




Am29000

User's Manual



ADVANCED MICRO DEVICES

Am29000



Streamlined Instruction Processor

User's Manual

© 1987 Advanced Micro Devices

Advanced Micro Devices reserves the right to make changes in its products without notice in order to improve design or performance characteristics.

This manual neither states nor implies any warranty of any kind, including but not limited to implied warranties of merchantability or fitness for a particular application. AMD assumes no responsibility for the use of any circuitry other than the circuitry embodied in an AMD product. The information in this publication is believed to be accurate in all respects at the time of publication, but is subject to change without notice. AMD assumes no responsibility for any errors or omissions, and disclaims responsibility for any consequences resulting from the use of the information included herein. Additionally, AMD assumes no responsibility for the functioning of undescribed features or parameters.

This Manual was written and edited by *Mike Johnson*, Manager of Microprocessor Product Planning.

Major contributions were provided by *Gigy Baror*, *Brian Case*, *Philip Freidin*, *Smeeta Gupta*, *Tim Olson* and *Dave Sorensen*.

Technical Writer:

Erlend Kyllonen, Senior Technical Writer, Technical Communications

TABLE OF CONTENTS

PREFACE	ix
1	FEATURES AND PERFORMANCE	1-1
1.1	DISTINCTIVE CHARACTERISTICS	1-1
1.2	INTRODUCTION	1-1
1.3	PERFORMANCE OVERVIEW	1-3
1.3.1	Cycle Time	1-3
1.3.2	Four-Stage Pipeline	1-3
1.3.3	System Interface	1-3
1.3.4	Register File	1-4
1.3.5	Instruction Execution	1-5
1.3.6	Branch Target Cache	1-5
1.3.7	Branching	1-6
1.3.8	Loads and Stores	1-6
1.3.9	Memory Management	1-8
1.3.10	Interrupts and Traps	1-8
1.3.11	Floating-Point Arithmetic Unit	1-9
1.4	OPTIMIZING COMPILERS	1-10
1.4.1	Optimizing-Compiler Overview	1-10
1.4.2	Optimizing-Compiler Operation	1-10
1.4.3	The Am29000 and Optimizing Compilers	1-11
2	ARCHITECTURE HIGHLIGHTS	2-1
2.1	PROGRAMMER'S OVERVIEW	2-1
2.1.1	Program Modes	2-1
2.1.2	Visible Registers	2-1
2.1.3	Instruction Set Overview	2-5
2.1.4	Data Formats and Handling	2-6
2.1.5	Interrupts and Traps	2-12
2.1.6	Memory Management	2-13
2.1.7	Coprocessor Programming	2-14
2.1.8	Timer Facility	2-14
2.1.9	Trace Facility	2-15
2.2	HARDWARE OVERVIEW	2-15
2.2.1	Four-Stage Pipeline	2-16
2.2.2	Instruction Fetch Unit	2-16
2.2.3	Execution Unit	2-17
2.2.4	Memory Management Unit	2-18
2.2.5	Processor Modes	2-18
2.3	SYSTEM INTERFACE OVERVIEW	2-20
2.3.1	Channel	2-20
2.3.2	Text/Development Interface	2-21
2.3.3	Clocks	2-22

	2.3.4	Master/Slave Operation	2-22
	2.3.5	Coprocessor Attachment	2-22
3		PROGRAMMER REFERENCE	3-1
	3.1	PROGRAM MODES	3-1
		3.1.1 Supervisor Mode	3-1
		3.1.2 User Mode	3-1
	3.2	VISIBLE REGISTERS	3-2
		3.2.1 General-Purpose Registers	3-2
		3.2.2 Special-Purpose Registers	3-7
		3.2.3 TLB Registers	3-26
	3.3	INSTRUCTION SET	3-30
		3.3.1 Integer Arithmetic	3-30
		3.3.2 Compare	3-30
		3.3.3 Logical	3-33
		3.3.4 Shift	3-34
		3.3.5 Data Movement	3-34
		3.3.6 Constant	3-36
		3.3.7 Floating-Point	3-36
		3.3.8 Branch	3-36
		3.3.9 Miscellaneous	3-36
	3.4	DATA FORMATS AND HANDLING	3-39
		3.4.1 Data Types	3-39
		3.4.2 External Data Accesses	3-40
		3.4.3 Addressing and Alignment	3-46
	3.5	INTERRUPTS AND TRAPS	3-51
		3.5.1 Interrupts	3-52
		3.5.2 Traps	3-52
		3.5.3 Wait Mode	3-52
		3.5.4 Interrupt Vectors	3-53
		3.5.5 Interrupt and Trap Handling	3-55
		3.5.6 *WARN Trap	3-59
		3.5.7 Sequencing of Interrupts and Traps	3-60
		3.5.8 Exception Reporting and Restarting	3-62
		3.5.9 Exceptions During Interrupt and Trap Handling	3-63
	3.6	MEMORY MANAGEMENT	3-64
		3.6.1 Translation Look-Aside Buffer	3-64
		3.6.2 Address Translation	3-65
		3.6.3 Reload	3-70
		3.6.4 Entry Invalidation	3-70
		3.6.5 Protection	3-71
	3.7	SERIALIZATION	3-72
	3.8	INITIALIZATION	3-72

4	HARDWARE FEATURES	4-1
4.1	FOUR-STAGE PIPELINE	4-1
4.2	INSTRUCTION FETCH UNIT	4-2
	4.2.1 Instruction Prefetch Buffer	4-2
	4.2.2 Branch Target Cache	4-6
	4.2.3 Non-Sequential Instruction Fetches	4-10
	4.2.4 Program Counter Unit	4-11
4.3	EXECUTION UNIT	4-12
	4.3.1 Register File	4-13
	4.3.2 Address Unit	4-15
	4.3.3 Arithmetic/Logic Unit	4-17
	4.3.4 Field Shift Unit	4-17
	4.3.5 Prioritizer	4-18
4.4	MEMORY MANAGEMENT UNIT	4-18
4.5	PIPELINE HOLD MODE	4-19
5	SYSTEM INTERFACES	5-1
5.1	SIGNAL DESCRIPTION	5-1
5.2	CHANNEL DESCRIPTION	5-6
	5.2.1 Channel Overview	5-7
	5.2.2 User-Defined Signals	5-7
	5.2.3 Instruction Accesses	5-8
	5.2.4 Data Accesses	5-8
	5.2.5 Reporting Errors	5-8
	5.2.6 Access Protocols	5-9
	5.2.7 Simple Accesses	5-11
	5.2.8 Pipelined Accesses	5-11
	5.2.9 Burst-Mode Accesses	5-13
	5.2.10 Arbitration	5-20
	5.2.11 Bus Sharing - Electrical Considerations	5-21
	5.2.12 Channel Behavior for Interrupts and Traps	5-22
	5.2.13 Effect of the *LOCK Output	5-22
5.3	TEST/DEVELOPMENT INTERFACE	5-23
	5.3.1 Processor Status Outputs	5-23
	5.3.2 CPU Control Inputs	5-24
	5.3.3 Hardware Development	5-25
	5.3.4 Hardware Testing	5-29
5.4	EXTERNAL INTERRUPTS AND TRAPS	5-30
5.5	PROCESSOR RESET	5-31
5.6	*WARN INPUT	5-32
5.7	CLOCKS	5-33
	5.7.1 Processor-Generated Clock	5-33
	5.7.2 System-Generated Clock	5-33
	5.7.3 Clock Synchronization	5-33
	5.7.4 Electrical Specifications	5-34

5.8	MASTER/SLAVE CHECKING	5-34
5.8.1	Master/Slave Operation	5-34
5.8.2	Preventing Spurious Errors	5-35
5.8.3	Switching Master and Slave Processors	5-35
6	COPROCESSOR INTERFACE	6-1
6.1	COPROCESSOR PROGRAMMING	6-1
6.1.1	Overview of Coprocessor Operations	6-1
6.1.2	Coprocessor Transfers	6-2
6.1.3	Coprocessor Exceptions	6-4
6.1.4	Coprocessor as a System Option	6-5
6.1.5	Interrupted Coprocessor Operations	6-5
6.2	COPROCESSOR ATTACHMENT	6-6
6.2.1	Signal Description	6-6
6.2.2	Coprocessor Communication	6-8
7	PROGRAMMING	7-1
7.1	APPLICATION-PROGRAMMING CONSIDERATIONS	7-1
7.1.1	Procedure Calls and Returns	7-1
7.1.2	Addressing General-Purpose Registers Indirectly	7-7
7.1.3	Run-Time Checking	7-8
7.1.4	Operating-System Calls	7-9
7.1.5	Multi-Precision Integer Addition and Subtraction	7-9
7.1.6	Integer Multiplication	7-10
7.1.7	Integer Division	7-12
7.1.8	Trapping Arithmetic Instructions	7-14
7.1.9	Complementing a Boolean	7-14
7.1.10	Generating Large Constants	7-15
7.1.11	Large Jump and Call Ranges	7-15
7.1.12	No-Ops	7-16
7.1.13	Character-String Operations	7-16
7.1.14	Movement of Large Data Blocks	7-17
7.2	SYSTEMS-PROGRAMMING CONSIDERATIONS	7-18
7.2.1	System Protection	7-18
7.2.2	Interrupts and Traps	7-19
7.2.3	Fast Context Switching	7-21
7.2.4	Memory Management	7-23
7.2.5	Restarting Faulting External Accesses	7-27
7.2.6	Multi-Processing	7-28
7.2.7	Timer Facility	7-29
7.2.8	Trace Facility	7-31
7.3	PIPELINE FEATURES EXPOSED TO SOFTWARE	7-32
7.3.1	Delayed Branch	7-32
7.3.2	Overlapped Loads and Stores	7-34
7.3.3	Delayed Effects of Registers	7-35

8	INSTRUCTION SET	8-1
8.1	INSTRUCTION-DESCRIPTION NOMENCLATURE	8-1
	8.1.1 Operand Notation and Symbols	8-1
	8.1.2 Operator Symbols	8-3
	8.1.3 Control-Flow Terminology	8-4
	8.1.4 Assembler Syntax	8-5
8.2	ARITHMETIC/LOGIC STATUS RESULTS OF INSTRUCTIONS	8-6
	8.2.1 Arithmetic/Logic Status Bits	8-6
	8.2.2 Arithmetic Operation Status Results	8-6
	8.2.3 Logical Operation Status Results	8-7
8.3	INSTRUCTION FORMATS	8-8
8.4	INSTRUCTION DESCRIPTIONS	8-10
8.5	INSTRUCTION INDEX BY OPERATION CODE	8-127
Appendix A.	Channel Operation Timing	A-1
Appendix B.	Register Summary	B-1

LIST OF FIGURES

1-1	Simplified System Diagram	1-2
2-1	Data-Unit Numbering Conventions	2-10
2-2	Am29000 Data Flow	2-15
3-1	General Purpose Register Organization	3-3
3-2	Register Bank Organization	3-6
3-3	Special Purpose Registers	3-8
3-4	Vector Area Base Address	3-9
3-5	Current Processor Status	3-10
3-6	Configuration Register	3-13
3-7	Channel Address Register	3-14
3-8	Channel Data Register	3-14
3-9	Channel Control Register	3-15
3-10	Register Bank Protect Register	3-16
3-11	Timer Counter Register	3-17
3-12	Timer Reload Register	3-18
3-13	Program Counter 0 Register	3-19
3-14	Program Counter 1 Register	3-19
3-15	Program Counter 2 Register	3-20
3-16	MMU Configuration Register	3-20
3-17	LRU Recommendation Register	3-21
3-18	Indirect Pointer C Register	3-22
3-19	Indirect Pointer A Register	3-22
3-20	Indirect Pointer B Register	3-23
3-21	Q Register	3-23
3-22	ALU Status Register	3-23
3-23	Byte Pointer	3-25
3-24	Funnel Shift Count	3-25
3-25	Load/Store Count Remaining	3-26
3-26	Translation Look-Aside Buffer Registers	3-27
3-27	TBL Entry Word 0	3-27
3-28	TBL Entry Word 1	3-29
3-29	Load/Store Instruction Format	3-41
3-30	Non-Coprocessor Load/Store Format	3-41
3-31	Byte and Half-Word Addressing With BO=0	3-48
3-32	Byte and Half-Word Addressing With BO=1	3-49
3-33	Vector Table Entry	3-53
3-34	Current Processor Status Register After an Interrupt or Trap	3-56
3-35	Current Processor Status Register Before Interrupt Return	3-57

3-36	Translation Look-Aside Buffer Organization	3-65
3-37	Virtual Address for 1,2,4, and 8 Kbyte Pages	3-66
3-38	Address Translation Process	3-68
3-39	Current Processor Status Register After Reset	3-73
4-1	Am29000 Data Flow	4-2
4-2	IPB State Transitions	4-4
4-3	Branch Target Cache Organization	4-6
4-4	Branch Target Cache Lookup Process	4-8
4-5	Program Counter Unit	4-12
4-6	Register File and Register Address Generation Unit	4-13
4-7	Address Unit	4-15
5-1	Channel Access Flow Diagram (Instruction and Data)	5-10
5-2	Processor Burst Mode Instruction Accesses: Control Flow	5-14
5-3	Slave Burst Mode Instruction Accesses: Control Flow	5-15
5-4	Processor Burst-Mode Data Accesses: Control Flow	5-16
5-5	Slave Burst-Mode Data Accesses: Control Flow	5-17
5-6	Valid Transitions on CNTL0 - CNTL1 Inputs	5-25
6-1	Coprocessor Load/Store Format	6-2
6-2	Coprocessor Attachment	6-7
7-1	Run-Time Stack Example	7-2
7-2	Procedure Prologue and Epilogue	7-5
7-3	Activation Record	7-6
8-1	Instruction Format	8-8
8-2	Frequently Occurring Instruction-Field Uses	8-10
8-3	Instruction Description Format	8-11
B-1	General-Purpose Register Organization	B-1
B-2	Register Bank Organization	B-2
B-3	Special-Purpose Registers	B-3
B-4	Translation Look-Aside Buffer Registers	B-7
B-5	Translation Look-Aside Buffer Entries	B-7

Tables

2-1	Am29000 Instruction Set	2-7
3-1	Integer Arithmetic Instructions	3-31
3-2	Compare Instructions	3-32
3-3	Logical Instructions	3-34
3-4	Shift Instructions	3-34
3-5	Data Movement Instructions	3-35
3-6	Constant Instructions	3-36
3-7	Floating-Point Instructions	3-37
3-8	Branch Instructions	3-38
3-9	Miscellaneous Instructions	3-38
3-10	Vector Number Assignments	3-54
3-11	Interrupt and Trap Priority Table	3-61
A-1	Signal Summary	A-1
B-6	Register Field Summary	B-8

PREFACE

Am29000 DESIGN PHILOSOPHY

The Am29000 Streamlined Instruction Processor is the result of a design philosophy which recognizes that processor performance must be considered in light of the processor's hardware and software environment. The key to maximizing performance lies in the realization that the processor is part of an integrated system, and is itself a collection of components which must be properly integrated.

Processor features must not be considered only on their own merits, but also in relation to other components of the system. A particular feature which—considered alone—increases one aspect of processor performance may actually decrease the performance of the total system, because of the burden which it places elsewhere in system. As an illustration, consider the factors involved in the execution time of any processor task:

$$\text{TASK TIME} = \text{INSTRUCTIONS/TASK} * \text{CYCLES/INSTRUCTION} * \text{TIME/CYCLE}$$

To minimize the time taken, it is necessary to minimize the above product. This is not equivalent to minimizing all of the terms which contribute to the product; this, in fact, is generally not possible due to the interaction of the terms.

As an example of the interaction of the above terms, consider the number of instructions required for a task. An attempt to minimize this number—a more or less traditional approach to processor architecture design—increases the average number of cycles required for the execution of an instruction, because of the increased number of operations performed by each instruction. In addition, cycle time is increased because of instruction-decode time.

A second example of the interaction in the above equation appears in an attempt to reduce the cycle time through the pipelining of operations. In theory, the cycle time can be made arbitrarily small by the definition of an arbitrarily large number of pipeline stages. In practice—at least in the case of general-purpose processors—pipelining rarely yields much of its potential benefit. This is due to situations where the pipeline cannot be kept fully occupied, such as when storage references and branches occur. In these situations, additional pipeline stages increase the number of cycles required for an operation, and thus affect the CYCLES/INSTRUCTION term.

OPTIMUM PERFORMANCE

Each of the terms in the above equation has some minimum bound for a given implementation technology and task. In general, this minimum bound cannot be approached without an offsetting increase in the other terms, making the overall product

less-than-optimum. The question then arises, what combination of terms does yield an optimum product? There are several things to note when answering this question.

The first observation is that the number of operations underlying a given task is more or less fixed. Any single processor ultimately limits the time required for a task because it has a single execution unit and a single instruction stream. The operations which must be performed are reflected in the INSTRUCTIONS/TASK and CYCLES/INSTRUCTION terms. These operations may be performed by relatively few instructions, where each instruction takes multiple cycles to execute, or by a larger number of instructions, where each takes a single cycle to execute. In the first case, the instructions are complex; in the second, they are simple.

The point is that the trade-off between simple and complex instructions is not one-to-one. For example, reducing the number of cycles per instruction by a factor of three does not increase the number of instructions per task by the same factor. There are two reasons for this. The first is that, even when an instruction set supports complex operations, a large proportion of the instructions which are executed perform operations which could be performed as well by simple instructions. The second is that simple instructions expose more of the internal processor operation to an optimizing compiler. This allows the compiler to tailor the organization and sequence of operations to the task at hand, thereby reducing the total number of instructions executed.

PERFORMANCE LEVERAGE

Another important observation is that there is a tremendous amount of leverage in the TIME/CYCLE and CYCLES/INSTRUCTION terms. As they are made smaller, they have a proportionately greater effect on performance.

For example, a reduction of 10 nsec in the cycle time of a processor operating with a 200 nsec cycle time yields an increase of 5% in the processor's performance. The same improvement in a processor operating with a 50 nsec cycle time yields a 20% increase in performance.

Correspondingly, a reduction of 0.2 in the number of cycles per instruction in a processor which averages 5 cycles per instruction yields a 4% increase in performance. However, the same reduction yields a 12.5% performance increase in a processor which averages 1.6 cycles per instruction.

CONCLUSION

The conclusion is that it is possible—and desirable—to yield somewhat in the number of instructions executed for a given task, and more than make up for the performance impact of this increase by reductions in the cycle time and in the number of cycles per instruction. For example, if both the cycle time and the number of cycles per instruction are reduced by a factor of three, while the number of instructions for a given task is allowed to grow by 50%, the resulting task time is reduced by a factor of 6.

The Am29000 architecture was designed with the above effects in mind. Maximum performance is obtained by the optimization of the product of the number of instructions per task, the number of cycles per instruction, and the cycle time, not by minimizing one factor at the expense of the others. This is accomplished by careful definition of all processor components. In particular:

- 1) The INSTRUCTIONS/TASK term is optimized by the definition of simple instructions. The processor provides an efficient instruction set and a large number of general-purpose registers to an optimizing, high-level language compiler. Most of the reductions in this term are accomplished by the compiler. The number of instructions for a given task may be greater than the number of instructions for processors with complex instruction sets. However, this increase is more than offset by other improvements in processor performance.
- 2) The CYCLES/INSTRUCTION term is optimized by the data-flow structure and performance-enhancing features of the processor. A large amount of processor hardware is dedicated to achieving an average instruction-execution rate which is close to single-cycle execution.
- 3) The TIME/CYCLE term is optimized by the implementation technology, the processor system interface, and judicious use of pipelining. The simplicity of the instruction set and processor features helps minimize the cycle time.

Am29000 USER MANUAL OVERVIEW

This Manual contains information on the Am29000 processor which is essential for computer hardware and software architects and system design engineers. Additional information is available in the form of Data Sheets, Application Notes, and other documentation which is provided with software products and hardware-development tools.

The information in this manual is organized into eight chapters, each viewing the processor from a different perspective, and each with a specific objective:

Chapter 1 introduces the features and performance aspects of the Am29000.

Chapter 2 contains brief technical descriptions of the processor architecture and implementation.

Chapter 3 describes the details of the Am29000 architecture.

Chapter 4 details the operation of the processor's internal functional units.

Chapter 5 describes the operation of the external interfaces of the Am29000.

Chapter 6 describes the attachment and use of coprocessors for the Am29000

Chapter 7 discusses the implementation of software systems for the processor, focusing on programming features which deserve more coverage than is provided by other chapters.

Chapter 8 specifies the instruction set of the Am29000. It describes the instruction formats in detail, and provides a detailed description of every instruction.

This Manual is organized around readers' concerns and objectives. Each chapter focuses on a particular aspect of the processor, and is organized so that it may be read independently, insofar as possible.

For those readers desiring only a brief overview of the Am29000, Chapters 1 and 2 identify the outstanding features of the processor, and give a brief overview of the processor. These chapters address both software and hardware concerns.

For software architects and system programmers interested mainly in software-related issues, Chapters 3, 7, and 8 provide the necessary information.

For hardware architects and systems hardware designers interested mainly in hardware-related issues, Chapters 4 and 5 provide most of the required information; Chapter 8 also provides some related information.

For those readers interested in the coprocessor interface, Chapter 6 describes the interface both from a software and hardware point-of-view.

CHAPTER 1

FEATURES AND PERFORMANCE

This chapter provides an evaluation of the Am29000 as an aid in considering a particular application. A detailed technical description of the Am29000 is contained in subsequent chapters. This chapter informally describes the features of the processor, concentrating on features which distinguish the Am29000 from other available processors.

1.1 DISTINCTIVE CHARACTERISTICS

- Full 32-bit architecture.
- CMOS technology / TTL-compatible interfaces.
- 25 MHz nominal operating frequency.
- 17 million instructions per second sustained.
- 1.5 clock cycles per instruction average.
- 4 giga-byte virtual address space.
- Double-precision, Floating-Point Arithmetic Unit (Am29027).
- 192 general-purpose registers.
- Three-address instruction architecture.
- Non-multiplexed, pipelined address, instruction and data buses.
- Concurrent instruction and data accesses.
- Burst-mode access support.
- 512-byte Branch Target Cache on-chip.
- 64-entry Memory Management Unit on-chip.
- Demand paging.
- Fully pipelined.
- On-chip Timer Facility.
- On-chip clock generation.
- On-chip debugging support.
- Master/slave chip output checking.

1.2 INTRODUCTION

The Am29000 Streamlined-Instruction Processor is a high-performance, general-purpose, 32-bit microprocessor implemented in complementary metal-oxide semiconductor (CMOS) technology. It supports a variety of applications, using a flexible architecture and rapid execution of simple instructions which are common to a wide range of tasks.

The Am29000 efficiently performs operations common to all systems, while deferring most decisions on system policies to the system architect. It is well-suited for application in high-performance workstations, general-purpose super-microcomputers, high-performance

real-time controllers, laser printer controllers, network protocol converters, and many other applications where high performance, flexibility, and the ability to program using standard software tools is important.

The Am29000 instruction set has been influenced by the results of high-level-language, optimizing-compiler research. It is appropriate for a variety of languages, because it efficiently executes operations which are common to all languages. Consequently, the Am29000 is an ideal target for high-level languages such as C, Fortran, Pascal, and Ada.

The Am29000 is packaged in a 169-pin, pin-grid-array (PGA) package, with 141 signal pins, 27 power and ground pins, and 1 alignment pin. D.C. power dissipation is 1.5 Watts. A representative system diagram is shown in Figure 1-1.

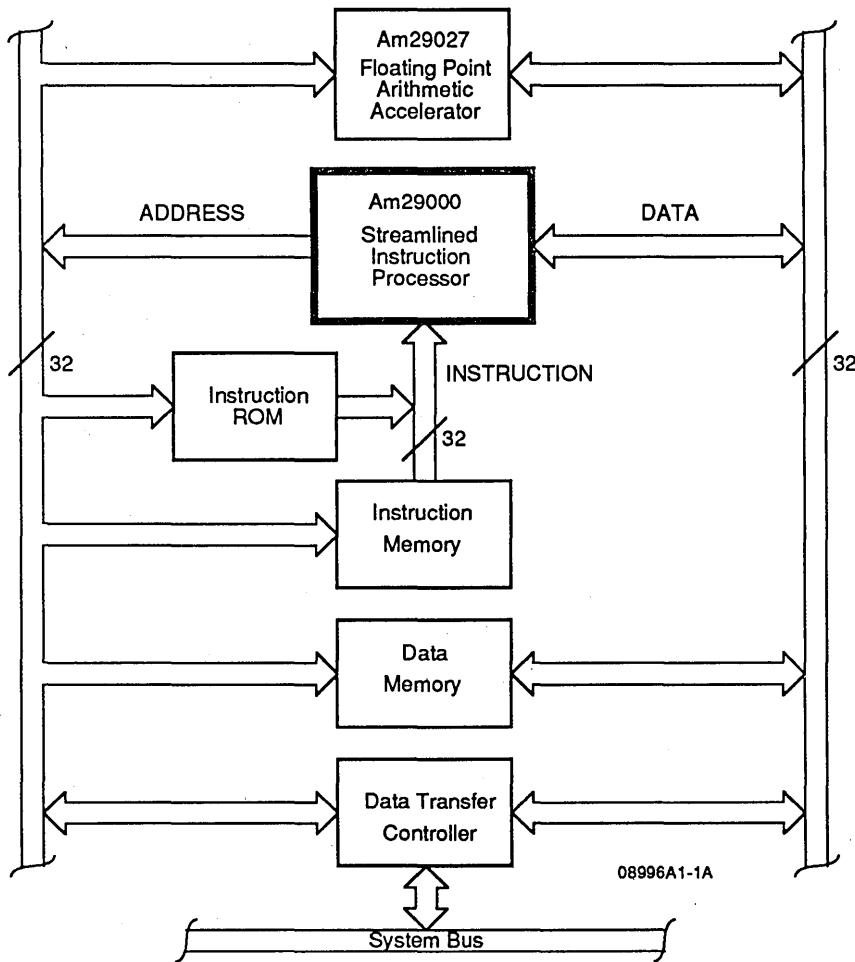


Figure 1-1. Simplified System Diagram

1.3 PERFORMANCE OVERVIEW

The Am29000 provides a significant margin of performance over other processors in its class, since the majority of processor features were defined with the maximum achievable performance in mind. This section describes the features of the Am29000 from the point-of-view of system performance.

1.3.1 CYCLE TIME

The Am29000 is implemented in CMOS technology, with a 1.2 micron effective transistor-channel length. This technology allows the processor to operate at a frequency of 25 MHz. The processor cycle time is a single, 40 ns clock period. The processor interface drivers can drive 80 pF loads at this frequency.

The Am29000 architecture and system interfaces are designed so that the processor cycle time can decrease with technology improvements.

1.3.2 FOUR-STAGE PIPELINE

The Am29000 utilizes a four-stage pipeline, allowing it to execute one instruction every clock cycle. The processor can complete an instruction on every cycle, even though four cycles are required from the beginning of an instruction to its completion.

At a 25 MHz operating frequency, the maximum instruction execution rate is 25 million instructions per second (MIPS). For most other processors, the maximum MIPS rate has little meaning, because it can be achieved only under special circumstances. However, the Am29000 pipeline is designed so that the Am29000 can operate at the maximum instruction-execution rate a significant portion of the time.

Pipeline interlocks are implemented by processor hardware. Except for a few special cases, it is not necessary to re-arrange programs to avoid pipeline dependencies.

1.3.3 SYSTEM INTERFACE

One of the most difficult tasks in the definition of a high-speed micro-processor is the definition of an off-chip interface which supports the operating frequency of the processor, and does not restrict the ability of the processor to fetch instructions and data. If the external interface of a microprocessor cannot support an instruction fetch rate of one instruction every cycle, there is little prospect that the processor will execute at this rate, even though it supports such a rate internally.

The Am29000 accesses external instructions and data using three non-multiplexed buses. These buses are collectively referred to as the channel. The channel protocol minimizes the logic chains involved in a transfer, and provides a maximum transfer rate of 200 Mbyte/sec.

Separate Address, Instruction, and Data Buses

The Am29000 incorporates two 32-bit buses for instruction and data transfers, and a third address bus which is shared between instruction and data accesses. This bus structure allows simultaneous instruction and data transfers, even though the address bus is shared. The channel achieves the performance of four separate 32-bit buses at a much reduced pin count.

Pipelined Addresses

The Am29000 address bus is pipelined, so that it can be released before an instruction or data transfer is completed. This allows a subsequent access to begin before the first has completed, and allows the processor to have two accesses in progress simultaneously.

Support of Burst Devices and Memories

Burst-mode accesses provide high transfer rates for instructions and data at sequential addresses. For such accesses, the address of the first instruction or datum is sent, and subsequent requests for instructions or data at sequential addresses do not require additional address transfers. These instructions or data are transferred until either party involved in the transfer terminates the access.

Burst-mode accesses can occur at the rate of one access per cycle after the first address has been processed. At 25 MHz, the maximum achievable transfer bandwidth for either instructions or data is 100 Mbyte/s.

Burst-mode accesses may occur to input/output devices, if the system design permits.

Interface to Fast Devices and Memories

The processor can be interfaced to devices and memories which complete accesses within one cycle. The channel protocol takes maximum advantage of such devices and memories, by allowing data to be returned to the processor during the cycle in which the address is transmitted. This allows a full range of memory-speed trade-offs to be made within a particular system.

1.3.4 REGISTER FILE

An on-chip Register File containing 192 general-purpose registers allows most instruction operands to be fetched without the delay of an external access. The Register File incorporates several features which aid the retention of data required by an executing program. Because of the number of general-purpose registers, the frequency of external

references for the Am29000 is significantly lower than the frequency of references in earlier processors having only 16 or 32 registers.

Triple-port access to the Register File allows two source-operands to be fetched, in one cycle, while a previously-computed result is written. Three, 32-bit internal buses prevent contention in the routing of operands. All operand fetches and result write-backs for instruction execution can be performed in a single cycle.

The registers allow efficient procedure linkage, by caching a portion of a compiler's run-time stack. On the average, procedure calls and returns can be executed 5 to 10 times faster (on a cycle-by-cycle basis) than in processors which require the implementation of a run-time stack in external memory (with the attendant loading and storing of registers on procedure call and return).

The registers can contain variables, constants, addresses, and operating-system values. In multi-tasking applications, they can be used to hold the processor status and variables for as many as 8 different tasks. A register-banking option allows the register file to be divided into segments which can be individually protected. In this configuration, a task switch can occur in as few as 17 cycles.

1.3.5 INSTRUCTION EXECUTION

The Am29000 uses an Arithmetic/Logic Unit, a Field Shift Unit, and a Prioritizer to execute most instructions. Each of these is organized to operate on 32-bit operands, and provide a 32-bit result. All operations are performed in a single cycle.

Instruction operations are overlapped with operand fetch and result write-back to the Register File. Pipeline forwarding logic detects pipeline dependencies and routes data as required, avoiding delays which might arise from these dependencies.

1.3.6 BRANCH TARGET CACHE

In general, the Am29000 meets its instruction bandwidth requirements via instruction prefetching. However, instruction prefetching is ineffective when a branch occurs. The Am29000, therefore, incorporates an on-chip Branch Target Cache to supply instructions for a branch, if this branch has been taken previously, while a new prefetch stream is established.

If branch-target instructions are in the Branch Target Cache, branches execute in a single cycle. This has a very positive effect on processor performance, due to the amount of time the processor could otherwise be idle waiting for the new instruction stream.

As an example, consider that successful branches are 20% of a dynamic instruction mix, and that 5 cycles are required to restart the processor pipeline after a branch. For 20% of the

instructions, the processor would take one cycle to execute the branch instruction and wait 5 cycles to refill the instruction pipeline. The overhead of branch instructions would be 6 cycles. If the remaining 80% of the instructions require a single cycle to execute, the latency involved in branching would reduce the average execution rate from one cycle per instruction to two, thus halving the performance.

The Branch Target Cache in the Am29000 has a average hit rate of 60%. In other words, it eliminates the branch latency for 60% of all successful branches on the average.

1.3.7 BRANCHING

Branch conditions in the Am29000 are based on Boolean data contained in general-purpose registers, rather than on arithmetic condition codes. Using a condition-code register for the purpose of branching—which is common in other processors—inhibits certain optimizations, because the condition-code register is modified by many different instructions. It is difficult for an optimizing compiler to schedule this shared use. By treating branch conditions as any other instruction operand, the Am29000 avoids this problem.

The Am29000 executes branches in a single cycle, for those cases where the target of the branch is in the Branch Target Cache. The single-cycle branch is unusual for a pipelined processor, and is due to processor hardware which allows much of the branch instruction operation to be performed early in the execution of the branch. Single-cycle branching has a dramatic effect on performance, since successful branches typically represent 15% to 25% of a processor's instruction mix.

The techniques used to achieve single-cycle branching also minimize the execution time of branches in those cases where the target is not in the Branch Target Cache. To keep the pipeline operating at the maximum rate, the instruction following the branch, referred to as the delay instruction, is executed regardless of the outcome of the branch. An optimizing compiler can define a useful instruction for the delay instruction in approximately 90% of branch instructions, thereby increasing the performance of branches.

1.3.8 LOADS AND STORES

The performance degradation of load and store operations is minimized in the Am29000 by overlapping them with instruction execution, by taking advantage of pipelining, and by organizing the flow of external data onto the processor so that the impact of external accesses is minimized.

Overlapped Loads and Stores

In the Am29000, a load or store is performed concurrently with execution of instructions which do not have dependencies on the load or store operation. An optimizing compiler can

schedule loads and stores in the instruction sequence so that, in most cases, data accesses are overlapped with instruction execution.

Overlapped load and store operations can achieve up to a 30% improvement in performance when data memory has a 2-cycle access time. Processor hardware detects dependencies while overlapped loads and stores are being performed, so that dependencies have no software implications.

A classical problem in the implementation of overlapped loads and stores is that of dealing with address-translation exceptions in a demand-paged environment. Overlap is not possible if any load or store which encounters an address-translation exception must be restarted by the re-execution of the initiating instruction. In this case, the processor would have to hold instruction execution until the success of every load or store were insured.

The Am29000 exception restart mechanism automatically saves information required to restart any load or store, until the operation successfully completes. Thus, it allows the overlapped execution of loads and stores while properly handling address-translation exceptions.

A second problem in the implementation of overlapped loads concerns the handling of data which is returned to the processor upon completion of the load. This data must be written to the register file, but it contends for register-file write-cycles with other instructions which are being overlapped with the load. This contention may be eliminated by adding a special write port to the register file. However, due to the size of the register file in the Am29000, a fourth port for writing incoming load data is not economical.

The Am29000 data-flow organization avoids the one-cycle penalty which would result from the contention between load data and the results of overlapped instruction execution. Load data is buffered in a latch while awaiting an opportunity to be written into the register file. This opportunity is guaranteed to arise before the next load is executed. While the data is buffered in this latch, it may be used as an instruction operand in place of the destination register for the load.

Load Multiple and Store Multiple

These instructions allow the transfer of the contents of multiple registers to or from external memories or devices. This transfer can occur at a rate of one register-content per cycle.

The advantage of Load Multiple and Store Multiple is best seen in task switching, register-file saving and restoring, and in block data moves. In many systems, such operations require a large percentage of execution time.

The load-multiple and store-multiple sequences are interruptible, so that they do not affect interrupt latency.

Forwarding of Load Data

Data which is sent to the processor at the completion of a load is forwarded directly to the appropriate execution unit if the data is required immediately by an instruction. This avoids the common one-cycle delay from bus transfer to use of data, and reduces the access latency of external data by one cycle.

1.3.9 MEMORY MANAGEMENT

A 64-entry Translation Look-Aside Buffer (TLB) on the Am29000 performs virtual-to-physical address translation, avoiding the cycle which would be required to transfer the virtual address to an external TLB. A number of enhancements improve the performance of address translation:

- 1) **Pipelining**—The operation of the TLB is pipelined with other processor operations.
- 2) **Early Address Translation**—Address translations for load, store, and branch instructions occur during the cycle in which these instructions are executed. This allows the physical address to be transferred externally in the next cycle.
- 3) **Task Identifiers**—Task Identifiers allow TLB entries to be matched to different processes, so that TLB invalidation is not required during task switches.
- 4) **Least-Recently-Used Hardware**—This hardware allows immediate selection of a TLB set to be replaced.
- 5) **Software Reload**—Software reload allows the operating system to use a page-mapping scheme which is best matched to its environment. Paged-segmented, one-level-page, and two-level-page mapping can be supported. Because Am29000 instructions execute at an average rate of nearly one instruction per cycle, software reload has a performance approaching that of hardware TLB reload.

1.3.10 INTERRUPTS AND TRAPS

When the Am29000 takes an interrupt or trap, it does not automatically save its current state information. This greatly improves the performance of temporary interruptions such as TLB reload, floating-point emulation, or other simple operating system calls which require no saving of state information.

In cases where the processor state must be saved, the saving and restoring of state information is under the control of software. The methods and data structures used to handle interrupts—and the amount of state saved—may be tailored to the needs of a particular system.

Interrupts and traps are dispatched through a 256-entry Vector Area, which directs the processor to a routine to handle a given interrupt or trap. The Vector Area may be relocated in memory by the modification of a processor register. There may be multiple Vector Areas in the system, though only one is active at any given time.

The Vector Area is either a table of pointers to the interrupt and trap handlers, or a segment of instruction memory (possibly read-only memory) containing the handlers themselves. The choice between the two possible Vector Area definitions is determined by the cost/performance trade-offs made for a particular system.

If the Vector Area is a table of vectors in data memory, it requires only 1 Kbyte of memory. However, this structure requires that the processor perform a vector fetch every time an interrupt or trap is taken. The vector fetch requires at least 3 cycles, in addition to the number of cycles required for the basic memory access.

If the Vector Area is a segment of instruction memory, it requires a maximum of 64 Kbytes of memory. The advantage of this structure is that the processor begins the execution of the interrupt or trap handler in the minimum amount of time.

1.3.11 FLOATING-POINT ARITHMETIC UNIT

The Am29027 is a Double-Precision, Floating-Point Arithmetic Unit for the Am29000. It can provide an order-of-magnitude performance increase over floating-point operations performed in software. It performs both single-precision and double-precision operations, using IEEE and other floating-point formats. The Am29027 also supports 32- and 64-bit integer functions.

The Am29027 performs floating-point operations using combinatorial—rather than sequential—logic, so that operations with the Am29027 require only five Am29000 cycles. Floating-point operations may be overlapped with other processor operations. Furthermore, the Am29027 incorporates pipeline registers and 8 operand registers, permitting very high throughput for certain types of operations (such as array computations).

The Am29027 attaches directly to the Am29000, using the coprocessor interface. The Am29000 can transfer two, 32-bit quantities to the Am29027 in one cycle.

The Am29027 is fully described in separate documentation.

1.4 OPTIMIZING COMPILERS

The number of instructions used to perform a given task is minimized by optimizing compilers which are supplied for the Am29000. A full discussion of optimizing-compiler technology is beyond the scope of this manual, but there are a few concepts which should be mentioned here, because the Am29000 was designed to be an excellent target for optimizing compilers.

1.4.1 OPTIMIZING-COMPILER OVERVIEW

In addition to performing the same tasks as any other compiler, an optimizing compiler re-arranges the generated code to minimize its size and execution time. This optimization occurs after the initial phases of code generation have been completed. The optimizer inspects large portions of the compiled program for frequently-occurring cases where the compiled results can be improved.

Many optimization opportunities arise precisely because the code is compiler-generated. Code translation is an automated process, so the initial phases of the compiler often generate code that is much less than optimum. However, the optimizer can produce results which are often better than those which can be produced by human assembly-language programmers, because it can deal with large portions of the program and an immense amount of data concerning program behavior.

1.4.2 OPTIMIZING-COMPILER OPERATION

Conceptually, the optimizer arranges program flow and the creation, modification, and use of program data to minimize the amount of time required to perform a given task. The reduction in program space is a normal side-benefit of the reduction in execution time. The optimizer is concerned not only with data explicit in the high-level program, but also with data created by other phases of the compiler in order to properly translate the program (for example, temporary values created during the evaluation of expressions). Optimization involves the following sorts of operations:

- 1) Reusing results rather than repeating computations. The optimizer attempts to eliminate redundant computations by performing a computation once, and saving the result for later use. Often these redundant computations are not apparent in the original program, but are created by the underlying definitions of high-level operations.
- 2) Reducing the amount of code executed within loops. In many cases, only a few computations change on different loop iterations. The optimizer attempts to reduce the amount of work performed within loops to a minimum, by moving loop-invariant computations outside of loops.

- 3) Replacing slow operations by faster ones. The optimizer can recognize special cases of multiply and divide, for example, and replace them with faster shift and add instructions. The slow operations, again, are often generated by earlier phases of the compiler because these operations are most general, and the early code-generation phases cannot recognize the special cases which allow the operations to be replaced with faster ones.
- 4) Allocating processor registers so that they contain frequently-used data. This reduces the number of relatively slow memory references, and replaces them by faster register references.
- 5) Scheduling the execution of instructions. The optimizer attempts to move instructions to a point in the program flow where they create fewer problems for the processor pipeline. For example, a register load may be moved to a point in the instruction sequence where its memory reference can be overlapped with other instructions.

Most optimizations performed rely heavily on two types of information collected by the optimizer: the first type deals with program flow, and the second with data dependencies which arise because of the program flow. The optimizer can tailor the code to the high-level task being compiled, not because it understands the task being performed by the high-level program, but because it understands the dependencies which arise in the generated code. As a result, it can adjust the instruction sequence to minimize the performance impact of these dependencies.

It is important to note that the optimizer does not directly optimize a given program, but rather optimizes a special representation of the program which is suitable for analysis and modification by the optimizer, which is, after all, just another program. The key to optimization is that this representation be easy to analyze for program- and data-flow information, and that it be easy to rearrange when optimizations are performed.

1.4.3 THE Am29000 AND OPTIMIZING COMPILERS

General Principles

The primary principle behind the Am29000 instruction set is that it matches the internal representation used by optimizing compilers to perform optimization. As discussed above, this representation is not arbitrary, but is rather strictly defined by the optimization algorithms.

It is important to realize that optimizations performed for the Am29000 would have limited effectiveness if applied to so-called complex-instruction processors. There are several fundamental problems which limit the effectiveness of optimizations for these other processors.

The first problem with complex instruction sets is that they normally provide a variety of instruction sequences which perform the same function as a sequence of instructions in the compiler's internal representation, but do not match it exactly. The trade-offs made by a compiler to decide among the available choices can get very complex.

In the first place, it is difficult for the compiler to determine the difference in execution time between multiple instruction sequences, because of the amount of information involved. For example, just changing the addressing mode of an instruction can change the execution time. This is further complicated in the cases where the compiled program is to be run on different implementations of the same processor, where execution times can depend on the implementation. If there is only one instruction sequence to choose from, and if all instructions execute in a single cycle, this problem is greatly reduced.

During the generation of code for a complex-instruction processor, it is nearly impossible to guarantee that the choice of a given code sequence will not force a less-than-optimum choice of code at some later point in the translation. Restrictions arise late in translation because of decisions made earlier. Often, these restrictions arise because of interactions between instructions; they are especially severe when instructions operate only on a specific register or group of registers.

An additional problem with complex instruction sets is that optimizations applied to them do not necessarily save execution time. An optimization may not be reflected in the final compiled code, because the instruction set may inhibit the realization of the optimization. However, in the case of the Am29000, an optimization is guaranteed to eliminate one or more execution cycles, due to the fact that all processor operations are exposed to the compiler.

The greatest benefit of exposing all processor operations to the compiler appears within loops, which is where processors spend a great deal of their execution time. The problem with complex instruction sets here is that, when an instruction set forces multiple operations with one instruction, the processor spends much time performing redundant computations within loops. Many times, the redundant computations are performed by microcode, which cannot detect that a computation is loop-invariant, because it knows nothing of loops. The compiler is in no position to do much about this, because it cannot remove the loop-invariant computations from the micro-sequence; it is forced to accept the definitions of the instructions as they are.

If an instruction set is defined so that all hardware-level operations are available to the compiler, the compiler is free to construct any sequence of these operations. This allows the movement of loop-invariant computations out of loops, which can result in tremendous performance improvements.

Special Am29000 Features

In addition to the above considerations, there are several other central principles behind the definition of the Am29000.

The Am29000 instruction set reduces the number of instructions required for most general-purpose tasks, by providing a complete set of operations. The instruction set is streamlined, but there is no attempt to minimize the number of instructions. Rather, the goal is to minimize the number of instructions required to execute most high-level language programs.

With a few minor exceptions, Am29000 instructions execute in a single cycle. As a result, the performance of an Am29000 instruction sequence is very easy to predict, simplifying the task of compiler instruction-selection. In addition, single-cycle instruction execution allows the Am29000 to take the maximum advantage of a high-performance system design. Instructions are executed at approximately the rate at which they are supplied to the processor. The Am29000 does not artificially constrain the instruction-execution rate by forcing instructions to require multiple cycles for execution.

The Am29000 contains a large number of registers which facilitate compiler optimizations. These registers allow frequently-used variables to be accessed quickly, provide a large number of temporary locations for the reuse of computational results, and simplify inter-procedural communication. The compiler is free to allocate these registers as required to improve performance. Register allocation is relatively simple, because there is such a large number of registers.

For other processors which have fixed register-addressing, a compiler has difficulties allocating the usage of registers, because registers must be allocated statically, at compile time. Procedure calls present the greatest difficulty. It is impossible for the compiler to determine exactly which procedures will be called during execution, and in what order they will be called. Thus, it is impossible to precisely allocate the usage of registers across procedure-call boundaries.

Since the Am29000 local registers are addressed relative to a Stack Pointer, compiler register-allocation is greatly simplified. The local registers are allocated dynamically, during execution. Thus, the compiler need not be concerned about the allocation of registers across procedure boundaries; this is handled automatically by the local-register addressing.

Am29000 pipelining is exposed to the compiler, in the form of delayed branches and overlapped loads and stores. The compiler is free to arrange instructions to reduce the performance impact of the processor pipeline. However, the compiler arranges instructions only because of the performance benefits. Pipeline interlocks in the Am29000 guarantee correct operation in any case.

CHAPTER 2

ARCHITECTURE HIGHLIGHTS

This chapter gives a brief overview of the Am29000 architecture, grouped into programming-related features, hardware features, and system interfaces. The technical information given in this chapter is also contained in subsequent chapters. Much of the detail is omitted here, since the objective is to provide a framework for understanding the information in later chapters.

Where appropriate, section titles in this chapter are followed by references to sections appearing in subsequent chapters. The referenced sections contain related detailed information.

2.1 PROGRAMMER REFERENCE OVERVIEW

This section gives a brief description of the Am29000 from a programmer's point of view. It introduces the processor's program modes, registers, and instructions. An overview of the processor's data formats and handling is given. This section also briefly describes interrupts and traps, memory management, and the coprocessor interface. Finally, the Timer Facility and Trace Facility are introduced.

2.1.1 PROGRAM MODES (see Section 3.1)

There are two mutually-exclusive modes of program execution; the Supervisor mode, and the User mode. In the Supervisor mode, executing programs have access to all processor resources. In the User mode, certain processor resources may not be accessed; any attempted access causes a trap.

2.1.2 VISIBLE REGISTERS (see Section 3.2)

The Am29000 incorporates three classes of registers which are accessed and manipulated by instructions: general-purpose registers, special-purpose registers, and Translation Look-Aside Buffer (TLB) registers.

General-Purpose Registers (see Section 3.2.1)

The Am29000 has 192 general-purpose registers. General-purpose registers are not dedicated to any special use, and are available for any appropriate program use.

Most processor instructions are three-address instructions. An instruction specifies any 3 of the 192 registers for use in instruction execution. Normally, two of these registers contain source-operands for the instruction, and a third stores the result of the instruction.

The 192 registers are divided into 64 global and 128 local registers. Global registers are addressed with absolute register-numbers, while local registers are addressed relative to an internal Stack Pointer.

For fast procedure calling, a portion of a compiler's run-time stack can be mapped into the local registers. Statically-allocated variables, temporary values, and operating-system parameters are kept in the global registers.

The Stack Pointer for local registers is mapped to Global Register 1. The Stack Pointer is a full 32-bit virtual address for the top of the run-time stack.

The general-purpose registers may be accessed indirectly, with the register-number specified by the content of a special-purpose register (see below) rather than by an instruction field. Three independent indirect register-numbers are contained in three separate special-purpose registers. The number for Global Register 0 specifies indirect register-addressing. An instruction can specify an indirect register access for any or all of the source operands or result.

General-purpose registers may be partitioned into segments of 16 registers for the purpose of access protection. A register in a protected segment may be accessed only by a program executing in the Supervisor mode. An attempted access (either read or write) by a User-mode program causes a trap to occur.

Special-Purpose Registers (see Section 3.2.2)

The Am29000 contains 23 special-purpose registers. These registers provide controls and data for certain processor functions.

Special-purpose registers are accessed by data movement only. Any special-purpose register can be written with the contents of any general-purpose register, and any general-purpose register can be written with the contents of any special-purpose register. Operations cannot be directly performed on the contents of special-purpose registers.

Some special-purpose registers are protected, and can be accessed only in the Supervisor mode. This restriction applies to both read and write accesses. An attempt by a User-mode program to access a protected register causes a trap to occur.

The protected special-purpose registers are defined as follows:

- 1) Vector Area Base Address—Defines the beginning of the interrupt/trap Vector Area.
- 2) Old Processor Status—Stores a copy of the Current Processor Status (see below) when an interrupt or trap is taken. It is later used to restore the Current Processor Status on an interrupt return.
- 3) Current Processor Status—Contains control information associated with the currently-executing process, such as interrupt disables and the Supervisor Mode bit.
- 4) Configuration—Contains control information which normally varies only from system to system, and is usually set only during system initialization.
- 5) Channel Address—Contains the address associated with an external access, and retains the address if the access does not complete successfully. The Channel Address Register, in conjunction with the Channel Data and Channel Control registers described below, allow the restarting of unsuccessful external accesses. This might be necessary for an access encountering a page fault in a demand-paged environment, for example.
- 6) Channel Data—Contains data associated with a store operation, and retains the data if the operation does not complete successfully.
- 7) Channel Control—Contains control information associated with a channel operation, and retains this information if the operation does not complete successfully.
- 8) Register Bank Protect—Restricts access of User-mode programs to specified groups of sixteen registers. This facilitates register banking for multi-tasking applications, and protects operating-system parameters kept in the global registers from corruption by User-mode programs.
- 9) Timer Counter—Supports real-time control and other timing-related functions.
- 10) Timer Reload—Maintains synchronization of the Timer Counter. It includes control bits for the Timer Facility.
- 11) Program Counter 0—Contains the address of the instruction being decoded when an interrupt or trap is taken. The processor restarts this instruction upon interrupt return.

- 12) Program Counter 1—Contains the address of the instruction being executed when an interrupt or trap is taken. The processor restarts this instruction upon interrupt return.
- 13) Program Counter 2—Contains the address of the instruction just completed when an interrupt or trap is taken. This address is provided for information only, and does not participate in an interrupt return.
- 14) MMU Configuration—Allows selection of various memory-management options, such as page size.
- 15) LRU Recommendation—Simplifies the reload of entries in the Translation Look-Aside Buffer (TLB) by providing information on the least-recently-used entry of the TLB when a TLB miss occurs (see Section 2.1.6).

The unprotected special-purpose registers are defined as follows:

- 1) Indirect Pointer C—Allows the indirect access of a general-purpose register.
- 2) Indirect Pointer A—Allows the indirect access of a general-purpose register.
- 3) Indirect Pointer B—Allows the indirect access of a general-purpose register.
- 4) Q—Provides additional operand bits for multiply and divide operations.
- 5) ALU Status—Contains information about the outcome of arithmetic and logical operations, and holds residual control for certain instruction operations.
- 6) Byte Pointer—Contains an index of a byte or half-word within a word. This register is also accessible via the ALU Status Register.
- 7) Funnel Shift Count—Provides a bit offset for the extraction of word-length fields from double-word operands. This register is also accessible via the ALU Status Register.
- 8) Load/Store Count Remaining—Maintains a count of the number of loads and stores remaining for load-multiple and store-multiple operations. The count is initialized to the total number of loads or stores to be performed before the operation is initiated. This register is also accessible via the Channel Control Register.

TLB Registers (see Section 3.2.3)

Translation Look-Aside Buffer (TLB) entries in the Am29000 Memory Management Unit are accessed via 128 TLB registers. A single TLB entry appears as two TLB registers; TLB registers are thus paired according to the corresponding TLB entry.

TLB registers are accessed by data movement only. Any TLB register can be written with the contents of any general-purpose register, and any general-purpose register can be written with the contents of any TLB register. Operations cannot be directly performed on the contents of TLB registers.

TLB registers can be accessed only in the Supervisor mode. This restriction applies to both read and write accesses. An attempt by a User-mode program to access a TLB register causes a trap to occur.

2.1.3 INSTRUCTION SET OVERVIEW (see Section 3.3 and Chapter 8)

The three-address architecture of the Am29000 instruction set allows a compiler or assembly-language programmer to prevent the destruction of operands, and aids register allocation and operand reuse. Instruction operands may be contained in any two of the 192 general-purpose registers, and instruction results may be stored in any of the 192 general-purpose registers.

The compiler or assembly-language programmer has complete freedom to allocate register usage. There is no dedication of a particular register or register group to a particular class of operations. The instruction set is designed to minimize the number of side-effects and implicit operations of instructions.

Most Am29000 instructions can specify an 8-bit constant as one of the source operands. Larger constants are constructed using one or two additional instructions and a general-purpose register. Relative branch instructions specify a 16-bit, signed, word offset. Absolute branches specify a 16-bit word address.

The Am29000 instruction set contains 115 instructions. These instructions are divided into 9 classes:

- 1) Integer Arithmetic—perform integer add, subtract, multiply, and divide operations.
- 2) Compare—perform arithmetic and logical comparisons. Some instructions in this class allow the generation of a trap if the comparison condition is not met.
- 3) Logical—perform a set of bit-wise Boolean operations.
- 4) Shift—perform arithmetic and logical shifts, and allow the extraction of 32-bit words from 64-bit double-words.

