc} 
Operand & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex \\
After
\end{tabular} \\
\hline\(-8(\) FP \()\) & 00001004 & 00001004 \\
PC & 0000909 A & 00001004
\end{tabular}

\section*{Load Floating-Point Status Register (FSR)}

\section*{Syntax: LFSR src}
gen
read.D


The LFSR instruction copies the double-word specified by src to the Floating Point Status register (FSR). See Section 2.4.2 for the format of the FSR.

Flags Affected: No PSR flags. All FSR flags are affected. All implemented FSR fields are loaded from the src operand.

Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

\section*{Example:}

LFSR R0
3E OF 00

This example copies the contents of register RO into the FSR.

Operand Values: Hex
\begin{tabular}{ccc} 
Operands & Before & After \\
\hline R0 & 00000028 & 00000028 \\
FSR & xxx00129 & xxx00028
\end{tabular}

\section*{LMR}

\section*{Load Memory Management Register}

\section*{Syntax:}

\section*{LMR mmureg, src}

\section*{short gen}
read.D


The LMR instruction copies the src operand to the Memory Management register specified by mmureg.

The LMR instruction may load the following registers. The short field of the basic instruction holds a 4-bit value which addresses the corresponding Memory Management register as shown below.
\begin{tabular}{lll} 
Register & mmureg & \begin{tabular}{l} 
short \\
field
\end{tabular} \\
Memory Management Control Register & MCR & 1000 \\
Memory Management Status Register & MSR & 1010 \\
Translation Exception Address Reg. & TEAR & 1011 \\
Page Table Base Register 0 & PTB0 & 1100 \\
Page Table Base Register 1 & PTB1 & 1101 \\
Invalidate Virtual Address 0 & IVAR0 & 1110 \\
Invalidate Virtual Address 1 & IVAR1 & 1111
\end{tabular}

Flags Affected: None.

Traps: Undefined Instruction Trap (UND) is activated if the M bit in the CFG register is clear. The instruction is not executed.

Illegal Instruction Trap (ILL) is activated if the \(U\) flag is set. The instruction is not executed.

\section*{Load Memory Management Register (continued)}

Example:
LMR PTB0, R0 1E OB 06

This example copies the contents of register R0 to the Page Table Register 0 .

Operand Values: Hex
\begin{tabular}{ccc} 
operands & Before & After \\
\hline R0 & 00009000 & 00009000 \\
PTB0 & AAAAAAAA & 00009000
\end{tabular}

\section*{LOGBf}

Logarithm Binary Floating


\section*{Example:}
```

LOGBF F3, F2
FE 94 18

```

This example computes the unbiased exponent from the number in the floating-point register F 3 and places the result in floating-point register F 2.

The instruction is illustrated below:
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: Hex \\
Before
\end{tabular} & \begin{tabular}{c} 
(Dec) \\
After
\end{tabular} \\
\hline & & \\
F2 & C2250000 & 40800000 \\
& \((-41.25)\) & \((+4.00)\) \\
F3 & & \\
& \(418 c 0000\) \\
\((+17.50)\) & 418 c 0000 \\
& & \((+17.50)\)
\end{tabular}

\section*{LPRi}

Load Processor Register

\section*{Syntax: LPRi procreg, src LPRB}
short gen LPRW
read.i


The LPRi instruction copies the src operand to the dedicated register specified by procreg. See Section 2.2 for the formats of these registers.

The src operand value is right-justified in the register. In registers other than PSR, the high-order bits are zero-filled if src is shorter than the register. The "LPRB PSR" form loads only the low-order byte of the PSR, but is privileged (see note 3 below). High-order src bits are discarded if src is longer than the register.

The Load Processor Register instruction may load the following registers. The specified procreg corresponds to the 4-bit short field in the basic instruction as shown below.
Register
User PSR
Debug Condition Register
Breakpoint Program Counter
Debug Status Register
Compare Address Register
Frame Pointer
Stack Pointer
Static Base Register
User Stack Pointer (SP1)
Configuration Register
Processor Status Register
Interrupt Base Register
Module Register
\begin{tabular}{lll} 
procreg & short & \\
\cline { 1 - 2 } & SPSR & 0000 \\
& & (Note 3) \\
DCR & 0001 & (Note 4) \\
BPC & 0010 & \\
(Note 4) \\
DSR & 0011 & (Note 4) \\
CAR & 0100 & (Note 4) \\
FP & 1000 & \\
SP & 1001 & (Note 1) \\
SB & 1010 & (Note 2) \\
USP & 1011 & (Note 4) \\
CFG & 1100 & (Note 4) \\
PSR & 1101 & (Note 4) \\
INTBASE & 1110 & (Note 4) \\
MOD & 1111 &
\end{tabular}

NOTES: 1. If \(S P\) is specified in the instruction and the \(S\) flag in the \(P S R\) is set, the instruction copies the src operand to the SP1 register. If the \(S\) flag is clear, the instruction copies the operand to the SPO register.
2. The SB register should not be loaded using LPRi except during system initialization after a Reset, because it is automatically reloaded from the current Module Table entry whenever an external procedure returns, or a trap or interrupt service procedure returns.
3. Specifying UPSR as the procreg operand causes only the low-order byte of the PSR to be affected, regardless of the operation length.
4. Specifying this register as the procreg operand is always a privileged operation, regardless of the operation length. See "Traps" below.

\section*{Load Processor Register (continued)}

Flags Affected: All PSR flags are affected if PSR is specified with operation length of \(W\) or \(D\). The \(N, ~ Z, ~ F, ~ L, ~ T ~ a n d ~ C ~ f l a g s ~ a r e ~ a f f e c t e d ~\) when UPSR is selected, or PSR is selected with operation length B. No flags are affected otherwise.

Traps: Illegal Operation Trap (ILL) is activated if the U flag is set and PSR, INTBASE, USP, CFG or a Debug Register is specified.

\section*{Examples:}


\section*{LSHi}

Logical Shift
\begin{tabular}{|c|c|c|c|}
\hline \multirow[t]{2}{*}{Syntax:} & \multirow[t]{3}{*}{LSHi} & count, & dest \\
\hline & & gen & gen \\
\hline & & read.B & \\
\hline
\end{tabular}


The LSHi instruction performs a logical shift on the dest operand in the manner specified by the count operand. The sign of count determines the direction of the shift. The absolute value of count gives the number of bit positions to shift the dest operand.

The count operand value must be within the range -7 to +7 for the LSHB form, -15 to +15 for the LSHW form, and -31 to +31 for the LSHD form. A positive count specifies a left shift; a negative count specifies a right shift. In a logical shift, all bits shifted out of dest are lost, and all bit positions emptied by the shift are zero-filled.

The count operand is interpreted as a signed integer. The dest operand is interpreted as an unsigned integer.

Flags Affected: None.

Traps: None.

\section*{Examples:}
1. LSHB 4, 8 (SB)
4E 94 A6 0408
2. LSHB \(-4(F P), 8(S B)\)
4 E 94 C6 7C 08

Example 1 shifts the 1-byte operand at address \(8(S B)\) four bit positions to the left.

Example 2 shifts the operand at address \(8(S B)\) according to the count given by the byte at address \(-4(F P)\). This value, -1, causes a 1-bit logical right shift.


Multiply Extended Integer

\section*{Syntax:}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline MEI & \[
\begin{aligned}
& \mathbf{s r} \\
& \text { ge } \\
& \text { re }
\end{aligned}
\] & & dest gen rmw. & & & & \\
\hline ! & src & ! & dest & ! & MEIi & & ! \\
\hline & gen & & gen & \(!1\) & \(!11\) & & \\
\hline 23 & & & 161 & & 87 & & 0 \\
\hline
\end{tabular}

The MEIi instruction multiplies the src operand and the low-order half of the dest operand and places the result in the entire dest operand location.

The src and dest operands are interpreted as unsigned integers.
The dest operand may be specified as an even-odd General-Purpose register pair. In such cases, the instruction reads the even-numbered register of the pair and places the low-order half (1, 2 or 4 bytes) of the result in the even register and the high-order half in the next consecutive register. The register pair must be specified in assembly language by the name of the even register of the pair.

If the Top of Stack (TOS) addressing mode is used for the dest operand, the Stack Pointer contents do not change. Note that this is not the same as popping a value of length "i" and pushing a result of length "2i". Space must already have been allocated on the stack to accommodate the entire result.

Flags Affected: None.
Traps: None.

MEIi
Multiply Extended Integer (continued)

\section*{Examples:}
1. MEIW R2, 10 (SB)
CE A5 16 0A
2. MEIW R2, R0
CE 2510

Example 1 multiplies the low-order word of register \(R 2\) and the word at the dest operand address \(10(S B)\) and places the double-word result at the dest operand address 10 (SB).

Example 2 multiplies the low-order word of register R2 and the low-order word of register R0. The result is a double-word. The low-order word of the result is written to the low-order word of register R0, the high-order word is written to the low-order word of register R1.

These instructions are illustrated below:
\begin{tabular}{ccc} 
& \begin{tabular}{c} 
Operand Values: \\
Operands \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex (Dec) \\
After
\end{tabular} \\
\hline
\end{tabular}
```

Syntax: MODi src, dest
gen gen MODW
read.i rmw.i MODD

```
\begin{tabular}{|c|c|c|c|}
\hline ! src & ! dest & ! & MODi \\
\hline ! gen & ! gen & & \\
\hline 23 & 161 & & \\
\hline
\end{tabular}

The MODi instruction places the value dest modulo src in the dest operand location. The modulus is the remainder after division is performed as per the DIVi instruction (q.v.). It is computed as
```

dest - ((dest DIV src) * src)

```
where dest DIV src is rounded to the next integer less than or equal to the exact quotient. The result of a MODi instruction always has the sign of the src operand (i.e., the divisor) unless the result is zero, which is always positive. Compare the REMi instruction (q.v.).

The src and dest operands are interpreted as signed integers.
Flags Affected: None.
Traps: Divide by Zero Trap (DVZ) is activated if src equals zero.

MODi
Modulus (continued)

\section*{Example:}


\section*{Move Floating Point}
```

Syntax: MOVf src, dest

```
                            gen gen
read.f write.f


The MOVf instruction copies the src operand to the dest operand location.
Flags Affected: No PSR Flags. The FSR TT field is set to all zeroes.

Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

\section*{Example:}
```

MOVF FO, 8(SB)

This example moves the single-precision number in register $F 0$ to the operand at address $8(\mathrm{SB})$.

| Operands | Operand Value: <br> Before | Hex (Dec) <br> After |
| :---: | :---: | ---: |
| F0 | 3F800000 |  |
| $(+1.0)$ | $3 F 800000$ |  |
| $(+1.0)$ |  |  |

MOVi
Move

Syntax: MOVi src, dest MOVB
$\begin{array}{lll}\text { gen } & \text { gen } & \text { MOVW } \\ \text { read.i } & \text { write.i } & \text { MOVD }\end{array}$


The MOVi instruction copies the src operand to the dest operand location.

Flags Affected: None.
Traps: None.

## Example:

MOVD R0, 8(SB)
970608

This example copies the contents of register $R 0$ to the double-word at the address specified by $8(S B)$.

Operand Values: Hex

| Operands | Before | After |
| :---: | :--- | :---: |
| R0 | 12345678 | 12345678 |
| $8(\mathrm{SB})$ | AAAAAAAA | 12345678 |


| Syntax: MOVif | sre, | dest | MOVBF |
| :--- | :--- | :--- | :--- |
|  | gen MOVBL |  |  |
|  | read.i | gen | write.f |



The MOVif instruction converts the integer operand src to a single- or doubleprecision floating-point number and places the result in the dest operand location.

Rounding, if required, is controlled by the Rounding Mode bits in the FSR.
Flags Affected: No PSR flags. FSR flags are affected as follows:
IF is set on an inexact result; unaffected otherwise.
$T T$ field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes.
See Sections 2.4.2 and 3.3 for details of exceptional conditions.

Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3. Particularly relevant is the Inexact Result trap if it is enabled in the FSR. It can occur in the MOVDF form, because in this case there are fewer significant bits in dest than in src. The smallest integer values for which this will happen are $+16,777,217$ (01000001 Hex) and -16,777,217 (FEFFFFFF Hex).

## MOVif

Move Converting Integer to Floating Point (continued)

## Examples:

1. MOVBF 2, F0
3E 04 A0 02
2. MOVDL $16(\mathrm{SB}), \mathrm{F} 2$
3E 83 D0 10

Example 1 converts the integer constant 2 to a single-precision number and copies the number to the register FO .

Example 2 converts the double-word integer at the address specified by 16 (SB) to a double-precision number and places the number into the register $F 2$.

| Operands | Operand Values: <br> Before |
| :---: | :---: |
| Hex (Dec) <br> After |  |

Ex. 1:
2
(immediate)

F0
AAAAAAAA
02
$(+2)$

$$
20401110
$$ (+541069584)

AAAAAAAAAAAAAAA

$$
16 \text { (SB) }
$$

$$
(+541069584)
$$

After--

Ex. 2:

$$
(+541069584)
$$

41C0200888000000 (+541069584.0)

## Move Floating to Long Floating

## Syntax: MOVFL src, dest

gen gen
read.F write.L


The MOVFL instruction converts the src operand to double-precision format and places the result in the dest operand location.

Flags Affected: No PSR Flags.
The FSR TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes.
See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3.

## Example:

MOVFL 8(SB), F0 3E 1B D0 08

This example converts the single-precision number at the address specified by $8(S B)$ to a double-precision number and places the number in the register $F 0$.

| Operands | Operand Values: <br> Before | Hex (Dec) <br> After |
| :---: | :---: | :---: |
| $8(\mathrm{SB})$ | 3 F 800000 <br> $(+1.0)$ | $3 F 800000$ <br> $(+1.0)$ |
| F0 | AAAAAAAAAAAAAAA | $3 F F 0000000000000$ |
|  | $(+1.0)$ |  |

## MOVLF

Move Long Floating to Floating

## Syntax: MOVLF src, dest

gen gen
read.L write.F


The MOVLF instruction converts the src operand to a single-precision number and places the result in the dest operand location.

Rounding is performed, if necessary, according to the rounding mode selected in the FSR. See Section 3.3 for details of rounding modes.

Flags Affected: No PSR flags. FSR flags are affected as follows:
UF is set if an underflow occurs; unaffected otherwise.
IF is set on an inexact result; unaffected otherwise.
TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes.
See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps:
Undefined Instruction Trap (UND) is activated if the $F$ bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3. Particularly relevant cases are:

* Overflow which occurs if the src operand is too great in absolute value to be represented as a single-precision number.
* Underflow which, if enabled in the FSR, occurs if the src operand is too small in absolute value to be represented as a normalized single-precision number.
* Inexact Result which, if enabled in the FSR, occurs if a loss of precision occurs in the conversion.


## Move Long Floating to Floating (continued)

## Example:

MOVLF FO, 12(SB) 3E 9606 0C

This example converts the double-precision number in register $F 0$ to a singleprecision number and places the result at address 12 (SB).

| Operands | Operand V Before | Hex (Dec) After |
| :---: | :---: | :---: |
| F0 | $\begin{gathered} 3 F F 0000000000000 \\ (+1.0) \end{gathered}$ | $\begin{gathered} 3 F F 0000000000000 \\ (+1.0) \end{gathered}$ |
| 12 (SB) | AAAAAAAA | $\begin{array}{r} 3 F 800000 \\ (+1.0) \end{array}$ |

## MOVMi

Move Multiple

| Syntax: | MOVMi | block1, | block2, | length | MOVMB |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | gen | gen | disp | MOVMW |
|  |  | addr | addr |  | MOVMD |



The MOVMi instruction copies the contents of block1 to block2. The instruction copies consecutive integers from block1 to consecutive integer locations in block2.

In assembly language, the length operand is specified as the number of integers in each block. In the machine instruction, however, the length operand is encoded according to the formula

```
(num - 1) * i
```

where num is the number of integers in each block, and $i$ is the number of bytes per integer. A block may not be greater than 16 bytes in length.

Flags Affected: None.
Traps: None.

## Example:

MOVMW 10 (R0), $16(\mathrm{R} 1), 4 \quad$ CE 4142 0A 1006
This instruction copies four word-long integers from the block starting at the address specified by $10(R 0)$ to the block starting at the address specified by 16 (R1).


## Move Quick Integer

## Syntax: MOVQi src

## dest

MOVQB
quick gen
MOVQW
write.i


The MOVQi instruction copies the src operand to the dest operand location. Before the copy operation, src is sign-extended to the length of dest.

Flags Affected: None.

Traps: None.

## Example:

MOVQW 7, TOS
DD BB

This example pushes the quick value 7 as a word onto the top of the stack. The high-order bits of the result are zero-filled due to sign-extension.

| Operands | Operand Values: <br> Before | Hex <br> After |
| :---: | :---: | :---: |
| 7 | $0007 *$ | -- |
| (quick) |  |  |
| SP | 0000 FFEE |  |
| Stack: |  |  |
| 0000 FFEC | XXXx | 000 FFEC |
| 0000 FFEE | AAAA | AAAA |

* This value is the internal representation of the Quick value 7, after signextension to Word length. The operand is encoded within the instruction as binary 0111.

MOVSi
MOVST
Move String

```
Syntax: MOVSi options


\section*{Syntax: MOVST options}
```

                                    MOVST
    +----------+---+-+-+---------+---+------------------
!0 0 0 0 0!UW !B!1!0 0 0 0 0!0 0!0 0 0 0 1 1 1 1 0!
!-++-+-+-++-+-+-+-!-+-++-+-++-+-+-++-!-+-++-+-+-++-+-+-!
23 16 15 8 7 0

```

Operands of the MOVSi and MOVST instructions are specified in General-Purpose registers:
```

R0 - Number of string elements to be processed.
R1 - Address of current String 1 element.
R2 - Address of current String 2 element.
R3 - Address of translation table (MOVST form only).
R4 - Match value (with Until Match or While Match option only).

```

The MOVSi instruction copies consecutive elements of String 1 (address in R1) to consecutive element locations in String 2 (address in R2). After an element is copied, the instruction sets register R1 to the address of the next element to copy, sets register \(R 2\) to the address of the next location to receive an element, and sets R0 to the number of elements remaining to be copied. See Section 3.7 for the exact sequences followed by String instructions.

The MOVST instruction copies one-byte elements from String 1, after translation, to String 2. The translated value to be copied is found by adding the current element from String 1 as an unsigned integer to the translation table address found in register R3. The instruction copies elements and sets registers as described above. See Section 3.7 for details of string translation.

Options may be specified by listing the letters B (Backward), U (Until Match) and \(W\) (While Match) as operands. The \(U\) and \(W\) options are mutually exclusive. See Section 3.7 for details of the options available in String instructions.

\section*{Move String (continued)}

In the machine instruction, the options are encoded in the \(B\) and UW fields as follows:

B field \(=0\)
1
UW field = \(00 \quad\) Neither Until Match nor While Match. 01 While Match.
10 (reserved)
11 Until Match.

String instructions are interruptible. See Section 3.7.
Flags Affected: \(F\) is set if the \(U\) or \(W\) option is specified and the corresponding Until/While condition is met, otherwise it is cleared.

Traps: None.

\section*{Example:}
MOVST
OE 8000

This example moves byte-long integers from the first string, after translation, to the second string.

The instruction is as follows:

\section*{MOVSi}

MOVST
Move String (continued)


This example translates 32 binary integers (in the range 0-31) into binary coded decimal (BCD) values. Each integer is read from String 1 (address given by R1) and used as an offset into the translation table at address 10000 (Hex). The BCD value found at that address in the translation table is then copied to the current location in String 2 (address in R2).

\section*{dest}

MOVSUB
MOVSUW
MOVSUD


The MOVSUi instruction copies the src operand in Supervisor space to the dest operand location in User space. User Mode protection is applied to the User space access.

Flags Affected: None.
Traps: Illegal Instruction Trap (ILL) is activated if the U flag is set.

\section*{Example:}

MOVSUB 5(SP), 9(SB) AE 8C CE 0509
This example copies the byte at the operand address \(5(\mathrm{SP})\) in supervisor space to the operand address \(9(S B)\) in user space.

The instruction is illustrated below:
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{l} 
Hex \\
After
\end{tabular} \\
\hline \begin{tabular}{c} 
(SP) \\
Supervisor
\end{tabular} & 10 & 10 \\
\(9(\mathrm{SB})\) & AA & 10 \\
User & &
\end{tabular}

\section*{MOVUSi}

Move Value from User to Supervisor Space
\begin{tabular}{llll} 
Syntax: MOVUSi & src, & dest & MOVUSB \\
& gen & gen & MOVUSW \\
& addr & addr & MOVUSD
\end{tabular}


The MOVUSi instruction moves the src operand in User space to the dest operand location in Supervisor space. User Mode protection is applied to the User space access.

Flags Affected: None.

Traps: Illegal Instruction Trap (ILL) is activated if the U flag is set.

\section*{Example:}
```

MOVUSB 9(SB), 5(SP) AE 5C D6 09 05

```

This example moves the byte at the address specified by \(9(\mathrm{SB})\) in user space to the address specified by \(5(S P)\) in supervisor space.
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex \\
After
\end{tabular} \\
\hline \(9(\mathrm{SB})\) & 10 & 10 \\
User & & \\
\(5(\mathrm{SP})\) & AA & 10 \\
Supervisor & &
\end{tabular}


\section*{Syntax: MOVXWD}
src,
gen
read.W write.D

\section*{dest}
gen

MOVXBD
MOVXWD
MOVXBW



Syntax: MOVXBW src, dest
gen gen
read.B write.W


The MOVX instructions convert signed integers to any greater length while preserving their signed values through sign-extension.

The MOVXBD instruction copies the byte-length src operand to the low-order byte of the double-word dest operand and extends the src operand's sign bit through the remaining high-order bits of the dest operand.

The MOVXWD instruction copies the word-length src operand to the low-order word of the double-word dest operand and extends the src operand's sign bit through the remaining high-order bits of the dest operand.

The MOVXBW instruction copies the byte-length src operand to the low-order byte of the word dest operand and extends the src operand's sign bit through the remaining high-order bits of the dest operand.

The src and dest operands are interpreted as signed integers.

MOVXii
Move with Sign-Extension (continued)

Flags Affected: None.
Traps: None.

\section*{Example:}

MOVXBW \(2(\mathrm{SB}), \mathrm{R} 0\)
CE 10 D0 02
This example copies the byte at the address specified by \(2(\mathrm{SB})\) to the low-order byte of register \(R 0\) and extends the sign bit of the byte through the next eight bits of RO. The instruction affects the low-order word of R0 only. The instruction (for two cases) is as follows:
\begin{tabular}{|c|c|c|c|}
\hline & Operands & Operand Values: Before & \[
\begin{aligned}
& \text { Hex } \text { (Dec) } \\
& \text { After } \\
& \hline
\end{aligned}
\] \\
\hline Case 1: & 2 (SB) & \[
\begin{gathered}
F 0 \\
(-16)
\end{gathered}
\] & \[
\begin{gathered}
\text { F0 } \\
(-16)
\end{gathered}
\] \\
\hline & R0 & AAAAAAAA & AAAAFFFO
\[
(-16)
\] \\
\hline Case 2: & 2 (SB) & \[
\begin{gathered}
70 \\
(+112)
\end{gathered}
\] & \[
\begin{gathered}
70 \\
(+112)
\end{gathered}
\] \\
\hline & R0 & AAAAAAAA & \[
\begin{aligned}
& \text { AAAAO } 070 \\
& (+112)
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{lll} 
sre, & dest & MOVZBD \\
gen & gen & MOVZWD \\
read.B & write.D & MOVZBW
\end{tabular}


Syntax.
\begin{tabular}{lll} 
MOVZWD & src, & dest \\
& gen & gen \\
& read.W & write.D
\end{tabular}


Syntax: MOVZBW src, dest
gen gen
read.B write.w


The MOVZ instructions convert unsigned integers to any greater length while preserving their unsigned values through zero-extension.

The MOVZBD instruction copies the byte-length src operand to the low-order byte of the double-word dest operand and zero-fills the remaining high-order bits of the dest operand.

The MOVZWD instruction copies the word-length src operand to the low-order word of the double-word dest operand and zero-fills the remaining high-order bits of the dest operand.

The MOVZBW instruction copies the src operand to the low-order byte of the dest operand and zero-fills the remaining eight bits of dest.

The src and dest operands are interpreted as unsigned integers.

MOVZii
Move with Zero-Extension (continued)

Flags Affected: None.
Traps: None.

\section*{Example:}

MOVZBW -4 (FP), R0
CE 14 CO 7C
This example copies the byte at the address specified by \(-4(F P)\) to the low-order byte of register R0 and sets the next eight bits of register R0 to zero. The instruction affects only the low-order word of R0.
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & Hex \begin{tabular}{c} 
(Unsigned Dec) \\
After
\end{tabular} \\
\hline\(-4(\mathrm{FP})\) & FF & FF \\
R0 & \((+255)\) & \((+255)\) \\
& AAAAAAAA & AAAA00FF \\
\((+255)\)
\end{tabular}
\begin{tabular}{lll} 
MULf & src,, & dest \\
& gen & gen
\end{tabular}
\begin{tabular}{|c|c|c|c|}
\hline ! src & ! dest & ! & MULf \\
\hline ! gen & gen & & \\
\hline 23 & 15 & & 7 \\
\hline
\end{tabular}

The MULf instruction multiplies the src and dest operands and places the result in the dest operand location. Results for normalized and zero operands are given in the table below. The symbols "+n" and "-n" represent non-zero normalized numbers, positive and negative, respectively. The symbols "+z" and "-z" represent positive and negative zero, respectively.
\begin{tabular}{cccccc}
\multicolumn{2}{l}{ dest: } & \(+n\) & \(-n\) & \(+z\) & \(-z\) \\
\hline src & \(!\) & & & & \\
\(+n\) & \(!\) & \(*\) & \(*\) & \(+z\) & \(-z\) \\
\(-n\) & \(!\) & \(*\) & \(*\) & \(-z\) & \(+z\) \\
\(+z\) & \(!\) & \(+z\) & \(-z\) & \(+z\) & \(-z\) \\
& \(!\) & & & & \\
\(-z\) & \(!\) & \(-z\) & \(+z\) & \(-z\) & \(+z\)
\end{tabular}
* The result in these cases is the product of the two operands.