- 5) **Data Movement**—perform movement of data fields between registers, and the movement of data to and from external devices and memories.
- 6) **Constant**—allow the generation of large constant values in registers.
- 7) **Floating-Point**—included for floating-point arithmetic, comparisons, and format conversions. These instructions are not currently implemented directly in processor hardware.
- 8) **Branch**—perform program jumps and subroutine calls.
- 9) **Miscellaneous**—perform miscellaneous control functions and operations not provided by other classes.

The Am29000 executes all instructions in a single cycle, except for interrupt returns, Load Multiple, and Store Multiple.

Table 2-1 shows a complete list of Am29000 instructions, listed alphabetically by instruction mnemonic. Table 2-1 is provided only to give a general overview of the instruction set. Section 3.3 defines the instructions grouped into classes, and Chapter 8 provides a detailed specification of the instruction set.

2.1.4 DATA FORMATS AND HANDLING (see Section 3.4)

This section introduces the data formats and data-manipulation mechanisms which are supported by the Am29000.

Data Types (see Section 3.4.1)

A word is defined as 32 bits of data. A half-word consists of 16 bits, and a double-word consists of 64 bits. Bytes are 8 bits in length. The Am29000 has direct support for word-integer (signed and unsigned), word-logical, word-Boolean, half-word integer (signed and unsigned), and character (byte) data.

Other data types, such as character strings, are supported with sequences of basic instructions and/or external hardware. Single- and double-precision floating-point types are defined for the Am29000, but are not directly supported by hardware.

The format for Boolean data used by the processor is such that the Boolean values TRUE and FALSE are represented by 1 and 0, respectively, in the most-significant bit of a word.

Table 2-1. Am29000 Instruction Set

Mnemonic	Instruction Name
ADD	Add
ADDC	Add with Carry
ADDCS	Add with Carry, Signed
ADDCU	Add with Carry, Unsigned
ADDS	Add, Signed
ADDU	Add, Unsigned
AND	AND Logical
ANDN	AND-NOT Logical
ASEQ	Assert Equal To
ASGE	Assert Greater Than or Equal To
ASGEU	Assert Greater Than or Equal To, Unsigned
ASGT	Assert Greater Than
ASGTU	Assert Greater Than, Unsigned
ASLE	Assert Less Than or Equal To
ASLEU	Assert Less Than or Equal To, Unsigned
ASLT	Assert Less Than
ASLTU	Assert Less Than, Unsigned
ASNEQ	Assert Not Equal To
CALL	Call Subroutine
CALLI	Call Subroutine, Indirect
CLZ	Count Leading Zeros
CONST	Constant
CONSTH	Constant, High
CONSTN	Constant, Negative
CPBYTE	Compare Bytes
CPEQ	Compare Equal To
CPGE	Compare Greater Than or Equal To
CPGEU	Compare Greater Than or Equal To, Unsigned
CPGT	Compare Greater Than
CPGTU	Compare Greater Than, Unsigned
CPLE	Compare Less Than or Equal To
CPLEU	Compare Less Than or Equal To, Unsigned
CPLT	Compare Less Than
CPLTU	Compare Less Than, Unsigned
CPNEQ	Compare Not Equal To
CVDF	Convert Floating-Point Double-Precision to Single-Precision
CVDINT	Convert Floating-Point Double-Precision to Integer
CVFD	Convert Floating-Point Single-Precision to Double-Precision

Table 2-1. Am29000 Instruction Set (Continued)


Mnemonic	Instruction Name
CVFINT	Convert Floating-Point Single-Precision to Integer
CVINTD	Convert Integer to Floating-Point Double-Precision
CVINTF	Convert Integer to Floating-Point Single-Precision
DADD	Floating-Point Add, Double-Precision
DDIV	Floating-Point Divide, Double-Precision
DEQ	Floating-Point Equal To, Double-Precision
DGT	Floating-Point Greater Than, Double-Precision
DIV	Divide Step
DIV0	Divide Initialize
DIVIDE	Integer Divide
DIVL	Divide Last Step
DIVREM	Divide Remainder
DLT	Floating-Point Less Than, Double-Precision
DMUL	Floating-Point Multiply, Double-Precision
DSUB	Floating-Point Subtract, Double-Precision
EMULATE	Trap to Software Emulation Routine
EXBYTE	Extract Byte
EXHW	Extract Half-Word
EXHWS	Extract Half-Word, Sign-Extended
EXTRACT	Extract Word, Bit-Aligned
FADD	Floating-Point Add, Single-Precision
FDIV	Floating-Point Divide, Single-Precision
FEQ	Floating-Point Equal To, Single-Precision
FGT	Floating-Point Greater Than, Single-Precision
FLT	Floating-Point Less Than, Single-Precision
FMUL	Floating-Point Multiply, Single-Precision
FSUB	Floating-Point Subtract, Single-Precision
HALT	Enter Halt Mode
INBYTE	Insert Byte
INHW	Insert Half-Word
INV	Invalidate
IRET	Interrupt Return
IRETINV	Interrupt Return and Invalidate
JMP	Jump
JMPF	Jump False
JMPFDEC	Jump False and Decrement
JMPFI	Jump False Indirect
JMPI	Jump Indirect

Table 2-1. Am29000 Instruction Set (Continued)


Mnemonic	Instruction Name
JMPT	Jump True
JMPTI	Jump True Indirect
LOAD	Load
LOADL	Load and Lock
LOADM	Load Multiple
LOADSET	Load and Set
MFSR	Move from Special Register
MFTLB	Move from Translation Look-Aside Buffer Register
MTSR	Move to Special Register
MTSRIM	Move to Special Register Immediate
MTTLB	Move to Translation Look-Aside Buffer Register
MUL	Multiply Step
MULL	Multiply Last Step
MULTIPLY	Integer Multiply
MULU	Multiply Step, Unsigned
NAND	NAND Logical
NOR	NOR Logical
OR	OR Logical
SETIP	Set Indirect Pointers
SLL	Shift Left Logical
SRA	Shift Right Arithmetic
SRL	Shift Right Logical
STORE	Store
STOREL	Store and Lock
STOREM	Store Multiple
SUB	Subtract
SUBC	Subtract with Carry
SUBCS	Subtract with Carry, Signed
SUBCU	Subtract with Carry, Unsigned
SUBR	Subtract Reverse
SUBRC	Subtract Reverse with Carry
SUBRCS	Subtract Reverse with Carry, Signed
SUBRCU	Subtract Reverse with Carry, Unsigned
SUBRS	Subtract Reverse, Signed
SUBRU	Subtract Reverse, Unsigned
SUBS	Subtract, Signed
SUBU	Subtract Unsigned
XNOR	Exclusive-NOR Logical
XOR	Exclusive-OR Logical

Figure 2-1 illustrates the numbering conventions for data units contained in a word. Within a word, bits are numbered in increasing order from right-to-left, starting with the number 0 for the least-significant bit. Bytes and half-words within a word are numbered in increasing order starting with the number 0. However, bytes and half-words may be numbered right-to-left or left-to-right, as controlled by the Configuration Register.


Bytes Within Words



OR



Half-Words Within Words:



OR




Figure 2-1. Data-Unit Numbering Conventions

Note that the numbering of bits within words is strictly for notational convenience. In contrast, the numbering conventions for bytes and half-words within words affect processor operations.

External Data Accesses (see Section 3.4.2)

External accesses move data between the processor and external devices and memories. These accesses occur only as a result of load and store instructions.

Load and store instructions move words of data to and from general-purpose registers. Each load and store instruction moves a single word. There are load and store instructions which support interlocking operations necessary for multi-processor exclusion, synchronization, and communication.

For the movement of multiple words, Load Multiple and Store Multiple instructions move the contents of sequentially-addressed external locations to or from sequentially-numbered general-purpose registers. The Load Multiple and Store Multiple allow the movement of up to 192 words at a maximum rate of one word per processor cycle. The multiple load and store sequences may be interrupted, and restarted at the point of interruption.

Load and store instructions provide no mechanism for computing the address associated with the external data access. All addresses are contained in a general-purpose register at the beginning of the access, or are given by an 8-bit instruction constant. Any address computation must be performed explicitly before the load or store instruction is executed. Since address computations are expressed directly, they are exposed for compiler optimizations as any other computations are.

External data accesses are overlapped with instruction execution. Processor performance is improved if instructions which follow loads do not immediately use externally-referenced data. In this manner, the time required to perform the external access is overlapped with subsequent instruction execution. Because of hardware interlocks, this concurrency has no effect on the logical behavior of an executing program.

Addressing and Alignment (see Section 3.4.3)

External instructions and data are contained in one of four, 32-bit address-spaces:

- 1) Instruction/Data Memory.
- 2) Input/Output.
- 3) Coprocessor.
- 4) Instruction Read-Only Memory (Instruction ROM).

An address in the Instruction/Data Memory address-space may be treated as virtual or physical, as determined by the Current Processor Status Register. Address translation for data accesses is enabled separately from address translation for instruction accesses. A program in the Supervisor mode may temporarily disable address translation for individual loads and stores; this permits load-real and store-real operations.

Bits contained within load and store instructions distinguish between the Instruction/Data Memory, Input/Output, and Coprocessor address-spaces. The Current Processor Status register determines whether instruction accesses are directed to the Instruction/Data Memory address-space or to the Instruction ROM address-space.

The Am29000 does not directly support data accesses to the Instruction ROM address-space. However, this capability is supported as a system option.

All addresses are interpreted as byte addresses, even though accesses are word-oriented. The number of a byte within a word is given by the two least-significant address bits. The number of a half-word within a word is given by the next-to-least-significant address bit.

The Am29000 supports the direct external access of bytes and half-words as a system option requiring external alignment hardware. However, the processor also supports mechanisms for the indirect external access of bytes and half-words. These mechanisms obviate the need for alignment hardware in the majority of systems.

The Am29000 sets a byte-position indicator in the ALU Status Register—as an option for load instructions—with the two least-significant bits of the address for the load. To load a byte or half-word, a word load is first performed. This load sets the byte-position indicator, and a subsequent instruction extracts the byte or half-word of interest from the accessed word. To store a byte or half-word, a load is also first performed; the byte or half-word of interest is inserted into the accessed word, and the resulting word is then stored.

Since only byte addressing is supported, it is possible that an address for the access of a word or half-word is not aligned to the desired word or half-word. For a word access, an unaligned address has a 1 in either or both of the two least-significant address bits. For a half-word access, an unaligned address has a 1 in the least-significant address bit.

In many systems, address alignment can simply be ignored, with addresses truncated to access the word or half-word of interest. However, as a user option, the Am29000 creates a trap when a non-aligned access is attempted. The trap allows software emulation of non-aligned accesses.

In the Am29000, all instructions are 32-bits in length, and are aligned on word-address boundaries.

2.1.5 INTERRUPTS AND TRAPS (see Section 3.5)

Normal program flow may be preempted by an interrupt or trap for which the processor is enabled. The effect on the processor is identical for interrupts and traps; the distinction is in the different mechanisms by which interrupts and traps are enabled. It is intended that interrupts be used for suspending current program execution and causing another program to execute, while traps are used to report errors and exceptional conditions.

The interrupt and trap mechanism supports high-speed, temporary context switching as well as user-defined interrupt-processing mechanisms.

Temporary Context Switching

The basic interrupt/trap mechanism of the Am29000 supports temporary context switching. During the temporary context switch, the interrupted context is held in processor registers. The interrupt or trap handler can return immediately to this context.

Temporary context switching is useful for instruction emulation, floating-point operations, TLB reload routines, and so forth. It is similar to an invocation of a microprogram. Many of its features are similar to microprogram execution: processor context does not have to be saved, interrupts are disabled for the duration of the program, and all processor resources are available, even if the context which was interrupted is in the User mode. The associated routine may execute from instruction/data memory or instruction ROM.

User-Defined Interrupt Processing

Since the basic interrupt/trap mechanism for the Am29000 keeps the interrupted context in the processor, dynamically nested interrupts are not directly supported. The context in the processor must be saved before another interrupt or trap can be taken.

The interrupt or trap handler executing during a temporary context switch is not required to return to the interrupted context. This routine may optionally save the interrupted context, load a new one, and return to the new context.

The implementation of the saving and restoring of contexts is completely user-defined. Thus, the context save/restore mechanism used (e.g. interrupt stack, program status word area, etc.) and the amount of context saved may be tailored to the needs of the system.

Vector Area (see Section 3.5.4)

Interrupt and trap dispatching occurs through a relocatable Vector Area which accommodates as many as 256 interrupt and trap handling routines. Entries into the Vector Area are associated with various sources of interrupts and traps; some are pre-defined, while others are user-defined.

The Vector Area is either a table of vectors in data memory, where each vector points to the beginning of an interrupt or trap handler, or it is a segment of instruction/data memory (or instruction ROM) containing the actual routines. The latter configuration for the Vector Area yields better interrupt performance, with the cost of additional memory.

2.1.6 MEMORY MANAGEMENT (see Section 3.6)

The Am29000 incorporates a Memory Management Unit (MMU) which accepts a 32-bit virtual byte-address and translates it to a 32-bit physical byte-address in a single cycle. The MMU is not dedicated to any particular address-translation architecture.

Address translation in the MMU is performed by a 64-entry Translation Look-Aside Buffer (TLB), an associative table which contains the most-recently-used address translations for the processor. If the translation for a given address cannot be performed by the TLB, a TLB miss occurs, and causes a trap which allows the required translation to be placed into the TLB.

Processor hardware maintains information for each TLB line indicating which entry was least recently used; when a TLB miss occurs, this information is used to indicate the TLB entry to be replaced. Software is responsible for searching system page tables and modifying the indicated TLB entry as appropriate. This allows the page tables to be defined according to the system environment.

TLB entries are modified directly by processor instructions. A TLB entry consists of 64 bits and appears as two word-length TLB registers which may be inspected and modified by instructions.

TLB entries are tagged with a Task Identifier field, which allows the operating system to create a unique 32-bit virtual address-space for each of 256 processes. In addition, TLB entries provide support for memory protection and user-defined control information.

2.1.7 COPROCESSOR PROGRAMMING (see Section 6.1)

The coprocessor interface for the Am29000 allows a program to communicate with an off-chip coprocessor for performing operations not directly supported by processor hardware.

The coprocessor interface allows the program to transfer operands and operation codes to the coprocessor, and then perform other operations while the coprocessor operation is in progress. The results of the operation are read from the coprocessor by a separate transfer. The processor may transfer multiple operands to the coprocessor without re-transferring operation codes or reading intermediate results. As many as 64 bits of information can be transferred to the coprocessor in a single cycle.

The Am29000 includes features which support the definition of the coprocessor as a system option. In this case, coprocessor operations are emulated by software when the coprocessor is not present in a system.

2.1.8 TIMER FACILITY (see Section 7.2.7)

The Timer Facility provides a counter for implementing a real-time clock or other software timing functions. This facility is comprised of two special-purpose registers: the Timer Counter Register, which decrements at a rate equal to the processor operating

frequency, and the Timer Reload Register, which re-initializes the Timer Counter Register when it decrements to zero. The Timer Facility may optionally create an interrupt when the Timer Counter decrements to zero.

2.1.9 TRACE FACILITY (see Section 7.2.8)

The Trace Facility allows a debug program to emulate single instruction stepping in a program under test. This facility allows a trap to be generated after the execution of any instruction in the program being tested.

Using the Trace Facility, the debug program can inspect and modify the state of the program at every instruction boundary. The Trace Facility is designed to work properly in the presence of normal system interrupts and traps.

2.2 HARDWARE OVERVIEW

This section briefly describes the operation of Am29000 hardware. It introduces the processor pipeline and the three major internal functional units: the Instruction Fetch Unit, the Execution Unit, and the Memory Management Unit. Finally, the processor's operational modes are described.

Figure 2-2 shows the Am29000 internal data-flow organization. The following sections refer to the various components on this data-flow diagram.

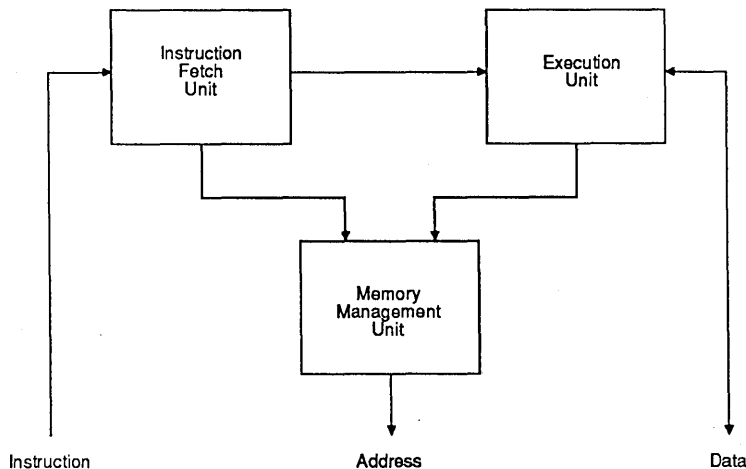


Figure 2-2. Am29000 Data Flow

2.2.1 FOUR-STAGE PIPELINE (see Section 4.1)

The Am29000 implements a four-stage pipeline for instruction execution. The four stages are: fetch, decode, execute, and write-back. The pipeline is organized so that the effective instruction-execution rate is as high as one instruction per cycle. Data forwarding and pipeline interlocks are handled by processor hardware.

2.2.2 INSTRUCTION FETCH UNIT (see Section 4.2)

The Instruction Fetch Unit fetches instructions, and supplies instructions to other functional units. It incorporates the Instruction Prefetch Buffer, the Branch Target Cache, and the Program Counter Unit. All components of the Instruction Fetch Unit operate during the fetch stage of the processor pipeline.

Instruction Prefetch Buffer (see Section 4.2.1)

Most instructions executed by the Am29000 are fetched from external instruction/data memory. The processor prefetches instructions so that they are requested at least four cycles before they are required for execution.

Prefetched instructions are stored in a four-word Instruction Prefetch Buffer while awaiting execution. An instruction-prefetch request occurs whenever there is a free location in this buffer (if the processor is otherwise enabled to fetch instructions). When a non-sequential instruction fetch occurs, prefetching is terminated, and then restarted for the new instruction stream.

Instruction prefetching de-couples the instruction-fetch rate from the instruction-access latency. For example, an instruction may be transferred to the processor two cycles after it is requested. However, as long as instructions are supplied to the processor at an average rate of one instruction per cycle, this latency has no effect on the instruction-execution rate.

Branch Target Cache (see Section 4.2.2)

The Am29000 incorporates a Branch Target Cache which contains as many as 128 instructions. The Branch Target Cache is a two-way, set-associative cache containing the first four target instructions of a number of recently-taken branches. Each of the two sets in the Branch Target Cache contains 64 instructions, and the 64 instructions are further divided into 16 blocks of 4 instructions each.

The purpose of the Branch Target Cache is to provide instructions for the beginning of a non-sequential instruction fetch sequence. This keeps the instruction pipeline full until the processor can establish a new instruction-prefetch stream from the external instruction/data memory.

The processor is organized so that branch instructions can execute in a single cycle if the target instruction sequence is present in the Branch Target Cache.

Program Counter Unit (see Section 4.2.4)

The Program Counter Unit creates and sequences addresses of instructions as they are executed by the processor.

2.2.3 EXECUTION UNIT (see Section 4.3)

The Execution Unit executes instructions. It incorporates the Register File, the Address Unit, the Arithmetic/Logic Unit, the Field Shift Unit, and the Prioritizer. The Register File and Address Unit operate during the decode stage of the pipeline. The Arithmetic/Logic Unit, Field Shift Unit, and Prioritizer operate during the execute stage of the pipeline. The Register File operates during the write-back stage.

Register File (see Section 4.3.1)

The general-purpose registers are implemented by a 192-location Register File. The Register File can perform two read accesses and one write access in a single cycle. Normally, two read accesses are performed during the decode pipeline-stage to fetch operands required by the instruction being decoded. The write access during the same cycle completes the write-back stage of a previously-executed instruction.

Addressing logic associated with the Register File distinguishes between the global and local general-purpose registers, and it performs the Stack-Pointer addressing for the local registers. Register File addressing functions are performed during the decode stage.

Address Unit (see Section 4.3.2)

The Address Unit evaluates addresses for branches, loads, and stores. It also assembles instruction-immediate data and computes addresses for load-multiple and store-multiple sequences.

Arithmetic/Logic Unit (see Section 4.3.3)

The ALU performs all logical, compare, and arithmetic operations (including multiply step and divide step).

Field Shift Unit (see Section 4.3.4)

The Field Shift Unit performs N-bit shifts. The Field Shift Unit also performs byte and half-word extract and insert operations, and it extracts words from double-words.

Prioritizer (see Section 4.3.5)

The Prioritizer provides a count of the number of leading zero bits in a 32-bit word; this is useful for performing floating-point normalization, for example.

2.2.4 MEMORY MANAGEMENT UNIT (see Section 4.4)

The Memory Management Unit (MMU) performs address translation and memory-protection functions for all branches, loads, and stores. The MMU operates during the execute stage of the pipeline, so the physical address which it generates is available at the beginning of the write-back stage.

All addresses for external accesses are physical addresses. MMU operation is pipelined with external accesses, so that an address translation can occur while a previous access completes.

Address translation is not performed for the addresses associated with instruction prefetching. Instead, these addresses are generated by an instruction prefetch pointer which is incremented by the processor. Address translation is performed only at the beginning of the prefetch sequence (as the result of a branch instruction), and when the prefetch pointer crosses a potential virtual-page boundary.

2.2.5 PROCESSOR MODES

The Am29000 operates in several different modes to accomplish various processor and system functions. All modes except for Pipeline Hold (see below) are under direct control of instructions and/or processor control inputs. The Pipeline Hold mode is normally determined by the relative timing between the processor and its external system for certain types of operations. The processor provides an external indication of its operational mode.

Executing

When the processor is in the Executing mode, it fetches and executes instructions as described in this manual. External accesses occur as required.

Wait (see Section 3.5.3)

When the processor is in the Wait mode, it does not execute instructions, and performs no external accesses. The Wait mode is controlled by the Current Processor Status Register. The processor leaves this mode when an interrupt or trap for which it is enabled occurs, or when a reset occurs.

Pipeline Hold (see Section 4.5)

Under certain conditions, processor pipelining might cause non-sequential instruction execution or timing-dependent results of execution. For example, the processor might attempt to execute an instruction which has not been fetched from instruction/data memory.

For such cases, pipeline interlock hardware detects the anomalous condition and suspends processor execution until execution can proceed properly. While execution is suspended by the interlock hardware, the processor is in the Pipeline Hold mode. The processor resumes execution when the pipeline interlock hardware determines that it is correct to do so.

Halt (see Section 5.3.3)

The Halt mode is provided so that the processor may be placed under the control of a hardware-development system for the purposes of hardware and software debug. The processor enters the Halt mode as the result of instruction execution, or as the result of external controls. In the Halt mode, the processor neither fetches nor executes instructions.

Step (see Section 5.3.3)

The Step mode allows a hardware-development system to step through processor pipeline operation on a stage-by-stage basis. The Step mode is nearly identical to the Halt mode, except that it enables the processor to enter the Executing mode while the pipeline advances by one stage.

Load Test Instruction (see Section 5.3.3)

The Load Test Instruction mode permits a hardware-development system to access data contained in the processor or system. This is accomplished by allowing the hardware-development system to supply the processor with instructions, instead of having the processor fetch instructions from instruction/data memory. The Load Test Instruction mode is defined so that, once the processor has completed the execution of instructions provided by the hardware-development system, it may resume the execution of its normal instruction sequence.

Test (see Section 5.3.4)

The Test mode facilitates testing of hardware associated with the processor, by disabling processor outputs so that they may be driven directly by test hardware. It also allows the addition of a second processor to a system, to monitor the outputs of the first and signal detected errors.

Reset (see Section 3.8 and Section 5.5)

The Reset mode provides initialization of certain processor registers and control state. This is used for power-on reset, for eliminating unrecoverable error conditions, and for supporting certain hardware debug functions.

2.3 SYSTEM INTERFACE OVERVIEW

This section briefly describes the features of the Am29000 which allow it to be connected to other system components.

The two major interfaces of the Am29000, introduced in this section, are the channel and the Test/Development interface. The other topics briefly described here are clock generation, master/slave checking, and coprocessor attachment.

Section 5.1 contains a complete pin description of the Am29000. Appendix A contains timing diagrams and related information.

2.3.1 CHANNEL (see Section 5.2)

The Am29000 channel consists of the following 32-bit buses and related controls:

- 1) An Instruction Bus, which transfers instructions into the processor.
- 2) A Data Bus, which transfers data to and from the processor.
- 3) An Address Bus, which provides addresses for both instruction and data accesses. The Address Bus is also used to transfer data to the coprocessor.

The channel performs accesses and data-transfers to all external devices and memories, including instruction/data memories, instruction caches, instruction read-only memories, data caches, input/output devices, bus converters, and coprocessors.

The channel defines three different access protocols. For simple accesses, the Am29000 holds the address valid throughout the entire access. This is appropriate for high-speed devices which can complete an access in one cycle, and for low-cost devices which are accessed infrequently (such as read-only memories containing initialization routines). For high performance with other types of devices and memories, the channel provides pipelined and burst-mode access protocols.

In the case of pipelined accesses, the address transfer is decoupled from the corresponding data or instruction transfer. After transmitting an address for a request, the processor may transmit one more address before receiving the reply to the first request. This allows address transfer and decoding to be overlapped with another access.

On the other hand, burst-mode accesses eliminate the address-transfer cycle completely. Burst-mode accesses are defined so that once an address is transferred for a given access, subsequent accesses to sequentially-increasing addresses may occur without re-transfer of the address. The burst may be terminated at any time by either the processor or responding device.

The Am29000 determines whether an access is simple, pipelined, or burst-mode on a transfer-by-transfer (i.e. generally device-by-device) basis. However, an access which begins as a simple access may be converted to a pipelined or burst-mode access at any time during the transfer. This relaxes the timing constraints on the channel-protocol implementation, since addressed devices do not have to respond immediately to a pipelined or burst-mode request.

Except for the shared Address Bus, the channel maintains a strict division between instruction and data accesses. In the most common situation, the system supplies the processor with instructions using burst-mode accesses, with instruction addresses transmitted to the system only when a branch occurs. Data accesses can occur simultaneously without interfering with instruction transfer.

The Am29000 contains arbitration logic to support other masters on the channel. A single external master can arbitrate directly for the channel, while multiple masters may arbitrate using a daisy-chain or other method which requires no additional arbitration logic. However, to increase arbitration performance in a multiple-master configuration, an external channel arbiter should be used. This arbiter works in conjunction with the processor's arbitration logic.

2.3.2 TEST/DEVELOPMENT INTERFACE (see Section 5.3)

The Am29000 supports the attachment of a general-purpose hardware-development system. This attachment is made directly to the processor in the system under development, without the removal of the processor from the system. The Test/Development Interface makes it possible for the hardware-development system to gain control over the Am29000, and inspect and modify its internal state (e.g. general-purpose register contents, TLB entries, etc.). In addition, the Am29000 may be used to access other system devices and memories on behalf of the hardware-development system.

The Test/Development Interface is comprised of controls and status signals provided on the Am29000, as well as the Instruction and Data buses. The Halt, Step, Reset, and Load Test Instruction modes allow the hardware-development system to control the operation of

the Am29000. The hardware-development system may supply the processor with instructions on the Instruction Bus using the Load Test Instruction mode. Internal processor state can be inspected and modified via the Data Bus.

2.3.3 CLOCKS (see Section 5.7)

The Am29000 generates and distributes a system clock at its operating frequency. This clock is specially designed to reduce skews between the system clock and the processor's internal clocks. The internal clock-generation circuitry requires a single-phase oscillator signal at twice the processor operating frequency.

For systems in which processor-generated clocks are not appropriate, the Am29000 also can accept a clock from an external clock-generator.

The processor decides between these two clocking arrangements, based on whether the power supply to the clock-output driver is tied to +5 volts or to GROUND.

2.3.4 MASTER/SLAVE OPERATION (see Section 5.8)

Each Am29000 output has associated logic which compares the signal on the output with the signal which the processor is providing internally to the output driver. The processor signals situations where the output of any enabled driver does not agree with its input.

For a single processor, the output comparison detects short circuits in output signals, but does not detect open circuits. It is possible to connect a second processor in parallel with the first, where the second processor has its outputs disabled due to the Test mode. The second processor detects open-circuit signals, as well as providing a check of the outputs of the first processor.

2.3.5 COPROCESSOR ATTACHMENT (see Section 6.2)

A coprocessor for the Am29000 attaches directly to the processor channel. However, this attachment has features which are different than those of other channel devices. The coprocessor interface is designed to support a high operand-transfer rate and to support the overlap of coprocessor operations with other processor operations, including other external accesses.

The coprocessor is assigned a special address-space on the channel. This permits the transfer of operands and other information on the Address Bus without interfering with normal addressing functions. Since both the Address Bus and Data Bus are used for data transfer, the Am29000 can transfer 64 bits of information to the coprocessor in one cycle.

CHAPTER 3

PROGRAMMER REFERENCE

This chapter contains a formal description of the Am29000 architecture. It concentrates on the features of the Am29000 and their logical behavior. Chapter 7 discusses the use of some of these features.

3.1 PROGRAM MODES

All system-protection features of the Am29000 are based on two mutually-exclusive program modes: the Supervisor mode, and the User mode.

3.1.1 SUPERVISOR MODE

The processor is in the Supervisor mode whenever the Supervisor Mode (SM) bit of the Current Processor Status Register (see Section 3.2.2) is 1. In this mode, executing programs have access to all processor resources.

During the address cycle of a channel request, the Supervisor mode is indicated by the SUP/*US output being High.

3.1.2 USER MODE

The processor is in the User mode whenever the SM bit in the Current Processor Status Register is 0. In this mode, any of the following actions by an executing program causes a Protection Violation trap to occur:

- 1) An attempted access of any TLB entry (see Section 3.2.3).
- 2) An attempted access of any general-purpose register for which a bit in the Register Bank Protect Register is 1 (see Section 3.2.1).
- 3) An attempted execution of a load or store instruction for which the PA bit is 1 (see Section 3.4.2).
- 4) An attempted execution of one of the following instructions: Interrupt Return, Interrupt Return and Invalidate, Invalidate, or Halt.

- 5) An attempted access of one of the following registers: Vector Area Base Address, Old Processor Status, Current Processor Status, Configuration, Channel Address, Channel Data, Channel Control, Register Bank Protect, Timer Counter, Timer Reload, Program Counter 0, Program Counter 1, Program Counter 2, MMU Configuration, LRU Recommendation (see Section 3.2.2).
- 6) An attempted execution of an assert or Emulate instruction which specifies a vector number between 0 and 63, inclusive (see Section 3.5.4).

Devices and memories on the channel can also generate traps based on the value of the SM bit. During the address cycle of a channel request, the User mode is indicated by the SUP/*US output being Low.

3.2 VISIBLE REGISTERS

The Am29000 has three classes of registers which are accessible by instructions. These are: general-purpose registers, special-purpose registers, and Translation Look-Aside Buffer (TLB) registers. Any operation available in the Am29000 can be performed on the general-purpose registers, while special-purpose registers and TLB registers are accessed only by explicit data movement to or from general-purpose registers. Various protection mechanisms prevent the access of some of these registers by User-mode programs.

A summary of the information in this section appears in Appendix B.

3.2.1 GENERAL-PURPOSE REGISTERS

The Am29000 incorporates 192 general-purpose registers. The organization of the general-purpose registers is diagrammed in Figure 3-1.

General-purpose registers hold the following types of operands for program use:

- 1) 32-bit data addresses
- 2) 32-bit signed or unsigned integers
- 3) 32-bit branch-target addresses
- 4) 32-bit logical bit strings
- 5) 8-bit characters
- 6) 16-bit signed or unsigned integers
- 7) word-length Booleans
- 8) single-precision floating-point numbers
- 9) double-precision floating-point numbers (in two register locations).

Because a large number of general-purpose registers is provided, a large amount of frequently-used data can be kept on-chip, where access time is fastest.




Figure 3-1. General-Purpose Register Organization

Am29000 instructions can specify two general-purpose registers for instruction source-operands, and one general-purpose register for storing the instruction result. These registers are specified by three 8-bit instruction fields containing register-numbers. A register may be specified directly by the instruction, or indirectly by one of three special-purpose registers.

Register Addressing

The general-purpose registers are partitioned into 64 global registers and 128 local registers, differentiated by the most-significant bit of the register-number. The distinction between global and local registers is the result of register-addressing considerations.

The following terminology is used to describe the addressing of general-purpose registers:

- 1) Register-number—this is a software-level number for a general-purpose register. For example, this is the number contained in an instruction field. Register-numbers range from 0 to 255.
- 2) Global register-number—this is a software-level number for a global register. Global register-numbers range from 0 to 127.
- 3) Local register-number—this is a software-level number for a local register. Local register-numbers range from 0 to 127.
- 4) Absolute register-number—this is a hardware-level number used to select a general-purpose register in the Register File. Absolute register-numbers range from 0 to 255.

Global Registers

When the most-significant bit of a register-number is 0, a global register is selected. The seven least-significant bits of the register-number give the global register-number. For global registers, the absolute register-number is equivalent to the register-number.

Global registers 2 through 63 are unimplemented. An attempt to access these registers yields unpredictable results; however, they may be protected from User-mode access by the Register Bank Protect Register (see below).

The register-numbers associated with Global Registers 0 and 1 have special meaning. The number for Global Register 0 specifies that an indirect pointer is to be used as the source of the register-number; there is an indirect pointer for each of the instruction operand/result registers. Global Register 1 contains the Stack Pointer, which is used in the addressing of local registers as explained below.

Local Register Stack Pointer

The Stack Pointer is a 32-bit register that may be an operand of an instruction as any other general-purpose register. However, a shadow-copy of Global Register 1 is maintained by processor hardware to be used in local-register addressing. This shadow-copy is set only with the results of Arithmetic and Logical instructions. If the Stack Pointer is set with the result of any other instruction class, local registers cannot be accessed predictably until the Stack Pointer is once again set with an Arithmetic or Logical instruction.

A modification of the Stack Pointer has a delayed effect on the addressing of local registers, as discussed in Section 7.3.3.

Local Registers

When the most-significant bit of a register-number is 1, a local register is selected. The seven least-significant bits of the register-number give the local register-number. For local registers, the absolute register-number is obtained by adding the local register-number to bits 8-2 of the Stack Pointer, truncating the result to seven bits; the most-significant bit of the original register-number is unchanged (i.e. it remains a 1).

The Stack Pointer addition applied to local register-numbers provides a limited form of base-plus-offset addressing within the local registers. The Stack Pointer contains the 32-bit base-address. This assists run-time storage management of variables for dynamically-nested procedures (see Section 7.1.1).

Register Banking

For the purpose of access restriction, the general-purpose registers are divided into register banks. Register banks consist of 16 registers (except for Bank 0, which contains unimplemented registers 2 through 15), and are partitioned according to absolute register-numbers, as shown in Figure 3-2.

The Register Bank Protect Register contains 16 protection bits, where each bit controls User-mode accesses (read or write) to a bank of registers. Bits 0–15 of the Register Bank Protect Register protect register banks 0 through 15, respectively.

When a bit in the Register Bank Protect Register is 1, and a register in the corresponding bank is specified as an operand-register or result-register by a User-mode instruction, a Protection Violation trap occurs. Note that protection is based on absolute register-numbers; in the case of local registers, Stack-Pointer addition is performed before protection checking.

When the processor is in Supervisor mode, the Register Bank Protect Register has no effect on general-purpose-register accesses.

Indirect Accesses

Specification of Global Register 0 as an instruction operand-register or result-register causes an indirect access to the general-purpose registers. In this case, the absolute register-number is provided by an indirect pointer contained in a special-purpose register.

Each of the three possible registers for instruction execution has an associated 8-bit indirect pointer. Indirect register-numbers can be selected independently for each of the three operands. Since the indirect pointers contain absolute register-numbers, the number in an indirect pointer is not added to the Stack Pointer when local registers are selected.

Register Bank Protect Register Bit	Absolute Register-Numbers	General-Purpose Registers
0	2 thru 15	Bank 0 (unimplemented)
1	16 thru 31	Bank 1 (unimplemented)
2	32 thru 47	Bank 2 (unimplemented)
3	48 thru 63	Bank 3 (unimplemented)
4	64 thru 79	Bank 4
5	80 thru 95	Bank 5
6	96 thru 111	Bank 6
7	112 thru 127	Bank 7
8	128 thru 143	Bank 8
9	144 thru 159	Bank 9
10	160 thru 175	Bank 10
11	176 thru 191	Bank 11
12	192 thru 207	Bank 12
13	208 thru 223	Bank 13
14	224 thru 239	Bank 14
15	240 thru 255	Bank 15

Figure 3-2. Register Bank Organization

The indirect pointers are set by the Move To Special Register, Floating-Point, MULTIPLY, DIVIDE, SETIP, and EMULATE instructions.

For a Move To Special Register instruction, an indirect pointer is set with bits 9-2 of the 32-bit source operand. This provides consistency between the addressing of words in general-purpose registers and the addressing of words in external devices or memories. A modification of an indirect pointer using a Move To Special Register has a delayed effect on the addressing of general-purpose registers, as discussed in Section 7.3.3.

For the remaining instructions, all three indirect pointers are set, simultaneously, with the absolute register-numbers derived from the register-numbers specified by the instruction. For any local registers selected by the instruction, the Stack-Pointer addition is applied to the register-numbers before the indirect pointers are set.

Register-numbers stored into the indirect pointers are checked for bank-protection violations, except when an indirect pointer is set by a Move-To-Special-Register instruction.

3.2.2 SPECIAL-PURPOSE REGISTERS

The Am29000 contains 23 special-purpose registers. The organization of the special-purpose registers is shown in Figure 3-3.

Special-purpose registers provide controls and data for certain processor operations. Some special-purpose registers are updated dynamically by the processor, independent of software controls. Because of this, a read of a special-purpose register following a write does not necessarily get the data which was written.

Some special-purpose registers have fields which are reserved for future processor implementations. When a special-purpose register is read, a bit in a reserved field is read as a 0. An attempt to write a reserved bit with a 1 has no effect; however, this should be avoided, because of upward-compatibility considerations.

The special-purpose registers are accessed by explicit data movement only. Instructions which move data to or from a special-purpose register specify the special-purpose register by an 8-bit field containing a special-purpose register-number. Register-numbers are specified directly by instructions.

An attempted read of an unimplemented special-purpose register yields an unpredictable value. An attempted write of an unimplemented special-purpose register has no effect; however, this should be avoided, because of upward-compatibility considerations.

The special-purpose registers are partitioned into protected and unprotected registers, differentiated by the most-significant bit of the register-number. If most-significant bit of the register-number is 0, the register is protected. If the most-significant bit of the register-number is 1, the register is unprotected.

**Special
Purpose
Reg. No.**

Protected Registers

0	Vector Area Base Address
1	Old Processor Status
2	Current Processor Status
3	Configuration
4	Channel Address
5	Channel Data
6	Channel Control
7	Register Bank Protect
8	Timer Counter
9	Timer Reload
10	Program Counter 0
11	Program Counter 1
12	Program Counter 2
13	MMU Configuration
14	LRU Recommendation

Unprotected Registers

128	Indirect Pointer C
129	Indirect Pointer A
130	Indirect Pointer B
131	Q
132	ALU Status
133	Byte Pointer
134	Funnel Shift Count
135	Load/Store Count Remaining

08996A 46

Figure 3-3. Special-Purpose Registers

Protected special-purpose registers are accessible only by programs executing in the Supervisor mode. An attempted read or write of a protected special-purpose register by a User-mode program causes a Protection Violation trap to occur.


Unprotected special-purpose registers are accessible by programs executing in both the User and Supervisor modes.

Vector Area Base Address (Register 0)

This protected special-purpose register specifies the beginning address of the interrupt/trap Vector Area. The Vector Area is either a table of 256 vectors which point to interrupt and trap handling routines, or a segment of 256, 64-instruction blocks which directly contain the interrupt and trap handling routines.

The organization of the Vector Area is determined by the Vector Fetch (VF) bit of the Configuration Register. If the VF bit is 1 when an interrupt or trap is taken, the vector number for the interrupt or trap (see Section 3.5.4) replaces bits 9-2 of the value in the Vector Area Base Address Register (Figure 3-4) to generate the physical address for a vector contained in instruction/data memory.

If the VF bit is 0, the vector number replaces bits 15-8 of the value in the Vector Area Base Address Register to generate the physical address of the first instruction of the interrupt or trap handler. The instruction fetch for this instruction is directed either to instruction memory or instruction read-only memory as determined by the ROM Vector Area (RV) bit of the Configuration Register.



08996A 1

Figure 3-4. Vector Area Base Address Register

Bits 31–16 : Vector Area Base (VAB)—The VAB field gives the beginning address of the Vector Area. This address is constrained to begin on a 64-Kbyte address-boundary in instruction/data memory or instruction read-only memory.

Bits 15–0 : Zeros—These bits force the alignment of the Vector Area.

Old Processor Status (Register 1)

This protected special-purpose register has the same format as the Current Processor Status described below. The Old Processor Status stores a copy of the Current Processor Status when an interrupt or trap is taken. This is required since the Current Processor

Status will be modified to reflect the status of the interrupt/trap handler.

During an interrupt return, the Old Processor Status is copied into the Current Processor Status. This allows the Current Processor Status to be set as required for the routine which is the target of the interrupt return.

Current Processor Status (Register 2)

This protected special-purpose register (see Figure 3-5) controls the behavior of the processor and its ability to recognize exceptional events.




Figure 3-5. Current Processor Status Register

Bits 31–16 : reserved.

Bit 15 : Coprocessor Active (CA)—The CA bit is set and reset under the control of load and store instructions which transfer information to and from a coprocessor. This bit indicates that the coprocessor is performing an operation at the time that an interrupt or trap is taken. This notifies the interrupt or trap handler that the coprocessor contains state information to be preserved. Note that this notification occurs because the CA bit of the Old Processor Status is 1 in this case, not because of the value of the CA bit of the Current Processor Status.

Bit 14 : Interrupt Pending (IP)—This bit allows software to detect the presence of external interrupts while they are disabled. The IP bit is set if one or more of the external signals *INTR0—*INTR3 is active, but the processor is disabled from taking the resulting interrupt due to the value of the DA, DI, or IM bits. If all external interrupt signals are subsequently de-asserted while still disabled, the IP bit is reset.

Bits 13–12 : Trace Enable, Trace Pending (TE, TP)—The TE and TP bits implement a software-controlled, instruction single-step facility. Single-stepping is not implemented directly, but rather emulated by trap sequences controlled by these bits. The value of the TE bit is copied to the TP bit whenever an instruction completes execution. When the TP bit is 1, a Trace trap occurs. Section 7.2.8 describes the use of these bits in more detail.

Bit 11 : Trap Unaligned Access (TU)—The TU bit enables checking of address alignment for external data-memory accesses. When this bit is 1, an Unaligned Access

trap occurs if the processor either generates an address for an external word which is not aligned on a word address-boundary (i.e. either of the least-significant two bits is 1), or generates an address for an external half-word which is not aligned on a half-word address-boundary (i.e. the least-significant address bit is 1). When the TU bit is 0, data-memory address alignment is ignored.

Alignment is ignored for input/output accesses and coprocessor transfers. The alignment of instruction addresses is also ignored (unaligned instruction addresses can be generated only by indirect jumps). Interrupt/trap vector addresses are always aligned properly.

Bit 10 : Freeze (FZ)—The FZ bit prevents certain registers from being updated during interrupt and trap processing, except by explicit data movement. The affected registers are: Channel Address, Channel Data, Channel Control, Program Counter 0, Program Counter 1, Program Counter 2, and the ALU Status Register.

When the FZ bit is 1, these registers hold their values. An affected register can be changed only by a Move To Special Register instruction. When the FZ bit is 0, there is no effect on these registers, and they are updated by processor instruction execution as described in this manual.

The FZ bit is set whenever an interrupt or trap is taken, holding critical state in the processor so that it is not used unintentionally by the interrupt or trap handler.

Bit 9 : Lock (LK)—The LK bit controls the value of the *LOCK external signal. If the LK bit is 1, the *LOCK signal is active. If the LK bit is 0, the *LOCK signal is controlled by the execution of the instructions Load and Set, Load and Lock, and Store and Lock. This bit is provided for the implementation of multi-processor synchronization protocols.

Bit 8 : ROM Enable (RE)—The RE bit enables instruction fetching from external instruction read-only memory (ROM). When this bit is 1, the IREQT signal directs all instruction requests to ROM. Instructions which are fetched from ROM are subject to capture and re-use by the Branch Target Cache when it is enabled; the Branch Target Cache distinguishes between instructions from ROM and those from non-ROM storage. When this bit is 0, off-chip requests for instructions are directed to instruction/data memory.

Bit 7 : WAIT Mode (WM)—The WM bit places the processor in the Wait mode. When this bit is 1, the processor performs no operations. The Wait mode is reset by an interrupt or trap for which the processor is enabled, or by the Reset mode.

Bit 6 : Physical Addressing/Data (PD)—The PD bit determines whether address translation is performed for load or store operations. Address translation is performed for an access only when this bit is 0, and the Physical Address (PA) bit in the load or store instruction causing the access is also 0.

Bit 5 : Physical Addressing/Instructions (PI)—The PI bit determines whether address translation is performed for external instruction accesses. Address translation is performed only when this bit is 0.

Bit 4 : Supervisor Mode (SM)—The SM bit protects certain processor state, such as protected special-purpose registers. When this bit is 1, the processor is in the Supervisor mode, and access to all state is allowed. When this bit is 0, the processor is in the User mode, and access to protected state is not allowed; an attempt to access (either read or write) protected processor state causes a Protection Violation trap.

Section 3.1 describes the processor state protected from User-mode access.

For an external access, the User Access (UA) bit in the load or store instruction also controls access to protected state. When the UA bit is 1, the Memory Management Unit and channel perform the access as if the program causing the access were in User mode.

Bits 3-2 : Interrupt Mask (IM)—The IM field is an encoding of the processor priority with respect to external interrupts. The interpretation of the interrupt mask is specified by the following table:

<u>IM Value</u>	<u>Result</u>
0 0	*INTR0 enabled
0 1	*INTR0—*INTR1 enabled
1 0	*INTR0—*INTR2 enabled
1 1	*INTR0—*INTR3 enabled

Bit 1 : Disable Interrupts (DI)—The DI bit prevents the processor from being interrupted by external interrupt requests *INTR0—*INTR3. When this bit is 1, the processor ignores all external interrupts. However, note that traps (both internal and external), Timer interrupts, and Trace traps will be taken. When this bit is 0, the processor will take any interrupt enabled by the IM field, unless the DA bit is 1 (see below).

Bit 0 : Disable All Interrupts and Traps (DA)—The DA bit prevents the processor from taking any interrupts and most traps. When this bit is 1, the processor ignores interrupts and traps, except for the *WARN, Instruction Access Exception, Data Access Exception, and Coprocessor Exception traps. When this bit is 0, all traps will be taken, and interrupts will be taken if otherwise enabled.

Configuration (Register 3)

This protected special-purpose register (see Figure 3-6) controls certain processor and system options. Most fields are normally modified only during system initialization. The Configuration Register is defined as follows:




Figure 3-6. Configuration Register

Bits 31–24: Processor Release Level (PRL)—The PRL field is an 8-bit, read-only identification number which specifies the processor version. The initial processor version has a PRL field of zero.

Bits 23–5 : reserved.

Bit 4 : Vector Fetch (VF)—The VF bit determines the structure of the interrupt/trap Vector Area. If this bit is 1, the Vector Area is defined as a block of 256 vectors which specify the beginning addresses of the interrupt and trap handling routines. If the VF bit is 0, the Vector Area is a segment of 256, 64-instruction blocks which contain the actual routines.

Bit 3 : ROM Vector Area (RV)—If the VF bit is 0, the RV bit specifies whether the Vector Area is contained in instruction memory (RV = 0) or instruction read-only memory (RV = 1). The value of the RV bit is irrelevant if the VF bit is 1.

Bit 2 : Byte Order (BO)—The BO bit determines the ordering of bytes and half-words within words. If the BO bit is 0, bytes and half-words are numbered left-to-right within a word. If the BO bit is 1, bytes and half-words are numbered right-to-left. Section 3.4.3 describes the interpretation of the BO bit in more detail.


Bit 1 : Coprocessor Present (CP)—The CP bit indicates the presence of a coprocessor which may be used by the processor. If this bit is 1, it enables the execution of load and store instructions which have a Coprocessor Enable (CE) bit of 1. If the CP bit is 0 and the processor attempts to execute a load or store instruction with a CE bit of 1, a Coprocessor Not Present trap occurs. This feature may be used to emulate coprocessor operations as well as to protect the state of a coprocessor shared between multiple processes.

Bit 0 : Branch Target Cache Disable (CD)—The CD bit determines whether or not the Branch Target Cache is used for non-sequential instruction references. When this bit is 1, all instruction references are directed to external instruction memory or instruction ROM, and the Branch Target Cache is not used. When this bit is 0, the targets of non-sequential instruction fetches are stored in the Branch Target Cache and re-used as described in Section 4.2.2. The value of the CD bit does not take effect until the execution of the next branch instruction.

Channel Address (Register 4)

This protected special-purpose register (Figure 3-7) is used to report exceptions during external accesses or coprocessor transfers. It is also used to restart interrupted load-multiple and store-multiple operations, and to restart other external accesses when possible (e.g., after TLB misses are serviced). The restarting of external accesses is described in Section 7.2.5.

The Channel Address Register is updated on the execution of every load or store instruction, and on every load or store in a load-multiple or store-multiple sequence, except when the Freeze (FZ) bit in the Current Processor Status Register is 1.



08996A 4


Figure 3-7. Channel Address Register

Bits 31-0 : Channel Address (CHA)—This field contains the address of the current channel transaction (if the FZ bit of the Current Processor Status Register is 0). For external data accesses, the address is virtual if address translation was enabled for the access, or physical if translation was disabled. For transfers to the coprocessor, the CHA field contains data transferred to the coprocessor.

Channel Data (Register 5)

This protected special-purpose register (Figure 3-8) is used to report exceptions during external accesses or coprocessor transfers. It is also used to restart the first store of an interrupted store-multiple operation and to restart other external accesses when possible (e.g., after TLB misses are serviced). The restarting of external accesses is described in Section 7.2.5.

The Channel Data Register is updated on the execution of every load or store instruction, and on every load or store in a load-multiple or store-multiple sequence, except when the Freeze (FZ) bit in the Current Processor Status Register is 1. When the Channel Data Register is updated for a load operation, the resulting value is unpredictable.



08996A 5

Figure 3-8. Channel Data Register

Bits 31-0 : Channel Data (CHD)—This field contains the data (if any) associated with the current channel transaction (if the FZ bit of the Current Processor Status Register is 0). If the current channel transaction is not a store or a transfer to the coprocessor, the value of this field is irrelevant.

Channel Control (Register 6)

This protected special-purpose register (Figure 3-9) is used to report exceptions during external accesses or coprocessor transfers. It is also used to restart interrupted load-multiple and store-multiple operations, and to restart other external accesses when possible (e.g., after TLB misses are serviced). The restarting of external accesses is described in Section 7.2.5.

The Channel Control Register is updated on the execution of every load or store instruction, and on every load or store in a load-multiple or store-multiple sequence, except when the Freeze (FZ) bit in the Current Processor Status Register is 1.




Figure 3-9. Channel Control Register

Bits 31-24:—These bits are a direct copy of bits 23-16 from the load or store instruction which started the current channel transaction (see Section 3.4.2 and Section 6.1.2).

Bits 23-16 : Load/Store Count Remaining (CR)—The CR field indicates the remaining number of transfers for a load-multiple or store-multiple operation which encountered an exception or was interrupted before completion. This number is zero-based; for example, a value of 28 in this field indicates that 29 transfers remain to be completed. If the fault or interrupt occurs on the last transaction, the CR field contains a value of 0, and the ML bit is 1 (see below).

Bit 15 : Load/Store (LS)—The LS bit is 0 if the channel transaction is a store operation, and 1 if it is a load operation.

Bit 14 : Multiple Operation (ML)—The ML bit is 1 if the current channel transaction is a partially-complete load-multiple or store-multiple operation; otherwise it is 0.

Bit 13 : Set (ST)—The ST bit is 1 if the current channel transaction is for a Load and Set instruction; otherwise it is 0.

Bit 12 : Lock Active (LA)—The LA bit is 1 if the current channel transaction is for a Load and Lock or Store and Lock instruction; otherwise it is 0. Note that this bit is not set as the result of the Lock (LK) bit in the Current Processor Status Register.

Bit 11 : reserved.

Bit 10 : Transaction Faulted (TF)—The TF bit indicates that the current channel transaction did not complete due to some exceptional circumstance. This bit is set only for exceptions reported via the *DERR input, and it causes a Data Access Exception or Coprocessor Exception trap to occur (depending on the value of the CE bit) when it is 1.

The TF bit allows the proper sequencing of externally-reported errors which get preempted by higher-priority traps (see Section 3.5.7). It is reset by software which handles the resulting trap.

Bits 9–2 : Target Register (TR)—The TR field indicates the absolute register-number of data operand for the current transaction (either a load target or store data-source). Since the register-number in this field is absolute, it reflects the Stack-Pointer addition when the indicated register is a local register.


Bit 1 : Not Needed (NN)—The NN bit indicates that, even though the Channel Address, Channel Data, and Channel Control registers contain a valid representation of an uncompleted load operation, the data requested is not needed. This situation arises when a load instruction is overlapped with an instruction which writes the load target-register.

Bit 0 : Contents Valid (CV)—The CV bit indicates that the contents of the Channel Address, Channel Data, and Channel Control registers are valid.

Register Bank Protect (Register 7)

This protected special-purpose register (Figure 3-10) protects banks of general-purpose registers from User-mode program accesses.

The general-purpose registers are partitioned into 16 banks of 16 registers each (except that Bank 0 contains 14 registers). The banks are organized as shown in Figure 3-2 of Section 3.2.1.



08996A 7

Figure 3-10. Register Bank Protect Register

Bits 31–16 : reserved.