Flags Affected: No PSR flags. FSR flags are affected as follows:
UF is set if an underflow occurs; unaffected otherwise.
IF is set on an inexact result; unaffected otherwise.
TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes.
See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3.

MULf
Multiply Floating (continued)

\section*{Examples:}
1. MULF F0, F7
BE F1 01
2. MULL \(-8(F P), 8(S B)\)
BE B0 C6 7808

Example 1 multiplies the single-precision numbers in registers \(F 0\) and \(F 7\) and places the result in register \(F 7\).

Example 2 multiplies the double-precision numbers at addresses -8(FP) and 8(SB) and places the double-precision result at address 8(SB).
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex (Dec) \\
After
\end{tabular} \\
\hline F0 & 42250000 \\
\((+41.25)\) & 42250000 \\
F7 & & \((+41.25)\) \\
& 40 A00000 & \((+5.0)\)
\end{tabular}
\begin{tabular}{ccc} 
Ex. 2: & \(-8(\mathrm{FP})\) & 409 F 440000000000 \\
& \((+2001.0)\) & 409 F 440000000000 \\
& \((+2001.0)\) \\
& \(8(\mathrm{SB})\) & 40 F 3218 E 00000000 \\
& \((+78360.875)\) & \(41 \mathrm{~A} 2 \mathrm{~B} 128 \mathrm{DDC00000}\) \\
& & \((+156800110.875)\)
\end{tabular}

Multiply

\section*{Syntax: MULi src, dest}
gen gen
read.i rmw.i


The MULi instruction multiplies the src and dest operands and places the product in the dest operand location. If the product is longer than dest, the high-order bits are truncated.

Flags Affected: None.

Traps: Integer Overflow Trap (OVF) is activated if the V flag is set and the result cannot be represented exactly in dest.

\section*{Examples:}
1. MULW 5, R0
2. MULD 4(-4(FP)), 3(SB)
CE 21 A0 0005
CE A3 86 7C 0403

Example 1 multiplies the constant 5 and the low-order word of register R0 and places the result in the low-order word of register R0.

Example 2 multiplies the double-word at the memory address specified by \(4(-4(F P))\) by the double-word at the address specified by \(3(\mathrm{SB})\). The instruction places the result in the double-word at address 3(SB).

These instructions are illustrated below:
\begin{tabular}{ll} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular}
\end{tabular} \begin{tabular}{c} 
Hex (Dec) \\
After
\end{tabular}

Ex. 1:
5
(immediate)

R0
BBBB000A
BBBB0032
(+10)
(+50)

Ex. 2:
\(4(-4(F P))\)

3 (SB)
FFFFFFF0
00000020
(+32)

NEGf
Negate Floating
```

Syntax: NEGf src, dest NEGF
gen gen NEGL

```
                                read.f write.f


The NEGf instruction complements the sign bit of the src operand and places the result in the dest operand location.

Flags Affected: No PSR flags.
The FSR TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes.
See sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3.

\section*{Example:}

NEGF F0, F2
BE 9500

This example complements the sign bit of the single-precision number in register FO and places the result in register \(F 2\).
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex (Dec) \\
After
\end{tabular} \\
\hline F0 & 42250000 & 42250000 \\
\((+41.25)\) & \((+41.25)\) \\
F2 & & C2250000 \\
& AAAAAAAA & \((-41.25)\)
\end{tabular}

\section*{Negate}
```

Syntax: NEGi src, dest NEGB
gen gen NEGW
read.i write.i NEGD

```


The NEGi instruction negates (takes the two's complement of) the src operand by subtracting src from zero and places the result in the dest operand location.

Flags Affected: C is set on a borrow from subtraction, cleared if no borrow. (A borrow will always occur in this instruction unless the src operand is zero.)
\(F\) is set on an overflow from subtraction, cleared if no overflow. This condition will occur if the src operand is the most negative number that can be represented in the operand length specified by the programmer. This value for bytes is -128 (Hex 80); for words it is -32768 (Hex 8000) and for double-words it is \(-2,147,483,648\) (Hex 80000000). These values have no corresponding positive values in the same operand length. The result returned on an overflow is the original src operand.

Integer borrow and overflow conditions are defined in Section 3.1.

Traps: Integer Overflow Trap (OVF) is activated if the V flag is set and the result cannot be represented exactly in dest.

NEGi
Negate (continued)

\section*{Examples:}
1. NEGB R5, R6
4E AO 29
2. NEGW 4 (SB) , 6 (SB)
4E A1 D6 0406

Example 1 negates the low-order byte of register R5 and places the result in the low-order byte of register R6. The remaining bytes of registers R5 and R6 are neither used nor affected.

Example 2 negates the word at the memory address specified by \(4(\mathrm{SB})\) and places the word result at the memory address specified by 6(SB).
\begin{tabular}{|c|c|c|c|}
\hline & Operands & Operand Values: Before & \[
\begin{aligned}
& \text { Hex } \text { (Dec) } \\
& \text { After }
\end{aligned}
\] \\
\hline \multirow[t]{3}{*}{Ex. 1:} & R5 & AAAAAAFF
\[
(-1)
\] & AAAAAAFF
\[
(-1)
\] \\
\hline & R6 & BBBBBBBB & BBBBBB01 (+1) \\
\hline & UPSR & \(n z f v x l t c\) & \(n z \underline{0 v x l t} \underline{1}\) \\
\hline \multirow[t]{3}{*}{Ex. 2:} & 4 (SB) & \[
\begin{aligned}
& 0041 \\
& (+65)
\end{aligned}
\] & \[
\begin{aligned}
& 0041 \\
& (+65)
\end{aligned}
\] \\
\hline & 6 (SB) & xxxx & \[
\begin{aligned}
& \text { FFBF } \\
& (-65)
\end{aligned}
\] \\
\hline & UPSR & \(n z f v x l t c\) & \(n \mathrm{z} \underline{0} \mathrm{vxlt} \underline{1}\) \\
\hline
\end{tabular}

\section*{No Operation}

Syntax: NOP
\begin{tabular}{|c|c|}
\hline ! & NOP \\
\hline & 0000101 \\
\hline & -+-+-+-+-! \\
\hline 7 & 0 \\
\hline
\end{tabular}

The NOP instruction passes control to the next sequential instruction. No operation is performed.

Flags Affected: None.
Traps: None.
Example:
NOP A2

\section*{NOTi}

Complement Boolean

\section*{Syntax:}


The NOTi instruction complements (inverts) the Boolean value of the src operand and places the result in the dest operand location. The complement of a Boolean value is that value with its least-significant bit complemented. The Boolean value "True" (the integer value 1) thus becomes "False" (the integer value 0) and vice versa. Boolean values are described in Section 3.4.

Flags Affected: None.
Traps: None.

\section*{Examples:}
\begin{tabular}{llllll} 
1. NOTB R0, R0 & 4 E & 24 & 00 \\
2. NOTW \(10(\mathrm{R} 1)\), TOS & 4 E & E 5 & 4 D & \(0 A\)
\end{tabular}

Example 1 complements the Boolean value in the low-order byte of register R0. The remaining bytes of \(R 0\) are neither used nor affected.

Example 2 complements the 1-word Boolean value at memory address 10(R1) and pushes the result as a word onto the top of the stack.
\begin{tabular}{|c|c|c|c|}
\hline & Operands & Operand Values: Hex Before & \begin{tabular}{l}
(Boolean) \\
After
\end{tabular} \\
\hline Ex. 1: & R0 & AAAAAA01 (True) & AAAAAA0 0 (False) \\
\hline Ex. 2: & 10 (R1) & \[
\begin{aligned}
& \text { AAAAO000 } \\
& \text { (False) }
\end{aligned}
\] & AAAA0000 (False) \\
\hline & \[
\begin{gathered}
\text { Stack: } \\
0000 \text { FFE0 } \\
0000 \text { FFE2 }
\end{gathered}
\] & \[
\begin{aligned}
& \text { XXXX } \\
& \text { AAAA }
\end{aligned}
\] & \[
\begin{aligned}
& 0001 \text { (True) } \\
& \text { AAAA }
\end{aligned}
\] \\
\hline & SP & 0000FFE2 & 0000FFE0 \\
\hline
\end{tabular}

\section*{Syntax:}


The ORi instruction performs a bit-wise logical OR operation between the src and dest operands and places the result in the dest operand location.

The instruction ORs each bit in src with the corresponding dest bit. If two corresponding bits are both "0", the dest bit is set to "0"; otherwise, the dest bit is set to "1".

Flags Affected: None.
Traps: None.

\section*{Example:}
```

ORB -6(FP), 11(SB) 98 C6 7A 0B

```

This example ORs the bytes at memory addresses \(-6(F P)\) and \(11(S B)\), placing the result at memory address \(11(\mathrm{SB})\).
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Binary \\
After
\end{tabular} \\
\hline\(-6(\mathrm{FP})\) & 11011000 & 11011000 \\
\(11(\mathrm{SB})\) & 00001011 & 11011011
\end{tabular}

\section*{POLYf}

Polynomial Floating
\begin{tabular}{|c|c|c|c|c|}
\hline Syntax: & POLY \({ }^{\text {S }}\) &  & & \\
\hline & src1 & 1 ! src2 & ! & POLYf \\
\hline & ! gen & ! gen & ! 0 & !f! 1 \\
\hline & ! -+-+-+ & -+-+-+-+-! & & -! -+ \\
\hline & 23 & 16 & & 87 \\
\hline
\end{tabular}

The POLYf instruction multiplies srcl and floating-point register \(F 0\) and than adds to the product src2. The result is placed in F0. ( F0 := (F0*src1) + src2 )

Flags Affected: No PSR flags.
The FSR TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes. See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap(UND) is activated if the \(F\) bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3.

\section*{Example:}
POLYF F2, F3
FE C9 10

This example multiplies the single-precision numbers in \(F 2\) and \(F 0\), places the result in register FO and finally adds the single-precision number in F 3 to FO .

The instruction is illustrated below:
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: Hex \\
Before
\end{tabular} & \begin{tabular}{c} 
(Dec) \\
After
\end{tabular} \\
\hline & & \\
F0 & C2250000 & C4323800 \\
& \((-41.250)\) & \((-712.875)\) \\
F2 & 418 C 0000 & 418 c 0000 \\
& \((+17.500)\) & \((+17.500)\) \\
F3 & 41100000 & 41100000 \\
& \((+9.000)\) & \((+9.000)\)
\end{tabular}

\section*{QUOi}

Quotient

\section*{Syntax:}
\begin{tabular}{lll} 
QUOi & src, & dest \\
& gen & gen \\
& read.i & rmw.i
\end{tabular}
read.i rmw.i


The QUOi instruction divides the dest operand by the src operand, returning as the result the nearest integer whose absolute value is less than or equal to the absolute value of the exact quotient. The result is placed in the dest operand location.

The src and dest operands are interpreted as signed integers.
Flags Affected: None.
Traps: Divide by Zero trap (DVZ) is activated if src equals zero.
Integer Overflow Trap (OVF) is activated if the V flag is set. It occurs only if the largest negative integer in a data format is divided by -1.

\section*{Examples:}
1. QUOB R0, R7
CE F0 01
2. QUOW \(4(S B), 8(S B)\)
CE B1 D6 0408

Example 1 divides the low-order byte of register \(R 7\) by the low-order byte of register R0, placing the result in the low-order byte of register R7.

Example 2 divides the word at address \(8(S B)\) by the word at address \(4(S B)\), placing the one-word result at address 8 (SB).
\begin{tabular}{|c|c|c|c|}
\hline & Operands & \begin{tabular}{l}
Operand \\
Before
\end{tabular} & Hex (Dec)
After \\
\hline \multirow[t]{2}{*}{Ex. 1:} & R0 & AAAAAA 05 (+5) & \[
\begin{array}{r}
\text { AAAAAA } 05 \\
(+5)
\end{array}
\] \\
\hline & R7 & \[
\begin{array}{r}
\text { BBBBBB22 } \\
(+34)
\end{array}
\] & BBBBBB 06
\((+6)\) \\
\hline \multirow[t]{2}{*}{Ex. 2:} & 4 (SB) & \[
\begin{aligned}
& 0010 \\
& (+16)
\end{aligned}
\] & \[
\begin{aligned}
& 0010 \\
& (+16)
\end{aligned}
\] \\
\hline & 8 (SB) & \[
\begin{aligned}
& \text { FFE1 } \\
& (-31)
\end{aligned}
\] & \[
\begin{gathered}
\text { FFFF } \\
(-1)
\end{gathered}
\] \\
\hline
\end{tabular}

\section*{Syntax: RDVAL loc}
gen
addr


The RDVAL instruction checks the protection level assigned to the user-mode virtual memory address specified as loc. If the address is allowed to be read while the CPU is in user mode, the F flag in the PSR is cleared. If the address is not allowed to be read, the \(F\) flag in the PSR is set. An address which is protected against reading is also protected against writing, and is therefore inaccessible for any use by a user-mode program.

NOTE: Although the final effective address of loc is interpreted as a user-mode virtual address, any memory references required in order to calculate that effective address are interpreted as using supervisor-mode addresses. This will occur in using the Memory Relative and External addressing modes for loc.

Flags Affected: \(F\) is set if loc is inaccessible in user mode, cleared otherwise.
Traps:
Undefined Instruction Trap (UND) is activated if the M bit in the CFG register is clear.

Illegal Operation Trap (ILL) is activated if this instruction is attempted while the PSR \(U\) bit is set.

Abort Trap (ABT) is activated if the Level 1 page table entry for loc is invalid (V bit \(=0\) ) and the Protection Level (PL) indicates that the access is allowed. No trap is issued for an invalid Level 2 page table entry, and the Protection Level field is assumed to be present regardless of the state of the \(V\) bit.

\section*{Example:}

RDVAL 512 (R0)
\[
\begin{array}{lllll}
1 \mathrm{E} & 03 & 40 & 82 & 00
\end{array}
\]

This example checks the protection level assigned to the address \(512(\mathrm{RO})\) and sets or clears the \(F\) flag to indicate the result.

REMi
Remainder

\section*{Syntax:}


The REMi instruction places the remainder from dividing dest by src into the dest operand location. The remainder is computed as
dest - src * (dest QUO src) ,
where dest QUO src is the result of dividing dest by src as per the QUOi instruction (q.v.).

The result of a REMi instruction always has the sign of the dest operand (i.e. the dividend) unless the result is zero, which is always positive. Compare the MODi instruction (q.v.).

The src and dest operands are interpreted as signed integers.
Flags Affected: None.
Traps: Divide by Zero trap (DVZ) is activated if src equals zero.

\section*{Remainder (continued)}

\section*{Example:}

REMB 4 (SB), 8 (SB)
CE B4 D6 0408

This example computes the remainder from dividing the 1-byte operand at address \(8(S B)\) by the 1-byte operand at address \(4(S B)\) and places the result as a byte at address \(8(\mathrm{SB})\).

The action of this instruction for four different cases is as follows:
\begin{tabular}{|c|c|c|c|}
\hline & Operands & Operand Values: Hex Before & \begin{tabular}{l}
(Dec) \\
After
\end{tabular} \\
\hline \multirow[t]{2}{*}{Case 1:} & 4 (SB) & \[
\begin{gathered}
0 F \\
(+15)
\end{gathered}
\] & \[
\begin{gathered}
0 F \\
(+15)
\end{gathered}
\] \\
\hline & 8 (SB) & \[
\begin{gathered}
21 \\
(+33)
\end{gathered}
\] & \[
\begin{gathered}
03 \\
(+3)
\end{gathered}
\] \\
\hline \multirow[t]{2}{*}{Case 2:} & 4 (SB) & \[
\begin{gathered}
F 1 \\
(-15)
\end{gathered}
\] & \[
\begin{gathered}
F 1 \\
(-15)
\end{gathered}
\] \\
\hline & 8 (SB) & \[
\begin{gathered}
21 \\
(+33)
\end{gathered}
\] & \[
\begin{gathered}
03 \\
(+3)
\end{gathered}
\] \\
\hline \multirow[t]{2}{*}{Case 3:} & 4 (SB) & \[
\begin{gathered}
F 1 \\
(-15)
\end{gathered}
\] & \[
\begin{gathered}
F 1 \\
(-15)
\end{gathered}
\] \\
\hline & 8 (SB) & \[
\begin{gathered}
D F \\
(-33)
\end{gathered}
\] & \[
\begin{gathered}
\text { FD } \\
(-3)
\end{gathered}
\] \\
\hline \multirow[t]{2}{*}{Case 4:} & 4 (SB) & \[
\begin{gathered}
0 F \\
(+15)
\end{gathered}
\] & \[
\begin{gathered}
0 \mathrm{~F} \\
(+15)
\end{gathered}
\] \\
\hline & 8 (SB) & \[
\begin{gathered}
\text { DF } \\
(-33)
\end{gathered}
\] & \[
\begin{aligned}
& \text { FD } \\
& (-3)
\end{aligned}
\] \\
\hline
\end{tabular}

\section*{RESTORE}

Restore General Purpose Registers

\section*{Syntax: RESTORE reglist}
imm
\begin{tabular}{|c|c|}
\hline ! & RESTORE \\
\hline & \(\begin{array}{llllll}1 & 1 & 0 & 0 & 1 & 0\end{array}\) \\
\hline & +-+-+-+-+- \\
\hline & 0 \\
\hline
\end{tabular}

The RESTORE registers instruction restores from the current stack the General Purpose registers specified by reglist.

In assembly language, the reglist operand is specified as a list of zero or more General-Purpose register names enclosed in brackets "[]". The instruction copies to each register in the list a double-word popped from the stack. Register names may appear in any order within reglist but must be separated by commas. Brackets are required even if no register names are given.

In the machine instruction, the reglist operand is encoded in an 8-bit field as shown below. Each bit in the field corresponds to one General-Purpose register. When the instruction is executed, the instruction reads the bits in the field from right to left beginning with bit 0 . If a bit is "0", the instruction ignores the corresponding register. If a bit is "1", it restores the corresponding register from the stack. Note that the binary format of the reglist operand is backward from the format of the reglist operand in the SAVE instruction; i.e., bit 0 corresponds to R7 instead of R0.
```

+--+--+--+--+--+--+--+---+
!R0!R1!R2!R3!R4!R5!R6!R7!
!--+--+--+--+--+--++--+---!
7
0

```

Flags Affected: None.
Traps: None

\section*{Example:}

RESTORE [R0, R2, R7]
72 A1

This instruction restores the contents of registers R0, R2, and R7 from the stack. The registers are restored in order beginning with register \(R 7\) and ending with RO.

\section*{Restore General Purpose Registers (continued)}

The action of the instruction is illustrated below.

Operand Values: Hex
\begin{tabular}{|c|c|c|}
\hline Operands & Before & After \\
\hline R0 & BBBBBBBB & 00000010 \\
\hline R2 & BBBBBBBBB & FFFFFFEF \\
\hline R7 & BBBBBBBBB & FFFFF9AB \\
\hline SP & 0000FFEO & 0000FFEC \\
\hline \multicolumn{3}{|l|}{Stack:} \\
\hline 0000FFE0 & FFFFF9AB & xxxxxxxx \\
\hline 0000 FFE 4 & FFFFFFEF & xxxxxxxx \\
\hline 0000 FFE 8 & 00000010 & xxxxxxxx \\
\hline 0000FFEC & AAAAAAAA & AAAAAAAA \\
\hline
\end{tabular}
* The RESTORE instruction does not itself change the contents of these memory locations. However, information that is outside the stack should be considered unpredictable for other reasons. See Section 2.8.1.

RET
Return from Subroutine

\section*{Syntax: RET constant}
disp
\begin{tabular}{|c|c|}
\hline & RET \\
\hline \multicolumn{2}{|l|}{} \\
\hline & -+-+-+-+- \\
\hline 7 & 0 \\
\hline
\end{tabular}

The RET instruction returns execution control from a local procedure and removes procedure parameters from the stack.

The instruction pops the return address as a 32-bit value from the currentlyselected stack. It then removes the number of bytes specified by the constant operand from the stack by adding the constant operand to the current stack pointer register. Finally, it transfers control by loading the return address into the PC register.

Flags Affected: None.

Traps: None

\section*{Return from Subroutine (continued)}

\section*{Example:}

RET \(16 \quad 1210\)

This example pops a new address from the currently-selected stack into the PC and adds 16 (H'10) to the stack pointer.
\begin{tabular}{|c|c|c|}
\hline Operand & Operand Before & \[
\begin{aligned}
& \text { Hex } \\
& \text { After } \\
& \hline
\end{aligned}
\] \\
\hline 16 & 10 & \\
\hline (disp) & (+16) & \\
\hline PC & 00009000 & 00009010 \\
\hline SP & 0000F000 & 0000F014 \\
\hline \multicolumn{3}{|l|}{Stack:} \\
\hline 0000F000 & 00009010 & mxxexxxy \\
\hline 0000F004 & BBBBBBBB & xxxxxxxx \\
\hline 0000F008 & BBBBBBBB & xxxxxxxx \\
\hline 0000F00C & BbBBBBBB & xxxxxxxx \\
\hline 0000 F 010 & BBBBBBBB & xxxxxxxx \\
\hline \(0000 \mathrm{F014}\) & AAAAAAAA & AAAAAAAA \\
\hline
\end{tabular}
* The RET instruction does not itself change the contents of these memory locations. However, information that is outside the stack should be considered unpredictable for other reasons. See Section 2.8.1.

\section*{RETI}

Return from Interrupt

\section*{Syntax: RETI}


The RETI instruction returns control from an interrupt service procedure to the program during which the interrupt was accepted, and informs any interrupt control circuitry present in the system that this is being done.

The RETI instruction does the following:
1. Performs either one or two "End of Interrupt" bus cycles in order to inform the appropriate Interrupt Controller(s) that this interrupt service procedure is ending. For details of this aspect of the RETI instruction, see the data sheets for the NS32202 Interrupt Control Unit and the appropriate CPU.
2. Pops a 32 -bit return address from the currently selected stack into the PC register.
3. Pops a 16-bit value from the currently-selected stack. If Direct-Exception mode is disabled the value is stored into the MOD register. Otherwise the value is discarded.
4. Pops a 16-bit PSR value from the currently selected stack into the PSR.
5. If Direct-Exception mode is disabled the instruction copies the doubleword from the address contained in the MOD register into the SB register.

Program execution continues at the new address placed in the PC register.
NOTE: The RETI instructon must not be used to return from the Non-Maskable or Non-Vectored interrupts or from any traps (including the Abort trap). Such use can cause anomalies in prioritization of interrupts by Interrupt Control circuits. For these use instead the Return from Trap instruction (RETT, q.v.).

Flags Affected: All flag states are restored from the stack.
Traps: Illegal Instruction Trap (ILL) is activated if this instruction is attempted while the \(U\) flag is set.

\section*{Return from Interrupt (continued)}

\section*{Example:}

RETI
52
This example returns control from an interrupt service procedure. Direct-Exception mode is disabled.

The action of this instruction is illustrated below. Note that the PSR \(S\) flag is assumed to be zero at the beginning of the instruction, thus selecting SPO as the current Stack Pointer. However, note also that after the instruction is completed the CPU is in User mode, the currently-selected Stack Pointer has become SP1, and interrupts are re-enabled.

Operand Values: Hex
\begin{tabular}{ccc} 
Operands & Before & After \\
\hline PC & & \\
SB & 0000 F 033 & 00009005 \\
MOD & 0000 F 100 & 00009080 \\
SP0 & 0020 & 0010 \\
PSR & 00001000 & \(00001008 *\) \\
& x000 & xB20 \\
& (xxxxipsu/nzfvxltc) & (xxxx1011/0010x000)
\end{tabular}
\begin{tabular}{cccc} 
Stack: & & & \\
00001000 & 00009005 & xxxxxxex & \(* *\) \\
00001004 & 0010 & xxxx & \(* *\) \\
00001006 & OB20 & xxxx & \(* *\) \\
00001008 & AAAA & & \\
Module & & & \\
Table: & & & \\
00000010 & 00009080 & (SB) & 00009080
\end{tabular}
* The final Stack Pointer value is the initial address plus 8 as follows: 4 (for double-word return address), 2 (for MOD address), and 2 (for PSR contents).
** The RETI instruction does not itself change the contents of these memory locations. However, information that is outside the stack should be considered unpredictable for other reasons. See Section 2.8.1.

\section*{RETT}

Return from Trap

\section*{Syntax: RETT constant}
disp
\begin{tabular}{|c|c|}
\hline ! & RETT \\
\hline \multicolumn{2}{|l|}{\(\begin{array}{llllllll}0 & 1 & 0 & 0 & 0 & 0 & 1 & 0\end{array}\)} \\
\hline \(1-\) & -+-+-+-+-+-! \\
\hline 7 & 0 \\
\hline
\end{tabular}

The RETT instruction returns control from a trap service procedure. It restores the PC, MOD and PSR. registers from the currently-selected stack, updates the SB register, and then removes any parameters passed by the procedure which caused the trap.

The instruction does the following:
1. Pops a 32-bit return address from the currently-selected stack into the PC register.
2. Pops a 16-bit value from the currently-selected stack. If Direct-Exception mode is disabled the value is stored into the MOD register. Otherwise the value is discarded.
3. Pops a 16-bit PSR value from the currently-selected stack into the PSR. Note that this may switch stack pointers by changing the PSR \(S\) bit.
4. If Direct-Exception mode is disabled the instruction copies the doubleword from the address contained in the MOD register into the SB register.
5. Adds the constant operand to the stack pointer newly selected in step 3.

Program execution continues at the new address placed in the PC register.
NOTE: When using the NS32202 Interrupt Control Unit, the RETT instruction must not be used to return from a vectored interrupt, since this instruction does not inform the Interrupt Control Unit that it is returning from an interrupt. To return properly from a vectored interrupt, use the RETI instruction.

Flags Affected: All flag states are restored from the stack.

Traps: Illegal Instruction Trap (ILL) is activated if the U flag is set.

\section*{Return from Trap (continued)}

\section*{Example:}

RETT 16
4210

This example returns control from a trap service procedure to a procedure which invoked the trap deliberately after pushing 16 bytes of parameters onto its stack. This instruction removes the 16 bytes from that stack as it returns control.

In the following illustration, it is assumed that the trap service routine is using the Interrupt Stack (with SPO as its stack pointer) and is returning to a procedure which is using the User Stack (with SP1). Direct-Exception mode is disabled.