Bits 15–0 : Bank 15 through Bank 0 Protection Bits (B15–B0)—In the Register Bank Protect Register, each bit is associated with a particular bank of registers, and the bit number gives the associated bank number (e.g., B11 determines the protection for Bank 11).


When a protection bit is 1, the corresponding bank is protected from access by programs executing in the User mode. A Protection Violation trap occurs when a User-mode program attempts to access (either read or write) a register in a protected bank. When a bit in this register is 0, the corresponding bank is available to programs executing in the User mode.

Supervisor-mode programs are not affected by the Register Bank Protect Register.

Register protection is based on absolute register-numbers. For local registers, the protection checking is performed after the Stack-Pointer addition is performed.

Timer Counter (Register 8)

This protected special-purpose register (Figure 3-11) contains the counter for the Timer Facility.



08996A 8

Figure 3-11. Timer Counter Register

Bits 31–24 : reserved.

Bits 23–0 : Timer Count Value (TCV)—The 24-bit TCV field decrements by one on each processor clock. When the TCV field decrements to zero, it is reloaded with the content of the Timer Reload Value field in the Timer Reload Register. At this time, the Interrupt bit in the Timer Reload Register is set.

Timer Reload (Register 9)

This protected special-purpose register (Figure 3-12) maintains synchronization of the Timer Counter Register, enables Timer interrupts, and maintains Timer Facility status information.




Figure 3-12. Timer Reload Register

Bits 31-27 : reserved.

Bit 26 : Overflow (OV)—The OV bit indicates that a Timer interrupt occurred before a previous Timer interrupt was serviced. It is set if the Interrupt (IN) bit is 1 (see below) when the Timer Count Value (TCV) field of the Timer Counter Register decrements to zero. In this case, a Timer interrupt caused by the IN bit has not been serviced when another interrupt is created.

Bit 25 : Interrupt (IN)—The IN bit is set whenever the TCV field decrements to zero. If this bit is 1 and the IE bit is also 1, a Timer interrupt occurs. Note that the IN bit is set when the TCV field decrements to zero, regardless of the value of the IE bit. The IN bit is reset by software which handles the Timer interrupt.


The TCV field is zero-based with respect to the Timer interrupt interval; for example, a value of 28 in the TCV field causes the IN bit to be set in the 29th subsequent processor cycle. The reason for this is that the TCV field is zero for a complete cycle before the IN bit is set.

Bit 24 : Interrupt Enable (IE)—When the IE bit is 1, the Timer interrupt is enabled, and the Timer interrupt occurs whenever the IN bit is 1. When this bit is 0, the Timer interrupt is disabled. Note that Timer interrupts may be disabled by the DA bit of the Current Processor Status Register regardless of the value of the IE bit.

Bits 23-0 : Timer Reload Value (TRV)—The value of this field is written into the Timer Count Value (TCV) field of the Timer Counter Register when the TCV field decrements to zero.

Program Counter 0 (Register 10)

This protected special-purpose register (Figure 3-13) is used, on an interrupt return, to restart the instruction which was in the decode stage when the original interrupt or trap was taken.



08996A 10

Figure 3-13. Program Counter 0 Register


Bits 31–2 : Program Counter 0 (PC0)—This field captures the word-address of an instruction as it enters the decode stage of the processor pipeline, unless the Freeze (FZ) bit of the Current Processor Status Register is 1. If the FZ bit is 1, PC0 holds its value.

When an interrupt or trap is taken, the PC0 field contains the word-address of the instruction in the decode stage; the interrupt or trap has prevented this instruction from executing. The processor uses the PC0 field to restart this instruction on an interrupt return.

Bits 1–0 : Zeros—These are zero-bits, since instruction addresses are always word-aligned.

Program Counter 1 (Register 11)

This protected special-purpose register (Figure 3-14) is used, on an interrupt return, to restart the instruction which was in the execute stage when the original interrupt or trap was taken.



08996A 11

Figure 3-14. Program Counter 1 Register


Bits 31–2 : Program Counter 1 (PC1)—This field captures the word-address of an instruction as it enters the execute stage of the processor pipeline, unless the Freeze (FZ) bit of the Current Processor Status Register is 1. If the FZ bit is 1, PC1 holds its value.

When an interrupt or trap is taken, the PC1 field contains the word-address of the instruction in the execute stage; the interrupt or trap has prevented this instruction from completing execution. The processor uses the PC1 field to restart this instruction on an interrupt return.

Bits 1–0 : Zeros—These are zero-bits, since instruction addresses are always word-aligned.

Program Counter 2 (Register 12)

This protected special-purpose register (Figure 3-15) reports the address of certain instructions causing traps.



08996A 12

Figure 3-15. Program Counter 2 Register


Bits 31–2 : Program Counter 2 (PC2)—This field captures the word-address of an instruction as it enters the write-back stage of the processor pipeline, unless the Freeze (FZ) bit of the Current Processor Status Register is 1. If the FZ bit is 1, PC2 holds its value.

When an interrupt or trap is taken, the PC2 field contains the word-address of the instruction in the write-back stage. In certain cases, as described in Section 3.5.8, PC2 contains the address of the instruction causing a trap. The PC2 field is used to report the address of this instruction, and has no other use in the processor.

Bits 1–0 : Zeros—These are zero-bits, since instruction addresses are always word-aligned.

MMU Configuration (Register 13)

This protected special-purpose register (Figure 3-16) specifies parameters associated with the Memory Management Unit (MMU).



08996A 13

Figure 3-16. MMU Configuration Register

Bits 31–10 : reserved.


Bits 9–8 : Page Size (PS)—The PS field specifies the page size for address translation. The page size affects translation as discussed in Section 3.6.2. The PS field has a delayed effect (see Section 7.3.3). At least one cycle of delay must separate an instruction which sets the PS field and an instruction that performs address translation. The PS field is encoded as follows:

<u>PS</u>	<u>Page Size</u>
0 0	1 Kbyte
0 1	2 Kbyte
1 0	4 Kbyte
1 1	8 Kbyte

Bits 7–0 : Process Identifier (PID)—This 8-bit field is compared to Task Identifier (TID) fields in Translation Look-aside Buffer entries when an address translation is performed. For the address translation to be valid, the PID field must match the TID field in an entry. This allows a separate 32-bit virtual-address space to be allocated to each of 256 active processes.

LRU Recommendation (Register 14)

This protected special-purpose register (Figure 3-17) assists Translation Look-aside Buffer (TLB) reloading by indicating the least-recently-used TLB entry in the required replacement line.



08996A 14

Figure 3-17. LRU Recommendation Register


Bits 31–7 : reserved.

Bits 6–1 : Least Recently Used Entry (LRU)—The LRU field is updated whenever a TLB miss occurs during an address translation. It gives the TLB register-number of the TLB entry selected for replacement. The LRU field is also updated whenever a memory-protection violation occurs; however, it has no interpretation in this case.

Bit 0 : Zero—The appended 0 serves to identify Word 0 of the TLB entry.

Indirect Pointer C (Register 128)

This unprotected special-purpose register (Figure 3-18) provides the RC-operand register-number (see Section 8.3) when an instruction RC field has the value zero (i.e. when Global Register 0 is specified).



08996A 15

Figure 3-18. Indirect Pointer C Register


Bits 31–10 : reserved.

Bits 9–2 : Indirect Pointer C (IPC)—The 8-bit IPC field contains an absolute register-number for a general-purpose register. This number directly selects a register (Stack-Pointer addition is not performed in the case of local registers).

Bits 1–0 : Zeros—The IPC field is aligned for compatibility with word-addresses.

Indirect Pointer A (Register 129)

This unprotected special-purpose register (Figure 3-19) provides the RA-operand register-number (see Section 8.3) when an instruction RA field has the value zero (i.e. when Global Register 0 is specified).



08996A 16

Figure 3-19. Indirect Pointer A Register


Bits 31–10 : reserved.

Bits 9–2 : Indirect Pointer A (IPA)—The 8-bit IPA field contains an absolute register-number for either a general-purpose register or a local register. This number directly selects a register (Stack-Pointer addition is not performed in the case of local registers).

Bits 1–0 : Zeros—The IPA field is aligned for compatibility with word-addresses.

Indirect Pointer B (Register 130)

This unprotected special-purpose register (Figure 3-20) provides the RB-operand register-number (see Section 8.3) when an instruction RB field has the value zero (i.e. when Global Register 0 is specified).



08996A 17

Figure 3-20. Indirect Pointer B Register


Bits 31-10 : reserved.

Bits 9-2 : Indirect Pointer B (IPB)—The 8-bit IPB field contains an absolute register-number for a general-purpose register. This number directly selects a register (Stack-Pointer addition is not performed in the case of local registers).

Bits 1-0 : Zeros—The IPB field is aligned for compatibility with word-addresses.

Q (Register 131)

The Q Register is an unprotected special-purpose register (Figure 3-21).




08996A 18

Figure 3-21. Q Register

Bits 31-0 : Quotient/Multiplier (Q)—During a divide operation, this field holds the low-order bits of the dividend; it contains the quotient at the end of the divide. During a multiply operation, this field holds the multiplier; it contains the low-order bits of the result at the end of the multiply.

ALU Status (Register 132)

This unprotected special-purpose register (Figure 3-22) holds information about the outcome of Arithmetic/Logic Unit (ALU) operations as well as control for certain operations performed by the Execution Unit.



DF

08996A 19

Figure 3-22. ALU Status Register

Bits 31–12 : reserved.

Bit 11 : Divide Flag (DF)—The DF bit is used by the instructions which implement division. This bit is set at the end of the division instructions either to 1 or to the complement of the 33rd bit of the ALU. When a Divide Step instruction is executed, then the DF bit determines whether an addition or subtraction operation is performed by the ALU.

Bit 10 : Overflow (V)—The V bit indicates that the result of a signed, two's-complement ALU operation required more than 32 bits to correctly represent the result. The value of this bit is determined by exclusive-ORing the ALU carry-out with the carry-in to the most-significant bit for signed, two's-complement operations. This bit is not used for any special purpose in the processor, and is provided for information only.

Bit 9 : Negative (N)—The N bit is set with the value of the most-significant bit of the result of an arithmetic or logical operation. If two's-complement overflow occurs, the N bit does not reflect the true sign of the result. This bit is used in divide operations.

Bit 8 : Zero (Z)—The Z bit indicates that the result of an arithmetic or logical operation is zero. This bit is not used for any special purpose in the processor, and is provided for information only.

Bit 7 : Carry (C)—The C bit stores the carry-out of the ALU for arithmetic operations. It is used by the add-with-carry and subtract-with-carry instructions to generate the carry into the Arithmetic/Logic Unit.

Bits 6–5 : Byte Pointer (BP)—The BP field holds a 2-bit pointer to a byte within a word. It is used by Insert Byte and Extract Byte instructions. The exact mapping of the pointer value to the byte position depends on the value of the Byte Order (BO) bit in the Configuration Register.

The most-significant bit of the BP field is used to determine the position of a half-word within a word for the Insert Half-Word, Extract Half-Word, and Extract Half-Word, Sign-Extended instructions. The exact mapping of the most-significant bit to the half-word position depends on the value of the BO bit in the Configuration Register.

The BP field is set by a Move To Special Register instruction with either the ALU Status Register or the Byte Pointer Register as the destination. It is also set with the low-order two bits of the address for a load or store instruction if the Set Byte Pointer (SB) bit in the instruction is 1.


Bits 4–0 : Funnel Shift Count (FC)—The FC field contains a 5-bit shift count for the Funnel Shifter. The Funnel Shifter concatenates two source-operands into a single, 64-bit operand and extracts a 32-bit result from this 64-bit operand; the FC field specifies the number of bit positions from the most-significant bit of the 64-bit operand

to the most-significant bit of the 32-bit result. The FC field is used by the EXTRACT instruction.

The FC field is set by a Move To Special Register instruction with either the ALU Status Register or the Funnel Shift Count Register as the destination.

Byte Pointer (Register 133)

This unprotected special-purpose register (Figure 3-23) provides an alternate access to the BP field in the ALU Status Register.



08996A 20


Figure 3-23. Byte Pointer

Bits 31–2 : Zeros.

Bits 1–0 : Byte Pointer (BP)—This field allows a program to change the BP field without affecting other fields in the ALU Status Register.

Funnel Shift Count (Register 134)

This unprotected special-purpose register (Figure 3-24) provides an alternate access to the FC field in the ALU Status Register.



08996A 21


Figure 3-24. Funnel Shift Count

Bits 31–5 : Zeros.

Bits 4–0 : Funnel Shift Count (FC)—This field allows a program to change the FC field without affecting other fields in the ALU Status Register.

Load/Store Count Remaining (Register 135)

This unprotected special-purpose register (Figure 3-25) provides alternate access to the CR field in the Channel Control Register.



08996A 22

Figure 3-25. Load/Store Count Remaining

Bits 31-8 : Zeros.

Bits 7-0 : Load/Store Count Remaining (CR)—This field allows a program to change the CR field without affecting other fields in the Channel Control Register, and is used to initialize the value before a Load Multiple or Store Multiple instruction is executed.

3.2.3 TLB REGISTERS

The Am29000 contains 128 Translation Look-Aside Buffer (TLB) registers. The organization of the TLB registers is shown in Figure 3-26.

The TLB registers comprise the TLB entries, and are provided so that programs may inspect and alter TLB entries. This allows the loading, invalidation, saving, and restoring of TLB entries.

TLB registers have fields which are reserved for future processor implementations. When a TLB register is read, a bit in a reserved field is read as a 0. An attempt to write a reserved bit with a 1 has no effect; however, this should be avoided, because of upward-compatibility considerations.

The Translation Look-aside Buffer (TLB) registers are accessed only by explicit data movement by Supervisor-mode programs. Instructions which move data to or from a TLB register specify a general-purpose register containing a TLB register-number. The TLB register-number is given by the contents of bits 6-0 of the general-purpose register. TLB register-numbers may only be specified indirectly by general-purpose registers.

TLB entries are accessed as registers numbered 0-127. Since two words are required to completely specify a TLB entry, two registers are required for each TLB entry. The words corresponding to an entry are paired as two sequentially-numbered registers starting on an even-numbered register. The word with the even register-number is called Word 0, and the word with the odd register-number is called Word 1. The entries for TLB Set 0 are in registers numbered 0-63, and the entries for TLB Set 1 are in registers numbered 64-127.




Figure 3-26. Translation Look-aside Buffer Registers

TLB Entry Word 0

The TLB Entry Word 0 register is shown in Figure 3-27.




Figure 3-27. TLB Entry Word 0

Bits 31–15 : Virtual Tag (VTAG)—When the TLB is searched for an address translation, the VTAG field of the TLB entry must match the most-significant 17, 16, 15, or 14 bits of the address being translated—for page sizes of 1, 2, 4, and 8 Kbytes, respectively—for the search to be successful.

When software loads a TLB entry with an address translation, the most-significant 14 bits of the Virtual Tag are set with the most-significant 14 bits of the virtual address whose translation is being loaded into the TLB. The remaining 3 bits of the Virtual Tag must be set either to the corresponding bits of the address, or to zeros, depending on the page size, as follows (“A” refers to corresponding address bits):

<u>Page Size</u>	<u>VTAG 2-0 (TLB Word 0 bits 17-15)</u>
1 Kbyte	AAA
2 Kbyte	AA0
4 Kbyte	A00
8 Kbyte	000

Bit 14 : Valid Entry (VE)—If this bit is 1, the associated TLB entry is valid; if it is 0, the entry is invalid.

Bit 13 : Supervisor Read (SR)—If the SR bit is 1, Supervisor-mode load operations from the virtual page are allowed; if it is 0, Supervisor-mode loads are not allowed.

Bit 12 : Supervisor Write (SW)—If the SW bit is 1, Supervisor-mode store operations to the virtual page are allowed; if it is 0, Supervisor-mode stores are not allowed.

Bit 11 : Supervisor Execute (SE)—If the SE bit is 1, Supervisor-mode instruction accesses to the virtual page are allowed; if it is 0, Supervisor-mode instruction accesses are not allowed.

Bit 10 : User Read (UR)—If the UR bit is 1, User-mode load operations from the virtual page are allowed; if it is 0, User-mode loads are not allowed.


Bit 9 : User Write (UW)—If the UW bit is 1, User-mode store operations to the virtual page are allowed; if it is 0, User-mode stores are not allowed.

Bit 8 : User Execute (UE)—If the UE bit is 1, User-mode instruction accesses to the virtual page are allowed; if it is 0, User-mode instruction accesses are not allowed.

Bits 7–0 : Task Identifier (TID)—When the TLB is searched for an address translation, the TID must match the Process Identifier (PID) in the MMU Configuration Register for the translation to be successful. This field allows the TLB entry to be associated with a particular process.

TLB Entry Word 1

The TLB Entry Word 1 register is shown in Figure 3-28.



08996A 25

Figure 3-28. TLB Entry Word 1

Bits 31–10 : Real Page Number (RPN)—The RPN field gives the most-significant 22, 21, 20, or 19 bits of the physical address of the page, for page sizes of 1, 2, 4, and 8 Kbytes, respectively. It is concatenated to bits 9–0, 10–0, 11–0, or 12–0 of the address being translated—for 1, 2, 4, and 8 Kbyte page sizes, respectively—to form the physical address for the access.

When software loads a TLB entry with an address translation, the most-significant 19 bits of the Real Page Number are set with the most-significant 19 bits of the physical address associated with the translation. The remaining three bits of the Real Page Number must be set either to the corresponding bits of the physical address, or to zeros, depending on the page size, as follows (“A” refers to corresponding address bits):

<u>Page Size</u>	<u>RPN 2-0</u> (TLB Word 1 bits 12-10)
1 Kbyte	AAA
2 Kbyte	AA0
4 Kbyte	A00
8 Kbyte	000

Bits 7–6 : User Programmable (PGM)—These bits are placed on the MPGM0–MPGM1 outputs when the address is transmitted for an access. They have no predefined effect on the access: any effect is defined by logic external to the processor.

Bit 1 : Usage (U)—This bit indicates which entry in a given TLB line was least recently used to perform an address translation. If this bit is a 0, then the entry in Set 0 in the line is least-recently-used; if it is 1, then the entry in Set 1 is least-recently-used. This bit has an equal value for both entries in a line. Whenever a TLB entry is used to translate an address, the Usage bit of both entries in the line used for translation are set according to the TLB set containing the translation. This bit is set whenever the translation is valid, regardless of the outcome of memory-protection checking.

Bit 0 : Flag (F)—The Flag bit has no effect on address translation, and is affected only by the MTTLB instruction. This bit is provided for software management of TLB entries.

3.3 INSTRUCTION SET

The Am29000 implements 115 instructions. All instructions execute in a single cycle, except for IRET, IRETINV, LOADM and STOREM.

Most instructions deal with general-purpose registers for operands and results; however, in most instructions, an 8-bit constant can be used in place of a register-based operand. Some instructions deal with special-purpose registers, TLB registers, external devices and memories, and coprocessors.

This section describes the 9 instruction classes in the Am29000, and provides a brief summary of instruction operations. A detailed instruction specification is contained in Chapter 8. Section 8.1 describes the nomenclature used here.

If the processor attempts to execute an instruction which is not implemented, an Illegal Opcode trap occurs.

3.3.1 INTEGER ARITHMETIC

The Integer Arithmetic instructions perform add, subtract, multiply, and divide operations on word-length integers. Certain instructions in this class cause traps if signed or unsigned overflow occurs during the execution of the instruction. There is support for multi-precision arithmetic on operands whose lengths are multiples of words. All instructions in this class set the ALU Status Register. The integer arithmetic instructions are shown in Table 3-1.

The MULTIPLY and DIVIDE instructions are not implemented directly by processor hardware, but cause MULTIPLY and DIVIDE traps.

3.3.2 COMPARE

The Compare instructions test for various relationships between two values. For all Compare instructions except the CPBYTE instruction, the comparisons are performed on word-length signed or unsigned integers. There are two types of Compare instructions. The first type places a Boolean value reflecting the outcome of the compare into a general-purpose register. For the second type (assert instructions), instruction execution continues only if the comparison is true; otherwise a trap occurs. The assert instructions specify a vector for the trap (see Section 3.5.4).

The assert instructions support run-time operand checking and operating-system calls. If the trap occurs in the User mode, and a trap number between 0 and 63 is specified by the instruction, a Protection Violation trap occurs. The Compare instructions are shown in Table 3-2.

Table 3-1 Integer Arithmetic Instructions

Mnemonic	Operation Description
ADD	$DEST \leftarrow SRCA + SRCB$
ADDS	$DEST \leftarrow SRCA + SRCB$ IF signed overflow THEN Trap (Out Of Range)
ADDU	$DEST \leftarrow SRCA + SRCB$ IF unsigned overflow THEN Trap (Out Of Range)
ADDC	$DEST \leftarrow SRCA + SRCB + C$
ADDCS	$DEST \leftarrow SRCA + SRCB + C$ IF signed overflow THEN Trap (Out Of Range)
ADDCU	$DEST \leftarrow SRCA + SRCB + C$ IF unsigned overflow THEN Trap (Out Of Range)
SUB	$DEST \leftarrow SRCA - SRCB$
SUBS	$DEST \leftarrow SRCA - SRCB$ IF signed overflow THEN Trap (Out Of Range)
SUBU	$DEST \leftarrow SRCA - SRCB$ IF unsigned underflow THEN Trap (Out Of Range)
SUBC	$DEST \leftarrow SRCA - SRCB - 1 + C$
SUBCS	$DEST \leftarrow SRCA - SRCB - 1 + C$ IF signed overflow THEN Trap (Out Of Range)
SUBCU	$DEST \leftarrow SRCA - SRCB - 1 + C$ IF unsigned underflow THEN Trap (Out Of Range)
SUBR	$DEST \leftarrow SRCB - SRCA$
SUBRS	$DEST \leftarrow SRCB - SRCA$ IF signed overflow THEN Trap (Out Of Range)
SUBRU	$DEST \leftarrow SRCB - SRCA$ IF unsigned underflow THEN Trap (Out Of Range)
SUBRC	$DEST \leftarrow SRCB - SRCA - 1 + C$
SUBRCS	$DEST \leftarrow SRCB - SRCA - 1 + C$ IF signed overflow THEN Trap (Out Of Range)
SUBRCU	$DEST \leftarrow SRCB - SRCA - 1 + C$ IF unsigned underflow THEN Trap (Out Of Range)
MULTIPLY	$DEST//Q \leftarrow SRCA \cdot SRCB$

(Continued)

08996A 26a

Table 3-1 Integer Arithmetic Instructions (Continued)

Mnemonic	Operation Description
MUL	Perform one-bit step of a multiply operation
MULL	Complete a sequence of multiply steps
MULU	Perform one-bit step of a multiply operation (unsigned)
DIVIDE	$DEST \leftarrow (SRCA/Q)/SRCB$ (unsigned)
DIVO	Intitialize for a sequence of divide steps (unsigned)
DIV	Perform one-bit step of a divide operation (unsigned)
DIVL	Complete a sequence of divide steps (unsigned)
DIVREM	Generate remainder for divide operation (unsigned)

08996A 26b

Table 3-2. Compare Instructions

Mnemonic	Operation Description
CPEQ	IF $SRCA = SRCB$ THEN $DEST \leftarrow TRUE$ ELSE $DEST \leftarrow FALSE$
CPNEQ	IF $SRCA \neq SRCB$ THEN $DEST \leftarrow TRUE$ ELSE $DEST \leftarrow FALSE$
CPLT	IF $SRCA < SRCB$ THEN $DEST \leftarrow TRUE$ ELSE $DEST \leftarrow FALSE$
CPLTU	IF $SRCA < SRCB$ (unsigned) THEN $DEST \leftarrow TRUE$ ELSE $DEST \leftarrow FALSE$
CPLE	IF $SRCA \leq SRCB$ THEN $DEST \leftarrow TRUE$ ELSE $DEST \leftarrow FALSE$
CPLEU	IF $SRCA \leq SRCB$ (unsigned) THEN $DEST \leftarrow TRUE$ ELSE $DEST \leftarrow FALSE$
CPGT	IF $SRCA > SRCB$ THEN $DEST \leftarrow TRUE$ ELSE $DEST \leftarrow FALSE$
CPGTU	IF $SRCA > SRCB$ (unsigned) THEN $DEST \leftarrow TRUE$ ELSE $DEST \leftarrow FALSE$

(Continued)

08996A 27a

Table 3-2. Compare Instructions (Continued)

Mnemonic	Operation Description
CPGE	IF SRCA >= SRCB THEN DEST ← TRUE ELSE DEST ← FALSE
CPGEU	IF SRCA >= SRCB (unsigned) THEN DEST ← TRUE ELSE DEST ← FALSE
CPBYTE	IF (SRCA.BYTE0 = SRCB.BYTE0) OR (SRCA.BYTE1 = SRCB.BYTE1) OR (SRCA.BYTE2 = SRCB.BYTE2) OR (SRCA.BYTE3 = SRCB.BYTE3) THEN DEST ← TRUE ELSE DEST ← FALSE
ASEQ	IF SRCA = SRCB THEN Continue ELSE Trap (VN)
ASNEQ	IF SRCA <> SRCB THEN Continue ELSE Trap (VN)
ASLT	IF SRCA < SRCB THEN Continue ELSE Trap (VN)
ASLTU	IF SRCA < SRCB (unsigned) THEN Continue ELSE Trap (VN)
ASLE	IF SRCA <= SRCB THEN Continue ELSE Trap (VN)
ASLEU	IF SRCA <= SRCB (unsigned) THEN Continue ELSE Trap (VN)
ASGT	IF SRCA > SRCB THEN Continue ELSE Trap (VN)
ASGTU	IF SRCA > SRCB (unsigned) THEN Continue ELSE Trap (VN)
ASGE	IF SRCA >= SRCB THEN Continue ELSE Trap (VN)
ASGEU	IF SRCA >= SRCB (unsigned) THEN Continue ELSE Trap (VN)

08996A 27b

3.3.3 LOGICAL

The Logical instructions perform a set of bit-by-bit Boolean functions on word-length bit strings. All instructions in this class set the ALU Status Register. These instructions are shown in Table 3-3.

Table 3-3. Logical Instructions

Mnemonic	Operation Description
AND	$DEST \leftarrow SRCA \& SRCB$
ANDN	$DEST \leftarrow SRCA \& \sim SRCB$
NAND	$DEST \leftarrow \sim(SRCA \& SRCB)$
OR	$DEST \leftarrow SRCA SRCB$
NOR	$DEST \leftarrow \sim(SRCA SRCB)$
XOR	$DEST \leftarrow SRCA \wedge SRCB$
XNOR	$DEST \leftarrow \sim(SRCA \wedge SRCB)$

08996A 28

3.3.4 SHIFT

The Shift instructions (Table 3-4) perform arithmetic and logical shifts. All but the EXTRACT instruction operate on word-length data and produce a word-length result. The EXTRACT instruction operates on double-word data and produces a word-length result. If both parts of the double-word for the EXTRACT instruction are from the same source, the EXTRACT operation is equivalent to a rotate operation. For each operation, the shift count is a 5-bit integer, specifying a shift amount in the range of 0 to 31 bits.

Table 3-4. Shift Instructions

Mnemonic	Operation Description
SLL	$DEST \leftarrow SRCA \ll SRCB$ (zero fill)
SRL	$DEST \leftarrow SRCA \gg SRCB$ (zero fill)
SRA	$DEST \leftarrow SRCA \ggg SRCB$ (sign fill)
EXTRACT	$DEST \leftarrow \text{high-order word of } (SRCA // SRCB \ll FC)$

08996A 29

3.3.5 DATA MOVEMENT

The Data Movement instructions (Table 3-5) move bytes, half-words, and words between processor registers. In addition, they move data between general-purpose registers and external devices, memories, and the coprocessor.

Table 3-5. Data Movement Instructions

Mnemonic	Operation Description
LOAD	DEST ← EXTERNAL WORD [SRCB]
LOADL	DEST ← EXTERNAL WORD [SRCB] assert *LOCK output during access
LOADSET	DEST ← EXTERNAL WORD [SRCB] EXTERNAL WORD [SRCB] ← h'FFFFFFF', assert *LOCK output during access
LOADM	DEST.. DEST + COUNT ← EXTERNAL WORD [SRCB] .. EXTERNAL WORD [SRCB + COUNT • 4]
STORE	EXTERNAL WORD [SRCB] ← SRCA
STOREL	EXTERNAL WORD [SRCB] ← SRCA assert *LOCK output during access
STOREM	EXTERNAL WORD [SRCB] .. EXTERNAL WORD [SRCB + COUNT • 4] ← DEST.. DEST + COUNT
EXBYTE	DEST ← SRCB, with low-order byte replaced by byte in SRCA selected by BP
EXHW	DEST ← SRCB, with low-order half-word replaced by half-word in SRCA selected by BP
EXHWS	DEST ← half-word in SRCA selected by BP, sign-extended to 32 bits
INBYTE	DEST ← SRCA, with byte selected by BP replaced by low-order byte of SRCB
INHW	DEST ← SRCA, with half-word selected by BP replaced by low-order half-word of SRCB
MFSR	DEST ← SPECIAL
MFTLB	DEST ← TLB [SRCA]
MTSR	SPDEST ← SRCB
MTSRIM	SPDEST ← 0H16
MTTLB	TLB [SRCA] ← SRCB

08996A 30

3.3.6 CONSTANT

The Constant instructions (Table 3-6) provide the ability to place half-word and word constants into registers. Most instructions in the instruction set allow an 8-bit constant as an operand. The Constant instructions allow the construction of larger constants.

Table 3-6. Constant Instructions

Mnemonic	Operation Description
CONST	DEST ← 0116
CONSTH	Replace high-order half-word of SRCA by 116
CONSTN	DEST ← 1116

08996A 31

3.3.7 FLOATING-POINT

The Floating-Point instructions (Table 3-7) provide operations on single-precision (32-bit) or double-precision (64-bit) floating-point data. In addition, they provide conversions between single-precision, double-precision, and integer number representations. In the current processor implementation, these instructions cause traps to routines which perform the floating-point operations.

3.3.8 BRANCH

The Branch instructions (Table 3-8) control the execution flow of instructions. Branch target addresses may be absolute, relative to the Program Counter (with the offset given by a signed instruction constant), or contained in a general-purpose register. For conditional jumps, the outcome of the jump is based on a Boolean value in a general-purpose register. Procedure calls are unconditional, and save the return address in a general-purpose register. All branches have a delayed effect; the instruction sequence following the branch is executed regardless of the outcome of the branch.

3.3.9 MISCELLANEOUS

The Miscellaneous instructions (Table 3-9) perform various operations which cannot be grouped into other instruction classes. In certain cases, these are control functions available only to Supervisor-mode programs.

Table 3-7. Floating-Point Instructions

Mnemonic	Operation Description
FADD	DEST (single-precision) \leftarrow SRCA (single-precision) + SRCB (single-precision)
DADD	DEST (double-precision) \leftarrow SRCA (double-precision) + SRCB (double-precision)
FSUB	DEST (single-precision) \leftarrow SRCA (single-precision) - SRCB (single-precision)
DSUB	DEST (double-precision) \leftarrow SRCA (double-precision) - SRCB (double-precision)
FMUL	DEST (single-precision) \leftarrow SRCA (single-precision) • SRCB (single-precision)
DMUL	DEST (double-precision) \leftarrow SRCA (double-precision) • SRCB (double-precision)
FDIV	DEST (single-precision) \leftarrow SRCA (single-precision)/ SRCB (single-precision)
DDIV	DEST (double-precision) \leftarrow SRCA (double-precision)/ SRCB (double-precision)
FEQ	IF SRCA (single-precision) = SRCB (single-precision) THEN DEST \leftarrow TRUE ELSE DEST \leftarrow FALSE
DEQ	IF SRCA (double-precision) = SRCB (double-precision) THEN DEST \leftarrow TRUE ELSE DEST \leftarrow FALSE
FLT	IF SRCA (single-precision) < SRCB (single-precision) THEN DEST \leftarrow TRUE ELSE DEST \leftarrow FALSE
DLT	IF SRCA (double-precision) < SRCB (double-precision) THEN DEST \leftarrow TRUE ELSE DEST \leftarrow FALSE
FGT	IF SRCA (single-precision) > SRCB (single-precision) THEN DEST \leftarrow TRUE ELSE DEST \leftarrow FALSE
DGT	IF SRCA (double-precision) > SRCB (double-precision) THEN DEST \leftarrow TRUE ELSE DEST \leftarrow FALSE
CVINTF	DEST (single-precision) \leftarrow SRCA (integer)
CVINTD	DEST (double-precision) \leftarrow SRCA (integer)
CVFINT	DEST (integer) \leftarrow SRCA (single-precision)
CVDINT	DEST (integer) \leftarrow SRCA (double-precision)
CVFD	DEST (double-precision) \leftarrow SRCA (single-precision)
CVDF	DEST (single-precision) \leftarrow SRCA (double-precision) 08996A 32

Table 3-8. Branch Instructions

Mnemonic	Operation Description
CALL	DEST ← PC//00 + 8 PC ← TARGET Execute delay instruction
CALLI	DEST ← PC//00 + 8 PC ← SRCB Execute delay instruction
JMP	PC ← TARGET Execute delay instruction
JMPI	PC ← SRCB Execute delay instruction
JMPT	IF SRCA = TRUE THEN PC ← TARGET Execute delay instruction
JMPTI	IF SRCA = TRUE THEN PC ← SRCB Execute delay instruction
JMPF	IF SRCA = FALSE THEN PC ← TARGET Execute delay instruction
JMPFI	IF SRCA = FALSE THEN PC ← SRCB Execute delay instruction
JMPFDEC	IF SRCA = FALSE THEN SRCA ← SRCA - 1 PC ← TARGET ELSE SRCA ← SRCA - 1 Execute delay instruction

08996A 33

Table 3-9. Miscellaneous Instructions

Mnemonic	Operation Description
CLZ	Determine number of leading zeros in a word
SETIP	Set IPA, IPB, and IPC with operand register-numbers
EMULATE	Load IPA and IPB with operand register-numbers, and Trap (VN)
INV	Reset all Valid bits in Branch Target Cache to zeros
IRET	Perform an interrupt return sequence
IRETINV	Perform an interrupt return sequence, and reset all Valid bits in Branch Target Cache to zeros
HALT	Enter Halt mode on next cycle

08996A 34

3.4 DATA FORMATS AND HANDLING

This section describes the various data types supported by the Am29000, and the mechanisms for accessing data in external devices and memories. The Am29000 includes provisions for the external access of bytes, half-words, unaligned words, and unaligned half-words. These accesses are also described in this section.

3.4.1 DATA TYPES

Most Am29000 instructions deal directly with word-length integer data; integers may be either signed or unsigned, depending on the instruction. Some instructions (e.g. AND) treat word-length operands as strings of bits. In addition, there is support for character, half-word, and Boolean data types. Single-precision and double-precision floating-point data types are defined, but not directly supported by processor hardware.

Byte Operations

The processor supports character data through extraction and insertion operations on word-length operands, and by a compare operation on byte-length fields within words. For sequences of characters within words, bytes are ordered either left-to-right or right-to-left, depending on the BO bit of the Configuration Register (see Section 3.4.3).

The Extract Byte (EXBYTE) instruction replaces the low-order character of a destination word with an arbitrary byte-aligned character from a source word. For the EXBYTE instruction, the destination word can be a zero word, which effectively zero-extends the character from the source operand.

The Insert Byte (INBYTE) instruction replaces an arbitrary byte-aligned character in a destination word with the low-order character of a source word. For the INBYTE instruction, the source operand can be a character constant specified by the instruction.

The Compare Bytes (CPBYTE) instruction compares two word-length operands and gives a result of TRUE if any corresponding bytes within the operands have equivalent values. This allows programs to detect characters within words without first having to extract individual characters, one-at-a-time, from the word of interest.

Half-word Operations

The processor supports half-word data through insertion and extraction operations on word-length operands. For sequences of half-words within words, half-words are ordered either left-to-right or right-to-left, depending on the Byte Order (BO) bit of the Configuration Register (see Section 3.4.3).

The Extract Half-Word (EXHW) instruction replaces the low-order half- word of a destination word with either the low-order or high-order half- word of a source word. For the EXHW instruction, the destination word can be a zero word, which effectively zero-extends the half-word from the source operand.

The Extract Half-Word, Sign-Extended (EXHWS) instruction is similar to the EXHW instruction, except that it sign-extends the half-word in the destination word (i.e. it replaces the most-significant 16 bits of the destination word with the most-significant bit of the source half-word).

The Insert Half-Word (INHW) instruction replaces either the low-order or high-order half-word in a destination word with the low-order half-word of a source word.

Boolean Data

Some instructions in the Compare class generate word-length Boolean results. Also, conditional branches are conditional upon Boolean operands. The Boolean format used by the processor is such that the Boolean values TRUE and FALSE are represented by a 1 or 0, respectively, in the most-significant bit of a word. The remaining bits are unimportant: for the compare instructions, they are reset. Note that two's-complement negative integers are indicated by the Boolean value TRUE in this encoding scheme.

Floating-point Data

The floating-point format defined for the Am29000 conforms to the IEEE Floating-point Standard P754.

Single-precision floating-point instructions operate on word-length floating-point operands. Double-precision floating-point instructions operate on double-word operands. The processor does not directly support mixed-mode floating-point operations, but provides for all possible conversions between single-precision floating-point, double-precision floating-point, and word-length integer data.

By convention, a double-precision floating-point operand is contained in two consecutive general-purpose registers, beginning on an even-numbered register. The processor does not enforce this restriction. However, it should be followed for compatibility with future processor versions.

3.4.2 EXTERNAL DATA ACCESSES

All processor external accesses occur between general-purpose registers and external devices and memories. Accesses occur as the result of the execution of load and store instructions. The load and store instructions specify which general-purpose register receives the data (for a load) or supplies the data (for a store). The format of the load and store instructions is shown in Figure 3-29.

Addresses for accesses are given either by the content of a general-purpose register or by a constant value specified by the load or store instruction. The load and store instructions do not perform address computation directly. Any required address computations are performed explicitly by other instructions.




Figure 3-29. Load/Store Instruction Format

In the load or store instruction, CE (Coprocessor Enable) bit (bit 23) determines whether or not the access is directed to the coprocessor. If the CE bit is 0, the access is directed to an external device or memory. If the CE bit is 1, data is transferred to or from the coprocessor. The CE bit affects the interpretation of the Control (CNTL) field as well as the channel protocol. Coprocessor accesses are discussed in Chapter 6. This section deals with all other external accesses.

The format of the instructions which do not perform coprocessor data-transfers (i.e. in which the CE bit is 0) is shown in Figure 3-30.




Figure 3-30. Non-Coprocessor Load/Store Format

In load and store instructions, the "RB or I" field specifies the address for access. The address is either the content of a general-purpose register, with register-number RB, or a constant of a value I (zero-extended to 32 bits). The M bit determines whether the register or the constant is used.

The data for the access is written into the general-purpose register RA for a load, and is supplied by register RA for a store.

The definitions for other fields in the load or store instruction are given below:

Bit 23 : Coprocessor Enable (CE)—The CE bit is 0 for a non-coprocessor load or store.

Bit 22 : Address Space (AS)—If the AS bit is 0, the access is directed to instruction/data memory. If the AS bit is 1, the access is directed to input/output.

Bit 21 : Physical Address (PA)—The PA bit may be used by a Supervisor-mode program to disable address translation for an access. If the PA bit is 1, then address translation is not performed for the access, regardless of the value of the Physical Addressing/Data (PD) bit in the Current Processor Status Register. If the PA bit is 0, address translation depends on the PD bit.

The PA bit may be 1 only for Supervisor-mode instructions. If it is 1 for a User-mode instruction, a Protection Violation trap occurs.

Bit 20 : Set Byte Pointer (SB)—If the SB bit is 1, the Byte Pointer Register is written with the two least-significant bits of the address for the access. These address bits can control subsequent character and half-word operations. If the BP bit is 0, the Byte Pointer Register is not affected.

Bit 19 : USER Access (UA)—The UA bit allows programs executing in the Supervisor mode to emulate User-mode accesses. This allows checking of the authorization of an access requested by a User-mode program.

If the UA bit is 1, the access associated with the instruction is performed in the User mode, regardless of the value of the Supervisor Mode (SM) bit in the Current Processor Status Register. In this case, the User mode affects only TLB protection-checking and the SUP/*US output; it has no effect on the registers which can be accessed by the instruction. If the UA bit is 0, the program mode for the access is controlled by the SM bit.

Bits 18–16 : Option (OPT)—This field is placed on the OPT0–OPT2 outputs during the address cycle of the access. There is a one-to-one correspondence between the OPT field and the OPT0–OPT2 outputs; that is, the most-significant OPT bit is placed on OPT2, and so on.

In a standard processor configuration, the OPT field controls the width of an external access, supports the detection of unaligned accesses, and allows the contents of the instruction read-only memory address-space to be accessed as data. In systems where such accesses are not important, the OPT field may control external hardware in a system-dependent fashion.

Bits 15–8 : (RA) The data for the access is written into the general-purpose register RA for a load, and is supplied by register RA for a store.

Bits 7–0 : (RB or I) In load and store instructions, the “RB or I” field specifies the address for the access. The address is either the content of a general-purpose register, with

register-number RB, or a constant value I (zero-extended to 32 bits). The M bit of the operation code (bit 24) determines whether the register or the constant is used.

Load and store operations are overlapped with the execution of instructions which follow the load or store instruction. Only one load or store may be in progress on any given cycle. If a load or store instruction is encountered while another load or store operation is in progress, the processor enters the Pipeline Hold mode until the first operation completes. However, the address for the second operation may appear on the Address Bus if the first operation is to a device or memory which supports pipelined operations (see Section 5.2.8).

Load Operations

The processor provides the following instructions for performing load operations: Load (LOAD), Load and Lock (LOADL), Load and Set (LOADSET), and Load Multiple (LOADM). All of these instructions transfer data from an external device or memory into one or more general-purpose registers.

The LOADL instruction supports the implementation of device and memory interlocks in a multi-processor configuration. It activates the *LOCK output during the address cycle of the access.

The LOADSET instruction implements a binary semaphore. It loads a general-purpose register and atomically writes the accessed location with a word which has 1 in every bit position (that is, the write is indivisible from the read). The *LOCK output is asserted during both the read and write access. Note that, if address translation is enabled for the LOADSET instruction, the TLB memory-protection bits must allow both the read and write access.

The LOADM loads a specified number of registers from sequential addresses, as explained below.

Load operations are overlapped with the execution of instructions which follow the load instruction. The processor detects any dependencies on the loaded data which subsequent instructions may have, and, if such a dependency is detected, enters the Pipeline Hold mode until the data is returned by the external device or memory. If a register which is the target of an incomplete load is written with the result of a subsequent instruction, the processor does not write the returning data into the register when the load completes; the Not Needed (NN) bit in the Channel Control Register is set in this case.

Store Operations

The processor provides the following instructions for performing store operations: Store (STORE), Store and Lock (STOREL), and Store Multiple (STOREM). All of these instructions transfer data from one or more general-purpose registers to an external device or memory.

The STOREL instruction supports the implementation of device and memory interlocks in a multi-processor configuration. It activates the *LOCK output during the address cycle of the access.

The STOREM instruction stores a specified number of registers to sequential addresses, as explained below.

Store operations are overlapped with the execution of instructions which follow the store instruction. However, no data dependencies exist, since the store prevents any subsequent accesses until it completes.

Multiple Accesses

Load Multiple (LOADM) and Store Multiple (STOREM) instructions move contiguous words of data between general-purpose registers and external devices and memories. The number of transfers is determined by the Load/Store Count Remaining Register.

The Load/Store Count Remaining (CR) field in the Load/Store Count Remaining Register specifies the number of transfers to be performed by the next LOADM or STOREM executed in the instruction sequence. The CR field is in the range of 0 to 255, and is zero-based: a count value of 0 represents one transfer, and a count value of 255 represents 256 transfers. The CR field also appears in the Channel Control Register.

Before a LOADM or STOREM is executed, the CR field is set by a Move To Special Register. A LOADM or STOREM uses the most recently written value of the CR field. If an attempt is made to alter the CR field, and the Channel Control Register contains information for an external access which has not yet completed, the processor enters the Pipeline Hold mode until the access completes. Since the CR is set independently of the LOADM and STOREM, the CR field may represent valid state of an interrupted program even if the Contents Valid (CV) bit of the Channel Control Register is 0.

Note: Because of the pipelined implementation of LOADM and STOREM, at least one instruction (e.g. the instruction which sets the CR field) must separate two successive LOADM and/or STOREM instructions.

After the CR field is set, the execution of a LOADM or STOREM begins the data transfer. As with any other load or store operation, the LOADM or STOREM waits until any pending load or store operation is complete before starting. The LOADM instruction specifies the starting address and starting destination general-purpose register. The STOREM instruction specifies the starting address and the starting source general-purpose register.

During the execution of the LOADM or STOREM instruction, the processor updates the address and register-number after every access, incrementing the address by four and the

register-number by 1. This continues until either all accesses are completed or an interrupt or trap is taken.

For a load-multiple or store-multiple address sequence, addresses wrap from the largest possible value (hexadecimal FFFFFFFF) to the smallest possible value (hexadecimal 00000000).

The processor increments absolute register-numbers during the load-multiple or store-multiple sequence. Absolute register-numbers wrap from 127 to 128, and from 255 to 128. Thus, a sequence which begins in the global registers may transition to the local registers, but a sequence which begins in the local registers remains in the local registers. Also, note that the local registers are addressed circularly.

The normal restrictions on register accesses apply for the load-multiple and store-multiple sequences. For example, if a protected general-purpose register is encountered in the sequence for a User-mode program, a Protection Violation trap occurs.

Intermediate addresses are stored in the Channel Address Register, and register-numbers are stored in the Target Register (TR) field of the Channel Control Register. For the STOREM instruction, the data for every access is stored in the Channel Data Register (this register is also set during the execution of the LOADM instruction, but has no interpretation in this case). The CR field is updated on the completion of every access, so that it indicates the number of accesses remaining in the sequence.

Load-multiple and store-multiple operations are indicated by the Multiple Operation (ML) bit in the Channel Control Register. This bit may be 1 even though the CR field has a value of zero (indicating that one transfer remains to be performed). The ML bit is used to restart a multiple operation on an interrupt return; if it is set independently by a Move To Special Register before a load or store instruction is executed, the results are unpredictable.

While a multiple load or store is executing, the processor is in the Pipeline Hold mode, suspending any subsequent instruction execution until the multiple access completes. If an interrupt or trap is taken, the Channel Address, Channel Data, and Channel Control registers contain the state of the multiple access at the point of interruption. The multiple access may be resumed at this point, at a later time, by an interrupt return.

The processor attempts to complete multiple accesses using the burst-mode capability of the channel (see Section 5.2.9). If the burst is preempted, the processor retransmits the address at the point of preemption. If the external device or memory cannot support burst-mode accesses, the processor transmits an address for every access. If the address sequence causes a virtual page-boundary crossing, the processor preempts the burst-mode access, translates the address for the new page, and re-establishes the burst-mode access using the new physical address.

Option Bits

The Option field in the load and store instructions supports system functions which have no direct hardware support in the processor, such as variable-length data accesses. The standard definition of this field for a load or store, depending on the AS bit of the instruction, is as follows:

<u>AS</u>	<u>OPT2</u>	<u>OPT1</u>	<u>OPT0</u>	<u>Meaning</u>
x	0	0	0	Word-length access
x	0	0	1	Byte access
x	0	1	0	Half-word access
x	0	1	1	24-bit access
0	1	0	0	Instruction ROM access (as data)
	— all others —			reserved

The processor enforces the above interpretation only if the Trap Unaligned Access (TU) bit of the Current Processor Status Register is 1. If unaligned accesses, or any other access defined by the OPT field, are not supported by the system, this field may have any user-defined interpretation in system hardware.

Note that non-standard uses of the OPT bits have an implication on the portability of software between different systems. Non-standard uses of this field should be restricted to control-program routines, to maintain application-software compatibility. In any event, a value of 000 for the OPT field should be defined to have no special effect on an access.

3.4.3 ADDRESSING AND ALIGNMENT

Address Spaces

External instructions and data are contained in one of four, 32-bit address-spaces:

- 1) Instruction/Data Memory.
- 2) Input/Output.
- 3) Coprocessor.
- 4) Instruction Read-Only Memory (Instruction ROM).

An address in the Instruction/Data Memory address-space may be treated as virtual or physical, as determined by the Current Processor Status Register. Address translation for data accesses is enabled separately from address translation for instruction accesses. A program in the Supervisor mode may temporarily disable address translation for individual loads and stores; this permits load-real and store-real operations.

It is possible to partition physical instruction and data addresses into two, separate physical address-spaces. However, virtual instruction and data addresses appear in the same virtual address-space (i.e instruction/data memory).

The Coprocessor address-space is not an address-space in the strictest sense. The Coprocessor address-space is defined so that transfers of operands and operation codes to the coprocessor do not interfere with other external devices and memories.

The processor does not directly support the access of the Instruction ROM address-space using loads and stores; this capability is defined as a system option.

For data accesses, bits contained in load and store instructions distinguish between the instruction/data memory, input/output and coprocessor address-spaces.

For instruction fetches, the ROM Enable (RE) bit of the Current Processor Status Register distinguishes between the instruction/data and instruction ROM address-spaces.

Byte and Half-word Addressing

The Am29000 generates word-oriented byte addresses for accesses to external devices and memories. Addresses are word-oriented because loads, stores, and instruction fetches access words. However, addresses are byte addresses because they are sufficient to select bytes within accessed words. For load and store operations, the processor provides means for using the least-significant address bits to access bytes and half-words within external words.

The selection of a byte within a word is determined by the two least-significant bits of an address, and the Byte Order (BO) bit of the Configuration Register. The selection of a half-word within a word is determined by the next-to-least significant bit of an address, and the BO bit. Figure 3-31 illustrates the addressing of bytes and half-words when the BO bit is 0, and Figure 3-32 illustrates the addressing of bytes and half-words when the BO bit is 1. In Figures 3-31 and 3-32, addresses are represented in hexadecimal notation.

In the processor, the two least-significant bits of an external address are reflected in the Byte Pointer (BP) field of the ALU Status Register. The BO bit affects only the interpretation of the BP field.

If the BO bit is 0, bytes are ordered within words such that a 00 in the BP field selects the high-order byte of a word, and a 11 selects the low-order byte. If the BO bit is 1, a 00 in the BP field selects the low-order byte of a word and a 11 selects the high-order byte.


If the BO bit is 0, half-words are ordered within words such that a 0 in the most-significant bit of the BP field selects the high-order half-word, and a 1 selects the low-order half-word. If the BO bit is 1, a 0 in the most-significant bit of the BP field selects the low-order half-word of a word, and a 1 selects the high-order half-word. Note that, since the least-significant bit of the BP field does not participate in the selection of

half-words, that the alignment of half-words is forced to half-word boundaries in this case.

Byte and Half-word Accesses

The processor allows the Byte Pointer Register to be set with the least- significant bits of an address specified by any load or store instruction, except those which transfer information to and from the coprocessor. The byte and half-word insert and extract instructions can then be used to manipulate the byte or half-word of interest, after the external word has been accessed. This provides a general-purpose mechanism for manipulating external byte and half-word quantities, without the need for external hardware support.

To load a byte or half-word, a load is first performed. This load sets the BP field with the two least-significant bits of the address. A subsequent EXBYTE, EXHW, or EXHWS instruction extracts the byte or half- word of interest from the accessed word.




08996A 47

Figure 3-31. Byte and Half-Word Addressing with BO = 0

To store a byte or half-word, a load is first performed, setting the BP field with the two least-significant bits of the address. A subsequent INBYTE or INHW instruction inserts the byte or half-word of interest into the accessed word, and the resulting word is then stored.

External Hardware Support for Byte and Half-word Accesses

Load and store instructions contain an Option (OPT) field which is transmitted on the OPT0–OPT2 outputs during the address cycle of a load or store. The OPT field has a standard definition which indicates the data length for an access. This field may be used in conjunction with the two least-significant address bits to implement byte and half-word manipulation in user-defined external hardware.



08996A 48

Figure 3-32. Byte and Half-Word Addressing with BO = 1

It is important to insure that the performance advantage of external hardware for byte and half-word accesses justifies the hardware and performance costs. Compared to the basic processor mechanism for byte and half-word accesses described above, external hardware can reduce the time required for byte and half-word loads by, at best, one cycle. However, there is most likely no reduction, because of the delay in the external alignment hardware. The improvement for byte and half-word stores is more significant, since external hardware can eliminate the extra load required by the basic processor mechanism. However, byte and half-word stores are relatively rare in many systems.

External hardware for byte and half-word accesses can be expected to slow all external accesses, since it appears in the critical memory-access path. Thus, the performance advantages of this hardware is at least partially offset by the negative impact on all accesses. When the hardware costs are also considered, it is likely that this hardware is justified only in special cases, where byte and/or half-word accesses are relatively frequent.

Alignment of Words and Half-words

Since only byte addressing is supported, it is possible that an address for the access of a word or half-word is not aligned to the desired word or half-word. The Am29000 either ignores or forces alignment in most cases. However, some systems may require that unaligned accesses be supported, for compatibility reasons. Because of this, the Am29000 provides an option which creates a trap when a non-aligned access is attempted. This trap allows software emulation of the non-aligned accesses, in a manner which is appropriate for the particular system.

The detection of unaligned accesses is activated by a 1 in the Trap Unaligned Access (TU) bit of the Current Processor Status Register. Unaligned-access detection is based on the data length as indicated by the OPT field of a load or store instruction, and on the two least-significant bits of the specified address. Only addresses for instruction/data memory accesses are checked; alignment is ignored for input/output accesses and coprocessor transfers.

An Unaligned Access trap occurs only if the TU bit is 1 and any of the following combinations of OPT field and address bits is detected for a load or store to instruction/data memory:

<u>OPT2</u>	<u>OPT1</u>	<u>OPT0</u>	<u>A1</u>	<u>A0</u>	
0	0	0	1	0	Unaligned word access
0	0	0	0	1	
0	0	0	1	1	
0	1	0	0	1	Unaligned half-word access
0	1	0	1	1	

The trap handler for the Unaligned Access trap is responsible for generating the correct sequence of aligned accesses and performing any necessary shifting, masking and/or merging. Note that virtual page-boundary crossing may have to be considered, also.

Alignment of Instructions

In the Am29000, all instructions are 32-bits in length, and are aligned on word-address boundaries. The processor's Program Counter is 30 bits in length, and the least-significant two bits are always 00 for processor-generated instruction addresses. An unaligned address can be generated by indirect jumps and calls. However, alignment is ignored by the processor in this case, and it expects the system to force alignment (i.e., by interpreting the two least-significant address bits as 00, regardless of their values).

Accessing Instructions as Data

To aid the external access of instructions and data on separate buses, the processor distinguishes between instruction and data accesses. However, it does not support a logical distinction between instruction and data address-spaces (except in the case of instruction read-only memory). In particular, address translation in the Memory Management Unit is in no way affected by this distinction (although memory protection is).

In systems where it is necessary to access instructions as data, this function should be performed via the shared address-space. The OPT field provides for accessing instructions in the instruction read-only memory address-space; however, this should be unnecessary in most systems.

3.5 INTERRUPTS AND TRAPS

Interrupts and traps cause the Am29000 to suspend the execution of an instruction sequence and begin the execution of a new sequence. The processor may or may not later resume the execution of the original instruction sequence.

The distinction between interrupts and traps is largely one of causation and enabling. Interrupts allow external devices and the Timer Facility to control processor execution, and are always asynchronous to program execution. Traps are intended to be used for certain exceptional events which occur during instruction execution, and are generally synchronous to program execution.

Throughout this manual a distinction is made between the point at which an interrupt or trap occurs and the point at which it is taken. An interrupt or trap is said to occur when all conditions which define the interrupt or trap are met. However, an interrupt or trap which occurs is not necessarily recognized by the processor, either because of various enables, or because of the processor's operational mode (e.g. Halt mode). An interrupt or trap is taken when the processor recognizes the interrupt or trap and alters its behavior accordingly.

3.5.1 INTERRUPTS

Interrupts are caused by signals applied to any of the external inputs *INTR0–*INTR3, or by the Timer Facility (see Section 7.2.7). The processor may be disabled from taking certain interrupts by the masking capability provided by the Disable All Interrupts and Traps (DA) bit, Disable Interrupts (DI) bit, and Interrupt Mask (IM) field in the Current Processor Status Register.

The DA bit disables all interrupts and most traps. The DI bit disables external interrupts without affecting the recognition of traps and Timer interrupts. The 2-bit IM field selectively enables external interrupts as follows:

<u>IM Value</u>	<u>Result</u>
0 0	*INTR0 enabled
0 1	*INTR0–*INTR1 enabled
1 0	*INTR0–*INTR2 enabled
1 1	*INTR0–*INTR3 enabled

Note that the *INTR0 interrupt cannot be disabled by the IM field. Also, note that no external interrupt is taken if either the DA or DI bit is 1. The Interrupt Pending bit in the Current Processor Status indicates that one or more of the signals *INTR0–*INTR3 is active, but that the corresponding interrupt is disabled due to the value of either DA, DI, or IM.

3.5.2 TRAPS

Traps are caused by signals applied to one of the inputs *TRAP0–*TRAP1, or by exceptional conditions such as protection violations. Except for the Instruction Access Exception, Data Access Exception, and Coprocessor Exception traps, traps are disabled by the DA bit in the Current Processor Status; a 1 in the DA bit disables traps, and a 0 enables traps. It is not possible to selectively disable individual traps.

3.5.3 WAIT MODE

A wait-for-interrupt capability is provided by the Wait mode. The processor is in the Wait mode whenever the Wait Mode (WM) bit of the Current Processor Status is 1. While in Wait mode, the processor neither fetches nor executes instructions, and performs no external accesses. The Wait mode is exited when an interrupt or trap is taken.

Note that the processor can take only those interrupts or traps for which it is enabled, even in the Wait mode. For example, if the processor is in the Wait mode with a DA bit of 1, it can leave the Wait mode only via the Reset mode (see Section 3.8) or a *WARN trap (see Section 3.5.6).


3.5.4 VECTOR AREA

Interrupt and trap processing relies on the existence of a user-managed Vector Area in external instruction/data memory or instruction read-only memory (instruction ROM). The Vector Area begins at an address specified by the Vector Area Base Address Register, and provides for as many as 256 different interrupt and trap handling routines. The processor reserves 32 routines for system operation and 32 routines for Floating-Point MULTIPLY and DIVIDE instruction emulation. The number and definition of the remaining 192 possible routines are system-dependent.

The Vector Area has one of two possible structures as determined by the Vector Fetch (VF) bit in the Configuration Register. The first structure, as described below, requires less external memory than the second, but imposes the performance penalty of the vector table lookup.

If the VF bit is 1, the structure of the Vector Area is a table of vectors in instruction/data memory. The layout of a single vector is shown in Figure 3-33. Each vector gives the beginning word-address of the associated interrupt or trap handling routine, and specifies, by the R bit, whether the routine is contained in instruction/data memory (R = 0) or instruction ROM (R = 1).

If the VF bit is 0, the structure of the Vector Area is a segment of contiguous blocks of instructions in instruction/data memory or instruction ROM. The ROM Vector Area (RV) bit of the Configuration Register determines whether the Vector Area is in instruction/data memory (RV = 0) or instruction ROM (RV = 1). A 64-instruction block contains exactly one interrupt or trap handling routine, and blocks are aligned on 64-instruction address boundaries.



08996A 37

Figure 3-33. Vector Table Entry

Vector Numbers

When an interrupt or trap is taken, the processor determines an 8-bit vector number associated with the interrupt or trap. The vector number gives either the number of a vector table entry or the number of an instruction block, depending on the value of the VF bit. If the VF bit is 1, the physical address of the vector table entry is generated by replacing bits 9-2 of the value in the Vector Area Base Address Register with the vector number. If the VF bit is 0, the physical address of the first instruction of the handling routine is generated by replacing bits 15-8 of the value in the Vector Table Base Address Register with the vector number.

Vector numbers are either predefined, or specified by an instruction causing the trap. The

assignment of vector numbers is shown in Table 3-10. Vector numbers 64 to 255 are for use by trapping instructions; the definition of the routines associated with these numbers is system-dependent.

Table 3-10. Vector Number Assignments

<u>Number</u>	<u>Type of Trap or Interrupt</u>
0	Illegal Opcode
1	Unaligned Access
2	Out of Range
3	Coprocessor Not Present
4	Coprocessor Exception
5	Protection Violation
6	Instruction Access Exception
7	Data Access Exception
8	User-Mode Instruction TLB Miss
9	User-Mode Data TLB Miss
10	Supervisor-Mode Instruction TLB Miss
11	Supervisor-Mode Data TLB Miss
12	Instruction TLB Protection Violation
13	Data TLB Protection Violation
14	Timer
15	Trace
16	*INTR0
17	*INTR1
18	*INTR2
19	*INTR3
20	*TRAP0
21	*TRAP1
22-31	reserved
32	MULTIPLY
33	DIVIDE
34	reserved
35	reserved
36	CVINTF
37	CVINTD
38	CVFINT
39	CVDINT
40	CVFD
41	CVDF
42	FEQ
43	DEQ
44	FGT
45	DGT

Table 3-10. Vector Number Assignments (Continued)

<u>Number</u>	<u>Type of Trap or Interrupt</u>
46	FLT
47	DLT
48	FADD
49	DADD
50	FSUB
51	DSUB
52	FMUL
53	DMUL
54	FDIV
55	DDIV
56-63	reserved
64-255	Assert and EMULATE instruction traps (vector number specified by instruction)

3.5.5 INTERRUPT AND TRAP HANDLING

Interrupt and trap handling consists of two distinct operations: taking the interrupt or trap, and returning from the interrupt or trap handler. If the interrupt or trap handler returns immediately to the interrupted routine, the interrupt or trap handler need not save and restore the processor state.

Taking An Interrupt or Trap

The following operations are performed in sequence by the processor when an interrupt or trap is taken:

- 1) Instruction execution is suspended.
- 2) Instruction fetching is suspended.
- 3) Any in-progress load or store operation is completed. Any additional operations are cancelled in the case of load-multiple and store-multiple.
- 4) The contents of the Current Processor Status Register are copied into the Old Processor Status Register.
- 5) The Current Processor Status register is modified as shown in Figure 3-34 (the value "u" means unaffected). Note that setting the Freeze (FZ) bit freezes the Channel Address, Channel Data, Channel Control, Program Counter 0, Program Counter 1, Program Counter 2, and ALU Status Registers.

- 6) The address of the first instruction of the interrupt or trap handler is determined. If the VF bit of the Configuration Register is 1, the address is obtained by fetching a vector from instruction/data memory, using the physical address obtained from the Vector Area Base Address Register and the vector number. If the VF bit is 0, the instruction address is given directly by the Vector Area Base Address Register and the vector number.
- 7) If the VF bit is 1, the R bit in the vector fetched in step 6 is copied into the RE bit of the Current Processor Status Register. If the VF bit is 0, the RV bit of the Configuration Register is copied into the RE bit. This step determines whether or not the first instruction of the interrupt handler is in instruction ROM.
- 8) An instruction fetch is initiated using the instruction address determined in step 6. At this point, normal instruction execution resumes.




Figure 3-34. Current Processor Status After an Interrupt or Trap

Note that the processor does not explicitly save the contents of any registers when an interrupt is taken. If register saving is required, it is the responsibility of the interrupt or trap handling routine. For proper operation, registers must be saved before any further interrupts or traps may be taken.

Returning From an Interrupt or Trap

Two instructions are used to resume the execution of an interrupted program: Interrupt Return (IRET), and Interrupt Return and Invalidate (IRETINV). These instructions are identical except in one respect: the IRETINV instruction resets all Valid bits in the Branch Target Cache, whereas the IRET instruction does not affect the Valid bits.

In some situations, the processor state must be properly set by software before the interrupt return is executed. The following is a list of operations normally performed in such cases:

- 1) The Current Processor Status is configured as shown in Figure 3-35 (the value "x" is a don't care). Note that setting the FZ bit freezes the registers listed below so that they may be set for the interrupt return.
- 2) The Old Processor Status is set to the value of the Current Processor Status for the target routine.
- 3) The Channel Address, Channel Data, and Channel Control registers are set to restart or resume uncompleted channel operations of the target routine.

- 4) The Program Counter 1 and Program Counter 0 registers are set to the addresses of the first and second instructions, respectively, to be executed in the target routine.
- 5) Other registers are set as required. These may include registers such as the ALU Status, Q, and so forth, depending on the particular situation. Some of these registers are not affected by the FZ bit, so they must be set in such a manner that they are not modified unintentionally before the interrupt return.




Figure 3-35. Current Processor Status Before Interrupt Return

Once the processor registers are properly configured, as described above, an interrupt return instruction performs the remaining steps necessary to return to the target routine. The following operations are performed by the interrupt return instruction:

- 1) Any in-progress load or store operation is completed. If a load-multiple or store-multiple sequence is in progress, the interrupt return is not executed until the sequence completes.
- 2) Interrupts and traps are disabled, regardless of the settings of the DA, DI, and IM fields of the Current Processor Status, for steps 3 through 10.
- 3) If the interrupt return instruction is an IRETINV, all Valid bits in the Branch Target Cache are reset.
- 4) The contents of the Old Processor Status Register are copied into the Current Processor Status Register. This normally resets the FZ bit allowing the Program Counter 0, 1, 2, Channel Address, Data, Control, and ALU Status registers to update normally. Since certain bits of the Current Processor Status Register are always updated by the processor, this copy operation may be irrelevant for certain bits (i.e., the Interrupt Pending bit).
- 5) If the Contents Valid (CV) bit of the Channel Control Register is 1, and the Not Needed (NN) and Multiple Operation (ML) bits are both 0, an external access is started. This operation is based on the contents of the Channel Address, Channel Data, and Channel Control registers. The Current Processor Status Register conditions the access—as is normally the case. Note that load-multiple and store-multiple operations are not restarted at this point.
- 6) The address in Program Counter 1 is used to fetch an instruction. The Current Processor Status Register conditions the fetch. This step is treated as a branch in

the sense that the processor searches the Branch Target Cache for the target of the fetch.

- 7) The instruction fetched in step 6 enters the decode stage of the pipeline.
- 8) The address in Program Counter 0 is used to fetch an instruction. The Current Processor Status Register conditions the fetch. This step is treated as a branch in the sense that the processor searches the Branch Target Cache for the target of the fetch.
- 9) The instruction fetched in step 6 enters the execute stage of the pipeline, and the instruction fetched in step 8 enters the decode stage.
- 10) If the CV bit in the Channel Control Register is a 1, the NN bit is 0, and the ML bit is 1, a load-multiple or store-multiple sequence is started, based on the contents of the Channel Address, Channel Data, and Channel Control registers.
- 11) Interrupts and traps are enabled per the appropriate bits in the Current Processor Status Register.
- 12) The processor resumes normal operation.

Fast Interrupt Processing

The registers affected by the FZ bit of the Current Processor Status Register are those which are modified by almost any usual sequence of instructions. Since the FZ bit is set by an interrupt or trap, the interrupt or trap handler is able to execute while not disturbing the state of the interrupted routine, though its execution is somewhat restricted. Thus, it is not necessary in many cases for the interrupt or trap handler to save the registers which are affected by the FZ bit.

The processor provides an additional benefit if the Program Counter 0 and Program Counter 1 Registers are not modified by the interrupt or trap handler. If Program Counters 0 and 1 contain the addresses of sequential instructions when an interrupt or trap is taken, and if they are not modified before an interrupt return is executed, step 8 of the interrupt return sequence above occurs as a sequential fetch—instead of a branch—for the interrupt return. The performance impact of a sequential fetch is normally less than that of a non-sequential fetch.

Because the registers affected by the FZ bit are sometimes required for instruction execution, it is not possible for the interrupt or trap handler to execute all instructions, unless the required registers are first saved elsewhere (e.g. in one or more global registers). Most of the restrictions due to register dependencies are obvious (e.g. the Byte Pointer for byte extracts), and will not be discussed here. Other, less obvious restrictions are listed below:

- 1) Load Multiple and Store Multiple. The Channel Address, Channel Data, and Channel Control registers are used to sequence load-multiple and store-multiple operations, so these instructions cannot be executed while the registers are frozen. However, note that other external accesses may occur; the Channel Address,

Channel Data, and Channel Control registers are required only to restart an access after an exception, and the interrupt or trap handler is not expected to encounter any exceptions.

- 2) Loads and stores which set the Byte Pointer. If the Set Byte Pointer (SB) of a load or store instruction is 1, and the FZ bit is also 1, there is no effect on the Byte Pointer. Thus, the execution of external byte and half-word accesses using this mechanism is not possible.
- 3) Extended arithmetic. The Carry bit of the ALU Status Register is not updated while the FZ bit is 1.
- 4) Divide instructions. The Divide Flag of the ALU Status Register is not updated when the FZ bit is 1.

If the interrupt or trap handler does not save the state of the interrupted routine, it cannot allow additional interrupts and traps. Also, the operation of the interrupt or trap handler cannot depend on any trapping instructions (e.g. Floating-Point instructions, illegal operation codes, arithmetic overflow, etc.), since these are disabled. There are certain cases, however, where traps are unavoidable; these are discussed in Section 3.5.9.

3.5.6 *WARN TRAP

The processor recognizes a special trap, caused by the activation of the *WARN input, which cannot be masked. The *WARN trap is intended to be used for severe system-error or deadlock conditions. It allows the processor to be placed in a known, operable state, while preserving much of its original state for error reporting and, possibly, recovery. Therefore, it shares some features in common with the Reset mode as well as features common to other traps described in this section.

The major differences between the *WARN trap and other traps are:

- 1) The processor does not wait for an in-progress external access to complete before taking the trap, since this access might not complete. However, the information related to any outstanding access is retained by the Channel Address, Channel Data, and Channel Control registers when the trap is taken.
- 2) The vector fetch operation is not performed, regardless of the VF bit of the Configuration Register, when the *WARN trap is taken. The ROM Enable (RE) bit in the Current Processor Status is set, and instruction fetching begins immediately at address 16 in the instruction ROM. The trap handler can execute directly from the instruction ROM without the need to access external (and possibly non-functional or invalid) instruction/data memory.
- 3) The *WARN trap sets the Current Processor Status Register as for the Reset Mode (see Figure 3-39, Section 3.8) rather than as for other traps (see Figure 3-34, Section 3.5.5). However, before the Current Processor Status Register is set, its contents are copied to the Old Processor Status Register, as for other traps; this is not the case for the Reset mode.

Note that *WARN trap may disrupt the state of the routine which is executing when it is taken, prohibiting this routine from being restarted.

3.5.7 SEQUENCING OF INTERRUPTS AND TRAPS

On every cycle, the processor decides either to execute instructions or to take an interrupt or trap. Since there are multiple sources of interrupts and traps, more than one interrupt or trap may be pending on a given cycle.

To resolve conflicts, interrupts and traps are taken according to the priority shown in Table 3-11. In this table, interrupts and traps are listed in order of decreasing priority. This section discusses the first three columns of Table 3-11. The last two columns are discussed in Section 3.5.8.

In Table 3-11, interrupts and traps fall into one of two categories depending on the timing of their occurrence relative to instruction execution. These categories are indicated in the third column of the table by the labels "inst" and "async." These labels have the following meaning:

- 1) Inst—generated by the execution or attempted execution of an instruction.
- 2) Async—generated asynchronous to and independent from the instruction being executed, although it may be a result of an instruction previously executed.

The principle for interrupt and trap sequencing is that the highest priority interrupt or trap is taken first. Other interrupts and traps remain active until they can be taken, or are regenerated when they can be taken. This is accomplished, depending on the type of interrupt or trap, as follows:

- 1) All traps in Table 3-11 with priority 13 or 14 are regenerated by the re-execution of the causing instruction.
- 2) Most of the interrupts and traps of priority 4 through 12 must be held by external hardware until they are taken. The exceptions to this are listed in 3) below.
- 3) The exceptions to 2) above are the Data Access Exception trap, the Coprocessor Exception trap, the Timer interrupt, and the Trace trap. These are caused by bits in various registers in the processor and are held by these registers until taken or cleared. The relevant bits are: the Transaction Faulted (TF) bit of the Channel Control Register for Data Access Exception and Coprocessor Exception traps, the Interrupt (IN) bit of the Timer Reload Register for Timer interrupts, and the Trace Pending (TP) bit of the Current Processor Status Register for Trace traps.
- 4) All traps of priority 2 and 3 in Table 3-11, except for the Unaligned Access trap, are not regenerated. These traps are mutually exclusive, and are given high priority because they cannot be regenerated; they must be taken if they occur. If

one of these traps occurs at the same time as a reset or *WARN trap, it is not taken, and its occurrence is lost.

- 5) The Unaligned Access trap is regenerated internally when an external access is restarted by the Channel Address, Channel Data, and Channel Control registers. Note that this trap is not necessarily exclusive to the traps discussed in 4) above.

Table 3-11. Interrupt and Trap Priority Table

PRIORITY	TYPE OF INTERRUPT OR TRAP	INST/ASYNC	PC1	Channel Regs
1 (highest)	*WARN	async	next	see Note 1
2	User-Mode Data TLB Miss	inst	next	all
	Supervisor-Mode Data TLB Miss	inst	next	all
	Data TLB Protection Violation	inst	next	all
3	Unaligned Access	inst	next	all
	Coprocessor Not Present	inst	next	all
	Out of Range	inst	next	N/A
	Assert Instructions	inst	next	N/A
	Floating-Point Instructions	inst	next	N/A
	MULTIPLY	inst	next	N/A
4	DIVIDE	inst	next	N/A
	EMULATE	inst	next	N/A
4	Data Access Exception	async	next	all
	Coprocessor Exception	async	next	all
5	*TRAP0	async	next	multiple
6	*TRAP1	async	next	multiple
7	*INTR0	async	next	multiple
8	*INTR1	async	next	multiple
9	*INTR2	async	next	multiple
10	*INTR3	async	next	multiple
11	Timer	async	next	multiple
12	Trace	async	next	multiple
13	User-Mode Instruction TLB Miss	inst	curr	N/A
	Supervisor-Mode Instr. TLB Miss	inst	curr	N/A
	Instruction TLB Protection Violation	inst	curr	N/A
	Instruction Access Violation	inst	curr	N/A
14 (lowest)	Illegal Opcode	inst	curr	N/A
	Protection Violation	inst	curr	N/A

08996A 40

Note1: The Channel Address, Channel Data, and Channel Control registers are set for a *WARN trap only if an external access is in progress when the trap is taken.

3.5.8 EXCEPTION REPORTING AND RESTARTING

When an instruction encounters an exceptional condition, the Program Counter 0, Program Counter 1, and Program Counter 2 registers report the relevant instruction address(es), and allow the instruction sequence to be restarted once the exceptional condition has been remedied (if possible). Similarly, when an external access or coprocessor transfer encounters an exceptional condition, the Channel Address, Channel Data, and Channel Control registers report information on the access or transfer, and allow it to be restarted. This section describes the interpretation and use of these registers.

Instruction Exceptions

The "PC1" column in Table 3-11 describes the value held in the Program Counter 1 Register (PC1) when the interrupt or trap is taken. For traps in the "inst" category, PC1 contains either the address of the instruction causing the trap, indicated by "curr," or the address of the instruction following the instruction causing the trap, indicated by "next".

For interrupts and traps in the "async" category, PC1 contains the address of the first instruction which was not executed due to the taking of the interrupt or trap. This is the next instruction to be executed upon interrupt return, as indicated by "next" in the PC1 column.

For traps caused by the execution of an instruction (for example, the Out of Range trap), the Program Counter 2 Register contains the address of the instruction causing the trap. In all of these cases, PC1 is in the "next" category.

The traps associated with instruction fetches (i.e. those of priority 13) occur only if the processor attempts the execution of the associated instruction. An exception may be detected during an instruction prefetch, but the associated trap does not occur if a non-sequential fetch occurs before the processor attempts the execution of the invalid instruction. This prevents the spurious indication of instruction exceptions.

Data Exceptions

The "Channel Regs" column of Table 3-11 indicates the cases for which the Channel Address, Channel Data, and Channel Control registers contain information related to an external access or coprocessor transfer (these registers are collectively termed "channel registers" in the following discussion). For the cases indicated, the access or transfer did not complete because of some exceptional condition. Note that the Channel Data Register contains relevant information only in the case of a store.

For the *WARN trap, the channel registers are valid only if a load or store were in progress when the trap was taken. Recall that the *WARN trap does not wait for any in-progress access to complete.

For the traps with an "all" in the "Channel Regs" column of Table 3-11, the channel registers contain information relevant to the trap in all cases. These traps are associated with exceptional events during external accesses or coprocessor transfers.

For the traps with a “multiple” in the “Channel Regs” column, the channel registers might contain information for restarting an interrupted load-multiple or store-multiple operation. In these cases, the operation did not encounter an exception, but was simply cancelled for latency considerations.

The information contained in the channel registers allows the processor to restart the related operation during an interrupt return sequence, without any special assistance by software. Software must only insure that the relevant information is retained in, or restored to, the channel registers before an interrupt return is executed.

3.5.9 EXCEPTIONS DURING INTERRUPT AND TRAP HANDLING

In most cases, interrupt and trap handling routines are executed with the DA bit in the Current Processor Status having a value of 1. It is assumed that these routines do not create many of the exceptions possible in most other processor routines, so most of these are ignored.

If the assumption of no exceptions is not valid for a particular interrupt or trap handler, it is important that the handler save the state of the processor and reset the FZ bit of the Current Processor Status, so that the handler itself may be restarted properly. This must be accomplished before any interrupts or traps can be taken. Of course, in this case, the state (or the state of some other process) must be restored before an interrupt return is executed.

It is possible that errors reported via the *IERR and *DERR signals are associated with hardware errors, independent of any routine being executed. For this reason, the Instruction Access Exception, Data Access Exception, and Coprocessor Exception traps cannot be disabled by the DA bit, and the processor may take one of these traps even while handling another interrupt or trap.

If the processor does take an unmaskable trap while handling another interrupt or trap, and the state of the interrupt or trap handler is not reflected in processor registers, it is not possible to return to the point at which the unmaskable trap is taken. When the unmaskable trap is taken, the processor state saved is that state associated with the original interrupt or trap, not with the unmaskable trap; however, the Old Processor Status Register is modified to reflect the Current Processor Status Register of the interrupt or trap handler. This situation, indicated by the DA bit being 1 in the Old Processor Status Register, may not be recoverable.

3.6 MEMORY MANAGEMENT

The Am29000 incorporates a Memory Management Unit (MMU) for performing virtual-to-physical address translation and memory access protection. This section describes the logical operation of the Memory Management Unit. Related issues are discussed in Sections 7.2.4 and 7.2.5.

Address translation can be performed only for instruction/data memory accesses. No address translation is performed for instruction ROM, input/output, coprocessor, or interrupt/trap vector accesses.

3.6.1 TRANSLATION LOOK-ASIDE BUFFER

The MMU stores the most recently performed address translations in a special cache, the Translation Look-Aside Buffer (TLB). All virtual addresses generated by the processor are translated by the TLB. Given a virtual address, the TLB determines the corresponding physical address.

The TLB reflects information in the processor system page tables, except that it specifies the translation for many fewer pages; this restriction allows the TLB to be incorporated on the processor chip where the performance of address translation is maximized.

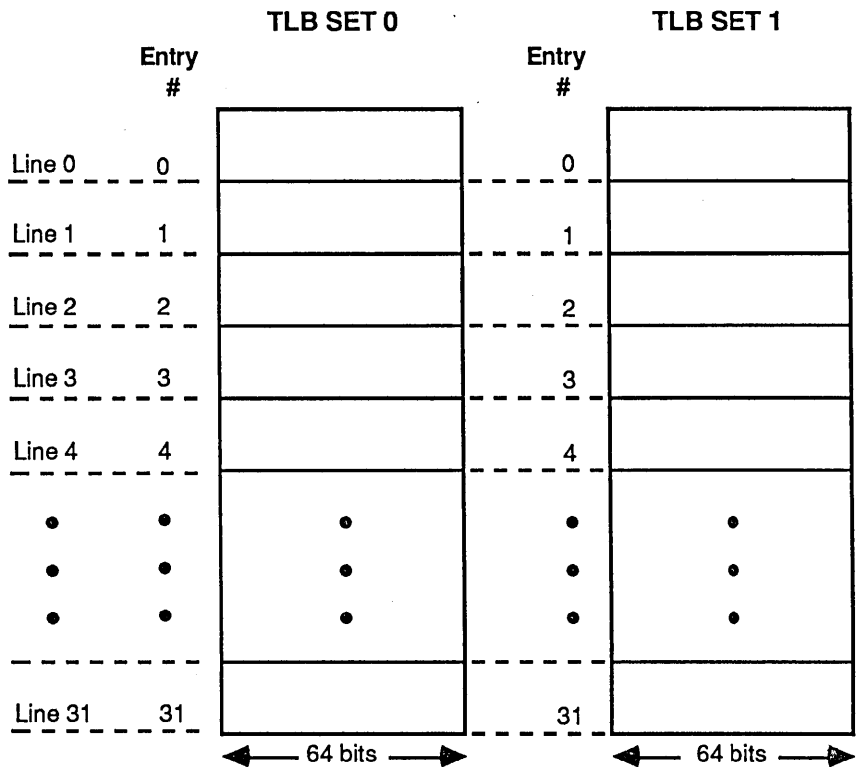
A diagram of the TLB is shown in Figure 3-36. The TLB is a table of 64 entries, divided into two equal sets, called Set 0 and Set 1. Within each set, entries are numbered 0 to 31. Entries in different sets which have equivalent entry-numbers are grouped into a unit called a line; there are thus 32 lines in the TLB, numbered 0 to 31.

Each TLB entry is 64 bits long, and contains mapping and protection information for a single virtual page. TLB entries may be inspected and modified by processor instructions executed in the Supervisor mode. The layout of TLB entries is described in Section 3.2.3.

The TLB stores information about the ownership of the TLB entries in an 8-bit Task Identifier (TID) field in each entry. This makes it possible for the TLB to be shared by several independent processes without the need for invalidation of the entire TLB as processes are activated. It also increases system performance by permitting processes to warm-start (i.e., to start execution on the processor with a certain number of TLB entries remaining in the TLB from a previous execution).

Each TLB entry contains two bits to assist management of the TLB entries. These are the Usage and Flag bits. The Usage bit indicates which set of the entry within a given line was least-recently used to perform an address translation. Usage bits for two entries in the same line are equivalent. The Flag bit has no effect on address translation, and is not affected by the processor except by explicit writes to the TLB. This bit is provided only for use by software.

The TLB contains other fields which are described in the following sections.



08996A3-36


Figure 3-36. Translation Look-Aside Buffer Organization

3.6.2 ADDRESS TRANSLATION


For the purpose of address translation, the virtual instruction/data address-space of a process is partitioned into regions of fixed size, called pages, which are mapped by the address-translation process into equivalent-sized regions of physical memory, called page frames. All accesses to instructions or data contained within a given page use the same virtual-to-physical address translation.

Pages may be of size 1, 2, 4, or 8 Kbytes, as specified by the MMU Configuration Register. Virtual addresses are partitioned into three fields for the address-translation process, as shown in Figure 3-37. The partitioning of the virtual address is based on the page size. The fields shown in Figure 3-37 are described in the following discussion.


1 Kbyte Page Size:




2 Kbyte Page Size:



4 Kbyte Page Size:



8 Kbyte Page Size:



08996A 41

Figure 3-37. Virtual Address for 1, 2, 4, and 8 Kbyte Pages

Address Translation Controls

The processor attempts to perform address translation for the following external accesses:

- 1) Instruction accesses, if the Physical Addressing/Instructions (PI) and ROM Enable (RE) bits of the Current Processor Status are both 0.
- 2) User-mode accesses to instruction/data memory if the Physical Addressing/Data (PD) bit of the Current Processor Status is 0.
- 3) Supervisor-mode accesses to instruction/data memory if the Physical Address (PA) bit of the load or store instruction performing the access is 0, and the PD bit of the Current Processor Status is 0.

Address translation is also controlled by the MMU Configuration Register. This register specifies the virtual page size, and contains an 8-bit Process Identifier (PID) field. The PID field specifies the process-number associated with the currently-running program. This value is compared with Task Identifier (TID) fields of the TLB entries during address translation. The TID field of a TLB entry must match the PID field for the translation to be valid.

Address Translation Process

The address-translation process is diagrammed in Figure 3-38. Address translation is performed by the following fields in the TLB entry: the Virtual Tag (VTAG), the Task Identifier (TID), the Valid Entry (VE) bit, and the Real Page Number (RPN). To perform an address translation, the processor accesses the TLB line whose number is given by certain bits in the virtual address. The bits used depend on the page size as follows:

<u>Page Size</u>	<u>Virtual Address Bits (for Line Access)</u>
1 Kbyte	14–10
2 Kbyte	15–11
4 Kbyte	16–12
8 Kbyte	17–13

The accessed line contains two TLB entries, which in turn contain two VTAG fields. The VTAG fields are both compared to bits in the virtual address. This comparison depends on the page size as follows (note that VTAG bit-numbers are relative to the VTAG field, not the TLB entry):

<u>Page Size</u>	<u>Virtual Address Bits</u>	<u>VTAG Bits</u>
1 Kbyte	31–15	16–0
2 Kbyte	31–16	16–1
4 Kbyte	31–17	16–2
8 Kbyte	31–18	16–3

Note: Certain bits of the VTAG field do not participate in the comparison for page sizes larger than 1 Kbyte. These bits of the VTAG field are required to be zero-bits.

For an address translation to be valid, the following conditions must be met:

- 1) The virtual address bits match corresponding bits of the VTAG field as specified above.
- 2) The TID field in the TLB entry matches the PID field in the MMU Configuration Register.
- 3) The VE bit in the TLB entry is 1.
- 4) Only one entry in the line meets conditions 1, 2, and 3 above. If this condition is not met, the results of the translation may be treated as valid by the processor, but the results are unpredictable.

If the address translation is valid for one TLB entry in the selected line, the RPN field in this entry is used to form the physical address of the access. The RPN field gives the portion of the physical address that depends on the translation; the remaining portion of the virtual address—called the Page Offset—is invariant with address translation.

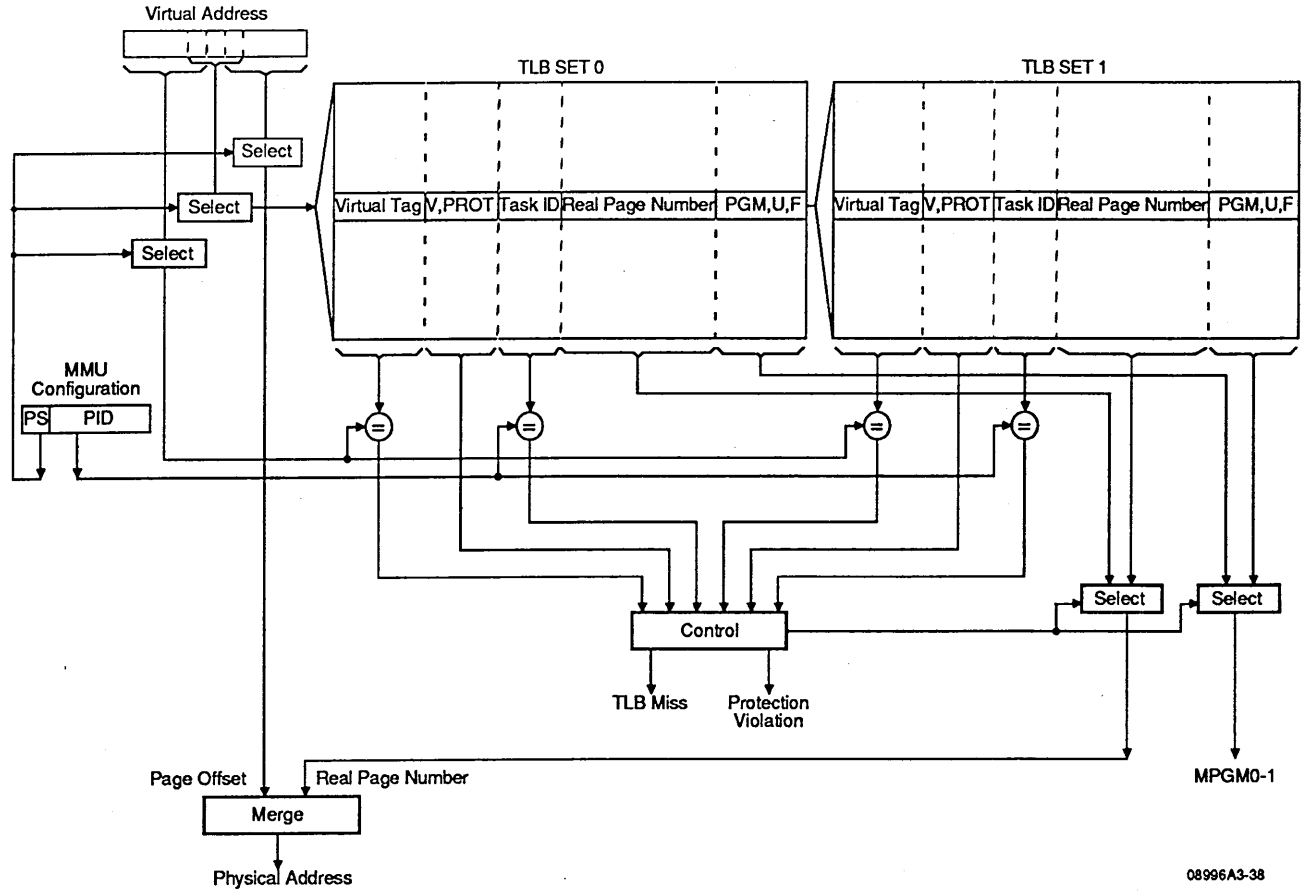


Figure 3-38. Address Translation Process

The Page Offset comprises the low-order bits of the virtual address, and gives the location of a byte (because of byte addressing) within the virtual page. This byte is located at the same position in the physical page frame, so the Page Offset also comprises the low-order bits of the physical address.

The 32-bit physical address is the concatenation of certain bits of the RPN field and Page Offset, where the bits from each depend on the page size as follows (note that RPN bit-numbers are relative to the RPN field, not the TLB entry):

<u>Page Size</u>	<u>RPN Bits</u>	<u>Virtual Address Bits for Page Offset</u>
1 Kbyte	21–0	9–0
2 Kbyte	21–1	10–0
4 Kbyte	21–2	11–0
8 Kbyte	21–3	12–0

Note: Certain bits of the RPN field are not used in forming the physical address for page sizes greater than 1 Kbyte. These bits of the RPN are required to be zero-bits. In addition, for certain instruction accesses, the Page Offset is incremented by 16 as described in Section 4.2.3.

Successful and Unsuccessful Translations

If an address translation is successful, the TLB entry is further used to perform protection checking for the access. Bits in the TLB make it possible to restrict accesses—independently for Supervisor-mode and User-mode accesses—to any combination of load, store, and instruction accesses, or to no access. Section 3.6.5 describes protection in more detail.

If the address translation is valid, and no protection violation is detected, the physical address from the translation is placed on the processor's Address Bus, and the access is initiated. If the translation is not valid, or a protection violation is detected, a trap occurs. Depending on the state of the channel interface, the access request may be placed on the Address Bus with the signal *BINV asserted, even though the trap occurs.

Also, if the address translation is successful, and there is no protection violation, the PGM bits from the TLB entry used for translation are placed on the MPGM0–MPGM1 outputs during the address cycle for the access. If address translation is not performed, these pins are both Low for the address cycle.

If the TLB cannot translate an address, a TLB miss occurs. The MMU causes a trap if either a TLB miss occurs, or the translation is successful and a protection violation is detected. The processor distinguishes between traps caused by instruction and data accesses, and between traps caused by User-and Supervisor-mode accesses, as follows:

<u>Trap Vector Number</u>	<u>Type of Trap</u>
8	User-Mode Instruction TLB Miss
9	User-Mode Data TLB Miss
10	Supervisor-Mode Instruction TLB Miss
11	Supervisor-Mode Data TLB Miss
12	Instruction TLB Protection Violation
13	Data TLB Protection Violation

The distinction between the above traps is made to assist trap handling, particularly the routines which load TLB entries.

3.6.3 RELOAD

So that the MMU may support a large variety of memory-management architectures, it does not directly load TLB entries which are required for address translation. It simply causes a TLB miss trap when an address translation is unsuccessful. The trap causes a program—called the TLB reload routine—to execute. The TLB reload routine is defined according to the structure and access method of the page table contained in an external device or memory.

When a TLB miss trap occurs, the LRU Recommendation Register is written with the TLB register-number for Word 0 of the TLB entry to be used by the TLB reload routine. For instruction accesses, the Program Counter 1 Register contains the instruction address which was not successfully translated. For data accesses, the Channel Address Register contains the data address which was not successfully translated.

The TLB reload routine determines the translation for the address given by the Program Counter 1 Register or Channel Address Register, as appropriate. The TLB reload routine uses an external page table to determine the required translation, and loads the TLB entry indicated by the LRU Recommendation Register so that it may perform this translation. In a demand-paged environment, the TLB reload routine may additionally invoke a page-fault handler when the translation cannot be performed.

TLB entries are written by the Move To TLB (MTTLB) instruction, which copies the contents of a general-purpose register into a TLB register. The TLB register-number is specified by bits 9-0 of a general purpose register. TLB entries are read by the Move From TLB (MFTLB) instruction, which copies the contents of a TLB register into a general-purpose register. Again, the TLB register-number is specified by a general purpose register.

3.6.4 ENTRY INVALIDATION

There are two methods for invalidating TLB entries which are no longer required at a given point in program execution. The first involves resetting the Valid Entry bit of a single

entry (this is done by a Move To TLB instruction). The second involves changing the value of the Process Identifier (PID) field of the MMU Configuration Register; this invalidates all entries whose Task Identifier (TID) fields do not match the new value.

If an entry is invalidated by changing the PID field, the TLB entry still remains valid in some sense. If the PID field is changed again to match the TID field, the entry may once again participate in address translation. This ability can be used to reduce the number of TLB misses in a system during process switching. However, it is important to manage TLB entries so that an invalid match cannot occur between the PID field and the TID field of an old TLB entry.

3.6.5 PROTECTION

If an address translation is performed successfully as described in Section 3.6.2, the TLB entry used in address translation is used to perform protection checking for the access. There are 6 bits in the TLB entry for this purpose: Supervisor Read (SR), Supervisor Write (SW), Supervisor Execute (SE), User Read (UR), User Write (UW), and User Execute (UE). These bits restrict accesses, depending on the program mode of the access, as follows (the value "x" is a don't care):

<u>SM</u>	<u>SR</u>	<u>SW</u>	<u>SE</u>	<u>UR</u>	<u>UW</u>	<u>UE</u>	<u>Type of Access Allowed</u>
0	x	x	x	0	0	0	No User access
0	x	x	x	0	0	1	User instruction
0	x	x	x	0	1	0	User store
0	x	x	x	0	1	1	User store or instruction
0	x	x	x	1	0	0	User load
0	x	x	x	1	0	1	User load or instruction
0	x	x	x	1	1	0	User load or store
0	x	x	x	1	1	1	Any User access
1	0	0	0	x	x	x	No Supervisor access
1	0	0	1	x	x	x	Supervisor instruction
1	0	1	0	x	x	x	Supervisor store
1	0	1	1	x	x	x	Supervisor store or instruction
1	1	0	0	x	x	x	Supervisor load
1	1	0	1	x	x	x	Supervisor load or instruction
1	1	1	0	x	x	x	Supervisor load or store
1	1	1	1	x	x	x	Any Supervisor access

Note that for the Load and Set (LOADSET) instruction, the protection bits must be set to allow both the load and store access.

If protection checking indicates that a given access is not allowed, a Data TLB Protection Violation or Instruction TLB Protection Violation trap occurs. The cause of the trap is determined by inspection of the Program Counter 1 Register for an Instruction TLB Protection Violation, or by inspection of the contents of the Channel Address and Channel Control registers for a Data TLB Protection Violation.

3.7 SERIALIZATION

Since the Am29000 overlaps external data references with other operations, it is necessary to restrict these other operations so that the data reference may be restarted if necessary.

The restriction is, in general, that the processor cannot perform any operation that changes the conditions under which the external access occurs. This is accomplished by serializing certain operations; that is, by having the processor enter the Pipeline Hold mode if an external access has not completed when the operation is attempted. Serialization is performed for the following operations:

- 1) The execution of one of the following instructions:
 - Move to Special Register
 - Move to Special Register Immediate
 - Move To TLB
 - Interrupt Return
 - Interrupt Return and Invalidate
 - Halt

- 2) The taking of an interrupt or trap, except for a *WARN trap.

If the processor is in the Pipeline Hold mode due to serialization, it enters the Executing mode once the external access completes. Note that the processor may immediately take a Data Access Exception or Coprocessor Exception trap instead of performing the operation which originally caused the serialization.

3.8 INITIALIZATION

When power is first applied to the processor, it is in an unknown state, and must be placed in a known state. Also, under certain circumstances, it may be necessary to place the processor in a defined state. This is accomplished by the Reset mode, which is invoked by activating the *RESET pin for the required duration. The Reset mode configures the processor state as follows:

- 1) Instruction execution is suspended.
- 2) Instruction fetching is suspended.
- 3) Any interrupt or trap conditions are ignored.
- 4) The Current Processor Status Register is set as shown in Figure 3-39.
- 5) The Cache Disable bit of the Configuration Register is set.
- 6) The Contents Valid bit of the Channel Control Register is reset.




Figure 3-39. Current Processor Status Register In Reset Mode

Except as previously noted, the contents of all general-purpose registers, special-purpose registers, and TLB registers are undefined. The contents of the Branch Target Cache are also undefined.

The Reset mode also configures the processor to initiate an instruction fetch using an address of 0. Since the ROM enable (RE) bit of the Current Processor Status is 1, this fetch is directed to external instruction read-only memory. This fetch occurs when the Reset mode is exited (i.e. when the *RESET input is de-asserted). Section 5.5 contains more information on this instruction fetch.

CHAPTER 4

HARDWARE FEATURES

This chapter describes the operation of the Am29000 pipeline, and the processor's three major functional units. The functional units are: the Instruction Fetch Unit, the Execution Unit, and the Memory Management Unit. These units, which were shown in abstract form in Figure 2-2, are shown in detail in Figure 4-1.

The operation of the functional units is coordinated by the Pipeline Hold mode, which insures that operations are performed in the proper order. This chapter also describes the Pipeline Hold mode.

Since this chapter describes the internal operation of the Am29000, it provides information which may not be required by some users. However, it aids an understanding of the behavior of the Am29000 under certain conditions, especially the behavior of the system interfaces described in Chapter 5.

4.1 FOUR-STAGE PIPELINE

The Am29000 implements a four-stage pipeline for instruction execution. The four stages are: fetch, decode, execute, and write-back. The pipeline is organized so that the effective instruction-execution rate may be as high as one instruction per cycle.

During the fetch stage, the Instruction Fetch Unit (Section 4.2) determines the location of the next processor instruction, and issues the instruction to the decode stage. The instruction is fetched either from the Instruction Prefetch Buffer, the Branch Target Cache, or an external instruction memory.

During the decode stage, the Execution Unit (Section 4.3) decodes the instruction selected during the fetch stage, and fetches and/or assembles the required operands. It also evaluates addresses for branches, loads, and stores.

During the execute stage, the Execution Unit performs the operation specified by the instruction. In the case of branches, loads, and stores, the Memory Management Unit (Section 4.4) performs address translation if required.

During the write-back stage, the results of the operation performed during the execute stage are stored. In the case of branches, loads, and stores, the physical address resulting from translation during the execute stage is transmitted to an external device or memory.

Most pipeline dependencies which are internal to the processor are handled by forwarding logic in the processor. For those dependencies which result from the external system, the Pipeline Hold mode insures proper operation.

In a few special cases (see Section 7.3) the processor pipeline is exposed to software executing on the Am29000.

4.2 INSTRUCTION FETCH UNIT

The Instruction Fetch Unit performs the functions required to keep the processor pipeline supplied with instructions. Since the processor can execute one instruction per cycle, instructions must be supplied at this rate if the execution stage is to perform at the maximum rate. To accomplish this, the Instruction Fetch Unit contains mechanisms for requesting instructions from instruction memory before they are required for execution, and for caching the most recently executed branch target instructions.

The Instruction Fetch Unit also incorporates the logic necessary to calculate and sequence instruction addresses. The processor is word-oriented, but generates byte addresses for all external accesses. Since all processor instructions are word-length, and are aligned on word-address boundaries, the Instruction Fetch Unit deals only with 30-bit addresses. For external instruction accesses, these addresses are appended with 00 in the two least-significant bits to

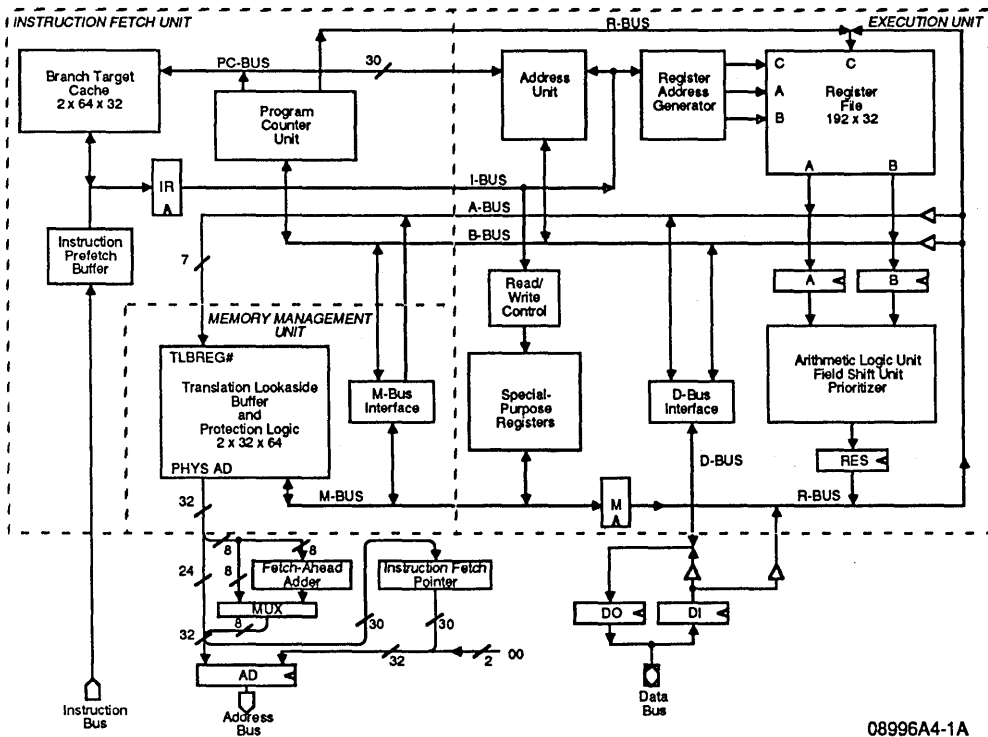


Figure 4-1. Am29000 Data Flow

form the required 32-bit address (note that the two least-significant bits of an external instruction address may not be 00 for indirect jumps).

4.2.1 INSTRUCTION PREFETCH BUFFER

All instructions executed by the processor are fetched either from the Branch Target Cache or from external instruction memory (i.e. instruction/data memory or instruction read-only memory). When instructions are fetched from the external memory, they are requested in advance, to assist the timing of instruction accesses. The processor attempts to initiate the fetch for any given instruction at least four cycles before it is required for execution.

Since instructions are requested in advance, based on a predicted need, it is possible that a prefetched instruction is not required immediately for execution when the prefetch completes. To accommodate this possibility, the Instruction Fetch Unit contains a four-word Instruction Prefetch Buffer (IPB), as shown in Figure 4-1. The IPB is a circularly-addressed buffer which acts as a first-in/first-out (FIFO) queue for instructions.

If instruction fetching is enabled, the processor requests an external instruction fetch on any cycle for which the IPB contains an available location. Instructions are stored in the IPB as they are returned from the external instruction memory. An instruction is stored into the IPB location whose number is given by bits 3-2 of the instruction address.

The instruction is held in the IPB until it is required for execution. When required, the instruction is sent to the decode stage, and the IPB location is freed to receive a subsequent instruction.

Instruction Prefetch Stream

An instruction prefetch stream is established whenever the processor performs a non-sequential instruction reference. Non-sequential references normally occur as the result of successful branches, but may also result from the taking of an interrupt or trap (including the *WARN trap), or an interrupt return.

The non-sequential instruction fetch is initiated by placing an instruction-fetch request on the Address Bus. Once an external instruction fetch has been initiated, the processor generates prefetches for subsequent instructions based on the availability of IPB locations, either by transmitting subsequent addresses, or by issuing burst-mode instruction requests.

The addresses for prefetched instructions are computed by a word-length register called the Instruction Fetch Pointer (IFP), which is maintained by the Instruction Fetch Unit. The IFP latches the physical instruction-address obtained from the Memory Management Unit whenever a non-sequential instruction reference occurs.

For instruction prefetches, an 8-bit incrementer associated with the IFP updates bits 9-2 of the IFP to point to sequential instructions in the prefetch stream. The incrementer is

limited to 8 bits because it increments physical addresses, and thus cannot increment beyond any possible virtual-page boundaries (recall that the minimum virtual page size is 1 Kbyte). If the incrementer overflows, as indicated by a carry-out, prefetching is preempted. The prefetch stream is later re-established, however, as described below.

The physical address in the IFP is always the address of the most-recently-prefetched instruction, even though this address may not appear on the Address Bus for burst-mode fetches. If the burst is externally preempted, the IFP is used to re-establish the burst at the point of preemption.

Instruction Prefetch Buffer States

Four states are associated with each Instruction Prefetch Buffer location. The state-transition diagram for these states is shown in Figure 4-2.

Available: The IPB location is free to receive a new fetch. It contains no valid instruction, and is not due to receive any requested instruction.

Allocated: The IPB location has been scheduled to receive a requested instruction which has not yet been returned from the external instruction memory.




Figure 4-2. IPB State Transitions

Valid: The IPB location contains a valid instruction.

Error: The IPB location contains an instruction which was returned from the external memory with an *IERR indication.

If all internal conditions are such that an instruction fetch can occur, the IPB location given by bits 3-2 of the instruction address is set to the Allocated state, and the instruction is requested externally. Once this instruction is returned to the processor, it is stored in the IPB location. The location is set to the Valid or Error state (based on the *IERR input), unless the instruction is immediately sent to the decode stage, in which case the buffer is set to the Available state.

The instruction remains in the buffer until it is required for execution. When the instruction is required, it is issued to the decode stage, and the IPB location is set to the Available state. If the buffer were in the Error state, it is still set to the Available state, but an Instruction Access Exception trap occurs.

It is possible for all IPB locations to be in the Available or Valid states, but only one is allowed to be in the Allocated state at any given time. This restricts the number of unsatisfied instruction prefetches to one, reducing the amount of logic required to keep track of external fetches. It additionally restricts the number of apparent pipeline stages in the external prefetch mechanism to one stage (the other stages involved in the four-stage prefetch pipeline are the request stage and the processor's fetch and decode stages). Larger external prefetch pipelines may be implemented, but they are required to appear as single-stage pipelines; at most one instruction can be returned to the processor from the old instruction prefetch stream after a non-sequential fetch occurs.

When a non-sequential fetch occurs, all buffer locations are set to the Available state during the execute stage of the non-sequential fetch. All instruction requesting for the previous prefetch stream is terminated at this time. There is at most one instruction which will be returned to the processor after instruction fetches are terminated; this instruction is returned before any instruction associated with the new instruction stream is requested externally.

The Error state is provided only to handle errors reported via the *IERR input. However, there are many other situations in which the IPB does not contain a valid instruction. These situations arise because of errors, such as memory-management protection violations, and because instruction fetching is sometimes preempted, such as is the case when the IFP adder overflows. All of these cases are indicated by the fact that the IPB location is in the Available state when the instruction is required for execution (note that the location should, normally, at least be in the Allocated state when the instruction is required).

If the processor requires an instruction from an IPB location which is in the Available state, it initiates the fetch for the instruction using the current value of the Program Counter. This fetch resolves the exceptional condition. It either performs an address translation with the proper address, eliminating page-boundary-crossing problems, or re-creates an error condition, in which case a trap occurs.