Operand Values: Hex
Operands Before After
\begin{tabular}{cc}
16 & 10 \\
(disp) &
\end{tabular}
\begin{tabular}{rrr} 
PC & 0000 F 033 & 00009005 \\
SB & 0000 F 100 & 00009080 \\
MOD & 0020 & 0010 \\
SP0 & 00001018 & 00001020 \\
SP1 & 0000 FFE & 0000 FFF \\
PSR & x 000 & x 320
\end{tabular}
(xxxxipsu/nzfvxltc) (xxxx0011/0010x000)
\begin{tabular}{|c|c|c|c|c|}
\hline Interrupt & 00001018 & 00009005 & xxxxxxxx & * \\
\hline \multirow[t]{2}{*}{Stack} & 0000101 C & 03200010 & xxxxxxix & * \\
\hline & 00001020 & AAAAAAAA & AAAAAAAA & \\
\hline \multirow[t]{5}{*}{User Stack} & O000FFE0 & BBBBBBBB & xxxxxxxx & * \\
\hline & 0000 FFE 4 & BBBBBBBB & xxxxxxxx & * \\
\hline & 0000 FFE 8 & BBBBBBBB & xxxxxxxx & * \\
\hline & 0000FFEC & BBBBBBBBB & xxxxxxxx & * \\
\hline & 0000FFFO & CCCCCCCC & CCCCCCCC & \\
\hline Module & 00000010 & 00009080 (SB) & 00009080 & (SB) \\
\hline \multirow[t]{2}{*}{Table} & 00000014 & 00002000 (LB) & 00002000 & (LB) \\
\hline & 00000018 & 00009000 (PB) & 00009000 & (PB) \\
\hline
\end{tabular}
* The RETT instruction does not itself change the contents of these memory locations. However, information that is outside the stack should be considered unpredictable for other reasons. See Section 2.8.1.

ROTi
Rotate
\begin{tabular}{lll} 
Syntax. ROTi & count, dest & ROTB \\
& gen & gen \\
& read.B rmw.i & ROTW \\
& & ROTD
\end{tabular}


The ROTi instruction performs a rotation shift on the dest operand in the manner specified by the count operand. The sign of count determines the direction of the shift. The absolute value of count gives the number of bit positions to shift the dest operand.

The count operand value must be within the range -7 to +7 for the ROTB form, -15 to +15 for the ROTW form, and -31 to +31 for the ROTD form. A positive count specifies a left shift; a negative count specifies a right shift. In a rotation, each bit rotated off one end of dest is moved to the emptied bit position at the other end of dest.

The count operand is interpreted as a signed integer. The dest operand is interpreted as an unsigned integer.

Flags Affected: None.

Traps: None.

\section*{Rotate (continued)}

\section*{Examples:}
1. ROTB 4, R5
4E 40 A1 04
2. ROTB \(-3,16\) (SP)
4E 40 A6 FD 10

Example 1 rotates the least-significant byte of register R5 four bit positions to the left. The remaining bytes of R5 are unaffected.

Example 2 rotates the operand at address \(16(S P)\) three bit positions to the right.
\begin{tabular}{|c|c|c|c|}
\hline & Operands & \begin{tabular}{l}
Operand Values: \\
Before
\end{tabular} & \[
\begin{gathered}
\text { Binary (Dec) } \\
\text { After } \\
\hline
\end{gathered}
\] \\
\hline Ex. 1: & \[
\begin{gathered}
4 \\
\text { (immediate) }
\end{gathered}
\] & \[
\begin{array}{r}
00000100 \\
(+4)
\end{array}
\] & -------- \\
\hline & \[
\begin{gathered}
\text { R5 } \\
\text { (low byte) }
\end{gathered}
\] & 00001111 & 11110000 \\
\hline Ex. 2: & \[
\begin{gathered}
-3 \\
\text { (immediate) }
\end{gathered}
\] & \[
\begin{array}{r}
11111101 \\
(-3)
\end{array}
\] & -------- \\
\hline & 16 (SP) & 00001111 & 11100001 \\
\hline
\end{tabular}

\section*{ROUNDfi}

Round Floating to Integer

\begin{tabular}{ll} 
ROUNDFB & ROUNDLB \\
ROUNDFW & ROUNDLW \\
ROUNDFD & ROUNDLD
\end{tabular}

The Round Floating to Integer instruction rounds the src operand to the nearest integer and places the result in the dest operand location as a signed integer. If src is exactly halfway between two integer values, it is rounded to the even value (i.e. the value that is exactly divisible by two).

Flags Affected: No PSR flags. FSR flags are affected as follows:
IF is set on an inexact result; unaffected otherwise.
TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes.
See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3. Particularly relevant to this instruction is the Overflow exception, which is caused by attempting to convert a floating-point number that is too great in absolute value to be held in a signed integer of the size specified for dest.

\section*{Round Floating to Integer (continued)}

\section*{Examples:}
\begin{tabular}{llll} 
1. & ROUNDFB & F 0, & R 0 \\
2. & ROUNDLD & F 2, & \(12(\mathrm{SB})\)
\end{tabular}

3E 2400
2. ROUNDLD \(\mathrm{F} 2,12\) (SB)

3E A3 16 0C

Example 1 rounds the single-precision number in register \(F 0\) to a 1 -byte integer and places the result in the low-order byte of register R0. The remaining bytes of RO are unaffected.

Example 2 rounds the double-precision number in register F 2 to a double-word integer and places the result at address \(12(\mathrm{SB})\).
\begin{tabular}{|c|c|c|c|}
\hline & Operands & Operand Va Before & s: Hex (Dec) After \\
\hline \multirow[t]{2}{*}{Ex. 1:} & F0 & \[
\begin{aligned}
& 40180000 \\
& (+2.375)
\end{aligned}
\] & \[
\begin{aligned}
& 40180000 \\
& (+2.375)
\end{aligned}
\] \\
\hline & R0 & AAAAAAAA & AAAAAA02 (+2) \\
\hline \multirow[t]{2}{*}{Ex. 2:} & F2 & \[
\begin{array}{r}
41 \mathrm{C0200888700000} \\
(+541069584.875)
\end{array}
\] & \[
\begin{array}{r}
41 \mathrm{C0200888700000} \\
(+541069584.875)
\end{array}
\] \\
\hline & 12 (SB) & AAAAAAAA & \[
\begin{gathered}
20401111 \\
(+541069585)
\end{gathered}
\] \\
\hline
\end{tabular}

\section*{RXP}

Return from External Procedure

\section*{Syntax: RXP constant}
disp
\begin{tabular}{|c|c|}
\hline ! & RXP \\
\hline & 10010 ! \\
\hline & -+-+-+-+-1 \\
\hline 7 & 0 \\
\hline
\end{tabular}

The RXP instruction returns control from an externally-called procedure and removes any procedure parameters from the stack.

The instruction does the following:
1. Pops a 32-bit return address from the currently-selected stack into the PC register.
2. Pops a 16-bit MOD address from the currently-selected stack to the MOD register and increments the stack pointer by two. The stack pointer is modified by a total of four in this step.
3. Copies the double-word at address MOD+0 to the SB register.
4. Adds the constant operand to the current stack pointer.

Program execution continues at the new address placed in the PC register.
Flags Affected: None.
Traps: None.

\section*{Return from External Procedure (continued)}

\section*{Example:}

RXP 16
3210

This example returns control from an externally-called procedure and removes 16 (H'10) bytes from the currently-selected stack.


Module
Table:
0000001000009080 (SB) 00009080 (SB)
0000001400002000 (LB) 00002000 (LB)

0000001800009000 (PB) 00009000 (PB)
* The final Stack Pointer content is the initial address plus 4 (for double-word return address), 2 (for word MOD address), 2 (additional from step 3), and 16 as specified by the constant operand.
** The RXP instruction does not itself change the contents of these memory locations. However, information that is outside the stack should be considered unpredictable for other reasons. See Section 2.8.1.

\section*{Scondi}

Save Condition as Boolean

\section*{Syntax: Scondi}
```

dest
gen
write.i

```

ScondB
ScondW
ScondD

The Scondi instruction sets the dest operand to the integer value "1" if the specified condition is true, and to "0" if false. These are the Boolean values "True" and "False", respectively.

Cond is a 2-character condition name that specifies the state of a flag or flags in the PSR. If the flag(s) have the specified state, the condition is true; otherwise, the condition is false.

The Save Condition as Boolean instruction may specify the following conditions:
\begin{tabular}{llll} 
Condition & \begin{tabular}{c} 
Condition \\
Name
\end{tabular} & \begin{tabular}{l} 
True \\
State
\end{tabular} & \begin{tabular}{c} 
Short \\
Field
\end{tabular} \\
\hline Equal & EQ & Z flag set & \\
Not Equal & NE & Z flag clear & 0000 \\
Carry Set & CS & C flag set & 0001 \\
Carry Clear & CC & C flag clear & 0010 \\
Higher & HI & L flag set & 0100 \\
Lower or Same & LS & L flag clear & 0101 \\
Greater Than & GT & N flag set & 0110 \\
Less Than or Equal & LE & N flag clear & 0111 \\
Flag Set & FS & F flag set & 1000 \\
Flag Clear & FC & F flag clear & 1001 \\
Lower & LO & Z and L flag clear & 1010 \\
Higher or Same & HS & Z or L flag set & 1011 \\
Less Than & Z and N flag clear & 1100 \\
Greater than or Equal & GE & Z or N flag set & 1101
\end{tabular}

The condition name must be embedded in the instruction mnemonic as illustrated in the examples below. The name is translated at assembly time to the corresponding Short Field of the basic instruction.

The interpretation of condition codes is such that the instruction sequence
\[
\begin{array}{ll}
\text { CMPB } & \text { A, B } \\
\text { SGTB } & \text { RESULT }
\end{array}
\]
will store "1" in RESULT if operand A is greater than operand B in the CMPB instruction.

\section*{Save Condition as Boolean (continued)}

Flags Affected: None.
Traps: None.

\section*{Examples:}
\begin{tabular}{llllll} 
1. & SEQB & R0 & 3C & 00 & \\
2. & SLOW & \(10(S B)\) & 3D & D5 & \(0 A\) \\
3. & SHID & TOS & \(3 F\) & BA &
\end{tabular}

Example 1 sets the low-order byte of register R0 to 1 if the \(Z\) flag is set, 0 if the Z flag is clear.

Example 2 sets the word at the operand address \(10(S B)\) to 1 if the \(Z\) and L flags are clear, 0 otherwise.

Example 3 pushes a double-word value onto the stack: 1 if the \(L\) flag is set, 0 otherwise.

In the following illustration, the \(Z\) and \(L\) flags are assumed to be set.
\begin{tabular}{|c|c|c|c|}
\hline & Operands & Operand Values: Before & \begin{tabular}{l}
Hex (Boolean) \\
After
\end{tabular} \\
\hline \multirow[t]{2}{*}{Ex. 1:} & R0 & AAAAAAAA & AAAAAA01 (True) \\
\hline & UPSR & n1fvx1tc & \(n 1 f v x 1 t c\) \\
\hline \multirow[t]{2}{*}{Ex. 2:} & 10 (SB) & AAAA & \[
\begin{gathered}
0000 \\
\text { (False) }
\end{gathered}
\] \\
\hline & UPSR & n1fvx1tc & n1fvx1tc \\
\hline \multirow[t]{3}{*}{Ex. 3:} & \begin{tabular}{l}
Stack: \\
00001000 \\
00001004
\end{tabular} & \[
\begin{aligned}
& \text { XXXXXXXX } \\
& \text { AAAAAAAA }
\end{aligned}
\] & 00000001 (True) AAAAAAAA \\
\hline & SP & 00001004 & 00001000 \\
\hline & UPSR & n1fvx1tc & \(n 1\) fvx1tc \\
\hline
\end{tabular}

\section*{SAVE}

Save General Purpose Registers

\section*{Syntax: SAVE reglist}
imm
\begin{tabular}{|c|c|}
\hline ! & SAVE \\
\hline & 00010 ! \\
\hline & +-+-+-+-+-! \\
\hline 7 & \\
\hline
\end{tabular}

The SAVE instruction saves the General-Purpose registers specified by reglist, pushing them onto the currently-selected stack.

The reglist operand is a list of zero or more general purpose register names, enclosed in brackets "[]". The instruction pushes the contents of each register in the list as a double-word onto the stack. Register names may appear in any order within reglist, but must be separated by commas. Brackets are required even if no register names are given.

In the machine instruction, the reglist operand is encoded in an 8-bit field as shown below. Each bit in the field corresponds to one general purpose register. When the instruction is executed, the instruction reads the bits in the field from right to left beginning with bit 0 . If a bit is "0", the instruction ignores the corresponding register. If a bit is "1", it pushes the corresponding register.
```

+------------------------+
!R7!R6!R5!R41R3!R2!R1!R0!
!--+--+---+--+---+--+--+---!
7 0

```

Flags Affected: None.

Traps: None.