4.2.2 BRANCH TARGET CACHE

The Branch Target Cache on the Am29000 allows fast access to instructions fetched non-sequentially. A branch instruction may execute in a single cycle, if the branch target is in the Branch Target Cache.

The target of a non-sequential fetch is in the Branch Target Cache if a similar fetch to the same target has occurred recently enough that it has neither been replaced by the target of another non-sequential fetch, nor invalidated by an INV or IRETINV instruction.

Branch Target Cache Organization

The organization of the Branch Target Cache is shown in Figure 4-3. To improve the ratio of the number of branch targets found in the cache, compared to the number of attempted cache accesses, two-way, set-associative mapping is used.

The Branch Target Cache is a 512-byte storage array divided into two sets each consisting of 64, 32-bit words (each instruction occupies a word). The sets are further divided into 16 blocks, numbered 0 to 15, which consist of 4 words each. Blocks in different sets with equivalent block-numbers are organized into a unit called a line.

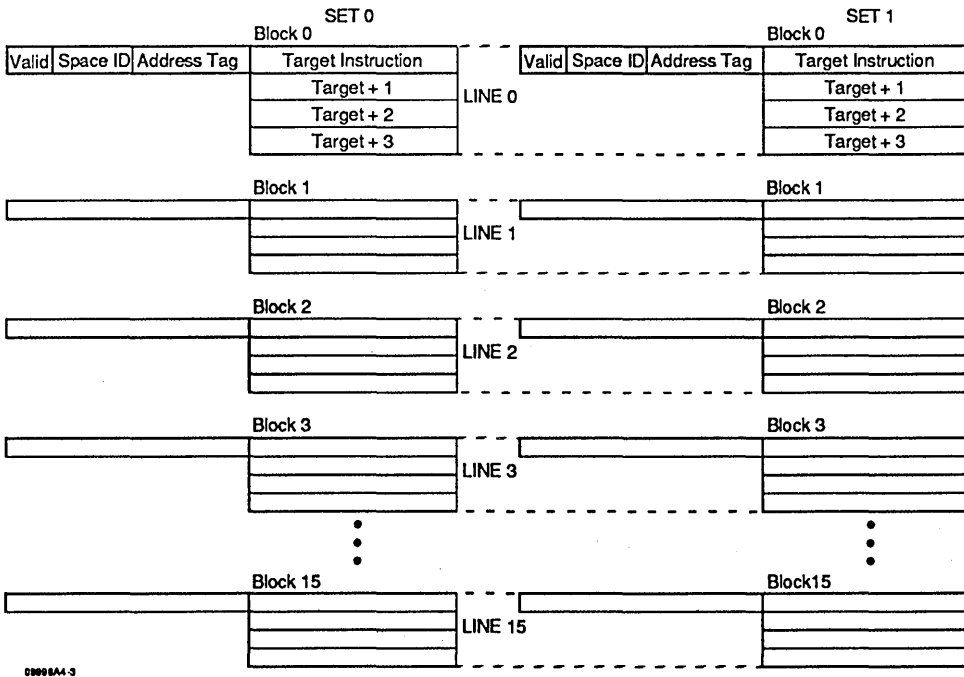


Figure 4-3. Branch Target Cache Organization

To eliminate fragmentation within the Branch Target Cache, each branch target entry is defined as a sequence of exactly four instructions, and is aligned on a cache-block boundary. A branch target sequence may occupy at most one block. This best utilizes the on-chip storage.

A 28-bit cache tag is associated with each four-word block. Of the 28 bits, 26 are derived from the address (possibly virtual) of the instructions in the block, and are called the Address Tag.

Note that the Address Tag is 26 bits in length, rather than 24 bits as might be implied by the organization of the Branch Target Cache. The reason for this is that branch target instruction-sequences are aligned on cache block boundaries. The result of this is that cache blocks are not aligned with respect to memory addresses. Thus, two additional bits are required in the Address Tag than would be required if cache locations were mapped one-to-one to memory locations.

Two additional bits in the cache tag, called the Space Identification field (Space ID), indicate the instruction memory from which the instructions were fetched (instruction/data or read-only memory) and the program mode under which the instructions were fetched (Supervisor or User). The encoding of these bits is described below:

<u>Space ID</u>	<u>Instruction Address Space</u>
0 0	User Instruction/Data Memory
0 1	User Instruction Read-Only Memory
1 0	Supervisor Instruction/Data Memory
1 1	Supervisor Instruction Read-Only Memory

A Valid bit associated with each cache word indicates that the word contains a valid instruction in the branch target sequence. There are thus four Valid bits for each cache block. Cache invalidation instructions make it possible to reset all Valid Bits in a single processor cycle. However, for the Invalidate instruction, the Valid bits are not reset until the next branch is executed.

Branch Target Cache Operation

It is possible to disable the operation of the Branch Target Cache via the Branch Target Cache Disable (CD) bit of the Configuration Register. If the CD bit is 1, all Branch Target Cache entries are made to appear invalid. If the CD bit is 0, there is no effect on Branch Target Cache entries. However, note that a change in the CD bit does not take effect until after the next non-sequential instruction fetch occurs.

When the Branch Target Cache is disabled, it continues to operate as described in this section. However, entries are made to appear invalid, even though they may be valid. If the Branch Target Cache is enabled after a period of being disabled, its contents reflect the most recent instruction execution, and it operates accordingly.

The Branch Target Cache lookup process is diagrammed in Figure 4-4. A given branch target sequence may be contained in one of two cache blocks, where these blocks are in the same line. The sequence is contained in the line whose number is given by bits 5-2 of the address of the first instruction of the sequence. A given branch target sequence is in a given cache block only if the following conditions are met:

- 1) Bits 31-6 of the address for the first instruction in the sequence match the corresponding bits in the Address Tag associated with the block.
- 2) The address of the first instruction in the block has a valid translation in the Memory Management Unit, if it is a virtual address.
- 3) The instruction address-space as indicated by the Current Processor Status Register matches the Space ID.
- 4) The CD bit of the Configuration Register was 0 for the previous non-sequential instruction fetch.

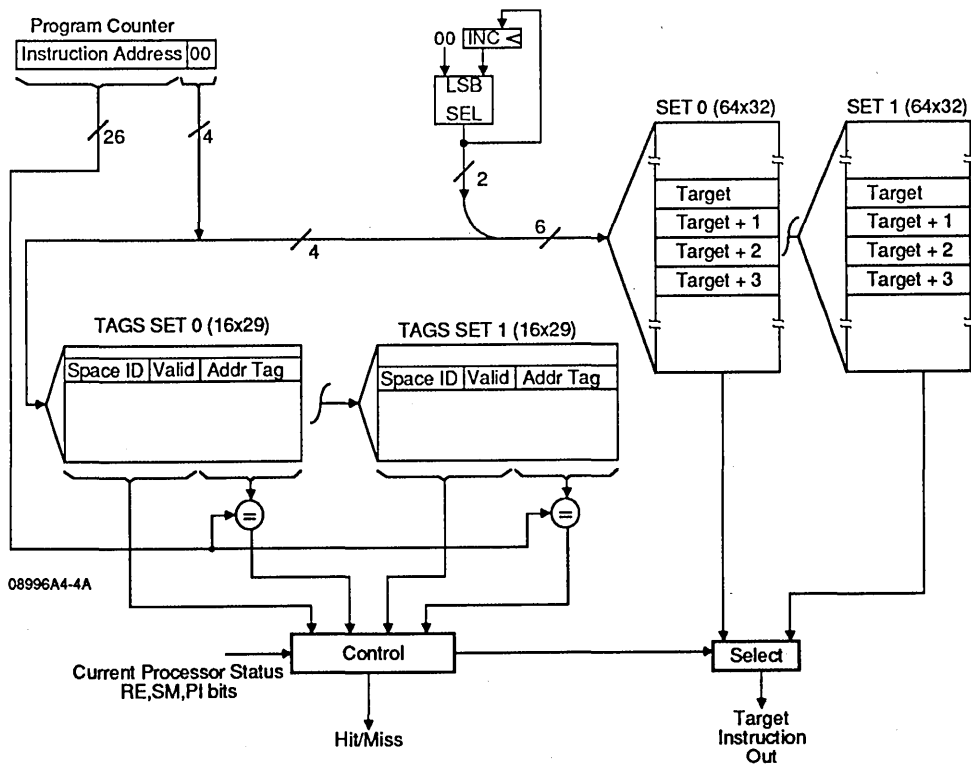


Figure 4-4. Branch Target Cache Lookup Process

In addition to the above requirements, the Valid bit must be 1 for any entry retrieved from the cache. Note that it is not required that all instructions in the sequence be present in the cache for the block to be considered valid.

Whenever a non-sequential fetch occurs (either for a branch instruction, an interrupt, or a trap), the address for the fetch is presented to the Branch Target Cache at the same time that the address is translated by the Memory Management Unit. If the target instruction for the non-sequential fetch is in the cache, it is presented for decoding in the next cycle. This instruction is always the first instruction of the cache block, and its address matches the cache tag. Subsequent instructions in the cache are presented for decoding as required in subsequent cycles. However, their addresses do not necessarily match the Address Tag.

Branch Target Cache Replacement

If, on a non-sequential fetch, the target instruction is not found in the Branch Target Cache, the address of the fetch selects a line to be used to store the instruction sequence of the new branch target. The replacement block within the line is selected at random, based on the processor clock. Random replacement has slightly better performance than least-recently-used replacement, and has a simpler implementation.

All Valid bits associated with the selected entry are reset, the Address Tag is set with the appropriate address bits of the first instruction in the sequence, and the Space ID bits are set according to the Current Processor Status Register.

Instructions from the new fetch stream are stored into the selected cache block as they are issued to the decode stage. The first instruction is stored into the first word of the block, the second instruction is stored into the second word, and so on up to a maximum of four instructions. The Valid bit for each word is set as the instruction is stored.

Special Cases of Branch Target Cache Entries

If a branch instruction appears as one of the first two instructions in a branch target sequence, the branch is executed before the Branch Target Cache block is filled. In this case, the cache block contains less than four valid instructions. The final valid instruction is the delay instruction of the branch.

When a block is only partially filled due to a branch within the block, the behavior of the cache during subsequent executions of the instructions in the block depends on the outcome of this branch.

If the branch is subsequently successful, then the instructions following the delay instruction of the branch are not needed, and the fact that they are not contained in the cache is irrelevant.

If the branch is subsequently unsuccessful, then the instructions following the delay instruction are required, and must be fetched externally. In this case, a required entry has a

Valid bit of 0. When the invalid entry is encountered, the Program Counter is used to create an external instruction fetch for the missing instruction. When the fetch completes, the instruction is stored in the cache location which was previously invalid, and the Valid bit for this entry is set.

Since an instruction sequence in the four-word cache block is not necessarily aligned on a four-word address-boundary, a virtual-page address-boundary may be crossed for the sequence in the cache. The processor does not prefetch instructions beyond this boundary, so the cache block is only partially filled in this case. If the processor requires instructions beyond the boundary, it creates a fetch for them as described above for the case of a branch instruction in the cache block.

When a fetch is created for a page-boundary crossing, this fetch is treated as a non-sequential fetch; a new cache block is allocated, and the first four instructions at the boundary are placed into the new cache block as they are returned by the instruction memory. Subsequent references to the original cache block also encounter an invalid instruction at the page boundary, and also create a special fetch for this instruction. However, since the instructions beyond this boundary are in the Branch Target Cache, subsequent boundary-crossings do not incur the instruction-fetch latency.

4.2.3 NON-SEQUENTIAL INSTRUCTION FETCHES

When a non-sequential instruction fetch occurs, the Memory Management Unit performs an address translation for target instruction, if address translation is enabled. If the address translation is valid, and the target of the fetch is not in the Branch Target Cache, an external instruction fetch is initiated. If there is a Translation Look-Aside Buffer (TLB) miss or memory-protection violation on this address, fetching is not initiated.

Instruction Fetch-Ahead

When a non-sequential fetch occurs, if the target of the fetch is found in the Branch Target Cache, the processor normally begins instruction fetching four instructions beyond the target. This behavior is termed fetch-ahead. The computation required to obtain the address for the fetch-ahead is performed in parallel with address translation, by a 6-bit adder called the Fetch-Ahead Adder (see Figure 4-1).

The Fetch-Ahead Adder is restricted to 6 bits so that the add cannot cause a page-boundary crossing (recall that the minimum virtual page size is 1 Kbyte and that all instructions are 32 bits in length). If the adder were larger, then the results of the add might affect the outcome of the address translation, and the add could not be performed in parallel with address translation.

Fetch-Ahead Disabling

When the target of a non-sequential fetch is in the Branch Target Cache, there are two cases

for which a fetch-ahead is not initiated.

The first case occurs when the Fetch-Ahead Adder overflows during the address computation for the fetch-ahead, as indicated by a carry out of the Fetch-Ahead Adder. Here, a page boundary may have been crossed, making the address translation—which is performed concurrently—invalid.

The second case occurs when the Branch Target Cache block containing the target instruction does not have Valid bits set for all entries within the block. In this case, the processor may have to fetch instructions for these entries, so it does not immediately initiate prefetching beyond the block.

If fetch-ahead is not initiated for an instruction which the processor eventually requires, this fetch is restarted on the cycle in which the missing instruction is required. The Program Counter is used in both of these cases, guaranteeing that the proper instruction address is used.

4.2.4 PROGRAM COUNTER UNIT

The Program Counter Unit, shown in Figure 4-5, forms and sequences instruction addresses for the Instruction Fetch Unit. It contains the Program Counter (PC), the Program-Counter Multiplexer (PC MUX), the Return Address Latch, and the Program-Counter Buffer (PC Buffer).

The PC forms addresses for sequential instructions executed by the processor. The master of the PC Register, PC L1, contains the address of the instruction being fetched in the Instruction Fetch Unit. The slave of the PC Register, PC L2, contains the next sequential address, which may be fetched by the Instruction Fetch Unit in the next cycle.

The Return Address Latch passes the address of the instruction following the delayed instruction of a call to the register file. This address is the return address of the call.

The PC Buffer stores the addresses of instructions in various stages of execution when an interrupt or trap is taken. The registers in this buffer—Program Counters 0, 1, and 2 (PC0, PC1, and PC2)—are normally updated from the PC as instructions flow through the processor pipeline.

When an interrupt or trap is taken, the Freeze (FZ) bit in the Current Processor Status is set, holding the quantities in the PC Buffer. When the FZ bit is set, PC0, PC1, and PC2 contain the addresses of the instructions in the decode, execute, and write-back stages of the pipeline, respectively.

Upon the execution of an interrupt return, the target instruction stream is restarted using the instruction addresses in PC0 and PC1. Two registers are required here because the processor implements delayed branches. An interrupt or trap may be taken when the processor is

executing the delay instruction of a branch and decoding the target of the branch. This discontinuous instruction sequence must be properly restarted upon an interrupt return. Restarting the instruction pipeline using two separate registers correctly handles this special case; in this case PC1 points to the delay instruction of the branch, and PC0 points to its target. PC2 does not participate in the interrupt return, but is included to report the addresses of instructions causing certain exceptions.

The PC is not defined as a special-purpose register. It cannot be modified or inspected by instructions. Instead, the interrupting and restarting of the pipeline is done by the PC Buffer registers PC0 and PC1.

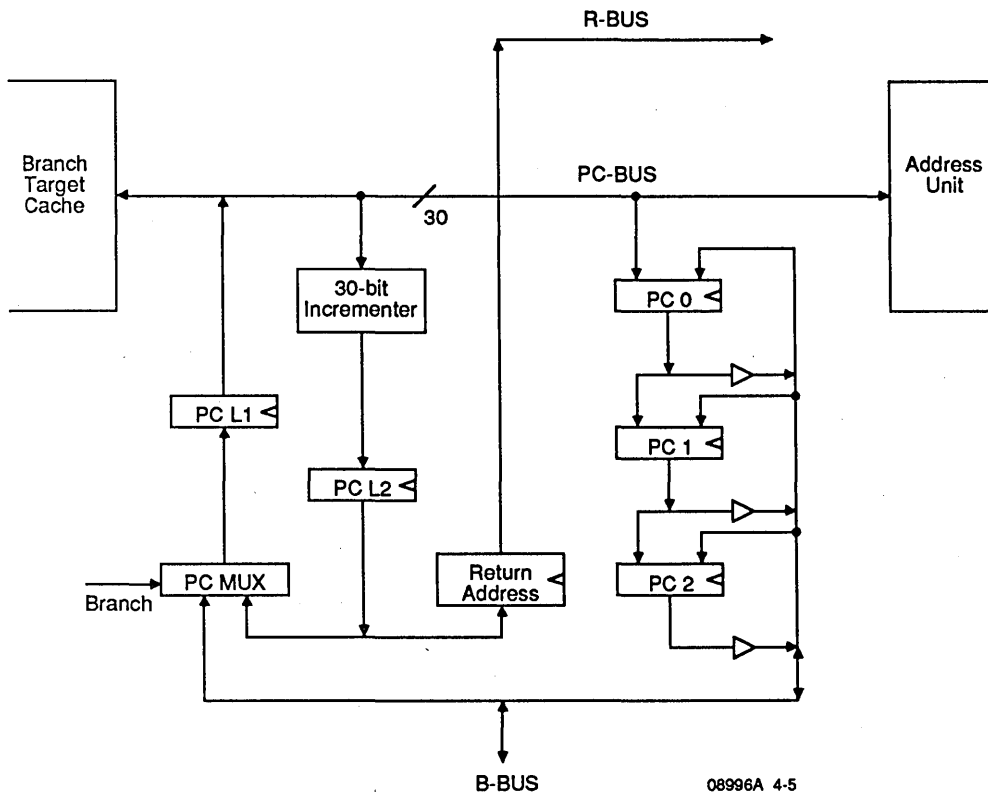


Figure 4-5. Program Counter Unit

4.3 EXECUTION UNIT

The Execution Unit performs most of the operations required for instruction execution. It incorporates the Register File, the Address Unit, the Arithmetic/Logic Unit, the Field Shift Unit, and the Prioritizer.

4.3.1 REGISTER FILE

The general-purpose registers are implemented by a triple-port, 192-location Register File. The Register File performs two read accesses and one write access in a single cycle. If a location is written and read in the same cycle, the data read is that written during the cycle.

The Register Address Generator, shown in Figure 4-6, computes register-numbers for operands, detects pipeline data-dependencies, and calculates register-number sequences for load-multiple and store-multiple operations.

Register Addressing

Register-numbers for instruction operands are computed during the decode stage. This computation is performed during the first half of a cycle, and the operands are read in the second half of a cycle. Three multiplexers select two source-operand register-numbers and a single destination register-number for any given instruction.

If the most-significant bit of a register-number is 0, the global registers are selected, and the register-number is used directly as a register address. If the most-significant bit of the register-number is 1, the local registers are selected, and the lower seven bits of the register-number are added to the Stack Pointer to form the desired local-register address.




Figure 4-6. Register File and Register Address Generator

The Stack Pointer is a hardware shadow-copy of bits 8-2 of Global Register 1, and is updated whenever Global Register 1 is written with the result of an Arithmetic or Logical instruction. Global Register 1 is implemented as a full, 32-bit register in the Register File; this register is distinct from the 192 locations which implement general-purpose registers.

If a register-number is zero (i.e. if Global Register 0 is specified as an operand), the Register Address Generator selects the content of an indirect pointer as the register-number. There are three indirect pointers, and each appears as a special-purpose register.

Pipeline Data-Dependencies

For the Register File, the pipeline delay in result write-back, compared to operand access, creates situations where a result from a previous operation may be required as an operand before it has been written into the register file. When one of these situations arises, a pipeline data-dependency is said to exist.

The register-numbers for the write-back of instruction results require two buffering registers, so that they are presented to the Register File during the write-back stage. In addition, the register-numbers for uncompleted load operations are held until the load completes (these register-numbers are held in the ETR Register shown in Figure 4-6).

Register read-address comparators detect pipeline data-dependencies, and activate multiplexers to forward data directly to the required functional unit, without waiting for the data to be written to the register file. The comparators activate the forwarding multiplexers if they detect one of the following situations:

- 1) One of the source register-numbers matches the destination register-number of the immediately-previous instruction.
- 2) One of the source register-numbers matches the target register-number (in the ETR) of an outstanding load.

In the first case listed above, the result of the execute stage is selected as an operand, instead of the output of the Register File port for which the forwarding condition is detected. In the second case, data from the channel is selected. The comparison may cause the processor to enter the Pipeline Hold mode if the load has not completed. However, data forwarding allows data from the Data Bus to be used immediately, in the cycle after it is returned on the Data Bus.

The content of the ETR is further compared to the register-numbers supplied to the write-back stage. If the target register for a load is written with the result of an overlapped instruction, the Not Needed (NN) bit in the Channel Control Register is set. If the comparators determine that the NN bit should be set, they also inhibit the write-back of load data on the completion of the load. The NN bit inhibits the restarting of the load operation if an exception occurs.

Load-Multiple and Store-Multiple Sequences

During load-multiple and store-multiple operations, sequential register-numbers are computed by an incrementer associated with the ETR/DTR pair shown in Figure 4-6. In the case of store-multiple, the register-numbers are supplied as read addresses to the Register File by the incrementer. The read addresses are latched by the DTR so that they may be incremented further. In the case of load-multiple, target register-numbers are held by the ETR as for any other load. However, the ETR is set with a sequence of incremented addresses in this case.

4.3.2 ADDRESS UNIT

The Address Unit, shown in Figure 4-7, computes addresses for branch target instructions, and load-multiple and store-multiple sequences. It also assembles instruction-immediate data and creates addresses for restarting terminated instruction prefetch streams.

The Address Unit consists of a 30-bit adder, the Decode PC Register, the ADRF Latch, and logic for formatting instruction-immediate data and generating the constants, zero and one. The Decode PC Register holds the address of the instruction in the decode stage of the pipeline.

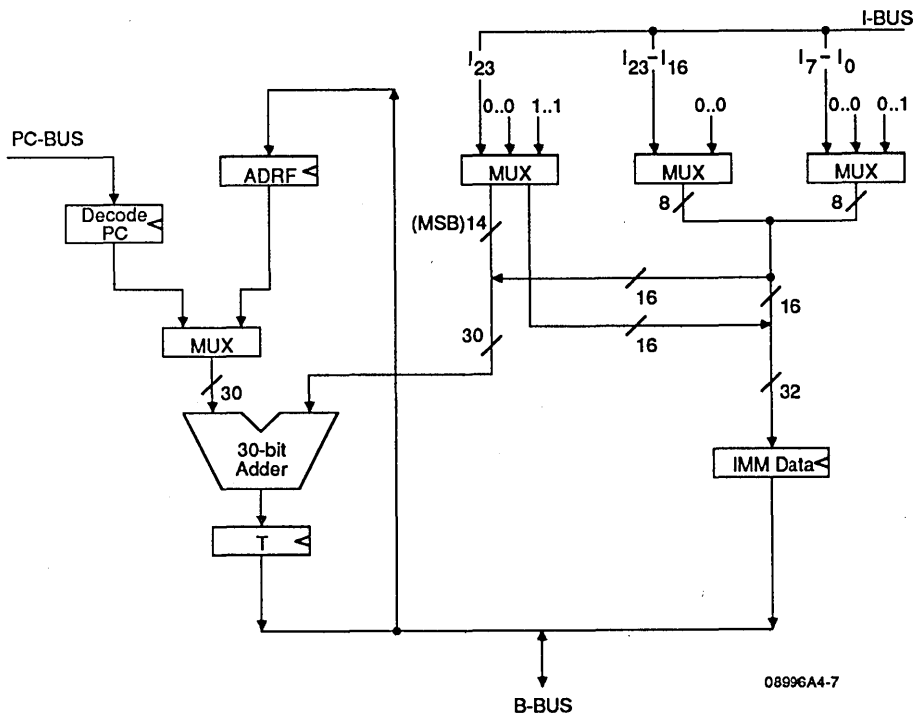


Figure 4-7. Address Unit

Branch Target Addresses

Branch target addresses are either fetched from the Register File or calculated by the Address Unit. The Address Unit calculates target addresses during the decode stage of branch instructions. These addresses are of two possible types:

- 1) **PC relative:** the current PC value is added to a sign-extended, 16-bit offset field from the branch instruction
- 2) **Absolute:** a zero-extended, 16-bit field of the branch instruction is used directly as an instruction address.

For each of the above types of addresses, the 16-bit instruction field is aligned on a word address-boundary (i.e. it is shifted left by two bits).

To calculate the branch target address, the Address Unit formats the 16-bit instruction field as required and presents it to the 30-bit adder. This adder adds the formatted field either to the contents of the Decode PC Register or to zero, as required for PC-relative and absolute addresses, respectively.

Load-Multiple and Store-Multiple Addresses

During the execution of Load Multiple and Store Multiple instructions, addresses for the access sequence are held in the ADRF Latch. An address in the ADRF Latch is updated, as required for an access in the sequence, by the 30-bit adder in the Address Unit. The formatting logic creates a constant offset of one for the update. The updated address is presented to the Memory Management Unit for translation and protection checking, and is placed into the ADRF Latch for further address computations.

For load-multiple and store-multiple operations performed using burst-mode accesses, the physical address for each access does not appear on the Address Bus, but the addresses are maintained in the processor so that they may be used to restart the burst-mode access upon preemption.

Special Instruction Fetches

As discussed in Section 4.2, the processor must create special instruction fetches when it encounters an invalid instruction in the middle of a Branch Target Cache block, or when it attempts to fetch an instruction from an Instruction Prefetch Buffer location which is in the Available state. The Address Unit routes the address for this fetch in a manner similar to the routing of a branch target address. It passes the contents of the Decode PC (containing the required instruction address) through the 30-bit adder, adding it to zero. This address is presented to the Memory Management Unit for translation, and is used in the Instruction Fetch Unit to complete the fetch.

4.3.3 ARITHMETIC/LOGIC UNIT

The Arithmetic/Logic Unit (ALU) performs 32-bit arithmetic and logical operations. The arithmetic operations consist of addition, subtraction, addition with carry-in, subtraction with carry-in, and primitives for multiplication and division. Instructions specify whether or not a trap is generated on signed or unsigned arithmetic overflow.

The A and B operands may be complemented independently in the ALU; complementors for data into the ALU are controlled by instructions. This allows subtraction and reverse subtraction to be formed from addition, and allows certain logical operations (e.g. XNOR) to be formed from other basic operations (e.g. XOR). The carry-in to the ALU can be 0, 1, or the value of the Carry bit in the ALU Status Register. The carry-out of the ALU is used in overflow detection, unsigned comparisons, multiplication, and division. It is stored in the ALU Status Register for multi-precision arithmetic.

The ALU also evaluates relational expressions with the operators equal, not equal, less-than, less-than-or-equal, greater-than, and greater-than-or-equal. Each comparison computes a Boolean corresponding to a relation between two integers, or creates a trap (possibly) based on this relation. The Boolean constants FALSE and TRUE are represented by a 0 and 1, respectively, in the most-significant bit of a word.

The relational operators may be applied to either signed or unsigned operands. For unsigned operands, these operators are implemented by recognizing that the ALU carry-out is the Boolean result of an unsigned comparison if the two numbers are subtracted and the carry-in is appropriately controlled. For comparison of signed numbers, the true sign of the result (i.e. the resulting sign exclusive-ORed with the overflow indication) gives the result of the compare.

The relational operators equal-to and not-equal-to are independent of the data type. These operators are implemented by a 32-bit equal-to-zero comparator.

The ALU also supports the 32-bit logical operations AND, OR, NAND, NOR, A AND-NOT B, XOR, and XNOR.

4.3.4 FIELD SHIFT UNIT

The Field Shift Unit contains a Funnel Shifter, logic for performing word extracts, and logic for performing byte and half-word extracts and inserts.

The Funnel Shifter is capable of performing N-bit shifts, where N is an integer between 0 and 31 inclusive given by a 5-bit shift count. The source of the shift count is specified by the shift instruction; the shift count is given either by a constant field in the shift instruction, bits 4-0 of a general-purpose register specified by the shift instruction, or by the 5-bit Funnel Shift Count field in the ALU Status Register.

Both arithmetic and logical shifts are supported, with the difference being the values stored into vacated bits: arithmetic shifts fill these bits with the sign bit of the operand, while logical shifts fill them with zero-bits. Arithmetic shifts are possible only for right shifts.

The Field Shift Unit operates on 32-bit words, 16-bit half-words, and 8-bit bytes. For byte operations, the position of a byte operand within a word is supplied by the 2-bit Byte Pointer (BP) field of the ALU Status Register. For half-word operations, the position of a half-word operand is given by the most-significant bit of the BP field: the least-significant bit is ignored. The processor supports either left-to-right or right-to-left byte and half-word ordering within a word.

4.3.5 PRIORITIZER

The prioritizer counts the number of leading zero bits in an operand. The count of the number of zero-bits up to the leading 1 is stored in the specified destination register. If the operand does not contain a 1, the value stored is 32.

4.4 MEMORY MANAGEMENT UNIT

The Memory Management Unit (MMU) performs all memory-management functions described in Section 3.6. Address translation is performed during the execute stage of any load, store, or branch instruction which requires address translation. Address translation is also performed whenever the processor requires an instruction which has not been prefetched; as discussed in Section 4.2, address translation is performed in this case to resolve certain exceptional events which occur during instruction prefetching.

Though the MMU is shared for instruction and data accesses, the processor pipeline is arranged so that there is no contention for the MMU. In general, this is the result of the instruction-set definition and the fact that instruction prefetch addresses are generated by the Instruction Fetch Pointer (see Section 4.2.1).

Instruction addresses are normally translated only when branches are executed. Since loads or stores cannot be executed at the same time, there is no contention for the MMU. If the Instruction Fetch Pointer overflows, the address translation is deferred until the Instruction Fetch Unit determines that the processor requires the associated instruction. Since instruction execution cannot occur at this time, the MMU cannot be required for the translation of a load or store address, and again there is no contention.

When the processor performs load-multiple and store-multiple operations, the MMU translates the address associated with every access. This allows the load-multiple and store-multiple address sequencing to be performed only in the virtual address space, rather than both the virtual and physical address-spaces. Since the execution of Load Multiple and Store Multiple instructions is not overlapped with the execution of other instructions, there is no penalty associated with using the MMU for every access.

The MMU performs address translation in a single cycle. If an address translation is valid, the results of the translation are placed on the Address Bus along with the instruction-access or data-access request. In many cases, the address appears on the Address Bus during the cycle immediately following address translation (it does not appear if the Address Bus is occupied with another access). This address appears regardless of the outcome of memory protection checking; this relaxes the timing constraints on protection checking, which can be performed only after address translation is complete. If a protection violation is detected, the processor activates the *BINV signal late in the first address cycle for the request.

4.5 PIPELINE HOLD MODE

The Pipeline Hold mode is activated whenever sequential processor operation cannot be guaranteed. When this mode is active, the pipeline stages do not advance, and most internal processor state is not modified. The processor places itself in the Pipeline Hold mode in the following situations:

- 1) The processor requires an instruction which has either not been fetched or not been returned by the external instruction memory.
- 2) The processor requires data from an in-progress load, and the access has not completed.
- 3) The processor attempts to execute a load or store instruction while another load or store is in progress.
- 4) The processor decodes an instruction which modifies any Translation Look-Aside Buffer entry or special-purpose register, and there is a load or store in progress. This is required for the serialization operation described in Section 3.7.
- 5) The processor is performing a sequence of load-multiple or store-multiple accesses. The Pipeline Hold mode in this case prevents further instruction execution until the completion of the load-multiple or store-multiple sequence.
- 6) The processor has taken an interrupt or trap, and the first instruction of the interrupt or trap handler has not entered the execute stage. The Pipeline Hold mode in this case prevents the processor pipeline from advancing until the interrupt or trap handler can begin execution.
- 7) The processor has executed an interrupt return, and the target instruction of the interrupt return has not entered the execute stage. The Pipeline Hold mode in this case prevents the processor pipeline from advancing until the interrupt return sequence is complete.

The Pipeline Hold mode is exited whenever the causing conditions no longer exist, or when the *WARN or *RESET input is asserted.

CHAPTER 5

SYSTEM INTERFACES

This chapter describes the attachment of the Am29000 to its hardware environment. It describes the channel, which allows the processor to communicate with external devices and memories. The Test/Development interface, provided for hardware development and testing, is also described. In addition, this chapter includes sections on external interrupts, traps, processor reset, clock generation, and master/slave checking.

In the signal descriptions of Section 5.1, certain outputs are described as being 3-state or bi-directional outputs. However, all outputs (except MSERR) may be placed in a high-impedance state by the Test mode. The 3-state and bi-directional terminology in this section is for those outputs (except SYSCLK) which are disabled when the processor grants the channel to another master.

5.1 SIGNAL DESCRIPTION

- A0-A31** **Address Bus (3-state output, synchronous)**
The Address Bus transfers the byte address for all accesses except burst-mode accesses. For burst-mode accesses, it transfers the address for the first access in the sequence.
- *BREQ** **Bus Request (input, synchronous)**
This input allows other masters to arbitrate for control of the processor channel.
- *BGRT** **Bus Grant (output, synchronous)**
This output signals to an external master that the processor is relinquishing control of the channel in response to *BREQ.
- *BINV** **Bus Invalid (output, synchronous)**
This output indicates that the Address Bus and related controls are invalid. It defines an idle cycle for the channel.
- R/*W** **Read/Write (3-state output, synchronous)**
This signal indicates whether data is being transferred from the processor to the system, or from the system to the processor.
- SUP/*US** **Supervisor/User Mode (3-state output, synchronous)**
This output indicates the program mode for an access.

***LOCK** **Lock (3-state output, synchronous)**
This output allows the implementation of various channel and device interlocks. It may be active only for the duration of an access, or active for an extended period of time under control of the Lock bit in the Current Processor Status.

The processor does not relinquish the channel (in response to *BREQ) when *LOCK is active.

MPGM0– **MMU Programmable (3-state output, synchronous)**
MPGM1 These outputs reflect the value of two PGM bits in the Translation Look-Aside Buffer entry associated with the access. If no address translation is performed, these signals are both Low.

***PEN** **Pipeline Enable (input, synchronous)**
This signal allows devices which can support pipelined accesses (i.e. which have input latches for the address and required controls) to signal that a second access may begin while the first completes.

I0–I31 **Instruction Bus (input, synchronous)**
The Instruction Bus transfers instructions to the processor.

***IREQ** **Instruction Request (3-state output, synchronous)**
This signal requests an instruction access. When it is active, the address for the access appears on the Address Bus.

IREQT **Instruction Request Type (3-state output, synchronous)**
This signal specifies the address-space of an instruction request, when *IREQ is active:

<u>IREQT</u>	<u>Meaning</u>
0	Instruction/data memory access
1	Instruction read-only memory access

***IRDY** **Instruction Ready (input, synchronous)**
This input indicates that a valid instruction is on the Instruction Bus. The processor ignores this signal if there is no pending instruction access.

***IERR** **Instruction Error (input, synchronous)**
This input indicates that an error occurred during the current instruction access. The processor ignores the content of the Instruction Bus, and an Instruction Access Exception trap occurs if the processor attempts to execute the invalid instruction. The processor ignores this signal if there is no pending instruction access.

- *IBREQ** **Instruction Burst Request (3-state output, synchronous)**
 This signal is used to establish a burst-mode instruction access and to request instruction transfers during a burst-mode instruction access. *IBREQ may be active even though the Address Bus is being used for a data access. This signal becomes valid late in the cycle, with respect to *IREQ.
- *IBACK** **Instruction Burst Acknowledge (input, synchronous)**
 This input is active whenever a burst-mode instruction access has been established. It may be active even though no instructions are currently being accessed.
- *PIA** **Pipelined Instruction Access (3-state output, synchronous)**
 If *IREQ is not active, this output indicates that an instruction access is pipelined with another, in-progress, instruction access. The indicated access cannot complete until the first access is complete. The completion of the first access is signalled by the assertion of *IREQ.
- D0–D31** **Data Bus (bi-directional, synchronous)**
 The Data Bus transfers data to and from the processor, for load and store operations.
- *DREQ** **Data Request (3-state output, synchronous)**
 This signal requests a data access. When it is active, the address for the access appears on the Address Bus.
- DREQT0–
DREQT1** **Data Request Type (3-state output, synchronous)**
 These signals specify the address-space of a data access, as follows (the value “x” is a don’t care):

DREQT1 DREQT0 Meaning

0	0	Instruction/data memory access
0	1	Input/output access
1	x	Coprocessor transfer

An interrupt/trap vector request is indicated as a data-memory read. If required, the system can identify the vector fetch by the STAT0–STAT2 outputs.

- *DRDY** **Data Ready (input, synchronous)**
 For loads, this input indicates that valid data is on the Data Bus. For stores, it indicates that the access is complete, and that data need no longer be driven on the Data Bus. The processor ignores this signal if there is no pending data access.

- *DERR** **Data Error (input, synchronous)**
 This input indicates that an error occurred during the current data access. For a load, the processor ignores the content of the Data Bus. For a store, the access is terminated. In either case, a Data Access Exception trap occurs. The processor ignores this signal if there is no pending data access.
- *DBREQ** **Data Burst Request (3-state output, synchronous)**
 This signal is used to establish a burst-mode data access and to request data transfers during a burst-mode data access. *DBREQ may be active even though the Address Bus is being used for an instruction access. This signal becomes valid late in the cycle, with respect to *DREQ.
- *DBACK** **Data Burst Acknowledge (input, synchronous)**
 This input is active whenever a burst-mode data access has been established. It may be active even though no data are currently being accessed.
- *PDA** **Pipelined Data Access (3-state output, synchronous)**
 If *DREQ is not active, this output indicates that a data access is pipelined with another, in-progress, data access. The indicated access cannot complete until the first access is complete. The completion of the first access is signalled by the assertion of *DREQ.
- OPT0–OPT2** **Option Control (3-state output, synchronous)**
 These outputs reflect the value of bits 18-16 of the load or store instruction which begins an access. Bit 18 of the instruction is reflected on OPT2, bit 17 on OPT1, and bit 16 on OPT0.

The standard definitions of these signals (based on DREQT) are as follows (the value “x” is a don’t care):

<u>DREQT1</u>	<u>DREQT0</u>	<u>OPT2</u>	<u>OPT1</u>	<u>OPT0</u>	<u>Meaning</u>
0	x	0	0	0	Word-length access
0	x	0	0	1	Byte access
0	x	0	1	0	Half-word access
0	x	0	1	1	24-bit access
0	0	1	0	0	Instruction ROM access (as data)
—all others—					reserved

If the interpretations above are irrelevant for a particular system, and compatibility issues are not important, other interpretations of the OPT0–OPT2 signals may be used.

- *CDA** **Coprocessor Data Accept (input, synchronous)**
 This signal allows the coprocessor to indicate the acceptance of operands or operation codes. For transfers to the coprocessor, the processor does not expect a *DRDY response; an active level on *CDA performs the function normally performed by *DRDY. *CDA may be active whenever the coprocessor is able to accept transfers.
- *WARN** **Warn (input, asynchronous, edge-sensitive)**
 A high-to-low transition on this input causes a non-maskable *WARN trap to occur. This trap bypasses the normal trap vector fetch sequence, and is useful in situations where the vector fetch may not work (e.g. when data memory is faulty).
- *INTR0-
*INTR3** **Interrupt Request (input, asynchronous)**
 These inputs generate prioritized interrupt requests. The interrupt caused by *INTR0 has the highest priority, and the interrupt caused by *INTR3 has the lowest priority. The interrupt requests are masked in prioritized order by the Interrupt Mask field in the Current Processor Status Register.
- *TRAP0-
TRAP1** **Trap Request (input, asynchronous)**
 These inputs generate prioritized trap requests. The trap caused by *TRAP0 has the highest priority. These trap requests are disabled by the DA bit of the Current Processor Status Register.
- STAT0-
STAT2** **CPU Status (output, synchronous)**
 These outputs indicate the state of the processor's execution stage on the previous cycle. They are encoded as follows:

<u>STAT2</u>	<u>STAT1</u>	<u>STAT0</u>	<u>Condition</u>
0	0	0	Halt or Step Modes
0	0	1	Pipeline Hold Mode
0	1	0	Load Test Instruction Mode
0	1	1	Wait Mode
1	0	0	Interrupt Return
1	0	1	Taking Interrupt or Trap
1	1	0	Non-sequential Instruction Fetch
1	1	1	Executing Mode

CNTL0– CPU Control (input, asynchronous)
CNTL1 These inputs control the processor mode:

<u>CNTL0</u>	<u>CNTL1</u>	<u>Mode</u>
0	0	Load Test Instruction
0	1	Step
1	0	Halt
1	1	Normal

***RESET** Reset (input, asynchronous)
 This input places the processor in the Reset mode.

***TEST** Test Mode (input, asynchronous)
 When this input is active, the processor is in Test mode. All outputs and bi-directional lines, except MSERR, are forced to the high-impedance state.

MSERR Master/Slave Error (output, synchronous)
 This output shows the result of the comparison of processor outputs with the signals provided internally to the off-chip drivers. If there is a difference for any enabled driver, this line is asserted.

SYCLK System Clock (bi-directional)
 This is either a clock output with a frequency which is half that of INCLK, or an input from an external clock generator at the processor's operating frequency.

INCLK Input Clock (input)
 When the processor generates the clock for the system, this is an oscillator input to the processor, at twice the processor's operating frequency. In systems where the clock is not generated by the processor, this signal must be tied High or Low, except in certain master/slave configurations as discussed in Section 5.8.

5.2 CHANNEL DESCRIPTION

The processor channel provides the bandwidth required for performance, while permitting the connection of many different types of devices. This section describes the channel, and methods of connecting devices and memories to the processor.

The channel is also used for transfers to and from the coprocessor. Coprocessor transfers are described in Section 6.2.

Timing diagrams for operations described in this chapter appear in Appendix A.

5.2.1 CHANNEL OVERVIEW

The channel consists of three, 32-bit synchronous buses with associated control and status signals: the Address Bus, Data Bus, and Instruction Bus. The Address Bus transfers addresses and control information to devices and memories. The Data Bus transfers data to and from devices and memories. The Instruction Bus transfers instructions to the processor from instruction memories. In addition, a set of signals allow control of the channel to be relinquished to an external master.

There are five logical groups of signals performing five distinct functions, as follows (since some signals perform more than one function, a signal may appear in more than one group):

- 1) Instruction Address Transfer and Instruction Access Requests: A0–A31, SUP/*US, MPGM0–MPGM1, *PEN, *IREQ, IREQT, *PIA, *BINV.
- 2) Instruction Transfer: I0–I31, *IBREQ, *IRDY, *IERR, *IBACK.
- 3) Data Address Transfer and Data Access Requests: A0–A31, R/*W, SUP/*US, *LOCK, MPGM0–MPGM1, *PEN, *DREQ, DREQT0–DREQT1, OPT0–OPT2, *PDA, *BINV.
- 4) Data Transfer: D0–D31, *DBREQ, *DRDY, *DERR, *DBACK, *CDA.
- 5) Arbitration: *BREQ, *BGRT, *BINV.

5.2.2 USER-DEFINED SIGNALS

Two types of user-defined outputs are available on the channel to control devices and memories directly in a system-dependent manner. Each of these outputs is valid simultaneously with—and for the same duration as—the address for an access.

The first set of user-defined signals, MPGM0–MPGM1, is determined by the PGM bits in the Translation Look-Aside Buffer entry used in translating the address for an access. If address translation is not performed, these outputs are both Low.

The second set of signals, OPT0–OPT2, are determined by bits 18-16 of the load or store instruction which initiates an access. These signals are valid only for data accesses, and have a pre-defined interpretation for coprocessor data-transfers.

Standard interpretations of OPT0–OPT2 are given in Section 5.1. Since the OPT0–OPT2 signals are determined by instructions, they may have an impact on application-software compatibility. If compatibility of application software among different systems is of concern, then system hardware should use the given definitions of OPT0–OPT2. In any event, a value of 000 on these signals should be defined to have no special effect on an access.

Note that the standard interpretations of OPT0–OPT2 apply only to accesses to

instruction/data memory and input/output. Other interpretations may be used for coprocessor transfers.

For interrupt and trap vector fetches, the MPGM0–MPGM1 and OPT0–OPT2 outputs are all Low.

5.2.3 INSTRUCTION ACCESSES

Instruction accesses occur to one of two address-spaces: instruction/data memory and instruction read-only memory (instruction ROM). The distinction between these address-spaces is made by the IREQT signal, which is in turn derived from the ROM Enable (RE) bit of the Current Processor Status Register. These are truly distinct address-spaces; each may be independently populated based on the needs of a particular system.

Instruction/data memory contains both instructions and data. Although the channel supports separate instruction and data memories, the Memory Management Unit does not. In certain systems, it may be required to access instructions via loads and stores, even though instructions may be contained in physically-separate memories. For example, this requirement might be imposed because of the need to load instructions into memory. Note also that the OPT signals may be used to allow the access of instructions in instruction ROM, using loads.

All processor instruction fetches are read-accesses, and the value of the R/*W signal is irrelevant for instruction fetches.

5.2.4 DATA ACCESSES

Data accesses occur to one of three address-spaces: instruction/data memory, input/output (I/O), and the coprocessor. The distinction between these spaces is made by the DREQT0–DREQT1 signals, which are in turn determined by the load or store instruction which initiates a data access. Each of these address-spaces is distinct from the others.

The protocol for data transfers to and from the coprocessor is slightly different than the protocol for instruction/data memory and I/O accesses. These transfers are described in Section 6.2.

Data accesses may occur either from a slave device or memory to the processor (for a load), or from the processor to a slave device or memory (for a store). The direction of transfer is determined by the R/*W signal. In the case of a load, the processor requires that data on the Data Bus be held valid only for a short time before the end of a cycle. In the case of a store, the processor drives the Data Bus as soon as it becomes available, and holds the data valid until the slave device or memory signals that the access is complete.

5.2.5 REPORTING ERRORS

The successful completion of an instruction access is indicated by an active level on the

*IRDY input, and the successful completion of a data access is indicated by an active level on the *DRDY input. If there are exceptional conditions for which an instruction or data access cannot complete successfully, the unsuccessful completion is indicated by an active level on the *IERR or *DERR input, as appropriate.

If the processor receives an *IERR or *DERR in response to an instruction or data access, it ignores the content of the Instruction or Data Bus and the value of *IRDY or *DRDY. An *IERR response causes an Instruction Access Exception trap, unless it is associated with an instruction which the processor does not ultimately execute (because of a non-sequential instruction fetch). A *DERR response always causes either a Data Access Exception trap or a Coprocessor Exception Trap.

The processor supports the restarting of unsuccessful accesses upon an interrupt return. In the case of an unsuccessful instruction access, the restart is performed by the Program Counter 0 and Program Counter 1 registers. In the case of an unsuccessful data access, the restart is performed by the Channel Address, Channel Data, and Channel Control registers. In any event, the control program must determine whether or not an access can and/or should be restarted.

The Instruction Access Exception and Data Access Exception traps cannot be masked. If one of these traps occurs within an interrupt or trap handler, the processor state may not be recoverable.

5.2.6 ACCESS PROTOCOLS

Figure 5-1 shows a control flow-chart for accesses performed by the Am29000. This control flow applies independently to both instruction and data accesses. Since the processor performs concurrent instruction and data accesses, these accesses may be at different points in the control flow at any given point in time.

Note that the items on the flow chart of Figure 5-1 do not represent actual states, and have no particular relationship to processor cycles. The flow chart only provides a high-level understanding of the control flow. Also, exceptions and error conditions are not shown.

The channel supports three protocols for accesses: simple, pipelined, and burst-mode. These are described in the following sections. The various protocols are defined to accommodate minimum-latency accesses as well as maximum-transfer-rate accesses. The protocols allow an access to complete in a single cycle, although they support accesses requiring arbitrary numbers of cycles. Address transfers for accesses may be independent of instruction or data transfers.



08996A5-1A

Figure 5-1. Channel Flow Chart

5.2.7 SIMPLE ACCESSES

For a simple access, the processor holds the address valid throughout the entire access. This protocol is used for single-cycle accesses, and for accesses to simple devices and memories.

On any cycle before the completion of the access, a simple access may be converted to a pipelined access (by the assertion of *PEN) or to a burst-mode access (by the assertion of *IBACK or *DBACK, if the processor is asserting *IBREQ or *DBREQ). Thus, the protocol for simple accesses may also be used during the initial cycles of pipelined and/or burst-mode accesses. This is advantageous, for example, in cases where the slave device or memory either requires the address to be held for multiple cycles at the beginning of the pipelined or burst-mode access, or cannot respond to the pipelined or burst-mode request within one cycle.

5.2.8 PIPELINED ACCESSES

A pipelined access is one which starts before an earlier, in-progress access completes. The in-progress access is called a primary access, and the second access is called a pipelined access. A pipelined access is of the same type as the primary access. For example, an instruction access which begins before the completion of a data access is not considered to be a pipelined access, whereas a second data access is.

The Am29000 allows only one pipelined access at any given time.

Tradeoffs

For accesses which require more than one cycle to complete, pipelined accesses perform better than simple accesses, because they allow the overlap of portions of two accesses. However, devices and memories which support pipelined accesses are somewhat more complex than devices and memories which support only simple accesses.

Support for pipelined operations is required for both the primary access and the pipelined access. The slave performing the primary access must contain some means for storing the address and other information about the access. The slave performing the pipelined access must be able to restrict its use of the Instruction Bus or Data Bus, and must be prepared to cancel the access (as explained below).

Pipelined Operation

Pipelined accesses are controlled by the signals *PEN, *PIA, and *PDA. Because of internal data-flow constraints, the Am29000 does not perform a pipelined store operation while a load is in progress. However, the protocol does not restrict pipelined operations. Other channel masters may perform a pipelined store during a load.

Except as noted above, the processor attempts to perform pipelining for every access; the

input *PEN indicates whether or not pipelining is supported for a given access. The *PEN input can be driven by individual devices, or can be tied active or inactive to enable or disable system-wide pipelined accesses. The processor ignores the value of *PEN unless it is performing an access.

The processor samples *PEN on every cycle during a primary access. If *PEN is active on any cycle, the processor ceases to drive the address and associated controls for the primary access in the next cycle. If the processor requires another access before the primary access completes, it drives the address and controls for the second access, asserting *PIA or *PDA to indicate that the second access is a pipelined access.

The output *IREQ or *DREQ, as appropriate, is not asserted for a pipelined access. Devices and memories which cannot support pipelined accesses should therefore ignore *PIA and/or *PDA, and base their operation upon *IREQ and/or *DREQ.

A device or memory which receives a request for a pipelined access may treat it as any other access, with one exception: the pipelined access cannot use the Instruction and Data buses nor the associated controls (e.g. *IRDY or *DRDY). In the case of a data read or instruction access, the results of the pipelined access cannot be driven on the appropriate bus. In the case of a data write, the data does not appear on the Data Bus. Any other operations for the access, such as address decoding, can occur.

When the primary access completes (as indicated by *IRDY or *DRDY), the pipelined access becomes a primary access. The processor indicates this by asserting *IREQ or *DREQ, depending on the type of access. The device or memory performing the pipelined access may complete the access as soon as *IREQ or *DREQ is asserted (possibly in the same cycle). When the access becomes a primary access, it controls the channel as any other primary access. For example, it may determine whether or not another pipelined access can be performed.

When the pipelined access becomes a primary access, the output *PIA or *PDA remains asserted for one cycle, to insure continuity of control within the slave device or memory. In the cycle after *IREQ or *DREQ is asserted, *PIA or *PDA is de-asserted, unless the processor initiates another pipelined access, in which case *PIA or *PDA remains asserted for the new access.

Cancellation of Pipelined Accesses

If the processor takes an interrupt or trap before a pipelined access becomes a primary access, the request for the pipelined access is removed from the channel. This may occur, for example, when *IERR or *DERR is signalled for the primary access.

If the pipelined access is removed from the channel, the slave device or memory does not receive an *IREQ or *DREQ for the pipelined access. Hence, the pipelined access does not become a primary access, and cannot complete. A pipelined access may be cancelled in this manner at any time before it becomes a primary access. Because of this, a pipelined access

should not change the state of a slave device or memory until the pipelined access becomes a primary access.

5.2.9 BURST-MODE ACCESSES

A burst-mode access allows multiple instructions or data words at sequential addresses to be accessed with a single address transfer. The number of accesses performed, and the timing of each access within the sequence, is controlled dynamically by the burst-mode protocol. Burst-mode accesses take advantage of sequential addressing patterns, and provide several benefits over simple and pipelined accesses:

- 1) Simultaneous instruction and data accesses. Burst-mode accesses reduce the utilization of the Address Bus. This is especially important for instruction accesses, which are normally sequential. Burst-mode instruction accesses eliminate most of the address transfers for instructions, allowing the Address Bus to be used for simultaneous data accesses.
- 2) Faster access times. By eliminating the address transfer cycle, burst-mode accesses allow addresses to be generated in a manner which improves access times.
- 3) Faster memory access modes. Many memories have special, high-bandwidth access modes (e.g. static-column page mode and nibble mode). These modes generally require a sequential addressing pattern, even though addresses may not be explicitly presented to the memory for all accesses. Burst-mode accesses allow the use of these access modes, without hardware to detect sequential addressing patterns.