\section*{Example:}

SAVE [R0, R2, R7] 6285

This instruction saves the contents of registers R0, R2, and \(R 7\) on the stack. The registers are stored in order beginning with register \(R 0\) and ending with \(R 7\).
\begin{tabular}{|c|c|c|}
\hline Operands & \begin{tabular}{l}
Operand Values: \\
Before
\end{tabular} & Hex After \\
\hline R0 & 00000010 & 00000010 \\
\hline R2 & FFFFFFEF & FFFFFFEF \\
\hline R7 & FFFFF9AB & FFFFF9AB \\
\hline SP & 0000FFEC & 0000 FFE 0 \\
\hline \multicolumn{3}{|l|}{Stack:} \\
\hline 0000FFE0 & xxxxxxxx & FFFFF9AB \\
\hline 0000 FFE 4 & mxxxxxmx & FFFFFFEF \\
\hline 0000 FFE 8 & SXXXXXXX & 00000010 \\
\hline 0000 FFEC & AAAAAAA & AAAAAAAA \\
\hline
\end{tabular}

SBITi
SBITIi
Set Bit, Set Bit Interlocked
Syntax: SBITi


Syntax: SBITIi offset, base
gen gen
read.i regaddr


The SBITi and SBITIi instructions set to 1 the register or memory bit specified by base and offset after copying the bit value to the \(F\) flag in the PSR.

The SBITIB, SBITIW, and SBITID instructions, in addition, activate the Interlocked Operation output pin on the CPU, which may be used in multiprocessor systems to interlock accesses to semaphore bits. See the applicable CPU data sheet for further details.

The location of the bit is determined from offset and base. Offset is a general operand, whose length is given by the operation length suffix. Base is an addressing expression giving a byte address from which offset specifies a bit position. See Section 3.5 for details of specifying bit positions.

If base is a register, then the bit is within that register, at the bit position given by the offset operand. If base is a memory location, then the bit is at bit position
```

Offset MOD 8

```
within the memory byte whose address is
EA(base) + (offset DIV 8),
where EA(base) is the effective address of base. See Section 3.5 for definitions of the operators MOD and DIV above, and for further details of bit instructions.

Offset is interpreted as a signed integer.

\section*{Set Bit (continued)}

Flags Affected: \(F\) is set to the original value of the specified bit.
Traps: None.

\section*{Example:}
```

SBITW R0, 1(R1)
4E 59 02 01

```

This example sets a bit in memory to 1 after copying the bit value to the \(F\) flag. This performs the basic operation of a semaphore "Test and Set". In a multiprocessor system, the SBITIW instruction would have been used instead. For designating the location of the target bit, the low-order word of register R0 supplies the bit offset, and \(1(R 1)\) is specified as the base address.

In the following illustration, the target bit is assumed to be 0 prior to instruction execution.
\begin{tabular}{|c|c|c|}
\hline Operands & Operand Values: Before & Hex (Dec) [Binary]
\(\qquad\) After \\
\hline \[
\begin{gathered}
\mathrm{R0} \\
(\mathrm{offset})
\end{gathered}
\] & \[
\begin{array}{r}
\text { AAAA004C } \\
(+76)
\end{array}
\] & \[
\begin{aligned}
& \text { AAAA004C } \\
& (+76)
\end{aligned}
\] \\
\hline R1 & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] \\
\hline base address 1 (R1) & \[
\begin{aligned}
& 00001001 \\
& \quad(+4097)
\end{aligned}
\] & -- \\
\hline \[
\begin{gathered}
0000100 \mathrm{~A} \text { * } \\
(+4106)
\end{gathered}
\] & \[
\begin{gathered}
\text { EF } \\
{[11101111]}
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{FF} \\
{[11111111]}
\end{gathered}
\] \\
\hline UPSR & \(n z f v x l t c\) & \(n z \underline{0} v x l t c\) \\
\hline
\end{tabular}
* The address 100A (Hex) is the effective address of the byte containing the desired bit. This address is computed from the offset and the base address as follows:
```

base address + (offset DIV 8)
4097 + 9
4106, or 100A (Hex) .

```

The bit number within this byte is calculated as:
```

offset MOD 8
76 MOD 8
4.

```

\section*{SCALBf}

Scale Binary Floating

\section*{Syntax:}
\begin{tabular}{lll} 
SCALBf & src, \\
& gen & dest \\
& gen
\end{tabular}

SCALBF
gen gen
read.f rmw.f


The SCALBf instruction multiplies the dest operand by 2 raised to the power of src. Only the integral part of src is used.

Flags Affected: No PSR flags.
The FSR TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes. See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap(UND) is activated if the \(F\) bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3.

\section*{Example:}

SCALBF F3, F2
FE 9018
This example multiplies the floating-point number in \(F 2\) by the unbiased exponent of F 3 raised to the power of 2 .

The instruction is illustrated below:
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: Hex \begin{tabular}{c}
\((\) Dec \()\) \\
Before
\end{tabular} \\
\hline
\end{tabular}\(\quad\)\begin{tabular}{cc} 
After
\end{tabular} \\
F2 & C2250000 \\
& \((-41.25)\) & \((-5406720.00)\) \\
F3 & \(418 C 0000\) & \(418 c 0000\) \\
& \((+17.50)\) & \((+17.50)\)
\end{tabular}

\section*{Syntax: SETCFG cfglist}

\author{
short
}


The SETCFG instruction loads the Configuration Register (CFG), enabling or disabling optional system features.

In assembly language, cfglist is a list of zero or more configuration bit names. The names are I, \(F\), \(M\) and \(C\). The names may appear in any order in the list, but must be separated by commas. The list itself must be enclosed in brackets. Brackets are required even if no bit name is given.

The cfglist operand is held in a 4-bit field in the basic instruction, as shown above. Each bit corresponds to one bit in the CFG register.

If I is specified, the I bit in the CFG register is set and vectored interrupt processing is enabled. Otherwise, the \(I\) bit is cleared and all maskable interrupts are serviced as non-vectored interrupts.

If \(F\) is specified, the \(F\) bit in the \(C F G\) register is set and Floating Point instructions are available. Otherwise, the \(F\) bit is cleared and Floating Point instructions activate the Undefined Instruction Trap (UND).

If \(M\) is specified, the \(M\) bit in the \(C F G\) register is set and Memory Management instructions are available. Otherwise, the M bit is cleared and Memory Management instructions activate the Undefined Instruction Trap (UND).

If C is specified, the C bit in the CFG register is set and Custom Slave instructions are available. (System-dependent Custom Slave hardware must be present to execute the instructions.) Otherwise, the \(C\) bit is cleared and Custom Slave instructions activate the Undefined Instruction Trap (UND).

See Section 2.3 for further information about the CFG register.

NOTE: The SETCFG instruction is not recommended for new systems. Please use the LPRi instruction instead.

NOTES: 1. A CFG bit name may only be specified if the corresponding device is present in the system.
2. The state of the \(M\) bit in the \(C F G\) register does not directly affect address translation hardware or bus timing. It only enables or disables the Memory Management instruction set.
3. When a Floating-Point, Memory Management, or Custom Slave instruction activates an Undefined Instruction Trap (UND) (i.e, when the corresponding CFG register bit is 0), it is possible to intercept the trap and simulate the instruction in software.

Flags Affected: None.
Traps: Illegal Instruction Trap (ILL) is activated if this instruction is attempted while the \(U\) flag is set.

\section*{Example:}

SETCFG [I,M,F] OE 8B 03

This instruction sets the \(I\), \(M\), and \(F\) bits in the CFG register, enabling vectored interrupt processing and the Memory Management and Floating Point instruction sets. The C bit is cleared, disabling the Custom Slave instruction set. The bits 8 through 13 are cleared. The bits 4 through 7 are always "1".

\section*{Bit Values}
\begin{tabular}{ccc} 
Operands & Before & After \\
\hline CFG & xxxxxx1111cmfi & \(\underline{00000011110111}\)
\end{tabular}

\section*{SFSR}

Store Floating-Point Status Register (FSR)

\section*{Syntax: \\ SFSR \\ dest}
gen
write.D


The SFSR instruction copies the contents of the Floating-Point Status Register (FSR) to the dest operand location. The FSR is treated as a 32-bit value. See Section 2.4.2 for the format of the FSR.

Flags Affected: None.
Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

\section*{Example:}
```

SFSR TOS
3E F7 05

```

This example pushes the contents of the FSR onto the top of the currently selected stack as a double word.
\begin{tabular}{clc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex \\
After
\end{tabular} \\
\hline FSR & xxx00028 & xxx00028 \\
SP & 0000 FFDE & 0000 FFDA \\
Stack: & & \\
0000 FFDA & XXXXXXXX & XXx00028 \\
0000 FFDE & AAAAAAAA & AAAAAAAA
\end{tabular}

\section*{Skip String}

\section*{Syntax:}

\section*{Syntax: SKPST options}


Operands of the SKPSi and SKPST instructions are specified in General-Purpose registers:
```

R0 - Number of string elements to be processed.
R1 - Address of current String 1 element.
R2 - (not used)
R3 - Address of translation table (SKPST form only).
R4 - Match value (with Until Match or While Match option only).

```

The SKPSi instruction examines and skips over consecutive elements in String 1 until either an Until/While condition is met or register RO is decremented to 0 (i.e., the string is exhausted). After each element is examined, the CPU sets register \(R 1\) to the address of the next element to be examined and register R0 to the number of integers remaining to be examined. Register R2 is not used or affected.

The SKPST instruction causes the CPU to internally replace the current String 1 element value with a corresponding translated value before performing its examination. The translated value to be examined is found by adding the current element from String 1 as an unsigned integer to the translation table address found in register R3. The instruction examines elements and sets registers as described above. The SKPST instruction operates on byte-long elements only.

Options may be specified by listing the letters B (Backward), U (Until Match) and \(W\) (While Match) as operands. The \(U\) and \(W\) options are mutually exclusive. See Section 3.7 for details of the options available in String instructions.

In the machine instruction, the options are encoded in the \(B\) and UW fields as follows:
```

    B field = 0 Forward direction.
    1 Backward direction.
    UW field = 00 Neither Until Match nor While Match.
0 1 ~ W h i l e ~ M a t c h .
10 (reserved)
11 Until Match.

```
String instructions are interruptible. See Section 3.7.

Flags Affected: \(F\) is set if the \(U\) or \(W\) option is specified and the corresponding Until/While condition is met, otherwise it is cleared.

Traps: None.

\section*{Example:}
```

SKPSB U
OE OC 06

```

This example examines and skips over byte-long elements in String 1 until the current integer and the contents of the low-order byte of register R4 are equal or until register RO contains zero.

In the following illustration, the underlined string element shows the point at which the instruction terminates.
\begin{tabular}{ccr} 
Operands & \begin{tabular}{c} 
Operand Values: Hex \begin{tabular}{c} 
(Dec) \\
Aefore
\end{tabular} \\
R0
\end{tabular} \begin{tabular}{r}
00000020 \\
\((+32)\)
\end{tabular} & \begin{tabular}{r}
00000016 \\
\((+22)\)
\end{tabular} \\
R1 & 00002000 & 0000200 A \\
R4 & \begin{tabular}{r} 
AAAAAA1F \\
\((+31)\)
\end{tabular} & \begin{tabular}{r} 
AAAAAA1F \\
\((+31)\)
\end{tabular} \\
UPSR & nzfvxltc & nz1vxltv
\end{tabular}
Starting String Contents

Address
```

2000 1E 04 05 1C 0A 14 0C 0B 09 07 1F 0F 17 01 00 11
1F 1D 1A 09 01 12 14 0E 1E 0A 00 03 09 06 16 18

```

\section*{Store Memory Management Register}

\section*{Syntax:}

\section*{SMR mmureg, dest}

\section*{short gen}
write.D


The SMR instruction copies the contents of the Memory Management register specified by mmureg to the dest operand location.

The Store MMU Register instruction may store the following registers. The short field of the basic instruction holds a 4-bit value which relates to the corresponding mmureg specifications as follows:
\begin{tabular}{lll} 
Register & mmureq & \begin{tabular}{l} 
short \\
field
\end{tabular} \\
Memory Management Control Register & MCR & 1000 \\
Memory Management Status Register & MSR & 1010 \\
Translation Exception Address Reg. & TEAR & 1011 \\
Page Table Base Register 0 & PTB0 & 1100 \\
Page Table Base Register 1 & PTB1 & 1101 \\
Invalidate Virtual Address 0 & IVAR0 & \(1110 *\) \\
Invalidate Virtual Address 1 & IVAR1 & \(1111 *\)
\end{tabular}
* This register are write-only. If they are read the result is undefined.

Flags Affected: None.

Traps: Undefined Instruction Trap (UND) is activated if the M bit in the CFG register is clear. The instruction is not executed.

Illegal Instruction Trap (ILL) is activated if the \(u\) flag is set. The instruction is not executed.

SMR
Store MMU Register (continued)
Example:
SMR PTB0, R0 1E 0F 06
This example copies the contents of the Page Table Register 0 in the MMU to register R0.
\begin{tabular}{clc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex \\
After
\end{tabular} \\
\hline PTB0 & 00009000 & 00009000 \\
R0 & AAAAAAAA & 00009000
\end{tabular}

\section*{Syntax:}

\section*{SPRi procreg, dest}
short gen
write.i
\begin{tabular}{|c|c|c|}
\hline dest & ! procreg! & SPRi \\
\hline gen & ! short ! 0 & 0111 i \\
\hline 15 & 87 & 0 \\
\hline
\end{tabular}

The SPRi instruction stores the contents of the dedicated processor register specified by procreg in the dest operand location.

The register contents are right-justified in dest. High-order dest bits are zero-filled if the register is shorter than dest. High-order register bits are discarded if the register is longer than dest. See Section 2.2 for the formats of the dedicated registers.