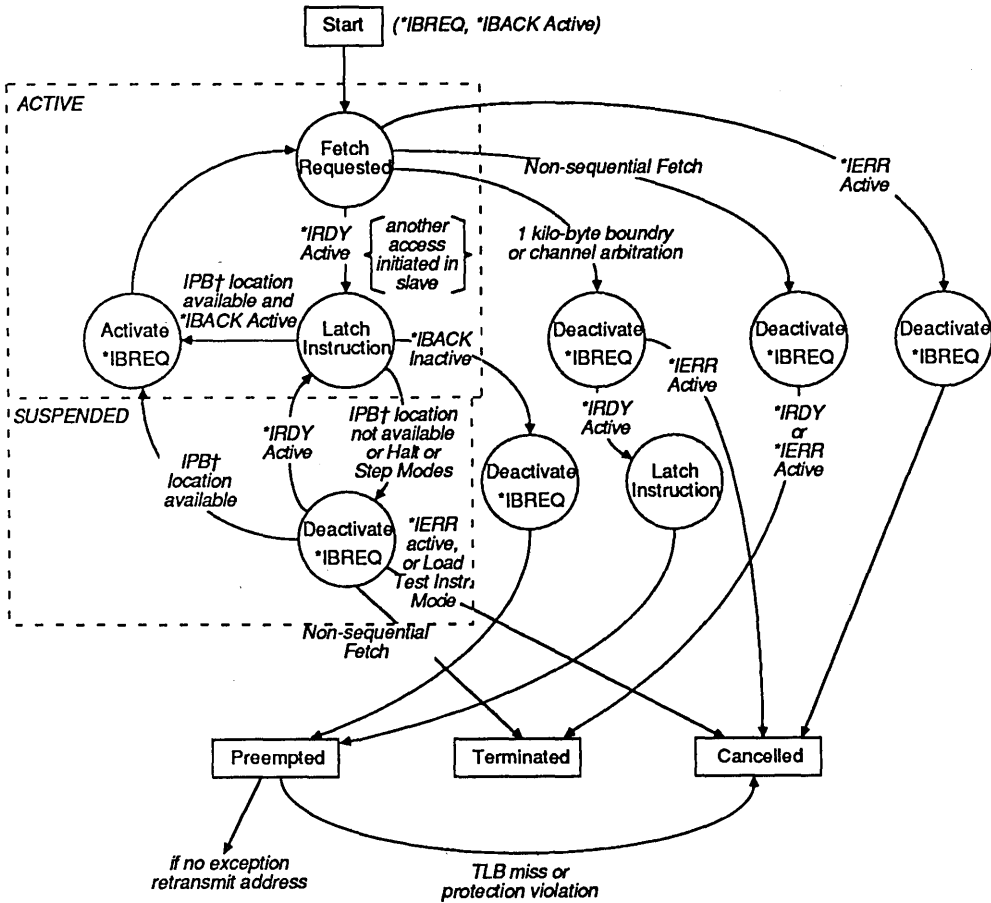
Burst-Mode Overview

The control-flow diagrams in Figures 5-2 and 5-3 illustrate the operation of the processor and an instruction memory during a burst-mode instruction access. The control-flow diagrams in Figures 5-4 and 5-5 illustrate the operation of the processor and a data memory or device during a burst-mode data access. Note that transitions on these diagrams do not necessarily correspond to processor cycles.

A burst-mode access is in one of the following operational conditions at any given time:

1. **Established:** The processor and slave device have successfully initiated the burst-mode access. A burst-mode access which has been established is either active or suspended. An established burst-mode access may become preempted, terminated, or cancelled.
2. **Active:** Instruction or data accesses and transfers are being performed as the result of the burst-mode access. An active burst-mode access may become suspended.

- 3. **Suspended:** No accesses or transfers are being performed as the result of the burst-mode access, but the burst-mode access remains established. Additional accesses and transfers may occur at some later time (i.e. the burst-mode access may become active) without the re-transmission of the address for the access.
- 4. **Preempted:** The burst-mode access can no longer continue because of some condition, but the burst-mode access can be re-established within a short amount of time.

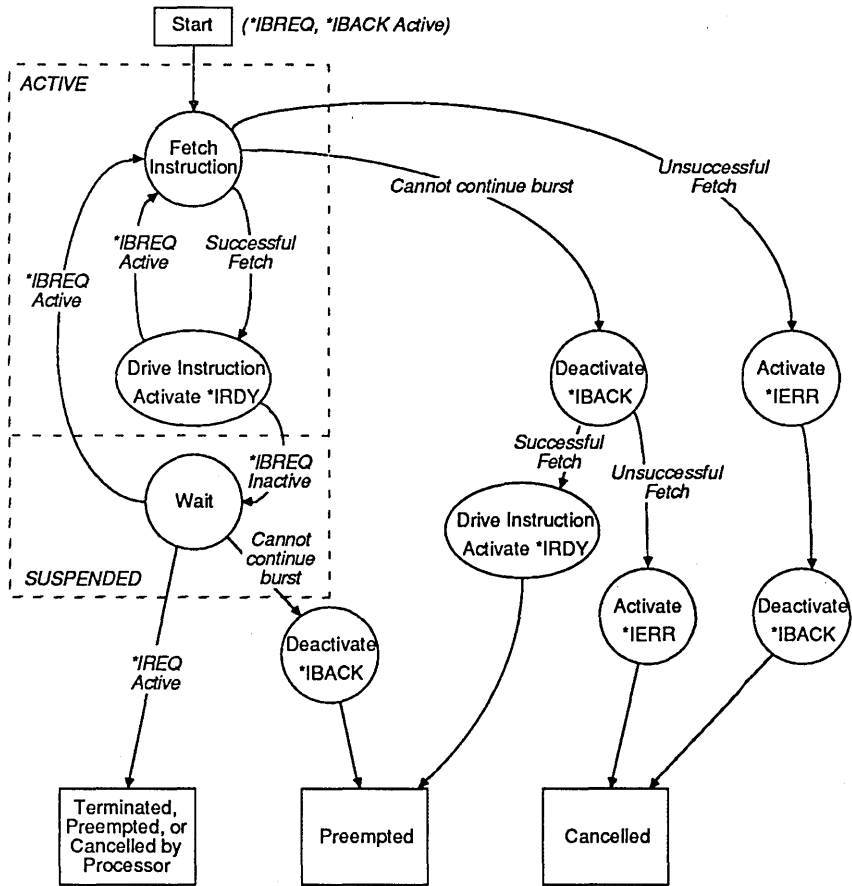


†IPB = Instruction Prefetch Buffer

08996A5-2A

Figure 5-2. Processor Burst-Mode Instruction Accesses: Control Flow

- 5. **Terminated:** All required accesses have been performed.
- 6. **Cancelled:** The burst-mode access can no longer continue because of some exceptional condition. The access may be re-established only after the exceptional condition has been corrected, if possible.

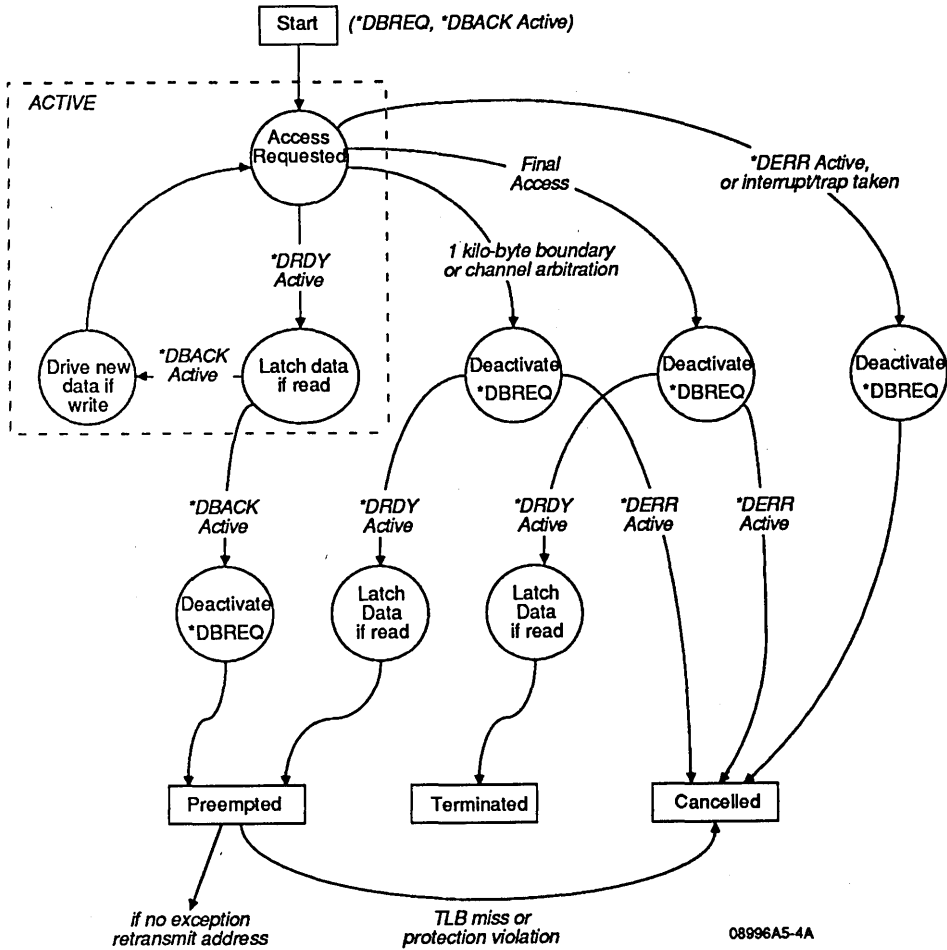


08996A5-3A

Note: A similar state transition may be used to support suspended burst-mode data accesses for a channel master other than the processor.

Figure 5-3. Slave Burst-Mode Instruction Accesses: Control Flow

Each of the above conditions, except for the terminated condition, is under the control of both the processor and slave device or memory. The terminated condition is determined by the processor, since only the processor can determine that all required accesses have been performed. The following sections discuss each of the above conditions with respect to the burst-mode protocol.



Note: The Am29000 does not suspend burst-mode data accesses.

Figure 5-4. Processor Burst-Mode Data Accesses: Control Flow

Establishing Burst-Mode Accesses

The Am29000 attempts to perform all instruction prefetches using burst-mode accesses. For data accesses, the processor attempts to perform load-multiple and store-multiple operations using burst-mode accesses. The processor indicates that it desires a burst-mode access by asserting **IBREQ* or **DBREQ* during the cycle in which the initial address is placed on the Address Bus (however, note that these signals become valid later in the cycle than the address).

The inputs **IBACK* and **DBACK* indicate that a requested burst-mode access is supported. The processor ignores the value of **IBACK* unless it is performing an instruction access, and it ignores the value of **DBACK* unless it is performing a load-multiple or store-multiple.

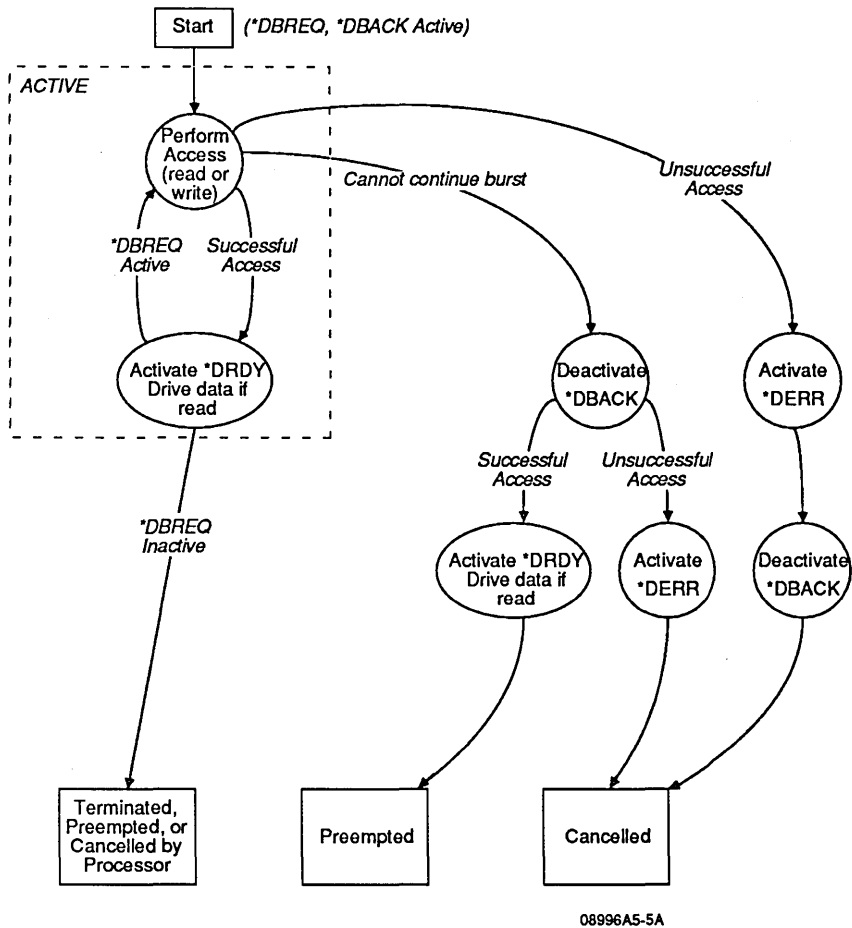


Figure 5-5. Slave Burst-Mode Data Accesses: Control Flow

When it desires a burst-mode access, the processor continues to drive *IBREQ or *DBREQ on every cycle for which the address is valid on the Address Bus. During this time, the device or memory involved in the access may assert *IBACK or *DBACK to indicate that it can perform the burst-mode access. If *IBACK or *DBACK (as appropriate) is asserted while the initial address appears on the Address Bus, the burst-mode access is established.

If the burst-mode access is not established on the first access, the processor attempts to establish a burst-mode access on each subsequent address transfer, as long as there are more accesses yet to be performed. During any subsequent access, the addressed device or memory may establish a burst-mode access by asserting *IBACK or *DBACK. If the burst-mode access is never established, the default behavior is to have the processor transmit an address for every access.

Active and Suspended Burst-Mode Accesses

After the burst-mode access is established, *IBREQ and *DBREQ are used during subsequent accesses to indicate that the processor requires at least one more access. If *IBREQ or *DBREQ is active at the end of the cycle in which an access successfully completes (i.e. when *IRDY or *DRDY is active), the processor requires another access. If the slave device or memory has not previously preempted the burst-mode access, and does not preempt or cancel (by asserting *IERR or *DERR) the burst-mode access in the cycle that the access completes, the additional access must be performed.

The execution rate of instructions is known only dynamically, so that, in certain situations, a burst-mode instruction access must be suspended. If *IBREQ is inactive during the cycle in which an instruction access completes, the burst-mode access is suspended (if it is neither preempted nor cancelled). The burst-mode access remains suspended unless the processor requests a new instruction access (in which case *IREQ is asserted), or unless the instruction memory preempts the burst-mode access.

A suspended burst-mode instruction access becomes active whenever the processor can accept more instructions. The processor activates the burst-mode access by asserting *IBREQ. If the instruction memory does not preempt the burst-mode access during this cycle, an instruction access must be performed.

When a suspended burst-mode instruction access is activated, the resulting instruction access is not permitted to complete in the cycle in which *IBREQ is asserted, but may complete in the next cycle. The reason for this restriction is that the burst-mode protocol is defined such that the combination of an active level on *IBREQ and *IRDY causes an instruction access (as previously discussed). If the instruction access completes immediately in the cycle that a suspended burst-mode access is activated, there is an ambiguity in the protocol: it is possible to interpret a single-cycle assertion of *IBREQ as a request for two instructions.

The above ambiguity is resolved by delaying the instruction access resulting from a re-activated burst-mode access for a cycle. Since this restriction applies only when the Instruction Prefetch Buffer is full and the instruction memory is capable of a very fast

access, the delayed instruction response has no performance impact.

The Am29000 does not suspend burst-mode data accesses, because the data transfers occur to and from general-purpose registers, which are always available. However, other channel masters may suspend burst-mode data accesses (during direct memory accesses, for example). The principles for suspending burst-mode accesses are the same as those for instruction accesses discussed above.

Processor Preemption, Termination, and Cancellation

The processor may preempt, terminate, or cancel a burst-mode access by de-asserting *IBREQ or *DBREQ, and asserting *IREQ or *DREQ at some later point. Normally, the processor receives one more instruction or data word after *IBREQ or *DBREQ is de-asserted. However, this access may complete in the same cycle that *IBREQ or *DBREQ is de-asserted. During the period after *IBREQ or *DBREQ is de-asserted and before *IREQ or *DREQ is asserted, the burst-mode access is in a suspended condition. The slave device or memory cannot—and need not—distinguish between preempted, terminated, and cancelled burst-mode accesses when they are caused by the processor.

The processor preempts a burst-mode access when an external channel master arbitrates for the channel, or when a burst-mode fetch crosses a potential virtual-page boundary. Since the minimum page size is 1 Kbyte, burst-mode instruction and data accesses are preempted whenever the address sequence crosses a 1 Kbyte address-boundary. The burst is re-established as soon as a new address translation is performed (if required). A new physical address is transmitted when the burst-mode access is re-established.

Note that the preemption resulting from page boundaries is advantageous for devices or memories which require counters to follow the burst-mode address sequence. Since all burst-mode accesses are word accesses, and the processor re-transmits an address at every 1 Kbyte address-boundary, an 8-bit counter in the slave device or memory is sufficient to follow the burst-mode address sequence. Additional address bits are simply latched.

The processor terminates a burst-mode access whenever all required instructions or data have been accessed. In the case of instruction accesses, the burst-mode access is terminated when a non-sequential fetch occurs. In the case of data accesses, the burst-mode access is terminated whenever the load-multiple or store-multiple sequence is complete (as determined by the Channel Control Register Load/Store Count Remaining field).

The processor cancels a burst-mode access when an interrupt or trap is taken. Note that a trap may be caused by the burst-mode access, for example when a Translation Look-Aside Buffer miss occurs on an address in the burst-mode sequence.

Cancelled burst-mode data accesses may be restarted at some (possibly much later) point in execution via the Channel Address, Channel Data, and Channel Control registers. In this case, the burst-mode access is restarted at the point at which it was cancelled, rather than at the beginning of the original address sequence.

Slave Preemption and Cancellation

The slave device or memory involved in a burst-mode access may preempt the access by de-asserting *IBACK or *DBACK. The processor samples *IBACK and *DBACK when *IRDY and *DRDY are active, so that *IBACK and *DBACK may be de-asserted as the last supported access is completed. However, *IBACK and *DBACK may also be de-asserted in any cycle before the access completes. If *IBACK or *DBACK is de-asserted when the processor is in a state where it expects an access, the access must be completed.

In general, the slave device or memory preempts the burst-mode access whenever it cannot support any further accesses in the burst-mode sequence. This normally occurs whenever an implementation-dependent address-boundary is encountered (for example, a cache-block boundary), but may occur for any reason. By preempting the burst-mode access, the slave receives a new request, with the address of the next instruction or data word required by the processor.

The slave device or memory may cancel a burst-mode access by asserting *IERR or *DERR in response to a requested access. The signals *IBACK or *DBACK need not be de-asserted at this time, but should be de-asserted in the next cycle.

Note that the *IERR and *DERR signals cause non-maskable traps, except in the case where *IERR is asserted for an instruction which the processor does not execute.

5.2.10 ARBITRATION

External masters can gain access to the Address, Data, and Instruction buses by asserting the *BREQ input. The processor completes any pending access, preempts any burst-mode access, and asserts the *BGRT output. At this time, the processor places all channel outputs associated with the Address, Data, and Instruction buses in the high-impedance state.

For the first cycle that *BGRT is asserted, the output *BINV is also asserted. If the external master cannot control the Address Bus and associated controls in the cycle that *BGRT is asserted, the active level on *BINV may be used to define an idle cycle for the channel (i.e. any spurious access requests are ignored). The *BINV signal is asserted only for a single cycle, so the external master must take control of the channel in the cycle after *BGRT is asserted.

While the *BREQ input remains asserted, the processor continues to assert *BGRT. The external master has control over the channel during this time.

To release the channel to the processor, the external master de-asserts *BREQ, but must continue to control the channel for the first cycle in which *BREQ is de-asserted. In the cycle after *BREQ is de-asserted, the processor asserts *BINV and de-asserts *BGRT; the external master should release control of the channel at this time. On the following cycle,

the processor de-asserts *BINV, and is able to use the channel. The processor re-establishes any burst-mode access preempted by arbitration.

The processor does not relinquish the channel when the *LOCK signal is active. This prevents external masters from interfering with exclusive accesses.

5.2.11 BUS SHARING—ELECTRICAL CONSIDERATIONS

When buses are shared among multiple masters and slaves, it is important to avoid situations where these devices are driving a bus at the same time. This may occur when more than one master or slave is allowed to drive a bus in the same cycle, because bus arbitration is incompletely or incorrectly performed. However, it also occurs when a master or slave releases a bus in the same cycle that another master or slave gains control, and the first master or slave is slow in disabling its bus drivers, compared to the point at which the second master or slave begins to drive the bus. The latter situation is called a bus collision in the following discussion.

In addition to the logical errors which can occur when multiple devices drive a bus simultaneously, such situations may cause bus drivers to carry large amounts of electrical current. This can have a significant impact on driver reliability and power dissipation. Since bus collisions usually occur for a small amount of time, they are of less concern, but may contribute to high-frequency electro-magnetic emissions.

The Am29000 channel is defined to prevent all situations where multiple drivers are driving a bus simultaneously. However, bus collisions may be allowed to occur, depending on the system design.

In the case of the Am29000 channel, arbitration for the channel prevents the processor from driving the Address and Data buses at the same time as another channel master. If there is more than one external master, the system design must include some means for insuring that only one external master gains control of the channel, and that no external master gains control of the channel at the same time as the processor.

When the processor relinquishes control of the channel to an external master, bus collisions may be prevented by not allowing the external master to drive any bus while *BINV is active. This insures that all processor outputs are disabled by the time the external master takes control of the channel. However, there is nothing in the channel protocol to prevent the external master from taking control as soon as *BGRT is asserted.

Slave devices and memories are prevented from simultaneously driving the Instruction Bus or Data Bus by allowing only the device or memory performing a primary access to drive the appropriate bus. When a pipelined access becomes a primary access, it may drive the Instruction or Data Bus immediately, so that there is a potential bus collision if the pipelined access is performed by a slave other than the slave performing the original primary access. This bus collision may be prevented by restricting all slaves to driving the Instruction and Data buses in the second half-cycle (using SYSCLK, for example). Since

the processor samples data only at the end of a cycle, this restriction does not affect performance.

When the processor performs a store immediately following a load, it drives the Data Bus for the store in the second cycle following the cycle in which the data for the load appears on the Data Bus. This provides a complete cycle for the slave involved in the load to disable its data drivers. The processor continues to drive the Data Bus until it receives a *DRDY or *DERR in response to the store; it ceases to drive the Data Bus in the cycle following the response.

5.2.12 CHANNEL BEHAVIOR FOR INTERRUPTS AND TRAPS

If an interrupt or trap is taken, any burst-mode accesses are cancelled. If a request for a pipelined access is on the Address Bus, this request is removed. Any other accesses are completed, and no new accesses are started, other than those required for the interrupt or trap.

When interrupt or trap processing is complete, any cancelled burst-mode accesses transactions are re-established, using the address of the access which was to be performed next when the interrupt or trap was taken. Uncompleted pipelined accesses are restarted, either by the interrupt return sequence in the case of an instruction access, or by restarting the initiating instruction in the case of a data access.

Note that the restarting of a pipelined access is not performed by the Channel Address, Channel Data, and Channel Control registers, since these registers may be required to restart the primary access. The instruction initiating the pipelined access is not allowed to complete until the primary access completes, so that the Program Counter 1 (PC1) Register contains the address of the initiating instruction when a pipelined access is cancelled. The address in PC1 can restart this instruction on interrupt return.

5.2.13 EFFECT OF THE *LOCK OUTPUT

The *LOCK output provides synchronization and exclusion of accesses in a multi-processor environment. *LOCK has no pre-defined effect for a system, other than the fact that the Am29000 does not grant the channel to an external master while *LOCK is active.

The *LOCK output is asserted for the address cycle of the Load-and-Lock and Store-and-Lock instructions, and is asserted for both the read and write accesses of a Load and Set instruction. *LOCK may also be active for an extended period of time, under control of the Lock bit in the Current Processor Status Register. (The latter capability is available only to Supervisor-mode programs.)

*LOCK may be defined to provide any level of resource locking for a particular system. For example, it may lock the channel, an individual device or memory, or a location within a device or memory.

When a resource is locked, it is available for access only by the processor with the appropriate access privilege. The mechanisms for restricting accesses, and the methods for reporting attempted violations of the restrictions, are system-dependent.

5.3 TEST/DEVELOPMENT INTERFACE

The Test/Development Interface consists of the inputs CNTL0–CNTL1 and *TEST, and the outputs STAT0–STAT2. The CNTL0–CNTL1 inputs provide control of processor operation, and the STAT0–STAT2 outputs provide information about processor operation for external monitoring.

An external hardware-development system uses CNTL0–CNTL1 and STAT0–STAT2 to control the processor for the purposes of processor and system debug.

A hardware tester uses the *TEST input to place all processor outputs in the high-impedance state. This allows the tester to check other system logic by driving processor outputs directly, without requiring that the processor be removed from the system.

5.3.1 PROCESSOR STATUS OUTPUTS

The STAT0–STAT2 outputs indicate certain information about processor modes, along with other information about processor operation. In addition to being used during normal processor operation, STAT0–STAT2 may be used to provide feedback of processor behavior when the processor is under the control of a hardware-development system.

The encoding of these signals is as follows:

<u>STAT2</u>	<u>STAT1</u>	<u>STAT0</u>	<u>Mode or Condition</u>
0	0	0	Halt or Step Modes
0	0	1	Pipeline Hold Mode
0	1	0	Load Test Instruction Mode
0	1	1	Wait Mode
1	0	0	Interrupt Return
1	0	1	Taking Interrupt or Trap
1	1	0	Non-sequential Instruction Fetch
1	1	1	Executing Mode

On any given cycle, the STAT0–STAT2 signals reflect the state of the processor's execute stage on the previous cycle. Where the conditions listed above are not mutually exclusive, the condition listed first is the one reflected on STAT0–STAT2.

For multiple-cycle operations (Load Multiple, Store Multiple, the taking of an interrupt or trap, and interrupt return), the first cycle is indicated appropriately, and additional cycles are indicated as Pipeline Hold cycles. The first cycle of the instructions Load Multiple, Store Multiple, Interrupt Return, and Interrupt Return and Invalidate are indicated as “Executing Mode” cycles. When an interrupt or trap is taken, the first cycle is indicated as “Taking Interrupt or Trap”.

A Low level on STAT2 indicates the processor is idle, and may be used as an indication of processor performance. Since most processor instructions execute in a single cycle, and since extra cycles spent executing multiple-cycle operations are counted as Pipeline Hold cycles, a count of the number of cycles, within a given time interval, that the processor is not idle (i.e. a count of the number of cycles for which STAT2 is High) is a close approximation to the number of instructions executed within that interval. This provides an approximation of the instruction execution rate. The only source of error in this approximation are the cycles in which the processor takes an interrupt or trap. If desired, this source of error can be eliminated by fully decoding the STAT0–STAT2 outputs.

The STAT2 output may also be used to implement processor timeouts for reliability. For example, a Low level on STAT2 may be used to start a hardware timeout counter, with a High level resetting and stopping the counter. If the counter exceeds a maximum expected count of idle cycles for a system, it is likely that an error has occurred. This error can be reported by the *WARN trap (see Section 3.5.6 and Section 5.6).

5.3.2 CPU CONTROL INPUTS

Certain processor operational modes are under the control of the CNTL0–CNTL1 inputs. These inputs have an effect on the processor mode as follows:

<u>CNTL0</u>	<u>CNTL1</u>	<u>Mode</u>
0	0	Load Test Instruction
0	1	Step
1	0	Halt
1	1	Normal

These inputs are asynchronous to the processor clock. In addition, changes on the CNTL0–CNTL1 inputs are restricted so that only CNTL0 or CNTL1, but not both, may change in any given processor cycle. The allowed transitions are shown in Figure 5-6.

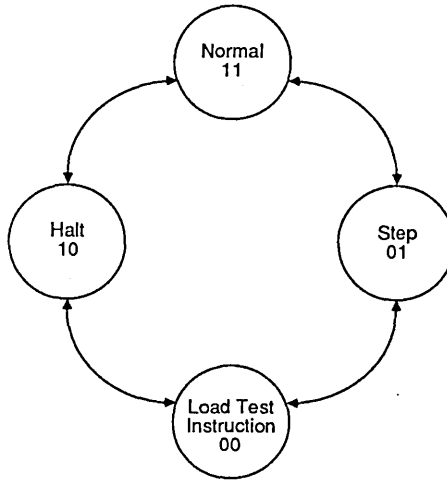


Figure 5-6. Valid Transitions on CNTL0-CNTL1 Inputs

The restriction on CNTL0–CNTL1 transitions allows these inputs to be driven directly by an external hardware-development system or tester, without any intervening logic. A violation of this restriction causes unpredictable processor operation. Proper operation is insured by making only single-input changes on CNTL0–CNTL1, and by restricting the interval between all changes to be greater than a processor cycle.

Note that, because of the restriction described above, it is not possible to transition directly between all possible modes which are controlled by these inputs. For example, the processor cannot go from the Load Test Instruction mode to Normal operation without first entering the Halt or Step modes.

5.3.3 HARDWARE DEVELOPMENT

The Halt, Step, and Load Test Instruction modes of operation are defined to support the debug of the processor system (both hardware and software) in a development environment. This section describes the use of these modes during debug, and describes the corresponding activity on the CNTL0–CNTL1 and STAT0–STAT2 lines.

Halt Mode

The Halt mode allows a hardware-development system to stop processor operation while preserving its internal state. The Halt mode is defined so that normal operation may resume from the point at which the processor enters the Halt mode. All external accesses are completed before the Halt mode is entered, so a minimum amount of system logic is required to support the Halt mode.

The Halt mode is invoked by the application of a value of 10 to the CNTL0–CNTL1 inputs. The processor enters the Halt mode within 2 or 3 cycles after the CNTL0–CNTL1 inputs are changed (depending on synchronization time), except that it first completes any external data access in progress.

The Halt mode is also entered as the result of the execution of a HALT instruction. When a HALT instruction is executed, the processor enters the Halt mode on the next cycle, except that it completes any external data accesses in progress. In this case, the processor remains in the Halt mode even though the CNTL0–CNTL1 inputs are 11. However, the processor cannot exit the Halt mode except as the result of the CNTL0–CNTL1 or *RESET inputs. If the instruction following the Halt instruction has an exception (e.g., TLB Miss), the trap associated with the exception is taken before the processor enters the Halt mode.

The STAT0–STAT2 lines have a value of 000 whenever the processor is in the Halt mode; these outputs can be used as a verification that the processor is in the Halt mode.

If a burst-mode instruction access is established before the processor enters the Halt mode, it remains established when the processor enters the Halt mode, but is suspended.

While in the Halt mode, the processor does not execute instructions, and performs no external accesses. The Timer Facility does not operate (i.e. the Timer Counter Register does not change).

The Halt mode is exited whenever the Reset mode is entered, or the CNTL0–CNTL1 lines place the processor into another mode. The only valid transitions on the CNTL0–CNTL1 lines from the value of 10 are to the value 00, which places the processor into the Load Test Instruction mode, and to the value 11, which causes the processor to resume normal execution.

Step Mode

The Step mode causes the Am29000 to execute at a rate determined by a hardware-development system, allowing the hardware-development system to control and monitor processor operation independent of speed mismatches. The Step mode is defined so that normal operation may resume after stepping is complete. Since all external accesses are completed during any step, a minimum amount of system logic is required to support the slower rate of execution.

The Step mode is invoked by the application of a value of 01 to the CNTL0–CNTL1 inputs. The processor enters the Step mode within 2 or 3 cycles after the CNTL0–CNTL1 inputs are changed (depending on synchronization time), except that it first completes any external data access in progress.

The STAT0–STAT2 lines have a value of 000 whenever the processor is in the Step mode; these outputs can be used as a verification that the processor is in the Step mode.

If a burst-mode instruction access is established before the processor enters the Step mode, it remains established when the processor enters the Step mode, but is suspended.

While in the Step mode, the processor does not execute instructions, and performs no external accesses. The Timer Facility does not operate (i.e. the Timer Counter Register does not change) while the processor is in the Step mode.

The Step mode is identical to the Halt mode in every respect except one. This difference is apparent on the transition of the CNTL0–CNTL1 lines from the value 01 (Step mode) to the value 11 (Normal). On this transition, the processor steps. That is, the processor state advances by one pipeline stage, and it completes any external access which is initiated by this state change.

If the processor immediately enters the Pipeline Hold mode on a step, the step may require multiple cycles to execute, since the processor pipeline cannot advance while the processor is in the Pipeline Hold mode. The STAT0–STAT2 lines reflect the state of the processor for every cycle of the step; STAT2 is High for one cycle, and only one cycle, before the step completes.

The Timer Counter decrements by one for every cycle of the step; if the Timer Counter decrements to zero, the usual Timer-Facility actions are performed, and a Timer interrupt may occur.

After the step is performed, the processor re-enters the Step mode, and remains in the Step mode even though the CNTL0–CNTL1 inputs have the value 11 (this prevents the need for a time-critical transition on the CNTL0–CNTL1 inputs). The processor remains in this condition until the CNTL0–CNTL1 inputs transition to 10 or 01 (or *RESET is asserted). The transition to 10 causes the processor to enter the Halt mode, and is used to clear the Step mode. The transition to 01 causes the processor to remain in the Step mode, so that it may perform additional steps.

Load Test Instruction Mode

The processor incorporates an Instruction Register (IR) which holds instructions while they are decoded. In the Load Test Instruction mode, the IR is enabled to receive the content of the Instruction Bus, regardless of the state of the processor's Instruction Fetch Unit. This allows a hardware-development system to directly provide instructions for execution, thereby providing means for the hardware-development system to examine and modify the internal state of the processor without altering the processor's instruction stream.

The hardware-development system can place an instruction in the IR by first placing 00 on CNTL0–CNTL1. The processor enters the Load Test Instruction mode within 2 or 3 cycles after the CNTL0–CNTL1 inputs are changed (depending on synchronization time), except that it first preempts any established burst-mode instruction access. The Load Test Instruction mode can be entered only from the Halt or Step modes. Note that the

burst-mode instruction access which is preempted here was previously suspended for the Halt or Step modes.

The STAT0–STAT2 lines have a value of 010 while the processor is in the Load Test Instruction mode; this may be used as a verification that the processor is loading the IR.

While the processor is in the Load Test Instruction mode, the IR is continually storing the value on the Instruction Bus; any change in the value on this bus is reflected in the IR on the next cycle. The hardware-development system can place a desired instruction into the IR by driving this instruction on the Instruction Bus. The value of *IRDY and *IERR are irrelevant.

The processor exits the Load Test Instruction mode in the second cycle following a change on the CNTL0–CNTL1 inputs. The only valid change here is either to the Halt mode (CNTL0–CNTL1 = 10) or the Step mode (CNTL0–CNTL1 = 01).

When the Load Test Instruction mode is exited, the most recent value stored into the IR is held. If the processor is placed in the Step mode, the IR is marked as having valid content, enabling the processor to decode and execute the instruction. If the processor is placed in the Halt mode, it ignores any instruction placed in the IR by the Load Test Instruction mode, and reverts to its normal instruction-fetch mechanism.

Once the IR has been set by the Load Test Instruction mode, the instruction in the IR may be executed via the Step mode as discussed in the previous sub-section. A single step is sufficient to cause the execution of this instruction. However, because of pipelining, multiple steps may be required before the instruction completes execution. If more than one step is performed, the processor executes the instruction in the IR on every step. If it is desired to step an instruction to completion without repeated execution, a NO-OP may be set into the IR (using the Load Test Instruction mode) after the first step.

The Load Test Instruction mode may be used to cause the execution of any processor instruction, except Load Multiple, Store Multiple, Interrupt Return, and Interrupt Return and Invalidate. This allows inspection and modification of processor state. For example, load and store instructions may be used to alter and inspect the contents of general-purpose registers; in this case, the hardware-development system supplies and reads register values on the Data Bus. Note that the external address for reading and writing registers in this manner should not be allowed to interfere with other system addresses.

The contents of the Program Counter 0, Program Counter 1, Program Counter 2, Channel Address, Channel Data, Channel Control, and ALU Status registers are not updated while instructions are executed via the Load Test Instruction mode, except explicitly by Move To Special Register instructions. Instructions executed using the Load Test Instruction mode may access protected processor state even though the processor is in the User mode.

Instructions executed via the Load Test Instruction mode may be used to access an external device or memory. Recall that the processor completes any data access before completing a

step. This allows the processor to access devices and memories on behalf of the hardware-development system, and simplifies the timing constraints on the hardware-development system.

During processor execution via the Load Test Instruction mode, the processor retains the information required to resume normal operation. If any processor state is modified by the hardware-development system, this state must be properly restored for normal operation to resume properly.

Once all instructions have been executed via the Load Test Instruction mode, the Halt mode (CNTL0–CNTL1 = 10) prepares the processor to resume normal operation. When the CNTL0–CNTL1 inputs transition to 11, the processor resumes normal operation, using a sequence very similar to that used for an interrupt return.

Summary of Development-System Operation

When the capabilities provided by the Halt, Step, and Load Test Instruction Register modes are combined, an extremely flexible test and development interface results. The following is an example sequence performed by a hardware-development system during debug:

- 1) Halt the processor either by a HALT instruction or by a 10 on the CNTL0–CNTL1 inputs. The HALT instruction may be used as a primitive in the implementation of a general instruction-breakpoint capability.
- 2) Load the IR with an instruction to inspect or alter the processor state. The hardware-development system should wait for the value 010 on STAT0–STAT2 (Load Test Instruction mode) before driving the Instruction Bus. After the IR is loaded, the hardware-development system sets CNTL0–CNTL1 to 01 (Step mode).
- 3) Step the processor by a transition of CNTL0–CNTL1 from 01 to 11 and back to 01. Data may be supplied on the Data Bus during one of the steps.
- 4) Repeat steps 2 and 3 as desired.
- 5) After the final step, enter the Halt mode by placing 10, instead of 01, on CNTL0–CNTL1.
- 6) Resume normal execution by placing 11 on CNTL0–CNTL1.

5.3.4 HARDWARE TESTING

The Test mode in the Am29000 allows processor outputs to be driven directly for testing or diagnostic purposes. The Test mode places all processor outputs (except MSERR) into the high-impedance state, so that they do not interfere electrically with externally-supplied signals. In all other respects, processor operation is unchanged.

The Test mode is invoked by an active level on the *TEST input, regardless of the processor's operational mode (for example, the Test mode is not affected by the Halt mode). The disabling of processor outputs is performed combinatorially. It occurs even though no clocks are applied to the processor.

For some outputs, the transition to the high-impedance state which results from the Test mode may occur at a much slower rate than applies during normal system operation (for example, when the processor relinquishes the channel to another master). For this reason, the Test mode may not be appropriate for special, user-defined purposes.

Note that SYSCLK is also placed in the high-impedance state by the Test mode. This allows the testing of external clock-distribution circuits, but care must be taken to insure that a high-impedance SYSCLK output does not have an adverse effect on the system. Furthermore, if SYSCLK is disabled, and a signal is not externally supplied, processor state may be lost.

5.4 EXTERNAL INTERRUPTS AND TRAPS

An external device causes an interrupt by asserting one of the *INTRO-*INTR3 inputs, and causes a trap by asserting one of the *TRAP0-*TRAP1 inputs. Transitions on each of these inputs may be asynchronous to the processor clock; they are protected against metastable states. For this reason, an assertion of one of these inputs which meets the proper set-up-time criteria does not cause the corresponding interrupt or trap until the second following cycle.

The *INTRO-*INTR3 inputs are prioritized with respect to each other and with respect to the processor. For resolving conflicts between these inputs which may arise, the inputs are prioritized in order, so that the interrupt caused by *INTRO has the highest priority, and the interrupt caused by *INTR3 has the lowest priority.

The interrupts caused by *INTRO-*INTR3 may be masked by the Disable Interrupts (DI) or Disable All Interrupts and Traps (DA) bits of the Current Processor Status Register. In addition, the Interrupt Mask (IM) field of the Current Processor Status Register sets the priority of the processor with respect to these inputs. The IM field enables the *INTRO-*INTR3 inputs as follows:

<u>IM Value</u>	<u>Result</u>
00	*INTRO enabled
01	*INTRO-*INTR1 enabled
10	*INTRO-*INTR2 enabled
11	*INTRO-*INTR3 enabled

Note that the interrupt caused by the *INTRO input cannot be disabled by the IM field.

If one of the *INTR0–*INTR3 inputs is active, and the resulting interrupt is disabled by the DA bit, DI bit, or IM field, the Interrupt Pending (IP) bit of the Current Processor Status Register is set. The IP bit is reset if the interrupt is enabled, or if all disabled external interrupts are de-asserted.

The *TRAP0–*TRAP1 inputs are prioritized with respect to each other, so that the trap caused by *TRAP0 has priority over the trap caused by *TRAP1 when a conflict occurs. Both *TRAP0 and *TRAP1 have priority over the *INTR0–*INTR3 inputs. The *TRAP0–*TRAP1 inputs cannot be selectively disabled. Both traps, however, can be disabled by the DA bit in the Current Processor Status Register.

The *INTR0–*INTR3 and *TRAP0–*TRAP1 inputs are level-sensitive. Once asserted, they must be held active until the corresponding interrupt or trap is acknowledged by the interrupt or trap handler (this acknowledgement is system-dependent, since there is no interrupt-acknowledge mechanism defined for the processor).

If any of these inputs is asserted, then de-asserted before it is acknowledged, it is not possible to predict (unless the interrupt or trap is masked) whether or not the processor has taken the corresponding interrupt or trap. During interrupt and trap processing, the vector number is determined in part by which of the *INTR0–*INTR3 and *TRAP0–*TRAP1 inputs is active. If the input causing an interrupt or trap is de-asserted before the vector number is determined, the vector number is unpredictable, with the result that processor operation is also unpredictable.

5.5 PROCESSOR RESET

When power is first applied to the processor, it is in an indeterminate state, and must be placed in a known state. Also, under certain circumstances, it may be necessary to place the processor in a defined state. This is accomplished by the RESET mode, which places the processor into a pre-defined state (see Section 3.8).

The Reset mode is invoked by asserting the *RESET input, and can be entered only if the SYSCLK pin is operating normally, whether or not the SYSCLK pin is being driven by the processor (see Section 5.7). The Reset mode is entered within 4 processor cycles after *RESET is asserted.

The Reset mode can be entered from any other processor mode (for example, the Reset mode can be entered from the Halt mode). If the *RESET input is asserted at the time that power is first applied to the processor, the processor enters the Reset mode only after four cycles have occurred on the SYSCLK pin.

The Reset mode is exited when the *RESET input is de-asserted. Either 3 or 4 cycles after *RESET is de-asserted (depending on internal synchronization time), the processor performs an initial instruction access on the channel. The initial instruction access is directed to

address 0 in the instruction read-only memory (instruction ROM). If instruction ROM is not implemented in a particular system, another device or memory must respond to this instruction fetch.

If the CNTL0–CNTL1 inputs are 10 or 01 when the initial instruction fetch completes, the processor enters the Halt or Step mode, respectively. Before completion of the initial instruction fetch, the CNTL0–CNTL1 inputs are irrelevant, except that the Load Test Instruction mode cannot be directly entered from the Reset mode. If the CNTL0–CNTL1 inputs are 00 immediately after *RESET is de-asserted, the effect on processor operation is unpredictable. If the CNTL0–CNTL1 inputs are 11, the processor enters the Executing mode.

The processor samples the STAT0-STAT1 output internally when *RESET is asserted. A High level on STAT0-STAT1 in this case is used to enable special test configurations, and may cause the processor to be inoperable. When *RESET is asserted, the processor drives STAT0-STAT1 Low in order to disable these considerations. However, if processor outputs are disabled by the Test mode, the processor is not able to drive STAT0-STAT1. Thus, if *RESET is asserted when the processor is in the Test mode, the STAT0-STAT1 pin must be driven Low externally. (In a master/slave configuration, as described in Section 5.8, STAT0-STAT1 is driven Low by the master processor when *RESET is asserted).

5.6 *WARN INPUT

An inactive-to-active transition on the *WARN input causes a *WARN trap to be taken by the processor. The *WARN trap cannot be disabled; the processor responds to the *WARN input regardless of its internal condition, unless the *RESET input is also asserted. This input is provided so that the system can gain control of the processor in emergency situations, such as when system power is about to be removed or when a severe, non-recoverable error occurs.

The *WARN input is edge-sensitive, so that an active level on the *WARN input for long intervals does not cause the processor to take multiple *WARN traps. However, *WARN must be held active for at least 4 cycles in order to be properly recognized by the processor. The processor still takes the *WARN trap if *WARN is de-asserted after 4 cycles. Another *WARN trap occurs if *WARN makes another inactive-to-active transition.

The processor enters the Executing mode when the *WARN input is asserted, regardless of its previous operational mode. Either 7 or 8 cycles after *WARN is asserted (depending on internal synchronization time), the processor performs a trap-handler instruction access on the channel. This instruction access is directed to address 16 in the instruction read-only memory (instruction ROM). If instruction ROM is not implemented in a particular system, another device or memory must respond to this instruction fetch.

If the CNTL0–CNTL1 inputs are 10 or 01 when the trap-handler instruction fetch completes, the processor enters the Halt or Step mode, respectively. Before the completion

of this instruction fetch, the CNTL0–CNTL1 inputs are irrelevant, except that the Load Test Instruction mode cannot be directly entered after a *WARN trap is taken. If the CNTL0–CNTL1 inputs are 00 immediately after *WARN is de-asserted, the effect on processor operation is unpredictable. If the CNTL0–CNTL1 inputs are 11, the processor remains in the Executing mode.

5.7 CLOCKS

The Am29000 supports two methods of system-clock generation and distribution. In one arrangement, the processor generates a clock for the system at its operating frequency; this clock appears on the SYSCLK pin, and may be distributed externally to other system components. In the second arrangement, the system provides its own clock generation and distribution; in this case, the processor receives the externally-generated clock on the SYSCLK pin.

In both arrangements, the circuits which generate and buffer SYSCLK are designed to minimize the apparent skew between internal processor clocks and external system clocks.

The processor provides a power-supply pin for the SYSCLK driver which is independent of all other chip power-distribution. This electrically isolates other processor circuits from noise which might be induced on the power supply by the SYSCLK driver. The separate power supply is also used to decide between the two possible clocking arrangements.

5.7.1 PROCESSOR-GENERATED CLOCK

If power (i.e. +5 volts) is applied to the SYSCLK power-supply pin, the processor is configured to generate clocks for the system. In this case, the SYSCLK pin is an output, and the signal on INCLK is used to generate the system clock. The processor divides the INCLK signal by two in the generation of SYSCLK, so INCLK should be driven at twice the processor's operating frequency.

5.7.2 SYSTEM-GENERATED CLOCK

If the SYSCLK power-supply pin is grounded, the processor is configured to receive an externally-generated clock. In this case, the SYSCLK pin is an input used directly as the processor clock. SYSCLK should be driven at the processor's operating frequency. In this configuration, the INCLK input should be tied High or Low, except in certain master/slave configurations as discussed in Section 5.8.

5.7.3 CLOCK SYNCHRONIZATION

The SYSCLK pin is at a High level during the first half of the processor cycle, and at a

Low level during the second half of the processor cycle. Thus, a processor cycle begins on a Low-to-High transition of SYSCLK. The definition of the beginning of the processor cycle is independent of the clocking arrangement chosen for a particular system.

In some systems, it might be desirable to have two or more processors operate in lock-step synchronization, with each processor driven by a common INCLK signal. In this case, synchronization of the processors is achieved by the *RESET input. If the de-assertion of *RESET meets a specified set-up time with respect to the Low-to-High transition of INCLK, the SYSCLK output is guaranteed to be Low in the next half-cycle. Thus, all processors may be synchronized as required.

5.7.4 ELECTRICAL SPECIFICATIONS

The electrical specifications for SYSCLK are different than the specifications for most other processor inputs and outputs. In order to reduce clock-skew effects, the SYSCLK pin is electrically compatible with the processor's CMOS circuits, rather than being compatible with transistor-transistor-logic (TTL) circuits.

Note that the SYSCLK pin is placed in the high-impedance state by the Test mode. If an externally-generated clock is not supplied in this case, processor state may be lost.

5.8 MASTER/SLAVE CHECKING

Each Am29000 output has associated logic which compares the signal on the output with the signal which the processor is providing internally to the output driver. The comparison between the two signals is made any time a given driver is enabled, and any time the driver is disabled only because of the Test mode. If, when the comparison is made, the output of a driver does not agree with its input, the processor asserts the MSERR output on the next cycle.

When the processor asserts MSERR, it takes no other actions with respect to the detected mis-comparison. In particular, no traps occur. However, the MSERR may be used externally to perform any system function, including the generation of a trap.

5.8.1 MASTER/SLAVE OPERATION

If there is a single processor in the system, the MSERR output indicates that a processor driver is faulty, or that there is a short-circuit in a processor output. However, a much higher level of fault detection is possible if a second processor (called a slave) is connected in parallel with the first (called a master), where the slave processor has outputs disabled by the Test mode.

The slave processor, by comparing its outputs to the outputs of the master processor,

performs a comprehensive check of the operation of the master processor. In addition, if the slave processor is connected at the proper position on the channel, it may detect open circuits and other faults in the electrical path between the master processor and its local devices and memories. Note that the master processor still performs the comparison on its outputs in this configuration.

5.8.2 PREVENTING SPURIOUS ERRORS

When two processors are connected in a master/slave configuration, it is necessary to prevent spurious assertions of MSERR. These result from situations where the outputs of the slave processor do not agree with the outputs of the master processor, but both processors are operating correctly.

One source of spurious errors can be unpredictable values for unimplemented bits in processor registers. This potential problem has been avoided by the Am29000 architecture; all unimplemented bits are read as 0.

Another source of spurious errors is a lack of synchronization between the master and slave processors. To maintain synchronization between the master and slave processors, it is first necessary that they operate with identical clocks. This is accomplished by having the master processor drive SYSCLK, with the slave processor receiving SYSCLK as an input, or by driving both processors' SYSCLK inputs with the same externally-generated clock.

However, the fact that both processors operate with the same clock is not sufficient to guarantee synchronization. Asynchronous processor inputs, if they are truly asynchronous to the operation of the master and slave processors, may affect the master processor a cycle sooner or later than they affect the slave processor. For this reason, the relevant asynchronous inputs (i.e. *WARN, *INTRO-*INTR3, *TRAP0-*TRAP1, CNTL0-CNTL1, and *RESET) must be externally synchronized to both the master and slave processors. Note that, in the case of *RESET, only the active-to-inactive transition must be synchronized.

5.8.3 SWITCHING MASTER AND SLAVE PROCESSORS

In some master/slave configurations, it might be desirable to give the slave processor control over the system when an error is isolated to the master processor. It is possible to grant control of the system to the slave processor by taking it out of the Test mode, and placing the master processor into the Test Mode. Note that synchronization must be maintained when this is accomplished (for example, using the Halt mode).

If the original master processor is configured to generate SYSCLK in this case, the slave processor must also generate SYSCLK when it becomes a master. Because of this, the INCLK signal must be supplied to both the master and slave processors, with both processors being configured to generate clocks.

In this master/slave configuration, the slave processor still receives SYSCLK from the master processor as described previously. The slave processor does not drive SYSCLK because of the Test mode. However, when the slave processor is taken out of the Test mode, it is able to drive SYSCLK as required.

Note that this processor-switching scheme may be generalized to more than two processors.

CHAPTER 6

COPROCESSOR INTERFACE

A coprocessor for the Am29000 is an off-chip extension of the processor's execution unit. The Am29000 communicates with the coprocessor using a mechanism which is very similar to the mechanism used to communicate with other external devices and memories.

However, because the coprocessor extends the instruction-execution capabilities of the processor, transfers to and from the coprocessor are in terms of operands, operation codes, results, and status information. This is in contrast to address and data transfers which occur for other types of external accesses. This chapter describes the coprocessor interface, both from a software and a hardware point-of-view.

6.1 COPROCESSOR PROGRAMMING

6.1.1 OVERVIEW OF COPROCESSOR OPERATIONS

A program executes the following steps to perform a coprocessor operation. This sequence is intended only as a guide, since there are many possible variations:

- 1) Send operands to the coprocessor. The number of transfers to the coprocessor depends on the number of operands, and the length of each operand. As many as 64 bits of information can be transferred in a single cycle.
- 2) Send an operation code and other operation information to the coprocessor. The operation can be specified by as many as 64 bits of information.
- 3) Start the coprocessor operation. This can occur simultaneously with the operation-code transfer of step 2.
- 4) Read the coprocessor results. The number of transfers from the coprocessor depends on the number of results, and the length of each result.

The above sequence is defined so that coprocessor operations may be concurrent with other processor operations, including external accesses. This is possible because coprocessor operations are decoupled from the transfer of information to and from the coprocessor. Once the operation is started, in step 3, the processor may continue further execution, overlapped with coprocessor execution, until the coprocessor results are read.

Because the Am29000 implements overlapped loads, it can continue execution after attempting to read a coprocessor result. However, if the processor attempts to use the result before the operation is complete, the processor enters the Pipeline Hold mode until the operation is complete.


In certain circumstances, it may be desired to perform multiple coprocessor operations before any results are read. For example, certain array computations form a single result from more than one operation. In this case, steps 1 through 3 above may be repeated—in any combination desired and as many times as desired—before results are read. The coprocessor interface allows the coprocessor to prevent the transfer of operands and/or operation codes if it is not prepared to receive them.

6.1.2 COPROCESSOR TRANSFERS

All coprocessor transfers occur between general-purpose registers and the coprocessor. The transfers occur as the result of the execution of load and store instructions for which the Coprocessor Enable (CE) bit has a value 1. For a store, the information transferred to the coprocessor is given either by the contents of two general-purpose registers, or by the contents of a general-purpose register and an 8-bit constant. For a load, information is transferred into a single general-purpose register in the Am29000.

The coprocessor model includes no provision for addressing. Although it is possible to extend the coprocessor interface to include addressing, addressing is more appropriately handled by normal external accesses defined for the processor (such as input/output).

The format of the instructions which transfer information to and from a coprocessor is shown in Figure 6-1.



08996A 43

Figure 6-1. Coprocessor Load/Store Format

For coprocessor stores, the RA and “RB or I” fields specify the source of data to be transferred to the coprocessor. The RA field specifies a general-purpose register whose contents are transferred to the coprocessor. The “RB or I” field specifies either a general-purpose register whose contents are transferred to the coprocessor, or a zero-extended constant which is transferred to the coprocessor. For the latter, the M bit of the operation code (bit 24) determines whether the register or the constant is used, as with most instructions. Note that as many as 64 bits of information may be transferred to the coprocessor by a single store instruction.

For coprocessor loads, the data transferred from the coprocessor is written to the general-purpose register given by RA; the “RB or I” field is unused in this case (however,

the contents of the specified register, or the zero-extended constant, does appear on the Address Bus). In contrast to the coprocessor store, a load transfers only 32 bits of information from the coprocessor.

Other bits in the coprocessor load and store instructions are defined as follows:

Bit 22: Transfer Control (TC)—This bit affects the behavior of the coprocessor for the transfer, depending on whether the transfer is for a load or store. The definition of this bit is by convention only, and is not enforced by the processor.

For transfers to the coprocessor (i.e. stores), a value of 1 for the TC bit causes a coprocessor operation to start. For transfers from the coprocessor (i.e. loads) a value of 1 for the TC bit causes the coprocessor to suppress exception-reporting. In either case, a value of 0 for the TC bit has no special effect on the coprocessor.

Bit 21: Set Coprocessor Active (SA)—This bit is provided to signal the beginning and end of a coprocessor operation, so that the proper action may be taken by software if the operation is interrupted.

An SA bit of 1 affects the Coprocessor Active (CA) bit in the Current Processor Status. If the SA bit is 1 for a store, the CA bit is set. If the SA bit is 1 for a load, the CA bit is reset. If the SA bit is 0, there is no effect on the CA bit.

Bit 20: Reserved

Bit 19: User Access (UA)—The UA bit allows programs executing in the Supervisor mode to emulate User-mode coprocessor transfers. This allows checking of the authorization of a transfer requested by a User-mode program. Note that this checking is performed externally, since the processor imposes no restriction on User-mode coprocessor transfers.

If the UA bit is 1, the coprocessor transfer is performed in the User mode, regardless of the value of the Supervisor Mode (SM) bit in the Current Processor Status. In this case, the User mode affects only the SUP/*US output; it has no effect on the registers which can be accessed by the instruction. If the UA bit is 0, the program mode for the transfer is controlled by the SM bit.

Bits 18-16: Option (OPT)—The OPT field is placed on the OPT0–OPT2 outputs during the coprocessor transfer. There is a one-to-one correspondence between the OPT field and the OPT0–OPT2 outputs; that is, the most-significant OPT bit is placed on OPT2, and so on.

The OPT bits define the quantities being transferred to or from the coprocessor. For example, they can specify whether operands or operation codes are being transferred. The interpretation of the OPT field depends on the definition of a given coprocessor.

The transfer of data to or from the coprocessor may be caused by any load or store instruction defined for the processor; the operation of coprocessor transfers is very similar to the operation of external accesses.

Coprocessor transfers are overlapped with the execution of instructions which sequentially follow the coprocessor load or store instruction. However, only one load or store may be in progress in any given cycle, whether or not the load or store is directed to a coprocessor. The pipeline interlocks which apply to external accesses also apply to coprocessor transfers, except that coprocessor-transfer interlocks are determined by the time taken by the coprocessor to perform an operation, rather than the time taken to perform an access.

Note that coprocessor transfers may be performed by Load Multiple and Store Multiple instructions. However, register RB has no defined interpretation for a Store Multiple to the coprocessor. For this reason, Store Multiple is defined to transfer multiple, 32-bit quantities to the coprocessor. Similarly, a Load Multiple transfers multiple, 32-bit quantities from the coprocessor. Note, however, that the incrementing address sequence defined for Load Multiple and Store Multiple still appears on the Address Bus for coprocessor transfers.

6.1.3 COPROCESSOR EXCEPTIONS

A Coprocessor Exception trap occurs if the coprocessor reports an exception (using the *DERR signal) during a coprocessor transfer. The Coprocessor Exception may occur either for a coprocessor load or store.

In the case of a load which reads a coprocessor result, the Coprocessor Exception can be used to indicate that the result is incorrect because of some exceptional condition. In some cases, the Am29000 might be able to correct the results of the operation.

In the case of a store to the coprocessor, the Coprocessor Exception can be used to indicate that the coprocessor cannot accept the transfer because of some exceptional condition. For example, it may indicate an error in a stream of calculations, where intermediate results are not being read. As with the load, the Am29000 may be able to correct the exceptional condition.

As noted above, the trap handler which executes as the result of the Coprocessor Exception trap may attempt to correct the exceptional condition. In many cases, the trap handler must be able to read the intermediate results of the operation from the coprocessor, along with other information about the operation. When this information is read, it may be necessary to suppress further exception-reporting, so that the trap handler does not create additional Coprocessor Exception traps. For this reason, the TC bit in the coprocessor load or store instruction allows the processor to read coprocessor results while suppressing exception-reporting.

Additionally, the TC bit allows a program to read the result of a coprocessor operation

regardless of any errors which may have occurred. This provides an optional trapping capability analogous to that provided for certain Am29000 arithmetic operations (for example, Am29000 instructions allow an optional trap on arithmetic overflow).

6.1.4 COPROCESSOR AS A SYSTEM OPTION

When the coprocessor is a system option, coprocessor operations are performed by the processor when the coprocessor is not present.

The coprocessor may be designed as a system option by use of the Coprocessor Present (CP) bit of the Configuration Register. The CP bit is set during system initialization, based on the presence (CP = 1) or absence (CP = 0) of the coprocessor. If the CP bit is 0 when the processor attempts to execute a coprocessor load or store instruction, a Coprocessor Not Present trap occurs.

When a Coprocessor Not Present trap is taken, the Channel Address, Channel Data, and Channel Control registers contain information related to the coprocessor transfer. This information may be used by the trap handler to emulate the operation of the coprocessor.

6.1.5 INTERRUPTED COPROCESSOR OPERATIONS

The Coprocessor Active (CA) bit of the Current Processor Status may be used to indicate the duration of a coprocessor operation. The value 1 in the CA bit indicates that the coprocessor has begun an operation which has not completed (i.e. the final results have not been read).

The CA bit is affected by the Set Coprocessor Active (SA) bit in the coprocessor load and store instructions. If the SA bit is 1 for a store, the CA bit is set; if the SA bit is 1 for a load, the CA bit is reset. The routine which accesses the coprocessor is responsible for setting and resetting the CA bit appropriately.

If an interrupt or trap is taken during a coprocessor operation, and the CA bit has been properly managed, the CA bit of the Old Processor Status signals to an interrupt or trap handler that the interrupted routine had begun a coprocessor operation, but had not completed the operation before the interrupt or trap was taken. In this case, the coprocessor contains state information which must be preserved. This information may be saved and restored across the interrupt or trap, or, alternatively, kept in the coprocessor.

Upon an interrupt or trap, the state information contained in the coprocessor depends on both the operation being performed and the definition of the coprocessor. The methods used to determine what state information must be saved, and the methods used to transfer this information, are also dependent on the definition of the coprocessor.

Due to interrupt-latency considerations, it may be desirable to leave state information in the

coprocessor upon interrupt, rather than require that it always be saved. A problem arises, however, when a routine other than the one which was originally interrupted attempts to use the coprocessor. The coprocessor may be protected from such use by resetting the CP bit in the Configuration Register. If another routine attempts to use the coprocessor in this case, a Coprocessor Not Present trap occurs. The trap handler for this trap may either save the coprocessor state and make the coprocessor available to the trapping routine, or return control to the routine which was originally using the coprocessor.

Certain coprocessor operations may not be interruptible. For these operations, interrupts may be disabled by the Disable Interrupts (DI) and/or Disable All Interrupts and Traps (DA) bits in the Current Processor Status Register. However, this disabling can be performed only by a program in the Supervisor mode. Any User-mode programs which perform non-interruptible coprocessor operations incur the overhead of a call to a Supervisor-mode program.

6.2 COPROCESSOR ATTACHMENT

Communication with the coprocessor occurs via the Am29000 channel. Figure 6-2 illustrates a typical coprocessor connection. For transfers to the coprocessor, 64 bits of data are transferred in a single cycle, using the Address Bus and Data Bus simultaneously. For transfers from the coprocessor, 32 bits of data are transferred in a cycle, using the Data Bus.

The width of transfers to the coprocessor is greater than the width of transfers from the coprocessor because the Am29000 is optimized for computations performed on two, word-length operands, with a single, word-length result. The operand/result data flow of the processor is reflected in the interface to the coprocessor.

The protocol for coprocessor transfers is nearly identical to the protocol for other external accesses on the channel. Minor differences result from the fact that there are no addresses for coprocessor transfers, and from the fact that the coprocessor is operation-oriented, rather than access-oriented.

6.2.1 SIGNAL DESCRIPTION

Coprocessor transfers are indicated on the channel by the DREQT1 output being High during a request. The DREQT0 output also affects the transfer, based on the R/*W signal, as follows:

<u>R/*W</u>	<u>DREQT1</u>	<u>DREQT0</u>	<u>Meaning</u>
0	1	0	Transfer to coprocessor
0	1	1	Transfer to coprocessor, start operation
1	1	0	Transfer from coprocessor
1	1	1	Transfer from coprocessor, suppress errors

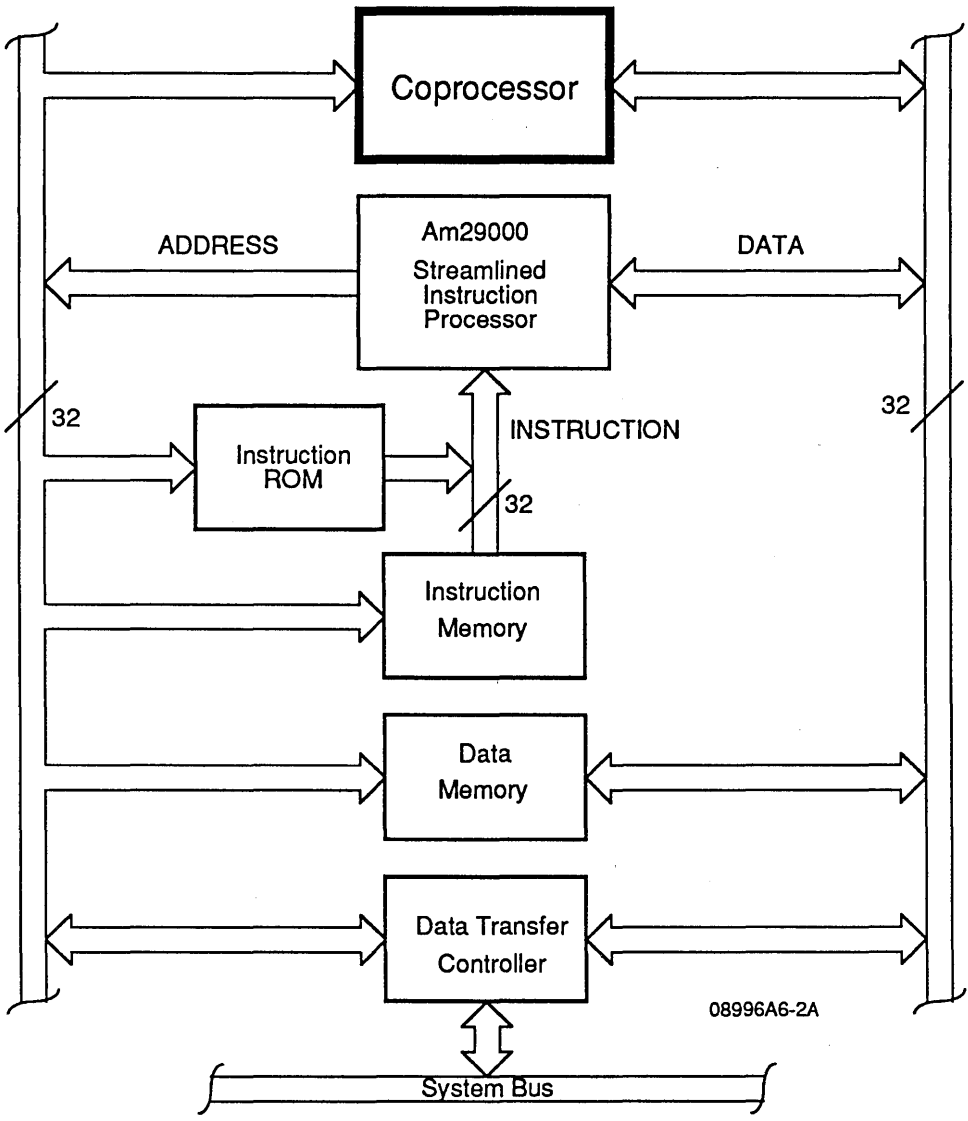


Figure 6-2. Coprocessor Attachment

The output DREQT1 is High only for coprocessor transfers. When the Address Bus is idle, the default value for DREQT1 is Low. As a result, the coprocessor can base its operation solely on DREQT1 and *BINV, without regard to *DREQ. Note that the interpretation of DREQT0 during a coprocessor transfer is by convention only.

The only signal unique to coprocessor transfers is the *CDA input. The coprocessor de-asserts this signal whenever it can accept no transfers from the processor (normally, this is because it is performing an operation).

The completion of a transfer to the coprocessor is indicated when the coprocessor asserts *CDA. The input *DRDY is not used in this case. The performance of transfers to the coprocessor is enhanced by the use of *CDA, since it eliminates the need for the coprocessor to decode a transfer request and respond with *DRDY, thereby eliminating the logic delay involved. Note that the coprocessor normally de-asserts *CDA when it starts an operation, so that *CDA can be independent of transfer requests.

6.2.2 COPROCESSOR COMMUNICATION

The Address Bus is used to transfer information to the coprocessor. Therefore, the addressing function of other devices and memories on the channel must be disabled during coprocessor transfers. Since DREQT1 is High for all coprocessor transfers, it should be used to inhibit the address-decoding function of channel devices and memories, as well as to indicate to the coprocessor that a transfer is occurring.

The OPT0–OPT2 outputs are used during coprocessor transfers to indicate the type of transfer, or to provide other controls for the coprocessor. The interpretation of the OPT0–OPT2 signals depends on the implementation of the coprocessor, and may also depend on the R/*W signal.

Coprocessor Transfer Protocols

The protocols available for coprocessor transfers are based on the protocols for simple, pipelined, and burst-mode data accesses discussed in Section 5.2. The protocols for write-accesses are used for transfers to the coprocessor, and the protocols for read-accesses are used for transfers from the coprocessor.

The coprocessor transfers differ in several respects from the protocol for external data accesses:

- 1) The *CDA signal consistently replaces the *DRDY for transfers to the coprocessor. An active level on *CDA, for transfers to the coprocessor, has an effect which is equivalent to the effect of an active level on *DRDY for normal store-operations. Note that *DRDY is still used for transfers from the coprocessor.

- 2) The Address Bus does not contain an address during a coprocessor transfer, but may contain data in the case of a transfer to the coprocessor. However, for transfers from the coprocessor, the Address Bus is still sequenced as described in Section 5.2, and the sequencing is determined by the same controls—except that *CDA replaces *DRDY for transfers to the coprocessor. The contents of the Address Bus are determined by the coprocessor load instruction, as for other load instructions.
- 3) For any coprocessor transfer, an active level on *DERR causes a Coprocessor Exception trap, rather than a Data Access Exception trap.
- 4) For burst-mode coprocessor transfers, the interpretation of sequential addressing is undefined. For this reason, burst-mode transfers are normally restricted to 32 bits of information for every transfer, regardless of whether the transfer is to or from the coprocessor. Note, however, that the incrementing address sequence is still present in the definition of a burst-mode coprocessor transfer, and may be useful in some cases.

Sequencing of *CDA

The coprocessor de-asserts *CDA whenever it cannot accept a transfer from the Am29000. An inactive level on *CDA prevents the Am29000 from transferring operands or operation codes to the coprocessor when these transfers might interfere with coprocessor operation.

Normally, the coprocessor de-asserts *CDA when it begins an operation. *CDA remains inactive until the coprocessor has completed the operation, and can accept further transfers from the processor. For some operations, a result may have to be read before the coprocessor can assert *CDA.

The coprocessor can acknowledge a transfer by asserting *CDA. However, it is generally more efficient for the coprocessor to hold *CDA active as long as it can accept transfers. In the latter case, multiple data-transfers can occur at a high rate, without involving long logic delays. *CDA is related to the operation of the coprocessor in this case, rather than to the transfer of data.

Exception Reporting

The coprocessor reports exceptions by the activation of *DERR during any coprocessor transfer. This causes a Coprocessor Exception trap to occur. However, if the DREQT0–DREQT1 signals have the value 11 for a transfer from the coprocessor, exception-reporting should be suppressed, and *DERR should not be asserted. Note, however, that the Am29000 does not enforce the suppression of exception reporting.

CHAPTER 7

PROGRAMMING

This chapter discusses programming topics as they relate to the Am29000. It focuses on the use of processor resources which were more formally described in Chapter 3. The presentation in this chapter is intended to be used as a guide in the implementation of software systems for the processor, not as a strict definition of how these systems should be implemented.

This chapter is organized into three sections. The first two sections discuss applications and systems programming for the processor, and the third discusses certain features of the processor pipeline which are exposed to—and must be properly handled by—software which executes on the processor.

7.1 APPLICATIONS—PROGRAMMING CONSIDERATIONS

This section discusses topics of general concern in the implementation of applications programs.

7.1.1 PROCEDURE CALLS AND RETURNS

The Am29000 is designed to minimize the overhead of calling a procedure. It allows the functions of passing parameters to a procedure and returning results from a procedure to be performed efficiently.

This efficiency is due largely to the definition of the local registers. The relative addressing of local registers greatly reduces the overhead of run-time storage management for variables, parameters, returned results, and other quantities required in procedure linkage.

Run-Time Stack Organization and Use

For programs written in a compiled, procedural language—such as C or Pascal—storage for certain program variables and compiler data is allocated, during the execution of the program, on a structure called a run-time stack. The compiler generates the instructions to create and manage the run-time stack, and compiler-generated instructions are based on its existence.

Figure 7-1 depicts part of a run-time stack as an example. The stack consists of consecutive, overlapping structures which are called activation records. An activation record contains dynamically-allocated information specific to a particular activation (i.e. call) of a procedure.

Because of recursion, multiple copies of a procedure may be active at any given time. Each active procedure has its own unique activation record, allocated somewhere on the run-time stack. The variables required by a particular procedure activation are contained only in the activation record associated with that activation. Thus, the variables for different activations do not interfere with one another.

There are three activation records in the example shown in Figure 7-1. This stack configuration was generated by procedure A calling procedure B, which in turn called procedure C. The fact that procedure C is the currently-active procedure is reflected by its activation record being on the top of the run-time stack. The Stack Pointer points to the top of C's activation record.

In Figure 7-1, the storage areas labelled "out args" and "in args" are shared between the called procedure and the caller for the communication of parameters and results. These are called the outgoing arguments area (for the caller) or the incoming arguments area (for the callee). The areas labelled "locals" contain storage for local variables, temporary variables (for example, for expression evaluation), and any other items required for the proper execution of the procedure.

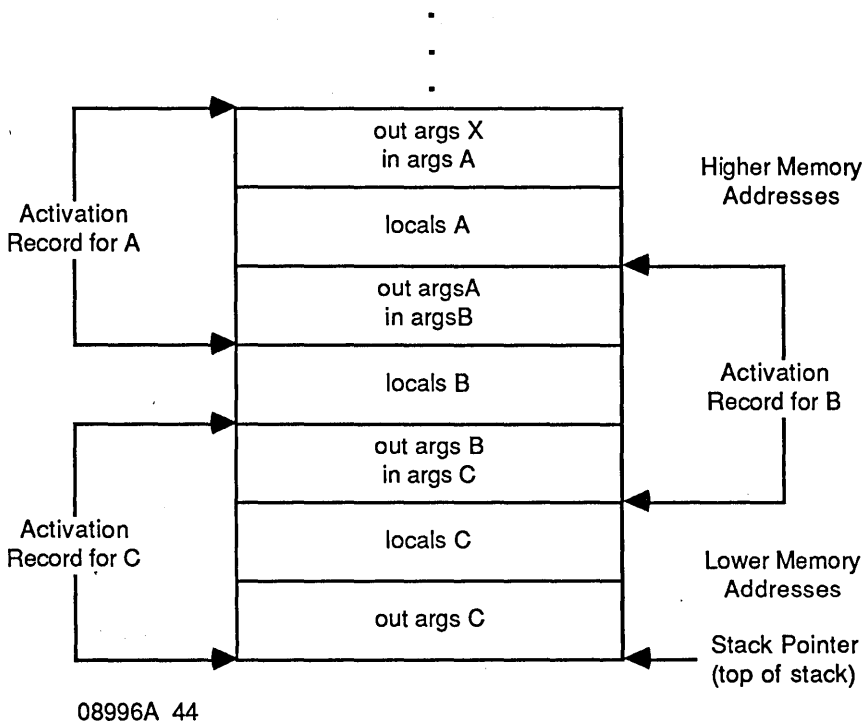


Figure 7-1. Run-time Stack Example

Management of the Run-Time Stack

When a procedure is called, a new activation record must be allocated on the run-time stack. An activation record is allocated by subtracting from the Stack Pointer the number of locations needed by the new activation record. The Stack Pointer is decremented so that variables referenced during procedure execution are referenced in terms of positive offsets from the Stack Pointer.

When storage for an activation record is allocated, the number of storage locations allocated is the sum of the number of locations needed for:

1. Local variables.
2. Temporary variables required by the compiler.
3. Restarting the caller, such as locations for return addresses.
4. Arguments of procedures which may be called in turn by the called procedure (the outgoing arguments area).

Note that, in some cases, no storage is required for one or more of the above items. Also, the incoming arguments area, though part of the activation record of the callee, is not allocated storage at this time, because this storage was allocated as the outgoing arguments area of the calling procedure.

An activation record is de-allocated by adding to the Stack Pointer the value which was subtracted during allocation.

More than one run-time stack may be used. In particular, it is possible to split activation records across multiple stacks. Storage is allocated and de-allocated on these stacks in synchronism. The reasons for such a split are explained below.

Am29000 Local Registers as a Stack Cache

A compiler targeted to the Am29000 should use two run-time stacks for activation records: one for often-used scalar data, and another for structured data and additional scalar data. The scalar portion of the activation record can then be mapped into the processor's local registers, because of the Stack-Pointer addressing which applies to the local registers.

Allocation and de-allocation of activation records can occur largely within the confines of the local registers. The result is that the currently-active portion of the scalar activation record is cached in high-speed local registers. The term "stack cache" in this section refers to the use of the local registers to cache a portion of the activation record stack.

The principle of locality of reference—which allows any cache to be effective—also applies to the stack cache. The entries in the stack cache are likely to remain there for re-use,

because the dynamic nesting-depth of activated procedures tends to remain near a given depth for long periods of time. As a result, the size of the run-time stack does not change very much over long intervals of program execution.

Since activation records are allocated and de-allocated within the local registers, most procedure linkage can occur without external references. Also, during procedure execution, most data accesses occur without external references, because the scalar data in an activation record is most frequently referenced. Activation records are typically small, so the 128 locations in the local register file can hold many activation records from the run-time stack.

Mapping of Activation Records to the Local Registers

Whenever a given location on the run-time stack is present in the local registers, it is mapped to the same register. The absolute number of the register that a given location occupies is given by bits 8-2 of the 32-bit memory address of the stack location. Thus, stack quantities whose addresses differ by 512 (since addresses are byte addresses) are mapped into the same local register.

Only one stack location can actually be mapped to a local register at any point in time. When the run-time stack grows beyond the 128-word capacity of the local registers, some movement of data between the stack cache and the scalar run-time stack in data memory must occur.

The terms “overflow” and “underflow” are used to describe the two kinds of conditions caused when the run-time stack cannot be contained completely within the stack cache. Both overflow and underflow can occur during the normal execution of a program.

An overflow occurs when a procedure is called, but the activation record of the callee requires more registers than can be allocated in the stack cache. In this case, the contents of a number of registers must be moved to data memory. The number of registers involved must be sufficient to allow the entire activation record of the callee to reside in the stack cache.

An underflow occurs when a procedure returns to the caller, and the entire activation record of the caller is not resident in the stack cache. In this case, there are enough unallocated registers to contain the activation record, since all locations below the activation record have been de-allocated and are no longer valid. However, the non-resident portion of the caller’s stack must be moved from the main-memory stack to the local registers. Underflow occurs because overflow occurred at some previous point during program execution, and the overflow caused part of the run-time stack to be moved to memory.

To use the stack cache properly, a compiler must not allow the size of an activation record on the scalar stack to exceed the size of the local register file (128 locations). This is required because the processor performs no dynamic management of the stack cache; management is performed by software. The processor cannot detect a reference to a quantity which is not in the stack cache. If the scalar portion of the activation record requires more than 128 location, the excess may be kept on the activation record used for structured data.

The software which performs procedure linkage must insure that the entire scalar activation record of the callee is in the local registers before the callee begins execution. The activation record must remain in the local registers as long as the callee is executing, since any quantity in the activation record may be referenced.

Implementation of the Stack Cache

The processor support for implementing a stack cache within the local register file consists of the 32-bit Stack Pointer (mapped to Global Register 1) and three 7-bit adders for computing local-register addresses. Overflow and underflow detection is performed by software when an activation record is allocated or de-allocated.

In this scheme, the value by which the Stack Pointer is adjusted for allocation and de-allocation of activation records must be a constant which can be determined by the compiler, since the compiler generates the instructions to do the adjustment and to check for overflow and underflow. The only problem in determining this constant is the outgoing arguments area: the size of the area required here varies with the number of parameters passed to procedures called in turn by the callee. The solution is to make the outgoing arguments area large enough to accommodate the call with the largest number of parameters.

The instruction sequence which executes as a result of the procedure call is the procedure prologue, and the sequence which executes as a result of the procedure return is the procedure epilogue. The procedure prologue allocates the activation record and checks for overflow, and the procedure epilogue de-allocates the activation record and checks for underflow. The two instruction sequences are shown in Figure 7-2.

The layout of an activation record is shown in Figure 7-3. Quantities on this diagram (for example, the size of the activation record "SIZE_A") relate to quantities shown in Figure 7-2.

```

*-----> A:  SUB      GR1,GR1,#ALLOC_A      (1)
procedure    ASGEU   #S_OVERFLOW,GR1,GR64  (2)
prologue     ADD     LR1,GR1,#SIZE_A       (3)
*----->      .
*----->      .
body of      .
procedure    .
*----->      .
*----->     ADD     GR1,GR1,#ALLOC_A      (4)
procedure    ASEQ   40#h,GR1,GR1         (5)
epilogue     JMPI   LRO                   (6)
*----->     ASLEU  #S_UNDERFLOW,LR1,GR65 (7)

```

Figure 7-2. Procedure Prologue and Epilogue

The procedure prologue consists of instructions (1), (2), and (3). Instruction (1) allocates the new activation record by decrementing the Stack Pointer (GR1) by the number of words in the new activation record. Instruction (2) compares the new value of the Stack Pointer (SP) with the stack-cache lower bound in GR64 (this value could be kept in any global register). If the new SP is below the lower bound of the stack cache, then an overflow has occurred, and the ASGEU instruction causes a trap.

If a trap occurs, the trap handler moves the contents of some of the local registers to instruction/data memory; the number of registers moved must be enough to allow the local registers to accommodate the new activation record. The trap handler also adjusts the lower bound of the stack cache (in this example, the lower bound is the value in GR64).

In the example of Figure 7-2, instruction (3) computes a pointer to the top of the new activation record and stores it in the location marked LR1(pro) in Figure 7-3. The use of this pointer is explained below. Note that LR1 is defined by the SP which was set by instruction (1).

The procedure epilogue consists of instructions (4), (5), (6), and (7). Instruction (4) de-allocates the activation record. Instruction (5) is a NO-OP (see Section 7.1.12), because a change in the value of the SP must be separated by at least one cycle from a use of the SP for local-register addressing. This restriction is caused by processor pipelining (see Section 7.3.3). The NO-OP can be replaced with a useful instruction, providing it does not reference any local registers.

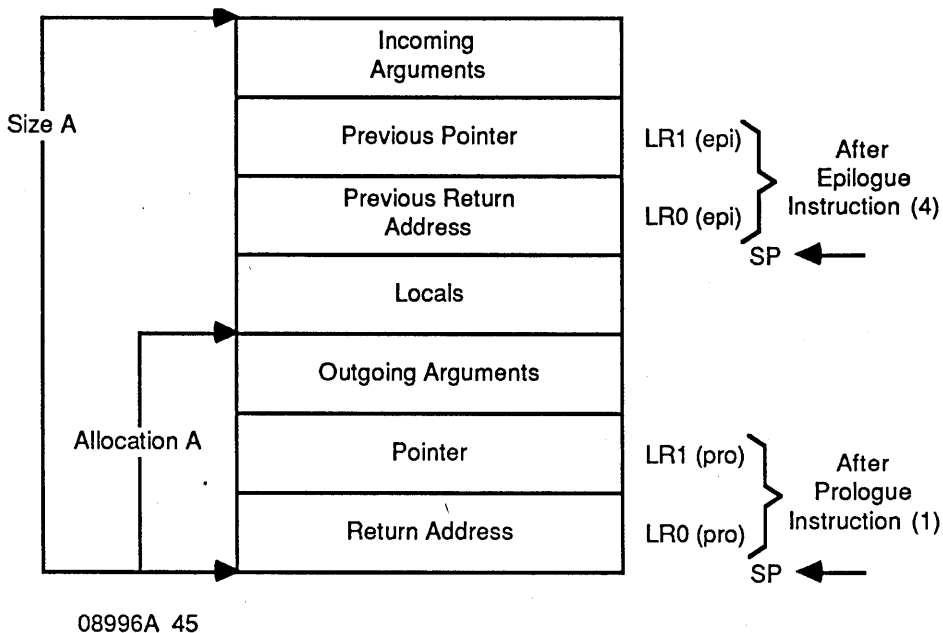


Figure 7-3. Activation Record

Instruction (6) is the return to the caller. The return address is in the location marked LR0(epi) in Figure 7-3. This location is defined by the SP set in instruction (4). With the new value of the SP, LR1(epi) contains the pointer to the top of the activation record of the caller, which was calculated during the procedure prologue of the caller. Instruction (7) compares this pointer to the upper bound of the stack cache in GR65. If this pointer is beyond the upper bound of the stack cache, then an underflow has occurred, and the ASLEU instruction causes a trap. The trap handler loads the required number of local registers from data memory, and adjusts the upper bound of the stack cache (in GR65).

The pointer to the top of the activation record computed in instruction (3) is needed to perform the compare in instruction (7). This pointer is used by the epilogue of every procedure that is called by the current procedure, to guarantee that the entire activation record of the caller is resident in the stack cache. All references to the scalar activation record are compiled as local-register references, and the entire activation record must be in the local register, as the compiler assumed it would be. Since the value of the pointer does not change during the execution of a procedure, it is computed only once—in the prologue.

The value in GR65 in this example is the virtual address of the boundary between the stack cache and the run-time stack in main memory. This address points to the highest memory location which is mapped to the local registers (e.g. adding 4 to this address gives the address of the first unmapped location, which is in instruction/data memory).

The value in GR64 is the virtual address of the lower bound of the stack cache. It may or may not point to a valid location; it is retained in a global register simply to make bounds-checking more efficient. Both of these values are software-defined, and can be kept in any global registers (except GR1). The values are maintained by the routines which handle the movement of data between the stack cache and main memory.

7.1.2 ADDRESSING GENERAL-PURPOSE REGISTERS INDIRECTLY

Registers in the processor are usually addressed directly by fields within instructions. However, indirect addressing of registers may be required in some situations, such as when a program pointer is known to point to a variable which is resident in the register file.

Three special registers—Indirect Pointers A, B, and C—are provided, so that separate indirect register-numbers can be set for each of the source and destination operands within an instruction. Indirect Pointer C corresponds to the destination register RC, Indirect Pointer A corresponds to the RA operand-register, and Indirect Pointer B corresponds to the RB operand-register.

A given indirect pointer (the value in the corresponding register) is used to address the register file whenever Global Register 0 is specified as a source or destination register. For example, a value of 0 in the RA field of an instruction causes the content of the Indirect Pointer A Register to be used to access the RA operand.

The indirect pointers can be set by the Move To Special Register and Floating-Point instructions, and by the instructions EMULATE, MULTIPLY, DIVIDE, and Set Indirect Pointers (SETIP). The Move To Special Register instructions set the indirect pointers individually as special-purpose registers. The Floating-Point, MULTIPLY, DIVIDE, and SETIP instructions set all three indirect pointers simultaneously, deriving the values which are written into the pointers from the instruction fields RC, RA, and RB. The EMULATE instruction sets all three indirect pointers, but only the Indirect Pointer A and Indirect Pointer B registers are written with meaningful values.

When an indirect pointer is set by a Move To Special Register, bits 9-2 of the source operand are copied to corresponding bits in the indirect pointer. This allows the addressing of general-purpose registers, via the indirect pointers, to be consistent with the addressing of words in external memories and devices.

When the indirect pointers are set from instruction fields, the resulting values reflect the Stack-Pointer addition which is performed on local registers. In addition, register bank-protection checking is performed on the values which are loaded. A Protection Violation trap occurs if the values represent registers which cannot be accessed.

The indirect pointers may thus be used to access exactly those operands which would be accessed by the instruction fields setting the indirect pointers. Consequently, a routine which emulates an instruction operation can access, with no overhead, the source and destination registers for the instruction being emulated. No copying of arguments and results needs to be done.

When using indirect register-addressing, at least one cycle of delay must separate any instruction which sets an indirect pointer and any instruction which de-references that pointer. (It can't access a general-purpose register using the indirect pointer.) This restriction is the result of processor pipelining (see Section 7.3.3).

7.1.3 RUN-TIME CHECKING

The assert instructions provide programs with an efficient means of comparing two values and causing a trap when a specified relation between the two values is not satisfied. Thus, the instructions assert that some specified relation is true, and trap if the relation is not true. This allows run-time checking—such as checking that a computed array index is within the boundaries of the storage for an array—to be performed with a minimum performance penalty.

Assert instructions are available for comparing two signed or unsigned operands. The following relations are supported: equal-to, not-equal-to, less-than, less-than or equal-to, greater-than, and greater-than-or-equal-to.

The assert instructions specify a vector number for the trap. However, only vector numbers 64 through 255 (inclusive) may be specified by User-mode programs. If a User-mode assert

instruction causes a trap, and the vector number is between 0 and 63 inclusive, a Protection Violation trap occurs, instead of the specified trap.

Since the assert instructions allow the specification of the vector number, several traps may be defined in the system, for different situations detected by the assert instructions.

7.1.4 OPERATING SYSTEM CALLS

An applications program can request a service from the operating system by using the following instruction:

```
ASNE SYSTEM_ROUTINE,GR1,GR1
```

This instruction always creates a trap, since it attempts to assert that the content of a register is not equal to itself (the register number used here is irrelevant, as long as the register is otherwise accessible).

The `SYSTEM_ROUTINE` vector number specified by the instruction invokes the execution of the operating system routine which provides the requested service. This vector number may have any value between 64 and 255, inclusive (vector numbers 0 through 63 are pre-defined or reserved). Thus, as many as 192 different operating-system routines may be invoked from the applications program.

In cases where the indirect pointers may be used, the `EMULATE` instruction allows two operand/result registers to be specified to the operating-system routine. The instruction is:

```
EMULATE SYSTEM_ROUTINE,LR3,LR6
```

In this case, the `SYSTEM_ROUTINE` vector number performs the same function as in the previous example. Here, however, `LR3` and `LR6` are specified as operand registers and/or result-registers (these particular registers are used only for illustration). The operating-system routine has access to these registers via the indirect pointers, allowing flexible communication.

7.1.5 MULTI-PRECISION INTEGER ADDITION AND SUBTRACTION

The processor allows the Carry (C) bit of the ALU Status Register to be used as an operand for add and subtract instructions. This provides for the addition and subtraction of operands which are greater than 32 bits in length. For example, the following code implements a 96-bit addition with signed overflow detection.

```
ADD      GR87,GR76,GR64
ADDC     GR88,GR77,GR65
ADDCS    GR89,GR78,GR66
```

Global registers GR76-GR78 contain the first operand, global registers GR64-GR66 contain the second operand, and global registers GR87-GR89 contain the result. The first two add instructions set the C bit, which is used by the second two instructions. If the addition causes a signed overflow, then an Out of Range trap occurs; overflow is detected by the final instruction.

7.1.6 INTEGER MULTIPLICATION

The processor performs integer multiplication by a series of multiply step instructions, rather than by a single instruction. Note that, when the product of a constant and a variable is to be computed, a more efficient sequence of shift and add instructions can usually be found.

If a program requires the multiplication of two variables, the required sequence of multiply steps may be executed in-line, or executed in a multiply routine called as a procedure. It may be beneficial to precede a full multiply procedure with a routine to discover whether or not the number of multiply steps may be reduced. This reduction is possible when the operands do not use all of the available 32 bits of precision

The following routine multiplies two, 32-bit, signed integers:

```

; Signed 32 bit multiply.

; multiplier in GR70
; multiplicand in GR71
; result MSW in GR72
; result LSW in GR73

MFSR   Q,GR70           ; multiplier to Q
MUL    GR72,GR71,00#h   ; step 1. no initial
                                ; partial product. Load GR72
                                ; with multiplicand (GR71)
                                ; if lsb of Q is 1, else
                                ; load it with zero. Then
                                ; down shift GR72 & Q by 1
                                ; bit.
MUL    GR72,GR71,GR72   ; step 2. conditional add
                                ; of GR71 to GR72 depending
                                ; on the least sig bit of Q.
                                ; Then down shift GR72 & Q
                                ; by 1 bit.
MUL    GR72,GR71,GR72   ; step 3
MUL    GR72,GR71,GR72   ; step 4
MUL    GR72,GR71,GR72   ; step 5
MUL    GR72,GR71,GR72   ; step 6
MUL    GR72,GR71,GR72   ; step 7

```

.
.
.
steps 8 thru 28
.
.
.

```
MUL    GR72,GR71,GR72    ; step 29
MUL    GR72,GR71,GR72    ; step 30
MUL    GR72,GR71,GR72    ; step 31
MULL   GR72,GR71,GR72    ; step 32. conditional
                          ; subtract of GR71 to GR72
                          ; depending on the least sig
                          ; bit of Q. Then down shift
                          ; GR72 & Q by 1 bit.
MFSR   GR73,Q            ; get LSW of result into
                          ; GR73.
```

The following routine multiplies two, 32-bit, unsigned integers:

```
; Unsigned 32 bit multiply.

; multiplier in GR70
; multiplicand in GR71
; result MSW in GR72
; result LSW in GR73

MTSR   Q,GR70            ; multiplier to Q
MULU   GR72,GR71,00#h    ; step 1. no initial
                          ; partial product. Load GR72
                          ; with multiplicand (GR71)
                          ; if lsb of Q is 1, else
                          ; load it with zero. Then
                          ; down shift GR72 & Q by 1
                          ; bit.
MULU   GR72,GR71,GR72    ; step 2. conditional add
                          ; of GR71 to GR72 depending
                          ; on the least sig bit of Q.
                          ; Then down shift GR72 & Q
                          ; by 1 bit.
MULU   GR72,GR71,GR72    ; step 3
MULU   GR72,GR71,GR72    ; step 4
MULU   GR72,GR71,GR72    ; step 5
MULU   GR72,GR71,GR72    ; step 6
MULU   GR72,GR71,GR72    ; step 7

.  
.
```

steps 8 thru 28

.
. .
. . .

```
MULU   GR72,GR71,GR72      ; step 29
MULU   GR72,GR71,GR72      ; step 30
MULU   GR72,GR71,GR72      ; step 31
MULU   GR72,GR71,GR72      ; step 32
MFSR   GR73,Q              ; get LSW of result into
                           ; GR73.
```

7.1.7 INTEGER DIVISION

The processor performs integer division by a series of divide step instructions, rather than by a single instruction. When the divisor is a power of 2, the divide should be accomplished by a right shift.

If a program requires the division of two integers, the required sequence of divide steps may be executed in-line, or executed in a divide routine called as a procedure. It may be beneficial to precede a full divide procedure with a routine to discover whether or not the number of divide steps may be reduced. This reduction is possible when the operands do not use all of the available 32 bits of precision

The following routine divides a 64-bit, unsigned dividend by a 32-bit, unsigned divisor:

```
;      64-bit dividend. Most-significant word in GR71,
;      Least-significant word in GR70.
;      32-bit divisor in GR72.
;
;
;      32-bit quotient in GR74
;      32-bit remainder in GR73
;
MTSR   Q,GR70              ; set Q to low half of 64-bit number
DIV0   GR73,GR71           ; do first step. GR73 & Q become
                           ; 64 bit shift area for divide.
DIV    GR73,GR73,GR72     ; divide step 1
.
.                          ; total of 31 DIV instructions
.
DIV    GR73,GR73,GR72     ; divide step 31
DIVL   GR73,GR73,GR72     ; last divide step
DIVREM GR73,GR73,GR72     ; remainder into GR73
MFSR   GR74,Q             ; resultant quotient into GR74
```

The following routine divides a 32-bit, unsigned dividend by a 32-bit, unsigned divisor:

```

;      32-bit dividend in GR70.
;      32-bit divisor in GR71.
;
;
;      32-bit quotient in GR73
;      32-bit remainder in GR72
;
MTSR   Q,GR70           ; set Q to 32-bit dividend
DIV0   GR72,00#h        ; do first step. GR72 & Q become
                        ; 64-bit shift area for divide.
DIV    GR72,GR72,GR71  ; divide step 1
.
.                        ; total of 31 DIV instructions
.
DIV    GR72,GR72,GR71  ; divide step 31
DIVL   GR72,GR72,GR71  ; last divide step
DIVREM GR72,GR72,GR71  ; remainder into GR72
MFSR   GR73,Q          ; resultant quotient into GR73

```

The following routine divides a 64-bit, signed dividend by a 32-bit, signed divisor:

```

;      64-bit dividend. Most-significant word in GR71,
;                          Least-significant word in GR70.
;      32-bit divisor in GR72.
;
;
;      32-bit quotient in GR74
;      32-bit remainder in GR73
;
ASNE   DIVBYZERO,GR72,00#h ; check for divide by zero
JMPF   GR71,SKIP1         ; jmp if dividend positive
CONST  GR75,0000#h        ; set flag to 0 for positive
CPEQ   GR75,GR75,00#h     ; set flag to TRUE if neg dividend
SUBR   GR70,GR70,00#h     ; negate low order word
SUBRC  GR71,GR71,00#h     ; negate high order word
SKIP1:
JMPF   GR72,SKIP2         ; jmp if divisor positive
OR     GR72,GR72,GR72     ; NOP
CPEQ   GR75,GR75,00#h     ; toggle flag
SUBR   GR72,GR72,00#h     ; negate divisor
SKIP2:
MTSR   Q,GR70            ; set Q to low half of 64-bit number
DIV0   GR73,GR71         ; do first step. GR73 & Q become
                        ; 64-bit shift area for divide.
DIV    GR73,GR73,GR72   ; divide step 1
.
.                        ; total of 31 DIV instructions
.

```

```

DIV      GR73,GR73,GR72 ; divide step 31
DIVL    GR73,GR73,GR72 ; last divide step
DIVREM  GR73,GR73,GR72 ; remainder into GR73
MFSR    GR74,Q          ; resultant quotient into GR74
CPLT    GR76,GR74,00#h ; set overflow flag if result is neg
JMPF    SKIP3,,GR76    ; jump if result is positive
CPEQ    GR77,GR77,GR77 ; set GR77 with constant 80000000
CPEQ    GR76,GR77,GR74 ; if result is equal to 80000000 and
                          ; need to negate answer, no overflow
CPNEQ   GR76,GR76,GR75 ; update overflow flag
SKIP3:
JMPF    POS,GR75        ; no correction if neg flag not set
ASEQ    DIVOVRFLOW,GR76,00#h ; if flag set, we have overflow
SUBR    GR74,GR74,00#h ; negate quotient
SUBR    GR73,GR73,00#h ; negate remainder
POS:

```

7.1.8 TRAPPING ARITHMETIC INSTRUCTIONS

The processor does not incorporate logic to directly support floating-point operations, nor does it directly support full multiply and divide operations. However, instructions to perform these operations are included in the instruction set. These instructions are included in the anticipation of future processor implementations which might include hardware to perform these operations.

In applications programs which must be fully object-code compatible with future processor versions—while taking advantage of the performance of future versions—these instructions should be used to perform floating-point, multiplication, and division operations.

In the Am29000, the Floating-Point, MULTIPLY, and DIVIDE instructions simply cause traps. The indirect pointers are set at the time the trap occurs, so that a trap handler can gain access to the operands of the instruction, and can determine where the result is to be stored. The trap handler can directly emulate the execution of the instruction, or can perform the instruction using an external coprocessor.

Note that interfacing to an external arithmetic coprocessor via the trapping arithmetic instructions simplifies the definition of the coprocessor as a system option. If the coprocessor is present, the trap handler uses the coprocessor to perform the arithmetic operations. If the coprocessor is not present, the trap handler emulates the operation by software.

7.1.9 COMPLEMENTING A BOOLEAN

To complement a Boolean in the processor's format, only the most-significant bit of the Boolean word should be considered, since the least-significant 31 bits may or may not be zeros. This is accomplished by the following instruction:

The Boolean is in GR64 in this example. This instruction is based on the observation that a Boolean TRUE is a negative integer, since the Boolean bit coincides with the integer sign bit. If the operand of this instruction is a negative integer (i.e. TRUE), the result is the Boolean FALSE. If the operand is non-negative (i.e. the Boolean FALSE), the result is TRUE.

7.1.10 GENERATING LARGE CONSTANTS

Eight-bit constants are directly available to most instructions. Larger constants must be generated explicitly by instructions and placed into registers before they can be used as operands. The processor has three instructions for the generation of large data constants: Constant (CONST); Constant, High (CONSTH); and Constant, Negative (CONSTN).

The CONST instruction sets the least-significant 16 bits of a register with a field in the instruction; the most-significant 16 bits are zero-bits. This instruction allows a 32-bit, positive constant to be generated with one instruction, when the constant lies in the range of 0 to 65535.

Any 32-bit constant may be generated with a combination of the CONST and CONSTH instructions. The CONSTH instruction sets the most-significant 16 bits of a register with a field in the instruction; the least-significant bits are set to the value of the corresponding bits in a source operand-register. Thus, to create a 32-bit constant in a register, the CONST instruction sets the least-significant 16 bits, and the CONSTH instruction sets the most-significant 16 bits.

The CONSTN instruction sets the least-significant 16 bits of a register with a field in the instruction; the most-significant 16 bits are one-bits. This instruction allows a 32-bit, negative constant to be generated with one instruction, when the constant lies in the range of -65536 to -1.

7.1.11 LARGE JUMP AND CALL RANGES

The 16-bit relative branch-displacement provided by processor instructions is sufficient in the majority of cases. However, addresses with a greater range are occasionally needed. In these cases, the CONST and CONSTH instructions generate the large branch-target address in a register. An indirect jump or call then uses this address to branch to the appropriate location.

When program modules are compiled separately, the compiler cannot determine whether or not the 16-bit displacement of a CALL instruction is sufficient to reach an external procedure, even though it is sufficient in most cases. Instead of generating instructions for the worst case (i.e. the CONST, CONSTH, and CALLI described above), it is more efficient

to generate a CALL as if it were appropriate, with the worst-case sequence (in this case, CONST, CONSTH, and JMPI) also appearing in the generated code somewhere (at the end of a compiled procedure, for example).

When the above scheme is used, the linker is able to determine whether or not the CALL is sufficient. If it is not, the CALL can be re-targeted to the worst-case sequence in the code. In other words, when the CALL is not sufficient, the linker causes the execution sequence to be:

```
*--CALL
|
|
|
*->CONST
    CONSTH
    JMPI
```

In this manner, the longer execution time for the call occurs only when necessary.

7.1.12 NO-OPS

When a NO-OP is required for proper operation (for example, as described in Section 7.3.3), it is important that the selected instruction not perform any operation, regardless of program operating conditions. For example, the NO-OP cannot access general-purpose registers, because a register may be protected from access in some situations. The suggested NO-OP is:

```
ASEQ 40#h,GR1,GR1
```

This instruction asserts that the Stack Pointer (GR1) is equal to itself. Since the assertion is always true, there is no trap. Note also that the Stack Pointer cannot be protected, and that the assert instruction cannot affect any processor state.

7.1.13 CHARACTER-STRING OPERATIONS

The need to perform operations on character strings arises frequently in many systems. The processor provides operations for manipulating character data, but these are frequently inefficient for dealing with character strings, since the processor is optimized for 32-bit data quantities.

It is much more efficient, in general, to perform character-string operations by operating on units of four bytes each. These four-byte units are more suited to the processor's data-flow organization. However, there are several things to be considered when dealing with four-byte units.

Alignment of Bytes Within Words

Character strings are normally not aligned with respect to 32-bit words. Thus, when word operations are used to perform character-string operations, alignment of the character strings must be taken into account.

For example, consider a character string aligned on the third byte of a word which is moved to a destination string aligned on the first byte of a word. If the movement is performed word-at-a-time, rather than byte-at-a-time, the move must involve shift and merge operations, since words in the destination character-string are split across word boundaries in the source character-string.

The processor's Funnel Shifter can be used to perform the alignment operations required when character operations are performed in four-byte units. Though the Funnel Shifter supports general, bit-aligned shift and merge operations, it is easily adapted to byte-aligned operations.

For byte-aligned shift and merge operations, it is only necessary to insure that the two most-significant bits of the Funnel Shift Count (FC) field of the ALU Status Register point to a byte within a word, and that the three least-significant bits of the FC field are 000.

Detection of Characters Within Words

Most character-string operations require the detection of a particular character within the string. For example, the end of a character string is identified by a special character in some character-string representations. In addition, character strings are often searched for a specific pattern. During such searches, the most-frequently executed operation is the search within the character string for the first character of the pattern.

The processor provides a Compare Bytes (CPBYTE) instruction, which directly supports the search for a character within a word. This instruction can provide a factor-of-four performance increase in character-search operations, since it allows a character string to be searched in four-byte units.

During the search, the words containing the character string are compared, a word at a time, to a search key. The search key has the character of interest in every byte position. The CPBYTE instruction then gives a result of TRUE if any character within the character-string word matches a byte in the search key.

7.1.14 MOVEMENT OF LARGE DATA BLOCKS

The movement of large blocks of data—for example, to perform a memory-to-memory move—can be performed by an alternating series of loads and stores. However, it is normally much more efficient to move large blocks of data by using an alternating series of Load Multiple and Store Multiple instructions. These instructions take better advantage of

the data-movement capabilities of the processor, though they require the use of a large number of registers.

During data movement, it is possible to perform alignment operations by a series of EXTRACT instructions between the Load Multiple and Store Multiple. Also, since the Load Multiple and Store Multiple are interruptible, these instructions may be used to move large amounts of data without affecting interrupt latency.

7.2 SYSTEMS-PROGRAMMING CONSIDERATIONS

This section discusses topics of general concern in the implementation of control programs and operating systems.

7.2.1 SYSTEM PROTECTION

The Am29000 provides protection of several different system resources. In general, this protection is based on the value of the Supervisor Mode (SM) bit in the Current Processor Status Register.

Memory Protection

Memory-access protection is provided by the Memory Management Unit. Each Translation Look-Aside Buffer entry in the MMU contains protection bits which determine whether or not a given routine can access the page associated with the entry.

There is a set of protection bits for Supervisor-mode programs, and a separate set for User-mode programs. Thus, for the same virtual page, the access authority of programs executing in the Supervisor mode can be different than the authority of programs executing in User mode.

A Data TLB Protection Violation or Instruction TLB Protection Violation trap occurs if a data or instruction access, respectively, is attempted, but is not allowed because of the value of the protection bits.

Register Protection

General-purpose registers are protected by the Register Bank Protection Register. The Register Bank Protection Register allows parameters for the operating system to be kept in general-purpose registers, protected from corruption by User-mode programs. Additionally, it allows processor registers to be partitioned among multiple tasks.

If a User-mode program attempts to access a protected general-purpose register, a Protection Violation trap occurs. Supervisor-mode programs may access any general-purpose register, regardless of protection.

The special-purpose registers numbered 0 to 127 (though not all are implemented) and all Translation Look-Aside Buffer registers are protected from User-mode accesses. Any attempted access of these registers by a User-mode program causes a Protection Violation trap.

External Access Protection

Other than the protection offered by the Memory Management Unit, the processor provides no specific protection for external devices and memories. However, the SUP/*US output reflects the value of the SM bit during the address cycle of an external access. This can signal external devices and memories to provide protection. Any protection violations can be reported via the *DERR input.

Note that loads and stores to input/output devices are not protected from execution by User-mode programs. Any protection of input/output devices must be performed by external hardware. This allows certain devices to be accessed by User-mode programs, without forcing the overhead of an operating-system call for all devices.

7.2.2 INTERRUPTS AND TRAPS

The Am29000 automatically saves only the Current Processor Status Register when an interrupt or trap is taken; it is saved in the Old Processor Status Register. The processor does not automatically save any other state when an interrupt or trap is taken, but rather freezes the contents of the following registers:

- 1) Program Counters 0, 1, and 2.
- 2) Channel Address, Channel Data, and Channel Control.
- 3) ALU Status.

When these registers are frozen, they are allowed to be updated only by Move To Special Register instructions. The frozen condition is directly controlled by the Freeze (FZ) bit in the Current Processor Status Register.

Since the Channel Address, Channel Data, and Channel Control registers are frozen when an interrupt or trap is taken, the interrupt handler may perform any single-word loads and stores without interfering with the restart state of a channel operation in the interrupted routine. However, load-multiple and store-multiple operations have unpredictable results if performed while the FZ bit is 1, since these operations are sequenced by the Channel Address, Channel Data, and Channel Control registers.

Vector Area

As discussed in Section 3.5.4, interrupts and traps are dispatched through a 256-entry Vector Area, which directs the processor to a routine to handle a given interrupt or trap. Only 64 entries of this area are required for basic processor operation (or 22, if Floating-Point, MULTIPLY and DIVIDE instructions are not used).