The Store Processor Register instruction may store the following registers. The specified procreg corresponds to the 4-bit short field in the basic instruction as shown below.
\begin{tabular}{llll} 
Register & procreg & \begin{tabular}{l} 
short \\
field
\end{tabular} & \\
& & & \\
User PSR & UPSR & 0000 & (Note 3) \\
Debug Condition Register & DCR & 0001 & (Note 2) \\
Breakpoint Program Counter & BPC & 0010 & (Note 2) \\
Debug Status Register & DSR & 0011 & (Note 2) \\
Compare Address Register & CAR & 0100 & (Note 2) \\
Frame Pointer & FP & 1000 & \\
Stack Pointer & SP & 1001 & (Note 1) \\
Static Base Register & SB & 1010 & \\
User Stack Pointer (SP1) & USP & 1011 & (Note 2) \\
Configuration Register & CFG & 1100 & (Note 2) \\
Processor Status Register & PSR & 1101 & (Note 2) \\
Interrupt Base Register & INTBASE & 1110 & (Note 2) \\
Module Register & MOD & 1111 &
\end{tabular}

NOTES: 1. If SP is specified in the instruction and the \(S\) flag in the PSR is set, the instruction copies the SP1 register to the dest operand location. If the \(S\) flag is clear, the instruction copies the SPO register to the dest operand location.
2. Specifying this register as the procreg operand is privileged, regardless of the operation length specified.
3. If UPSR is specified only the low-order eight bits of the PSR are stored, regardless of the operation length specified. It is zero-extended as necessary to fill the dest operand.

Flags Affected: None.
Traps: Illegal Instruction Trap (ILL) is activated if the U flag is set and PSR, INTBASE, USP, CFG or a Debug Register is specified.

\section*{SPRi}

\section*{Store Processor Register (continued)}

\section*{Examples:}
1. SPRD FP, R0
2F 04
2. SPRW MOD, 4 (SB)
AD D7 04

Example 1 copies the entire \(F P\) register to register \(R 0\).

Example 2 copies the contents of the MOD register to a word at the address specified by 4 (SB).
\begin{tabular}{ccc} 
& \multicolumn{2}{c}{\begin{tabular}{c} 
Operand Values: Hex \\
Operands \\
Ex. 1: After
\end{tabular}} \\
\cline { 2 - 3 } & FP & 00000021
\end{tabular}

\section*{Subtract Floating}

\section*{Syntax:}


The SUBf instruction subtracts the src operand from the dest operand and places the result in the dest operand location. Subtraction can be modelled as negating the src operand and adding the result to the dest operand. For details of the addition step see the ADDf instruction.

Flags Affected: No PSR flags. FSR flags are affected as follows:
UF is set if an underflow occurs; unaffected otherwise.
IF is set on an inexact result; unaffected otherwise.
TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes.
See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3.

\section*{SUBf}

Subtract Floating (continued)

\section*{Examples:}
1. SUBF F0, F7
BE D1 01
2. SUBL F2, 16 (SB)
BE 901610

Example 1 subtracts the single-precision number in register \(F 0\) from the singleprecision number in register \(F 7\) and places the result in register F 7 .

Example 2 subtracts the double-precision number in register F2 from the doubleprecision number at the address \(16(S B)\) and places the double-precision result at the address 16(SB).
\begin{tabular}{|c|c|c|c|}
\hline & Operands & Operand V Before & \begin{tabular}{l}
Hex (Dec) \\
After
\end{tabular} \\
\hline \multirow[t]{2}{*}{Ex. 1:} & F0 & \[
\begin{aligned}
& 40840000 \\
& (+4.125)
\end{aligned}
\] & \[
\begin{array}{r}
40840000 \\
(+4.125)
\end{array}
\] \\
\hline & F7 & \[
\begin{aligned}
& 41 \text { F50000 } \\
& (+30.625)
\end{aligned}
\] & \[
\begin{aligned}
& 41 D 40000 \\
& (+26.500)
\end{aligned}
\] \\
\hline \multirow[t]{2}{*}{Ex. 2:} & F2 & \[
\begin{gathered}
4114 C 86300000000 \\
(+340504.750)
\end{gathered}
\] & \[
\begin{gathered}
4114 C 86300000000 \\
(+340504.750)
\end{gathered}
\] \\
\hline & 16 (SB) & \[
\begin{aligned}
& \text { 41C022A194900000 } \\
& (+541410089.125)
\end{aligned}
\] & \[
\begin{aligned}
& 41 \mathrm{C0200888300000} \\
& (+541069584.375)
\end{aligned}
\] \\
\hline
\end{tabular}

\section*{Subtract}

\section*{Syntax:}
\begin{tabular}{llllll} 
SUBi src, & dest \\
gen & gen \\
read.i & rmw.i
\end{tabular}

The SUBi instruction subtracts the src operand from the dest operand and places the result in the dest operand location.

Flags Affected: \(C\) is set on a borrow from subtraction, cleared if no borrow. \(F\) is set on an overflow from subtraction, cleared if no overflow.

Integer carry and overflow conditions are defined in Section 3.1.

Traps:
Integer Overflow Trap (OVF) is activated if the \(V\) flag is set and the result cannot be represented exactly in dest.

\section*{SUBi}

Subtract (continued)

\section*{Examples:}
1. \(\operatorname{SUBB}\) R0, R1
2. SUBD \(4(S B), 20(S B)\)

6000
2. SUBD \(4(\mathrm{SB}), 20(\mathrm{SB})\)

A3 D6 0414

Example 1 subtracts the low-order byte of register R0 from the low-order byte of register R1 and places the result in the low-order byte of register R1. The remaining bytes of R1 are not affected.

Example 2 subtracts the double-word at the memory address specified by 4 (SB) from the double-word at the memory address specified by \(20(\mathrm{SB})\). The instruction places the result at memory address \(20(\mathrm{SB})\).
\begin{tabular}{|c|c|c|c|}
\hline & Operands & Operand Values Before & \[
\begin{aligned}
& \text { (Dec) } \\
& \text { After } \\
& \hline
\end{aligned}
\] \\
\hline \multirow[t]{3}{*}{Ex. 1:} & R0 & AAAAAA01 (+1) & AAAAAA01 (+1) \\
\hline & R1 & \[
\begin{aligned}
& \text { BBBBBB7F } \\
& (+127)
\end{aligned}
\] & \[
\begin{aligned}
& \text { BBBBBB7E } \\
& (+126)
\end{aligned}
\] \\
\hline & UPSR & \(n z f v x l t c\) & \(n z \underline{0} \mathrm{vxlt} \underline{1}\) \\
\hline \multirow[t]{3}{*}{Ex. 2:} & 4 (SB) & FFFFFFFE
\[
(-2)
\] & FFFFFFFE
\[
(-2)
\] \\
\hline & 20 (SB) & \[
\begin{aligned}
& 00010000 \\
& (+65536)
\end{aligned}
\] & \[
\begin{aligned}
& 00010002 \\
& (+65538)
\end{aligned}
\] \\
\hline & UPSR & \(n z f v x l t c\) & \(n z \underline{0} \mathrm{vxlt} \underline{1}\) \\
\hline
\end{tabular}

\section*{Subtract with Carry [Borrow]}
```

Syntax: SUBCi src, dest
gen gen
read.i rmw.i
SUBCD

```


The SUBCi instruction subtracts the sum of the src operand and the \(C\) flag from the dest operand and places the result in the dest operand location.

Flags Affected: C is set on a borrow from subtraction, cleared if no borrow. \(F\) is set on an overflow from subtraction, cleared if no overflow.

Integer carry and overflow conditions are defined in Section 3.1.

Traps: Integer Overflow Trap (OVF) is activated if the V flag is set and the result cannot be represented exactly in dest.

SUBCi
Subtract with Carry [Borrow] (continued)

\section*{Examples:}
1. SUBCB 32, R1
70 A0 20
2. SUBCW TOS, -8 (FP)
31 BE 78

Example 1 subtracts the sum of 32 and the \(C\) flag value from the low-order byte of register R1 and places the result in the low-order byte of register R1. The remaining bytes of R1 are not affected.

Example 2 subtracts the sum of the word at the top of the stack and the \(C\) flag value from the word at the memory address specified by -8(FP). The instruction then places the 2-byte result at the memory address specified as -8(FP).

In the following illustration, the \(C\) flag value is assumed to be 1.

* The instruction has not itself changed the contents of these memory locations. However, information that is outside the stack should be considered unpredictable for other reasons. See Section 2.8.1.

\section*{Subtract Packed Decimal}

\section*{Syntax:}


The SUBPi instruction subtracts the src operand from the dest operand and then subtracts the C flag. The instruction places the result in the dest operand location as a packed decimal (BCD) integer.

The src and dest operands are interpreted as unsigned packed decimal (BCD) integers. If either operand contains invalid digits, the result is undefined. See Section 3.2 for details of packed decimal arithmetic.

Flags Affected: C is set on a borrow from subtraction, cleared if no borrow. F is cleared.

Traps: None.

\section*{SUBPi}

\section*{Subtract Packed Decimal (continued)}

\section*{Examples:}
```

1. SUBPD H'99, R1 4E 6F A0 00 00 00 99
2. SUBPB -8(FP), 16(FP)
4E 2C C6 78 10
```

Example 1 subtracts the packed decimal integer 99 from the contents of register R1 and then subtracts the \(C\) flag. The result is placed in register R1.

Example 2 subtracts the packed decimal integer at memory address -8(FP) from the packed decimal integer at memory address \(16(\mathrm{FP})\) and then subtracts the C Flag. The instruction places the result at memory address 16(FP).

In the following illustration, the \(C\) flag value is assumed to be 0 .
\begin{tabular}{|c|c|c|c|}
\hline & Operands & \begin{tabular}{l}
Operan \\
Before
\end{tabular} & \[
\begin{aligned}
& \text { Hex }{ }^{*} \\
& \text { After } \\
& \hline
\end{aligned}
\] \\
\hline \multirow[t]{3}{*}{Ex. 1:} & \[
\begin{gathered}
\text { H'99 } \\
\text { (immediate) }
\end{gathered}
\] & 00000099 & -- \\
\hline & R1 & 00000187 & 00000088 \\
\hline & UPSR & \(n z f v x l t 0\) & \(n z \underline{0} \mathrm{vxlt} \underline{0}\) \\
\hline \multirow[t]{3}{*}{Ex. 2:} & -8(FP) & 10 & 10 \\
\hline & 16 (FP) & 01 & 91 ** \\
\hline & UPSR & \(n z f v x l t 0\) & \(n z \underline{0} \mathrm{vxlt} \underline{1}\) \\
\hline \multicolumn{4}{|l|}{* The hexadecimal representation also expresses the decimal the value.} \\
\hline
\end{tabular}

\section*{Supervisor Call}

\section*{Syntax: SVC}


The SVC instruction activates the Supervisor Call Trap (SVC). The Supervisor Call Trap passes program execution control to the SVC service procedure. The return address pushed onto the Interrupt Stack is the address of the SVC instruction itself.

Flags Affected: None.
Traps: Supervisor Call Trap (SVC) is activated.

\section*{Example:}

SVC
E2

TBITi
Test Bit

\section*{Syntax:}


The TBITi instruction copies the register or memory bit specified by base and offset to the F flag.

The location of the bit is determined from offset and base. Offset is a general operand, whose length is given by the operation length suffix. Base is an addressing expression giving a byte address from which offset specifies a bit position. See Section 3.5 for details of specifying bit positions.

If base is a register, then the bit is within that register, at the bit position given by the offset operand. If base is a memory location, then the bit is at bit position
offset MOD 8
within the memory byte whose address is
EA (base) + (offset DIV 8),
where EA(base) is the effective address of base. See Section 3.5 for definitions of the operators MOD and DIV above, and for further details of bit instructions.

Offset is interpreted as a signed integer.
Flags Affected: \(F\) is set to the value of the specified bit.
Traps: None.

\section*{Example:}
```

TBITW R0, 0(R1)
750200

```

This example copies a bit from memory to the \(F\) flag. The low-order word of register RO supplies the bit offset, and \(0(R 1)\) is specified as the base address.