The total number of Vector Area entries required is system-dependent, as determined by the vector numbers which are specified in the assert and EMULATE instructions. The number of entries can be restricted to reduce the memory requirements for the Vector Area, especially when the Vector Area is organized as a sequence of 64-instruction blocks. However, there is nothing to prevent an instruction from specifying a vector number in the range 64 to 255. For this reason, it may not be possible to reduce the size of the Vector Area, since erroneous instruction vector numbers might cause unpredictable results.

The Vector Area may be relocated by the Vector Area Base Address Register, and there may be multiple Vector Areas in the system, with the Vector Area Base Address Register pointing to the one which is currently active.

Interrupt Handling

For temporary program interruptions, such as for Translation Look-Aside Buffer reload, the basic processor interrupt mechanism is sufficient to eliminate the need for the interrupt or trap handler to save any state for the interrupted routine. This state may be left in the appropriate registers while the handler executes. An interrupt return returns immediately to the interrupted program.

Besides the direct performance advantage which results from not saving state for temporary program interruptions, there is an additional advantage provided by the processor. When the state of the interrupted routine remains in the appropriate registers, the processor can detect that the Program Counter 0 and Program Counter 1 registers contain sequential addresses. Instead of performing two non-sequential instruction fetches for the interrupt return in this case, the processor initiates only a single non-sequential fetch (the second fetch is performed as a sequential fetch). This reduces the overhead of the interrupt return for these routines.

Note that, when the state of an interrupted program remains in the processor, the processor cannot be enabled to take any further interrupts until an interrupt return is executed. Therefore, this capability should be restricted to time-critical routines, where the execution time of the routine does not interfere with interrupt-latency considerations. (Note that the Interrupt Pending bit of the Current Processor Status Register may be used to detect the presence of external interrupts while these interrupts are disabled).

To support dynamically-nested interrupts and traps, the interrupt or trap handler must save state as necessary for the application, using an appropriate data structure (such as an interrupt stack or program status area). Once the state has been saved (or, alternately, while it is being saved), the handler can load the state for a new program to be executed. An

interrupt return then initiates the execution of the new program.

Interrupt Return

An interrupt return resumes the execution of a program whose processor state is contained in the following registers:

- 1) Old Processor Status.
- 2) Program Counters 0 and 1.
- 3) Channel Address, Channel Data, and Channel Control.

This state is most likely different from the state of the program executing the interrupt return. These registers must be set appropriately before an interrupt return is executed. Note that the instruction sequence which sets these registers must have a Current Processor Status which is equivalent to that of an interrupt or trap handler; the FZ bit must be 1, and interrupts and traps must be disabled.

Simulation of Interrupts and Traps

Assert instructions may be used by a Supervisor-mode program to simulate the occurrence of various interrupts and traps defined for the processor. A Supervisor-mode assert instruction can specify a vector number between 0 and 63. If this instruction causes a trap, the effect is to create an interrupt or trap which is similar to that associated with the specified vector number.

Thus, the interrupt and trap routines defined for basic processor operation can be invoked without creating any particular hardware condition. For example, an *INTR1 interrupt may be simulated by an assert instruction which specifies a vector number of 17, without the activation of the *INTR1 signal.

7.2.3 FAST CONTEXT SWITCHING

The Am29000 allows general-purpose registers to be partitioned among multiple tasks, so that context switching can be very fast. However, in this configuration, fewer registers are available to each task, and the local registers cannot be used as a stack cache (see Section 7.1.1). Thus, task-switch time is minimized at the possible expense of an increase in procedure-call and procedure-execution time. Even so, this trade-off is appropriate in many real-time applications.

The 128 local registers may be partitioned into 8 banks of 16 registers each, with a each bank of registers allocated to one of 8 tasks resident in the processor. Partitioning can be made transparent to resident tasks (except that each task has a small number of registers), because the Stack Pointer can be set so that the first register in each bank is addressed as

Local Register 0. Even when two or more banks of local registers are combined into larger banks, each of the larger banks can still start with Local Register 0. Registers within a given bank can be protected from access by other tasks by the Register Bank Protect Register.

Since the Stack Pointer does not affect the addressing of global registers, global registers cannot be partitioned among multiple tasks. A task cannot be made to reference the proper global registers unless the registers allocated to a given task are known before execution; this restriction is too severe in most cases. Because of this, the global registers should be restricted for use by the operating system, and protected from access by resident tasks. Given this restriction, it is best to use the global registers to contain the processor state of resident tasks.

The processor state which normally must be saved on a task switch consists of the contents of the following registers:

- 1) Old Processor Status
- 2) Channel Address
- 3) Channel Data
- 4) Channel Control
- 5) Program Counter 0
- 6) Program Counter 1
- 7) Q
- 8) ALU Status

Thus, the processor state can be saved in 8 general-purpose registers, and the state for 8 resident tasks can be saved in the 64 global registers. This state is protected from access (by resident tasks) by the Register Bank Protect Register.

In summary, the general-purpose registers can accommodate as many as 8 tasks resident in the processor simultaneously. In this configuration, the global registers contain the processor state for the tasks, and the local registers contain the program state (e.g. variables) for the tasks. Each task has 16 general-purpose registers, numbered 0 to 15, which are used in the normal way.

The above configuration allows a complete context switch to be performed by the adjustment of the Stack Pointer and the movement of processor state to and from general-purpose registers. These operations can be completed within 17 processor cycles.

7.2.4 MEMORY MANAGEMENT

This section discusses various issues involved in memory management as they relate to an operating system. The focus is on virtual-addressing issues.

Virtual Page Size

The MMU Configurations Register determines the size of a virtual page mapped by the Memory Management Unit. The choices for page size are 1, 2, 4, and 8 Kbytes. The selection of page size is based on several considerations:

- 1) For a given page size, any allocation of pages to a process will, on average, waste half of one page. With smaller page sizes, the waste is smaller. In systems with a large number of processes, each with a small amount of memory, small page sizes can reduce waste significantly.
- 2) Smaller page sizes allow finer memory-protection granularity.
- 3) The maximum amount of memory that can be referenced by Translation Look-Aside Buffer (TLB) entries is set by the number of TLB entries and the page size. Larger page sizes allow the fixed number of TLB entries to address more memory, and generally reduce the number of TLB misses. For example, with 1-Kbyte pages, a process requiring 8 Kbytes of contiguous memory would create eight TLB misses; with 8-Kbyte pages, the process would create only one TLB miss.
- 4) The page is usually the unit of memory moved between memory and backing storage. The design of the backing storage sub-system may also influence the choice of page size, because of transfer-efficiency considerations. For example, if the backing storage is a disk, the disk seek time is large compared to transfer time. Thus, it is more efficient to transfer large amounts of data with a single seek. Efficiency may also depend on disk organization (i.e. the number of seeks possibly required to transfer a page).

Page Reference and Change Information

In a demand-paged environment, it is important to be able to collect information on the use and modification of pages. The processor does not collect this information directly, but the information may be collected by the operating system, without requiring hardware support.

Each TLB entry contains 6 bits which specify the type of accesses which are permitted for the corresponding page. When a TLB entry is loaded, the TLB reload routine can set the protection bits so that an access to the corresponding page is not allowed. If an access is attempted, a TLB protection violation traps occurs. This trap may be used to signal that the page is being referenced. After noting this fact, the trap handler may set the protection bits to allow the access, and return to the trapping routine.

A technique similar to the one just described can be used to collect information on the modification of a page. However, in this case, the TLB protection bits are initially set so that a store is not allowed.

It is also possible to create reference information by noting references during TLB reload. For example, reference bits are normally reset periodically, so that they reflect current references. When reference bits are reset, the entire TLB may be invalidated. Reference bits are then set as TLB entries are loaded. Note that this scheme relies on the fact that a TLB miss implies a reference to the corresponding page. Also, this scheme does not account for page change information.

The disadvantage of the above schemes is one of possible performance loss. This is the result of the additional traps required to monitor page references and changes. If the performance impact is unacceptable, references and changes can be easily monitored by hardware which detects reads and writes to page frames in instruction/data memory.

Monitoring Critical Areas of Memory

In certain fault-tolerant systems, it is necessary to detect changes to critical areas of memory, so that these changes may be immediately reflected on a non-volatile storage device. To monitor critical memory areas, the TLB protection bits can be set so that any change to the area causes a Data TLB Protection Violation trap. This trap signals that the area is being modified.

In this use of the protection bits, the trap handler does not set the bits to allow the access. Rather, the trap handler must emulate the access, using the Channel Address, Channel Data, and Channel Control registers. The Contents Valid (CV) bit of the Channel Control Register is reset before the trapping routine is restarted, so that the trap does not re-occur.

TLB Miss Handling

The address translation performed by the MMU is ultimately determined by routines which place entries into the Translation Look-Aside Buffer (TLB). TLB entries are normally based on system page tables, which give the translation for a large number of pages. The TLB simply caches the currently-needed translations, so that system page tables do not have to be accessed for every translation.

If a required address translation cannot be performed by any entry in the TLB, a TLB miss trap occurs. The trap handling routine—called the TLB reload routine—accesses the system page tables to determine the required translation, and sets the appropriate TLB entry. Note that the access requiring this translation can be restarted by the interrupt return at the end of the TLB reload routine (see Section 7.2.5).

A large number of different page-table organizations are possible. Since the TLB reload routine is a sequence of processor instructions, the page tables may have a structure and

access method which satisfies trade-offs of page table size, translation lookup time, and memory-allocation strategies.

Another possibility supported by the TLB reload mechanism is that of a second-level TLB. The TLB reload routine is not required to access the system page tables immediately upon a TLB miss, but may access an external TLB which can be much larger than the processor's TLB. The amount of time required to access the external TLB is normally much smaller than the amount of time required to access the page tables, leading to an overall improvement in performance. Of course, if a translation is not in the external TLB, a page table lookup must still be performed.

Because the TLB reload routine may depend on the type of access causing the TLB miss, the processor differentiates between misses on instruction and data accesses by Supervisor-mode and User-mode programs. This eliminates any time which might be spent by the TLB reload routine in making the same determination. Performance is also enhanced by the LRU Recommendation Register, which gives the TLB register-number for Word 0 of the TLB entry to be replaced by the TLB reload routine (the least-recently-used entry).

Warm Start

When a process switch occurs, there is a high probability that most of the TLB entries of the old process will not be used by the new process. Thus, the new process most likely creates many TLB miss traps early in its execution. This is unavoidable on the first initiation of a process, but may be prevented on subsequent initiations.

When a given process is suspended, the operating system can save a copy of its TLB contents. When the task is restarted, the copy can be loaded back into the TLB. This warm start prevents many of the process' initial TLB misses, at the expense of the time required to save and restore the copy of the TLB entries. However, this time may be much shorter than the time required to perform all TLB reloads individually.

Note that, if this warm-start strategy is adopted, any change in address translation must be reflected in all copies of TLB entries for all affected processes. If address translation is often changed so that it affects more than one process, warm start may not be advantageous.

Minimum Number of Resident Pages

In any processor which supports demand-paging, there is a minimum number of pages which must be resident for any active process. This minimum is determined by the maximum number of pages which might be referenced by an atomic operation in the processor's architecture (e.g. an instruction, normally). If this maximum number is not guaranteed to be resident in memory, some operations might never complete, since they may never have all of the required pages resident in memory at one time.

For the Am29000, two pages are required for a process to make progress through the system. The reason for this requirement is that the Am29000, on interrupt return, restarts an interrupted Load Multiple or Store Multiple only after fetching two instructions (see Section 3.5.5). The first of these instructions must be resident in memory—and mapped by the TLB—and the page required to complete the Load Multiple or Store Multiple must also be resident—and mapped by the TLB—for the interrupt return to complete successfully.

Branch Target Cache Considerations

The Branch Target Cache is accessed with virtual as well as physical addresses, depending on whether address translation is enabled for instruction accesses. Because of this, the Branch Target Cache may contain entries which might be considered valid, even though they are not.

For example, address translation may be changed by a change in the Process Identifier of the MMU Configuration Register. This change is not reflected in the Branch Target Cache tags, so they do not necessarily perform valid comparisons. Also, the Branch Target Cache does not differentiate between virtual and physical addresses, so that it may perform an invalid comparison after address translation for instructions is enabled or disabled.

If a TLB miss occurs during the address translation for a branch target instruction, the processor considers the contents of the Branch Target Cache to be invalid. This is required to properly sequence the LRU Recommendation Register, and does not solve the problem just described. If the TLB is changed at some point, so that the TLB miss does not occur, the Branch Target Cache may still perform an invalid comparison.

To avoid the above problem, the contents of the Branch Target Cache must be explicitly invalidated. This can be accomplished by executing an Invalidate (INV) instruction whenever an address translation is changed. The INV instruction causes all entries of the Branch Target Cache to become invalid (after the next successful branch). However, since the change in address translation rarely affects the program performing the change, the INV may unnecessarily affect the performance of this program.

The IRETINV instruction has the same effect on the Branch Target Cache as the INV instruction, but can reduce the performance impact. The IRETINV delays invalidation until an interrupt return is executed, eliminating the need to disrupt an operating-system routine when it changes address translation. At the point of interrupt return, the contents of the Branch Target Cache are most likely not of much use anyway.

Note that the Branch Target Cache is not invalidated when the Cache Disable (CD) bit of the Configuration Register is set. When the CD bit is 1, the Branch Target Cache continues to operate, but the processor considers its contents to be invalid. Thus, the CD bit cannot be used to invalidate the cache, and, furthermore, the Branch Target Cache may have to be invalidated whenever the CD bit is to be reset (i.e. when the cache is to be enabled).

The Branch Target Cache does not distinguish between virtual and physical addresses, but does distinguish between the instruction/data memory and instruction read-only memory (ROM) address-spaces, and between User-mode and Supervisor-mode addresses. Thus, the Branch Target Cache does not have to be invalidated on transitions between these address-spaces. This improves the performance of applications which make heavy use of ROM-based and/or operating-system routines.

7.2.5 RESTARTING FAULTING EXTERNAL ACCESSES

In a demand-paged system environment, virtual pages and their associated virtual-to-physical mappings are made available to programs on demand. In other words, the memory-management routines generally execute only when a given page or mapping is needed by a program. This need is signalled by a page fault trap caused by a program access (normally, the page fault occurs during a TLB reload).

Since the page fault trap is part of normal system operation, and does not represent an error, the access which causes the trap must be restarted—once the trapping condition is remedied—in a manner that is not detectable to the program causing the trap.

Additionally, in the Am29000, the TLB reload mechanism relies on the ability to restart an access which causes a TLB miss trap. This restart, also, must be accomplished in a manner which cannot be detected by the trapping program.

The Am29000 overlaps external accesses with the execution of instructions. Thus, traps caused by accesses are imprecise: that is, the address of the instruction which initiated the access cannot be determined by the trap handler. Since the address of the initiating instruction is unknown, the access cannot be restarted by re-executing this instruction. Even if the address could be determined, the instruction might not be restartable, since an instruction executed before the trap occurred may have altered the conditions of the access, such as by altering the address source-register.

In order to provide for the restarting of loads and stores which cause exceptions, the processor saves all information required to restart these accesses in the Channel Address, Channel Data, and Channel Control registers. The Contents Valid (CV) and Not Needed (NN) bits in the Channel Control Register indicate that the information contained in these registers represents an access which must be restarted. The CV bit indicates that the access did not complete, and the NN bit indicates whether or not the data from the access is required by the processor.

Note that, since instruction execution is overlapped with external accesses, an instruction which executes after a load may alter the destination-register for the load. If a trap occurs in this situation, the access information in the Channel Address, Data, and Control registers is correct, but the load cannot be restarted. The NN bit provides correct operation in this case.

When an interrupt or trap is taken, the handling routine has access to the Channel Address, Data, and Control registers; the contents of these registers may contain information relevant to an incomplete access, and can be preserved for restarting this access. Note that, since these registers are frozen (due to the FZ bit of the Current Processor Status), they are not available to monitor any external accesses in the interrupt or trap handler until their contents are saved, and the FZ bit is reset.

The processor restarts an access, using the Channel Address, Channel Data, and Channel Control registers, upon an interrupt return (IRET or IRETINV). The access is initiated if the CV bit of the Channel Control Register is 1, and the NN bit is 0. The restart cannot be detected in the logical operation of the restarted routine, although the timing of its execution is altered.

The mechanism used to restart faulting accesses has the additional benefit of allowing a fast interrupt-response time when the processor is performing a load-multiple or store-multiple operation. Interrupted load-multiple and store-multiple operations are restarted as if they had faulted. In this case, the operation resumes from the point of interruption, not the beginning of the sequence.

7.2.6 MULTI-PROCESSING

The Am29000 provides several facilities for the implementation of multi-programming and multi-processing systems. These facilities help provide mutual exclusion, synchronization, and communication between multiple processes, whether these processes execute on a single processor or multiple processors.

Binary semaphores are supported by the Load and Set (LOADSET) instruction. This instruction loads the contents of an external location into a register and atomically sets the contents of the location to the integer -1. This instruction requires no special hardware support in the system, since all sequencing is performed by the processor. Also, the LOADSET is available to User-mode programs. This eliminates the overhead of an operating-system call in the use of binary semaphores.

The instructions Load and Lock (LOADL) and Store and Lock (STOREL) support the locking of external devices and memories, or the locking of particular locations within an external device or memory. This prevents access by any process or processor other than the one which performed the lock, and provides the flexibility of locking in a manner appropriate to the system and application. The LOADL and STOREL instructions are available to User-mode programs.

To indicate that a LOADL or STOREL is being executed, the processor asserts the *LOCK output during the external access. Since the processor cannot directly control the behavior of external devices and memories, system hardware must support locking, if required.

Note also that the protocol for the locking and unlocking of devices and memories must be defined by the system. For example, the protocol may be defined such that a LOADL locks the device or memory, and a STOREL unlocks the device or memory. Between the execution of the LOADL and the STOREL, the device can be accessed by the locking process, with any combination of normal loads and stores.

For the implementation of a general-purpose exclusion, synchronization, and/or communication scheme, the processor allows Supervisor-mode programs to set the Lock (LK) bit in the Current Processor Status. This bit activates the *LOCK pin, and prevents the processor from relinquishing the channel to another channel master. (If another master already has control of the channel when the LK bit is set, the LK bit does not take effect until control of the channel is returned to the processor).

The LK bit allows a Supervisor-mode program to execute with mutual exclusion for any sequence of instructions. However, because interrupts must also be disabled for true exclusion, this may have a negative impact on system performance if used improperly.

7.2.7 TIMER FACILITY

The processor has a built-in Timer Facility which can be configured to cause periodic interrupts. The Timer Facility consists of 2 special-purpose registers—the Timer Counter and the Timer Reload registers—which are accessible only to Supervisor-mode programs. These registers implement timing functions independent of program execution.

Timer Facility Operation

The Timer Counter Register has a 24-bit Timer Count Value (TCV) field which decrements by one on every processor cycle. If the TCV field decrements to zero, it is written with the Timer Reload Value (TRV) field of the Timer Reload Register on the next cycle; the Interrupt (IN) bit of the Timer Reload register is set at the same time. The reloading of the TCV field by the TRV field maintains the accuracy of the Timer Facility.

The Timer Reload Register contains the 24-bit TRV field and the control bits Overflow (OV), Interrupt (IN), and Interrupt Enable (IE). The TCV field and IN bit were described above. If the IN bit is 1 and the IE bit also 1, a Timer interrupt occurs. If the IN bit is 1 when the TCV field decrements to zero, the OV bit is also set. The OV bit indicates that a Timer interrupt may have occurred before a previous interrupt was serviced.

Timer Facility Initialization

To initialize the Timer Facility, the following steps should be taken in the specified order (it is assumed that Timer interrupts are disabled by the DA bit of the Current Processor Status Register during the following steps):

- 1) Set the TCV field with the desired interval count for the first timing interval. Note that this interval must be sufficiently large to allow the execution of the next step before the TCV field decrements to zero (this is normally the case).
- 2) Set the TRV field with the desired interval count for the second timing interval. The OV and IN bits are reset, and the IE bit is set as desired. Note that the second timing interval may be equivalent to the first timing interval.

Handling Timer Interrupts

The following is a suggested list of actions to be taken to handle a Timer interrupt:

- 1) Read the Timer Reload register into a general-purpose register.
- 2) Reset the IN bit in the general-purpose register.
- 3) Set the TRV field in the general-purpose register to the desired value for the next timing interval. Note that, at this time, the Timer Counter is timing the current interval. Also, this step may be omitted, if all intervals are equivalent.
- 4) Write the contents of the general-purpose register back into the Timer Reload register.
- 5) Test the general-purpose-register copy of the OV bit, and, if it is set, report the error as appropriate.
- 6) Perform any system operations required for the Timer interrupt.
- 7) Execute an interrupt return.

Timer Facility Uses

Since the Timer Facility has a resolution of a single processor cycle, it may be used to perform precise timing of system events. For example, it may be used to determine an exact measurement of the number of cycles between two events in the system, or to perform precise, time-critical control functions. Note that the Timer interrupt is enabled and disabled separately from other processor interrupts, so that its priority can be separately specified.

The Timer Facility can be used to generate time intervals for collecting virtual page usage information (see Section 7.2.4). For example, if memory management relies on a working-set page-replacement algorithm, the Timer Facility can establish the working set window.

The Timer Facility can be shared among multiple processes. This sharing is accomplished by the implementation of a queue for timer events, which are sorted in order of increasing event time. On each occurrence of a Timer interrupt, the TRV field is set for the interval

between the next two events in the queue, while the Timer Counter Register is counting the current interval (because of a previous setting of the TRV field). The event at the beginning of the queue identifies other system actions to be taken for the Timer interrupt. This event is removed from the queue after the appropriate actions are taken.

7.2.8 TRACE FACILITY

Software debug is supported by the Trace Facility. The Trace Facility guarantees exactly one trap after the execution of any instruction in a program being tested. This allows a debug routine to follow the execution of instructions, and to determine the state of the processor and system at the end of each instruction.

Tracing is controlled by the Trace Enable (TE) and Trace Pending (TP) bits of the Current Processor Status. The value of the TE bit is always copied into the TP bit when an instruction enters the write-back stage. A Trace trap occurs whenever the TP bit is 1. As with most traps, the Trace trap can be disabled only by the DA bit of the Current Processor Status.

In order to trace the execution of a program, the debug routine performs an interrupt return to cause the program to begin or resume execution. However, before the interrupt return is executed, the TE and TP bits of the Old Processor Status are set with the values 1 and 0, respectively. The interrupt return causes these bits to be copied into the TE and TP bits of the Current Processor Status.

When the target of the interrupt return (whose address is contained in the Program Counter 1 Register when the interrupt return is executed) enters the write-back stage, the processor copies the value of the TE bit into the TP bit. Since the TP bit is a 1, a Trace trap occurs. This trap prevents any further instruction execution in the target routine until the interrupt is taken and the routine is resumed with an interrupt return. When the Trace trap is taken, the TE and TP bits are both reset automatically, preventing any further Trace traps.

Since the Trace Facility is managed by the Old and Current Processor Status registers, it operates properly in the event that the processor takes an interrupt or trap—which is unrelated to the Trace Facility—before the above trace sequence completes. When the unrelated interrupt or trap is taken, the state of the Trace Facility (i.e. the values of the TE and TP bits) is copied into the Old Processor Status from the Current Processor Status. The Trace Facility then resumes operation when the interrupted routine is restarted by an interrupt return.

Note that it is possible to cause a Trace trap by directly setting the TP and/or TE bits in the Current Processor Status Register. This may be accomplished only by a Supervisor-mode program.

7.3 PIPELINE FEATURES EXPOSED TO SOFTWARE

In certain cases, the Am29000 pipeline is exposed during instruction execution, in that the execution of certain instructions are dependent on the execution of previous instructions. This section discusses the cases where the pipeline is exposed to software, and the resulting effect on instruction execution.

7.3.1 DELAYED BRANCH

The effect of jump and call instructions is delayed by one cycle, to allow the processor pipeline to achieve maximum throughput. When one of these branches is successful, the instruction immediately following the jump or call is executed before the target instruction of the jump or call is executed. Jump and call instructions are collectively referred to as delayed branches, and the immediately-following instruction is called the delay instruction.

For example, in the following code fragment:

```

      .
      .
      .
      CPEQ      GR88,GR96,GR97   (1)
      JMPF     LBL,GR88         (2)
      SUB      GR96,GR96,01#h   (3)
      CONST   GR96,00#h        (4)
      .
      .
LBL:   CALL    SORT,LR0         (5)
      ADD     LR2,LR5,00#h     (6)
      CPNEQ   LR3,LR2,00#h     (7)
      .
      .
      .
```

The SUB instruction (3) is executed regardless of the outcome of the JMPF instruction (2). Of course, if the JMPF is not successful, the CONST instruction (4) is also executed. If the JMPF is successful, then the instruction sequence is: (3), (5), (6), and then the first instruction of the SORT procedure. Note that the CALL instruction (5) is also a delayed branch, so the instruction immediately following it, (6), is always executed. After the SORT procedure executes the return sequence, the CPNEQ instruction (7) is the next instruction executed.

The benefit of delayed branches is improved performance and a simplified processor implementation. Performance is improved because the processor pipeline executes useful instructions in a larger number of cycles, compared to an implementation without delayed branches.

For example, ignoring all other effects on performance, and assuming that 15% of all instructions are branches, then a processor without delayed branches would take at least 2 cycles for 15% of its instructions, leading to $0.85(1) + 0.15(2) = 1.15$ cycles per instruction, on average. This represents a 15% performance degradation compared to a processor with delayed branches (assuming, for this simple example, that the delay instruction is always useful).

The cost of having delayed branches is either the extra effort required when the compiler takes advantage of delayed branches (by re-organizing code), or the extra NO-OP instruction which the compiler inserts after every branch to guarantee correct program operation. Since the compiler expends only a small amount of effort to avoid wasting time and space with NO-OPs, and since the performance improvement resulting from this effort is significant, delayed branches are beneficial overall.

When two immediately-adjacent branches are taken, the target of the first branch preempts execution of the delay cycle of the second branch, and the target of the second branch then follows the target of the first branch. For example, in the following code fragment:

```

      .
      .
      .
      JMP    L1                (1)
      JMP    L2                (2)
      ADD    GR68,GR68,01#h    (3)
      .
      .
L1:   SUB    GR77,GR77,01#h    (4)
      SUBC   GR68,GR68,00#h    (5)
      .
      .
L2:   CONST  GR77,ff0f#h      (6)
      SUBR   GR68,GR68,01#h    (7)
      OR     GR77,GR77,GR68    (8)
      .
      .
      .

```

an unconditional JMP instruction (1) is followed immediately by another unconditional JMP instruction (2). (In this example, unconditional JMPs are used; however, any two immediately-adjacent taken branches exhibit the same behavior.) The sequence of executed instructions in this case is: JMP instruction (1), JMP instruction (2), SUB instruction (4), CONST instruction (6), SUBR instruction (7), OR instruction (8), and so on. Note that the ADD instruction (3) is not executed. Also, the target of the first JMP instruction (1) was merely visited; control did not continue sequentially from L1 but rather continued from L2.

7.3.2 OVERLAPPED LOADS AND STORES

The processor allows an external access to be overlapped with instruction execution. This means that, while the access is being performed, the processor continues to execute instructions, as long as the instructions and data required for execution are available.

In order to make full use of overlapped storage accesses, some instruction reorganization may be necessary. For example, in the following sequence:

```
LOOP: .  
      .  
      .  
      CONST   LR3, #ARRAYBASE      (1)  
      ADD     LR5, LR3, LR4        (2)  
      LOAD    LR6, LR5             (3)  
      ADD     LR7, LR6, LR7        (4)  
      SUB     LR4, LR4, 01#h       (5)  
      CPEQ    LR8, LR4, 00#h       (6)  
      JMPF    LOOP, LR8            (7)  
      ASEQ    40#h, GR1, GR1       (8)  
      .  
      .  
      .
```

the ADD instruction (4) uses the result of the LOAD instruction (3). However, the following four instructions do not depend on the result of the LOAD. Therefore, the ADD instruction (4) can be moved past the JMPF (7)—since it will always be executed even if the JMPF is taken—and replace the ASEQ instruction (8), used as a NO-OP. The resulting sequence is:

```
LOOP: .  
      .  
      .  
      CONST   LR3, #ARRAYBASE      (1)  
      ADD     LR5, LR3, LR4        (2)  
      LOAD    LR6, LR5             (3)  
      SUB     LR4, LR4, 01#h       (4)  
      CPEQ    LR8, LR4, 00#h       (5)  
      JMPF    LOOP, LR8            (6)  
      ADD     LR7, LR6, LR7        (7)  
      .  
      .  
      .
```

The instructions (4) through (6) are likely to be executed while external memory satisfies the load request, resulting in improved throughput. The processor thus allows parallelism to be exploited by instruction reordering.

The overlapped load feature may be used to improve processor performance, but imposes no constraints on instruction sequences, as delayed branches do. The processor implements the proper pipeline interlocks to make this parallelism transparent to a running program.

7.3.3 DELAYED EFFECTS OF REGISTERS

The modification of some registers has a delayed effect on processor behavior, because of the processor pipeline. The affected registers are the Stack Pointer (Global Register 1), Indirect Pointers A, B, and C, and the MMU Configuration Register.

An instruction which writes to the Stack Pointer can be followed immediately by an instruction which reads the Stack Pointer. However, any instruction which references a local register also uses the value of the Stack Pointer to calculate an absolute register-number.

At least one cycle of delay must separate an instruction which updates the Stack Pointer and an instruction which references a local register. In most systems, this affects procedure call and return only (see Section 7.1.1). In general, though, an instruction which immediately follows a change to the Stack Pointer should not reference a local register (however, note that this restriction does not apply to a reference of a local register via an indirect pointer).

The indirect pointers have an implementation similar to the Stack Pointer, and exhibit similar behavior. At least one cycle of delay must separate an instruction which modifies an indirect pointer and an instruction which uses that indirect pointer to access a register.

At least one cycle of delay must separate a Move To Special Register which modifies the Page Size (PS) field of the MMU Configuration Register and an instruction which performs address translation. The latter instruction includes successful branches, loads, and stores.

Note that it is normally not possible to guarantee that the delayed effect of the Stack Pointer and indirect pointers is visible to a program. If an interrupt or trap is taken immediately after one of these registers is set, then the interrupted routine sees the effect in the following instruction, because many cycles elapse between the two instructions. For this reason, a program should not be written in a manner which relies on the delayed effect; the results of this practice may be unpredictable.

CHAPTER 8

INSTRUCTION SET

This chapter provides a specification of the Am29000 instruction set. Sections 8.1 through 8.3 describe the terminology used, the setting of the ALU Status Register by instructions, and the instruction formats. Section 8.4 describes each instruction in detail; instructions are presented alphabetically, by assembler mnemonic. Finally, Section 8.5 gives an index of instructions by operation code.

8.1 INSTRUCTION-DESCRIPTION NOMENCLATURE

In order to simplify the specification of the instruction set, special terminology is used throughout this chapter. This section defines the terminology and symbols used to describe instruction operands, operations, and the assembly-language syntax.

This section does not describe all terminology used. It excludes certain descriptive terms which have an obvious meaning.

8.1.1 OPERAND NOTATION AND SYMBOLS

Throughout this chapter, instruction operands are signed, two's-complement, word integers, unless otherwise noted. The term "register" is used consistently to denote a general-purpose register; other types of registers are explicitly described.

The following notation is used in the description of instruction operands:

0I16	16-bit immediate data, zero-extended to 32 bits.
1I16	16-bit immediate data, ones-extended to 32 bits
BP	The Byte Pointer (BP) field of the ALU Status Register. The BP field selects a byte or half-word within a word, and is interpreted according to the Byte Order bit of the Configuration Register.
C	The Carry (C) bit of the ALU Status Register. The C bit is logically zero-extended to 32 bits when it is involved in a word operation.
COUNT	The value of the Count Remaining field of the Channel Control Register. Note that COUNT does not refer to this field directly, but rather to the value of the field at the beginning of a LOADM or STOREM instruction.

DEST	The general-purpose register which is the destination of an instruction, i.e. the register used to store the result.
EXTERNAL WORD[n]	The word in an external device or memory with address n. This terminology is also used for coprocessor words, except that the address n either has no pre-defined interpretation or is a data item transferred to the coprocessor.
FALSE	The Boolean constant FALSE.
FC	The Funnel Shift Count (FC) field of the ALU Status Register.
h'n'	The hexadecimal constant n.
I16	16-bit immediate data.
IPA	Indirect Pointer A Register.
IPB	Indirect Pointer B Register.
IPC	Indirect Pointer C Register.
PC	The Program Counter Register. This register is not explicitly accessible by instructions, but does appear as an operand for certain instructions. The Program Counter always contains the word address of the instruction being executed, and is 30 bits in length.
Q	The Q Register.
register RA register RB register RC	These designate the general-purpose registers specified by the instruction fields RA, RB, and RC (see Section 8.3).
SPDEST	The special-purpose register which is the destination of an instruction.
SPECIAL	The content of a special-purpose register, used as an instruction operand.
special- purpose register SA	Designates the special-purpose register specified by the instruction field SA (see Section 8.3).

SRCA	
SRCB	The contents of general-purpose registers, used as instruction operands.
SRCA.BYTE _n	
SRCB.BYTE _n	Designate the byte numbered n within the SRCA or SRCB operand.
TARGET	The target-instruction address specified by a jump or call instruction. This address is either absolute, or Program-Counter relative.
TLB[n]	The Translation Look-aside Buffer Register with register-number n.
TRUE	The Boolean constant TRUE.
TWIN	The twin of a general-purpose register is the odd-numbered register whose register is one greater than the register number for a given even-numbered register. For example, Local Register 5 is the twin of Local Register 4.

8.1.2 OPERATOR SYMBOLS

The following symbols are used to describe instruction operations:

A << B	Left shift of the A operand by the shift amount given by the B operand.
A >> B	Right shift of the A operand by the shift amount given by the B operand.
A // B	Concatentation. The B operand is appended to the A operand. In the resulting quantity, the A operand makes up the high-order part, and the B operand makes up the low-order part.
A & B	Bitwise AND.
A B	Bitwise OR.
A ^ B	Bitwise exclusive-OR.
~ A	One's-complement.
A ← exp	Assignment of the A location by the result of the expression on the right side.
A = B	Equal to.
A <> B	Not equal to.

$A > B$	Greater than.
$A \geq B$	Greater than or equal to.
$A < B$	Less than.
$A \leq B$	Less than or equal to.
$A + B$	Addition.
$A - B$	Subtraction.
$A \cdot B$	Multiplication.
A / B	Division.
$A..B$	A subrange which includes the A operand and the B operand. This symbol is used for subranges of bits as well as subranges of words.
$A \text{ OR } B$	Logical OR of two Boolean conditions.

8.1.3 CONTROL-FLOW TERMINOLOGY

The following terminology is used to describe the control functions performed during the execution of various instructions:

Continue	Continue execution of the current instruction sequence.
IF condition THEN operations ELSE operations	The condition following the IF is tested. If the condition holds, the operations following the THEN are performed. If the condition does not hold, the operations following the ELSE are performed. If the ELSE is not present and the condition does not hold, no operation is performed.
signed overflow	This condition is present when the result of an add or subtract of two's-complement operands cannot be represented by a signed, word integer.
Trap(n)	Specifies a trap with vector number n. The vector number n may be specified indirectly, e.g. Trap (VN), or explicitly by symbolic name, e.g. Trap (Out of Range).

unsigned overflow	This condition is present when the result of an add of unsigned operands cannot be represented by an unsigned, word integer.
unsigned underflow	This condition is present when the result of a subtract of unsigned operands cannot be represented by an unsigned integer, i.e. when the result is less than zero.
VN	Designates the trap vector number specified by the instruction field VN (see Section 8.3).

8.1.4 ASSEMBLER SYNTAX

This chapter does not contain a full description of the instruction assembler, but provides a rudimentary description of the assembler syntax. The following notation is used to describe assembler tokens:

ce	Determines the Coprocessor Enable (CE) bit of a load or store instruction.
cntl	Determines the 7-bit control field in a load or store instruction.
const8	Specifies a constant which can be expressed by 8 bits.
const16	Specifies a constant which can be expressed by 16 bits.
ra rb rc	These tokens name general-purpose registers. In a formal sense, these represent the same token, since the name of a register does not depend on its instruction use. However, three distinct tokens are used to clarify the relationship between the assembler syntax, instruction operands, and instruction fields.
spid	A symbolic identifier for a special-purpose register.
target	A symbolic label for the target of a jump or call instruction.
vn	Specifies a trap vector number.

8.2 ARITHMETIC/LOGIC STATUS RESULTS OF INSTRUCTIONS

8.2.1 ARITHMETIC/LOGIC STATUS BITS

The arithmetic/logic status bits of the ALU Status Register are:

V	Overflow
N	Negative
Z	Zero
C	Carry

The C bit is used in extended arithmetic operations (i.e. on operands greater than 32 bits in length), and the N bit is used in divide step operations. Other than these uses, the status bits are not involved in instruction operations. In particular, they are not used to determine the outcome of conditional jump instructions: Boolean values in registers are used instead for this purpose. The status bits are primarily informational.

Except for instructions which explicitly modify the ALU Status Register, the status bits are modified only by the execution of instructions in the Arithmetic and Logical classes. The Arithmetic and Logical instructions affect the status bits differently. The following two sections describe the setting of the status bits by Arithmetic and Logical instructions.

When the Freeze (FZ) bit of the Current Processor Status Register is 1, the ALU Status Register is not modified except by the Move To Special Register instruction.

8.2.2 ARITHMETIC OPERATION STATUS RESULTS

The Arithmetic instructions modify the V, N, Z, and C bits. These bits are set according to the result of the operation performed by the instruction.

All instructions in the Arithmetic class—except for MULTIPLY and DIVIDE—perform an add. In the case of subtraction, the subtract is performed by adding the two's-complement or one's-complement of an operand to the other operand. The multiply step and divide step operations also perform adds, again possibly complementing one of the operands before the operation is performed. In general, the status bits are based on the results of the add.

If two's complement overflow occurs during the add, the V bit of the ALU Status Register is set; otherwise it is reset. Two's complement overflow occurs when the carry-in to the most-significant bit of the intermediate result differs from the carry-out. When this occurs, the result cannot be represented by a signed, word integer. Note that the V bit is always set in this manner, even when the result is unsigned.

The N bit of the ALU Status Register is set to the value of the most-significant bit of the result of the add. Note that the divide step and multiply step operations may shift the result after the operation is performed. In the cases where shifting occurs, the N bit may not agree with the result which is written into a general-purpose register, since the N bit is based only on the result of the add, not on the shift.

If the result of the add causes a zero word to be written to a general-purpose register, the Z bit of the ALU Status Register is set; otherwise, it is reset. Note that the Z bit always reflects the result written into a general-purpose register; if shifting is performed by a multiply or divide step, the Z bit reflects the shifted value.

If there is a carry out of the add operation, the C bit is set; otherwise it is reset.

Correcting Out-of-Range Results

Some Arithmetic instructions cause an Out of Range trap if the arithmetic operation causes an overflow or underflow. When an Out of Range trap occurs, the result of the operation—though incorrect—is written into the destination register.

Furthermore, the Program Counter 2 Register contains the address of the trapping instruction, and the ALU Status Register contains an indication of the cause of the trap. It is possible, if required, for the trap handler to use this information to form the correct result.

The ALU Status indicates the cause of the Out of Range trap, based on the operation performed, as follows:

- 1) Signed overflow. If the Out of Range trap is caused by signed, two's-complement overflow (this can occur for both signed adds and subtracts), the V bit is 1.
- 2) Unsigned overflow. If the Out of Range trap is caused by unsigned overflow (this can occur only for unsigned adds), the C bit is 1.
- 3) Unsigned underflow. If the Out of Range trap is caused by unsigned underflow (this can occur only for unsigned subtracts), the C bit is 0.

8.2.3 LOGICAL OPERATION STATUS RESULTS

The Logical instructions modify the N and Z bits. These bits are set according the result of the instruction. The V and C bits are meaningless in regard to the logical instructions, so they are not modified.

The N bit of the ALU Status Register is set to the value of the most-significant bit of the result of the logical operation.

If the result of the logical operation is a zero word, the Z bit of the ALU Status Register is set; otherwise, it is reset.

8.3 INSTRUCTION FORMATS

All instructions for the Am29000 are 32 bits in length, and are divided into four fields, as shown in Figure 8-1. These fields have several alternative definitions, as discussed below. In certain instructions, one or more fields are not used, and are reserved for future use. Even though they have no effect on processor operation, bits in reserved fields should be 0, to insure compatibility with future processor versions.




Figure 8-1. Instruction Format

The instruction fields are defined as follows:

BITS 31-24

OP This field contains an operation code, defining the operation to be performed. In some instructions, the least-significant bit of the operation code selects between two possible operands. For this reason, the least-significant bit is sometimes labelled "A" or "M", with the following interpretations:

A (Absolute) : The A bit is used to differentiate between Program-Counter relative (A = 0) and absolute (A = 1) instruction addresses, when these addresses appear within instructions.

M (IMmediate) : The M bit selects between a register operand (M = 0) and an immediate operand (M = 1), when the alternative is allowed by an instruction.

BITS 23-16

RC	The RC field contains a global or local register-number.
I17..I10	This field contains the most-significant 8 bits of a 16-bit instruction address. This is a word address, and may be Program-Counter relative or absolute, depending on the A bit of the operation code.
I15..I8	This field contains the most-significant 8 bits of a 16-bit instruction constant.
VN	This field contains an 8-bit trap vector number.
CE//CNTL	This field controls a load or store access, as described in Sections 3.4.2 and 6.1.2.

BITS 15-8


RA	The RA field contains a global or local register-number.
SA	The SA field contains a special-purpose register-number.

BITS 7-0


RB	The RB field contains a global or local register-number.
RB or I	This field contains either a global or local register-number, or an 8-bit instruction constant, depending on the value of the M bit of the operation code.
I9..I2	This field contains the least-significant 8 bits of a 16-bit instruction address. This is a word address, and may be Program-Counter relative, or absolute, depending on the A bit of the operation code.
I7..I0	This field contains the least-significant 8 bits of a 16-bit instruction constant.

The fields described above may appear in many combinations. However, certain combinations which appear frequently are shown in Figure 8-2.


Three operands, with possible 8-bit constant:




Three operands, without constant:




One register operand, with 16-bit constant:




Jumps and calls with 16-bit instruction address:



Two operands with trap vector number:



Loads and stores:



CE

08996A 8-2A

Figure 8-2. Frequently Occurring Instruction-Field Uses

8.4 INSTRUCTION DESCRIPTION

This section describes each Am29000 Instruction in detail. Figure 8-3 illustrates the layout of the information given for each description.

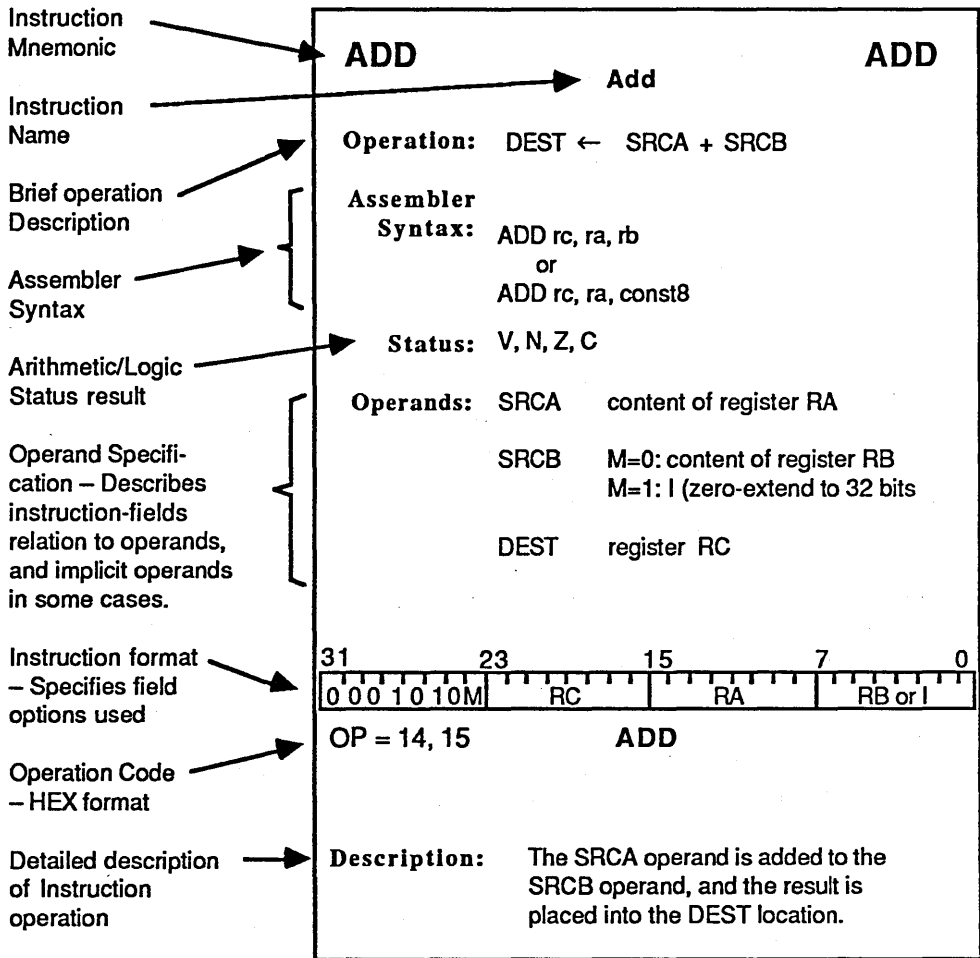


Figure 8-3. Instruction-Description Format

ADDC

ADDC

Add with Carry


Operation: $DEST \leftarrow SRCA + SRCB + C$

Assembler

Syntax: `ADDC rc, ra, rb`
or
`ADDC rc, ra, const8`

Status: V, N, Z, C

Operands: `SRCA` content of register RA
`SRCB` M = 0 : content of register RB
M = 1 : 1 (zero-extended to 32 bits)
`DEST` register RC



Description: The SRCA operand is added to the SRCB operand and the value of the ALU Status Carry bit, and the result is placed into the DEST location.

ADDCS

ADDCS

Add with Carry, Signed

Operation: $DEST \leftarrow SRCA + SRCB + C,$
IF signed overflow THEN Trap (Out of Range)

Assembler


Syntax: ADDCS rc, ra, rb
 or
 ADDCS rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA

 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)

 DEST register RC



OP = 18, 19

ADDCS

Description: The SRCA operand is added to the SRCB operand and the value of the ALU Status Carry bit, and the result is placed into the DEST location. If the add operation causes a two's-complement signed overflow, an Out of Range trap occurs.

Note that the DEST location is altered whether or not an overflow occurs.

ADDCU

ADDCU

Add with Carry, Unsigned

Operation: $DEST \leftarrow SRCA + SRCB + C$,
IF unsigned overflow THEN Trap (Out of Range)

Assembler


Syntax: `ADDCU rc, ra, rb`
 or
`ADDCU rc, ra, const8`

Status: V, N, Z, C

Operands: `SRCA` content of register RA

`SRCB` M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)

`DEST` register RC



Description: The SRCA operand is added to the SRCB operand and the value of the ALU Status Carry bit, and the result is placed into the DEST location. If the add operation causes an unsigned overflow, an Out of Range trap occurs.

Note that the DEST location is altered whether or not an overflow occurs.

ADDS

ADDS

Add, Signed

Operation: DEST ← SRCA + SRCB
 IF signed overflow THEN Trap (Out of Range)

Assembler


Syntax: ADDS rc, ra, rb
 or
 ADDS rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA

 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)

 DEST register RC



Description: The SRCA operand is added to the SRCB operand, and the result is placed into the DEST location. If the add operation causes a two's-complement signed overflow, an Out of Range trap occurs.

Note that the DEST location is altered whether or not an overflow occurs.

ADDU

ADDU

Add, Unsigned


Operation: $DEST \leftarrow SRCA + SRCB$
IF unsigned overflow THEN Trap (Out of Range)

Assembler

Syntax: ADDU rc, ra, rb
 or
 ADDU rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA
 SRCB M = 0 : content of register RB
 M = 1 : 1 (zero-extended to 32 bits)
 DEST register RC



OP = 12, 13

ADDU

Description: The SRCA operand is added to the SRCB operand, and the result is placed into the DEST location. If the add operation causes an unsigned overflow, an Out of Range trap occurs.

Note that the DEST location is altered whether or not an overflow occurs.

AND

AND

AND Logical


Operation: DEST ← SRCA & SRCB

Assembler

Syntax: AND rc, ra, rb
 or
 AND rc, ra, const8

Status: N, Z

Operands: SRCA content of register RA
 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)
 DEST register RC



Description: The SRCA operand is logically ANDed, bit-by-bit, with the SRCB operand, and the result is placed into the DEST location.

ANDN

ANDN

AND-NOT Logical


Operation: DEST ← SRCA & ~SRCB

Assembler

Syntax: ANDN rc, ra, rb
 or
 ANDN rc, ra, const8

Status: N, Z

Operands: SRCA content of register RA
 SRCB M = 0 : content of register RB
 M = 1 : 1 (zero-extended to 32 bits)
 DEST register RC



OP = 9C, 9D

ANDN

Description: The SRCA operand is logically ANDed, bit-by-bit, with the one's-complement of the SRCB operand, and the result is placed into the DEST location.

ASEQ

ASEQ

Assert Equal To


Operation: IF SRCA = SRCB THEN Continue
 ELSE Trap (VN)

Assembler

Syntax: ASEQ vn, ra, rb
 or
 ASEQ vn, ra, const8

Status: Not affected

Operands: SRCA content of register RA
 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)
 VN Trap vector number



Description: If the SRCA operand is equal to the SRCB operand, instruction execution continues; otherwise, a trap with the specified vector number occurs.

For programs in the User mode, a Protection Violation trap occurs—instead of the assert trap—if a vector number between 0 and 63 is specified and the assert condition is not met.

ASGE

ASGE

Assert Greater Than or Equal To

Operation: IF SRCA >= SRCB THEN Continue
ELSE Trap (VN)

Assembler


Syntax: ASGE vn, ra, rb
 or
 ASGE vn, ra, const8

Status: Not affected

Operands: SRCA content of register RA

 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)

 VN Trap vector number



Description: If the value of the SRCA operand is greater than or equal to the value of the SRCB operand, instruction execution continues; otherwise, a trap with the specified vector number occurs.

For programs in the User mode, a Protection Violation trap occurs—instead of the assert trap—if a vector number between 0 and 63 is specified and the assert condition is not met.

ASGEU

ASGEU

Assert Greater Than or Equal To, Unsigned


Operation: IF SRCA >= SRCB (unsigned) THEN Continue
 ELSE Trap (VN)

Assembler

Syntax: ASGEU vn, ra, rb
 or
 ASGEU vn, ra, const8

Status: Not affected

Operands: SRCA content of register RA
 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)
 VN Trap vector number



OP = 5E, 5F

ASGEU

Description: If the value of the SRCA operand is greater than or equal to the value of the SRCB operand, instruction execution continues; otherwise, a trap with the specified vector number occurs. For the comparison, both operands are treated as unsigned integers.

For programs in the User mode, a Protection Violation trap occurs—instead of the assert trap—if a vector number between 0 and 63 is specified and the assert condition is not met.

ASGT

ASGT

Assert Greater Than

Operation: IF SRCA > SRCB THEN Continue
ELSE Trap (VN)

Assembler


Syntax: ASGT vn, ra, rb
 or
 ASGT vn, ra, const8

Status: Not affected

Operands: SRCA content of register RA

 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)

 VN Trap vector number



Description: If the value of the SRCA operand is greater than the value of the SRCB operand, instruction execution continues; otherwise, a trap with the specified vector number occurs.

For programs in the User mode, a Protection Violation trap occurs—instead of the assert trap—if a vector number between 0 and 63 is specified and the assert condition is not met.

ASGTU

ASGTU

Assert Greater Than, Unsigned

Operation: IF SRCA > SRCB (unsigned) THEN Continue
ELSE Trap (VN)

Assembler


Syntax: ASGTU vn, ra, rb
 or
 ASGTU vn, ra, const8

Status: Not affected

Operands: SRCA content of register RA

 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)

 VN Trap vector number



OP = 5A, 5B

ASGTU

Description: If the value of the SRCA operand is greater than the value of the SRCB operand, instruction execution continues; otherwise, a trap with the specified vector number occurs. For the comparison, both operands are treated as unsigned integers.

For programs in the User mode, a Protection Violation trap occurs—instead of the assert trap—if a vector number between 0 and 63 is specified and the assert condition is not met.

ASLE

ASLE

Assert Less Than or Equal To

Operation: IF SRCA <= SRCB THEN Continue
 ELSE Trap (VN)

Assembler


Syntax: ASLE vn, ra, rb
 or
 ASLE vn, ra, const8

Status: Not affected

Operands: SRCA content of register RA

 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)

 VN Trap vector number



Description: If the value of the SRCA operand is less than or equal to the value of the SRCB operand, instruction execution continues; otherwise, a trap with the specified vector number occurs.

For programs in the User mode, a Protection Violation trap occurs—instead of the assert trap—if a vector number between 0 and 63 is specified and the assert condition is not met.

ASLEU

ASLEU

Assert Less Than or Equal To, Unsigned


Operation: IF SRCA <= SRCB (unsigned) THEN Continue
 ELSE Trap (VN)

Assembler

Syntax: ASLEU vn, ra, rb
 or
 ASLEU vn, ra, const8

Status: Not affected

Operands: SRCA content of register RA
 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)
 VN Trap vector number



Description: If the value of the SRCA operand is less than or equal to the value of the SRCB operand, instruction execution continues; otherwise, a trap with the specified vector number occurs. For the comparison, both operands are treated as unsigned integers.

For programs in the User mode, a Protection Violation trap occurs— instead of the assert trap—if a vector number between 0 and 63 is specified and the assert condition is not met.

ASLT

ASLT

Assert Less Than

Operation: IF SRCA < SRCB THEN Continue
ELSE Trap(VN)

Assembler


Syntax: ASLT vn, ra, rb
 or
 ASLT vn, ra, const8

Status: Not affected

Operands: SRCA content of register RA

 SRCB M = 0 : content of register RB
 M = 1 : 1 (zero-extended to 32 bits)

 VN