In the following illustration, the target bit is assumed to be 1.
\begin{tabular}{|c|c|c|}
\hline Operands & Operand Values: Before & \[
\begin{gathered}
\text { Hex (Dec) } \begin{array}{c}
\text { [Binary] } \\
\text { After }
\end{array} \\
\hline
\end{gathered}
\] \\
\hline \[
\begin{gathered}
\mathrm{R} 0 \\
\text { (offset) }
\end{gathered}
\] & \[
\begin{aligned}
& \text { AAAA004C } \\
& (+76)
\end{aligned}
\] & \[
\begin{aligned}
& \text { AAAA } 004 \mathrm{C} \\
& (+76)
\end{aligned}
\] \\
\hline R1 & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] \\
\hline base address 0 (R1) & \[
\begin{aligned}
& 00001000 \\
& \quad(+4096)
\end{aligned}
\] & -- \\
\hline \[
\begin{gathered}
00001009 \text { * } \\
(+4105)
\end{gathered}
\] & \[
\begin{gathered}
10 \\
{[00010000]}
\end{gathered}
\] & \[
\begin{gathered}
10 \\
{[00010000]}
\end{gathered}
\] \\
\hline UPSR & \(n z f v x l t c\) & \(n z 1 v x l t v\) \\
\hline
\end{tabular}
* The address 1009 (Hex) is the effective address of the byte containing the desired bit. This address is computed from the offset and the base address as follows:
```

    base address + (offset DIV 8)
    4096 + 9
    4105, or 1009 (Hex) .
    ```

The bit number within this byte is calculated as: offset MOD 8

76 MOD 8
4 .

\section*{Syntax: TRUNCfi src, dest \\ gen gen}
read.f write.i
TRUNCFB TRUNCLB
TRUNCFW TRUNCLW
TRUNCFD TRUNCLD


The TRUNCfi instruction truncates the src operand to the nearest integer which is less than or equal to it in absolute value and places the result in the dest operand location as a signed integer.

Flags Affected: No PSR flags. FSR flags are affected as follows: IF is set on an inexact result; unaffected otherwise. TT field is set to reflect any exceptional conditions encountered in executing the instruction. If none is encountered, TT is set to all zeroes.
See Sections 2.4.2 and 3.3 for details of exceptional conditions and reporting.

Traps: Undefined Instruction Trap (UND) is activated if the F bit in the CFG register is clear.

Floating-Point Trap (FPU) is activated if a floating-point exception is detected. See Section 3.3. Particularly relevant to this instruction is the Overflow exception, which is caused by attempting to convert a floating-point number which is too great in absolute value to be held in a signed integer of the size specified for dest.

\section*{Examples:}
1. TRUNCFB F0, R0
3E 2C 00
2. TRUNCLD F2, 8 (SB)
3E AB 1608

Example 1 truncates the single-precision number in register \(F 0\) to a one-byte integer and copies the integer to the low-order byte of register R0.

Example 2 truncates the double-precision number in register \(F 2\) to a double-word integer and copies the integer to address 8(SB).
\begin{tabular}{ccc} 
Operands & \begin{tabular}{c} 
Operand Values: \\
Before
\end{tabular} & \begin{tabular}{c} 
Hex (Dec) \\
After
\end{tabular} \\
\hline F0 & C0280000 \\
\((-2.625)\) & C0280000 \\
R0 & AAAAAAAA & AAAAAAFE \\
& & \((-2)\)
\end{tabular}

Ex. 2:

\section*{F2}
\(41 C 0200888700000\)
\((+541069584.875\)
41C0200888700000 \((+541069584.875) \quad(+541069584.875)\)

8 (SB)
AAAAAAA
20401110
(+541069584)

WAIT
Wait

\section*{Syntax: WAIT}


The WAIT instruction suspends program execution until an interrupt occurs. An interrupt restores program execution by passing it to an interrupt service procedure. When the WAIT instruction is interrupted, the return address saved is the address of the instruction following the WAIT instruction.

Flag Affected: None.
Traps: None.

\section*{Example:}
WAIT
B2

\section*{Syntax: WRVAL loc}
gen
addr


The WRVAL instruction checks the protection level assigned to the user-mode virtual memory address specified as loc. If the address can be written to while in user mode, the \(F\) flag in the \(P S R\) is cleared. If the address cannot be written to (i.e., if loc is write-protected), the \(F\) flag in the PSR is set. See Section 3.12 for details of Memory Management instructions.

NOTE: Although the final effective address of loc is interpreted as a user-mode virtual address, any memory references required in order to calculate that effective address are interpreted as using supervisor-mode addresses. This will occur in using the Memory Relative and External addressing modes for loc.

Flags Affected: \(F\) is set if loc is write-protected, cleared otherwise.
Traps: Undefined Instruction Trap (UND) is activated if the M bit in the CFG register is clear.

Illegal Operation Trap (ILL) is activated if this instruction is attempted while the PSR U flag is set.

Abort Trap (ABT) is activated if the Level 1 page table entry for loc is invalid (V bit \(=0\) ) and the Protection Level (PL) indicates that the access is allowed. No trap is issued for an invalid Level 2 page table entry, and the Protection Level field is assumed to be present regardless of the state of the \(V\) bit.

\section*{Example:}

WRVAL 512 (R0)
1E 07408200

This example checks the protection level assigned to the user-mode virtual address \(512(R 0)\) and sets or clears the \(F\) flag to indicate the result.

\section*{XORi}

Exclusive Or

\section*{Syntax: XORi src, dest}
```

gen gen
read.i rmw.i


The XORi instruction performs a bit-wise Exclusive-OR operation on the src and dest operands and places the result in the dest operand location.

The instruction XORs each bit of src with the corresponding dest bit. If two corresponding bits are equal, the dest bit is set to "0"; otherwise, the dest bit is set to "1".

Flags Affected: None.
Traps: None.

## Example:

$$
\text { XORB -8(FP), -4(FP) } 38 \text { C6 } 7870
$$

This example XORs the bytes at the addresses specified by $-8(F P)$ and $-4(F P)$ and places the result in the byte at address -4 (FP).

| Operands | Operand Values: <br> Before | Binary <br> After |
| :--- | :--- | :---: |
| $-8(\mathrm{FP})$ | 11110000 | 11110000 |
| $-4(\mathrm{FP})$ | 10010101 | 01100101 |

INSTRUCTION SET LISTED BY FUNCTIONAL GROUPS

## Instruction

INTEGER

Arithmetic
Add
Add Quick
Add with Carry
Subtract
Subtract with Carry [Borrow]
Negate
Absolute Value
Multiply
Multiply Extended Integer
Divide
Modulus
Quotient
Remainder
Divide Extended Integer
Movement and Conversion

```
Move
Move Quick
Move with Sign-Extension
Move with Zero-Extension
```

Comparison

Compare
Compare Quick

PACKED DECIMAL

Add Packed Decimal

Subtract Packed Decimal

## Mnemonic Forms

## Index

| ADDB, ADDW, ADDD | ADDi |
| :--- | :--- |
| ADDQB, ADDQW, ADDQD | ADDQi |
| ADDCB, ADDCW, ADDCD | ADDCi |
| SUBB, SUBW, SUBD | SUBi |
| SUBCB, SUBCW, SUBCD | SUBCi |
| NEGB, NEGW, NEGD | NEGi |
| ABSB, ABSW, ABSD | ABSi |
| MULB, MULW, MULD | MULi |
| MEIB, MEIW, MEID | MEIi |
| DIVB, DIVW, DIVD | DIVi |
| MODB, MODW, MODD | MODi |
| QUOB, QUOW, QUOD | QUOi |
| REMB, REMW, REMD | REMi |
| DEIB, DEIW, DEID | DEIi |

MOVB, MOVW, MOVD MOVi
MOVQB, MOVQW, MOVQD MOVQi MOVXBD, MOVXWD, MOVXBW MOVXii MOVZBD, MOVZWD, MOVZBW MOVZii

CMPB, CMPW, CMPD
CMPi
CMPQB, CMPQW, CMPQD
CMPQi

## Instruction

## FLOATING POINT

Add Floating
Subtract Floating
Multiply Floating
Divide Floating
Dot Product Floating
Polynomial Floating
Negate Floating
Absolute Value Floating
Compare Floating
Move Floating
Logarithm Binary Floating
Scale Binary Floating
Move Long Floating to Floating
Move Floating to Long Floating
Move Integer to Floating
Round Floating to Integer
Truncate Floating to Integer
Floor Floating to Integer

Load FSR
Store FSR

LOGICAL

Arithmetic

Logical AND
Logical OR
Bit Clear
Exclusive OR
Complement

Shift
Arithmetic Shift
Logical Shift
Rotate

Boolean

Complement Boolean
Save Condition as Boolean

Mnemonic Forms

ADDF, ADDL ADDf
SUBF, SUBL SUBf
MULF, MULL MULf
DIVF, DIVL DIVf
DOTF, DOTL DOTf
POLYF, POLYL POLYf
NEGF, NEGL NEGf
ABSF, ABSL ABSf
CMPF, CMPL CMPf

MOVF, MOVL MOVf

LOGBF, LOGBL LOGBf
SCALBF, SCALBL SCALBf
MOVLF
MOVFL MOVFL

MOVBF, MOVWF, MOVDF, MOVif
MOVBL, MOVWL, MOVDL
ROUNDFB, ROUNDFW, ROUNDFD, ROUNDfi
ROUNDLB, ROUNDLW, ROUNDLD
TRUNCFB, TRUNCFW, TRUNCFD, TRUNCfi
TRUNCLB, TRUNCLW, TRUNCLD
FLOORFB, FLOORFW, FLOORFD, FLOORfi
FLOORLB, FLOORLW, FLOORLD
LFSR LFSR
SFSR SFSR

ANDB, ANDW, ANDD ANDi
ORB, ORW, ORD ORi
BICB, BICW, BICD BICi
XORB, XORW, XORD XORi
COMB, COMW, COMD COMi

| ASHB, ASHW, ASHD | ASHi |
| :--- | :--- |
| LSHB, LSHW, LSHD | LSHi |
| ROTB, ROTW, ROTD | ROTi |

NOTB, NOTW, NOTD NOTi
ScondB, ScondW, ScondD Scondi

Instruction

BIT

Test Bit

Set Bit

Clear Bit

Invert Bit

Find First Set Bit

Convert to Bit Pointer

BIT FIELD

Extract Field

Extract Field Short

Insert Field

Insert Field Short

STRING
Move String
Move String, Translating
Compare Strings
Compare Strings, Translating
Skip String
Skip String, Translating

| MOVSB, MOVSW, MOVSD | MOVSi |
| :--- | :--- |
| MOVST |  |
| MOVST |  |
| CMPSB, CMPSW, CMPSD | CMPSi |
| CMPST |  |
| CMPST |  |
| SKPSB, SKPSW, SKPSD |  |
| SKPST | SKPSi |
|  |  |

BLOCK
Move Multiple
Compare Multiple

ARRAY

Bounds Check

Calculate Index

| TBITB, TBITW, TBITD | TBITi |
| :--- | :--- |
| SBITB, SBITW, SBITD, | SBITi, |
| SBITIB, SBITIW, SBITID | SBITIi |
| CBITB, CBITW, CBITD, | CBITi, |
| CBITIB, CBITIW, CBITID | CBITIi |
| IBITB, IBITW, IBITD | IBITi |
| FFSB, FFSW, FFSD | FFSi |
| CVTP | CVTP |

EXTB, EXTW, EXTD EXTi

EXTSB, EXTSW, EXTSD EXTSi

INSB, INSW, INSD INSi

INSSB, INSSW, INSSD
INSSi

MOVMB, MOVMW, MOVMD MOVMi

CMPMB, CMPMW, CMPMD

CHECKB, CHECKW, CHECKD

INDEXB, INDEXW, INDEXD
CHECKi

INDEXi

## PROCESSOR CONTROL

## Branches

```
Jump
Conditional Branch
Unconditional Branch
Case Branch (Multiway)
Add, Compare and Branch
```

Local Procedure Calls/Returns

Jump to Subroutine
Branch to Subroutine
Return from Subroutine

External Procedure Calls/Returns
Call External Procedure
Call External Procedure with Descriptor
Return from
External Procedure

Explicit Trap Instructions
Breakpoint Trap
Trap on Flag (conditional)
Supervisor Call Trap
Trap/Interrupt Returns

| Return from Trap* | RETT | RETT |
| :--- | :--- | :--- |
| Return from Interrupt* | RETI | RETI |

* Privileged instruction.

| JUMP | JUMP |
| :--- | :--- |
| Bcond | Bcond |
| BR | BR |
| CASEB, CASEW, CASED | CASEi |
| ACBB, ACBW, ACBD | ACBi |

JSR
JSR
BSR BSR
RET RET

CXP
CXP
CXPD
CXPD

RXP

| BPT | BPT |
| :--- | :--- |
| FLAG | FLAG |
| SVC | SVC |

RETI

## Instruction

## PROCESSOR SERVICE

## Effective Address

Calculate Effective Address

Context Instructions

Save General Purpose Registers
Restore General Purpose Registers
Enter New Procedure Context
Exit Procedure Context

Register/Stack Manipulation

Adjust Stack Pointer
Bit Clear in PSR*

Bit Set in PSR*

Load Processor Register*
Store Processor Register*
Set Configuration Register*

Miscellaneous

No Operation
Wait for Interrupt
Diagnose
Cache Invalidate*

* Privileged, or having privileged forms.

MEMORY MANAGEMENT

Load Memory Management
Register
Store Memory Management
Register

Validate Address for Reading
Validate Address for Writing

Move Value from Supervisor
to User Space
Move Value from User
to Supervisor Space

MOVSUB, MOVSUW, MOVSUD
MOVSUi
ADDR

| SAVE | SAVE |
| :--- | :--- |
| RESTORE | RESTORE |
| ENTER | ENTER |
| EXIT | EXIT |

ADJSPB, ADJSPW, ADJSPD
ADJSPi
BICPSRB, BICPSRW

BISPSRB, BISPSRW

LPRB, LPRW, LPRD BICPSRB BICPSRW BISPSRB BISPSRW LPRi
SPRi
SETCFG
SETCFG

| NOP | NOP |
| :--- | :--- |
| WAIT | WAIT |
| DIA | DIA |
| CINV | CINV |

LMR

SMR
SMR

RDVAL
RDVAL
WRVAL
WRVAL

MOVUSB, MOVUSW, MOVUSD
MOVUSi


[^0]:    * This shows the internal format of the quick operand after sign-extension to byte length. The operand is encoded within the instruction as binary 1000.

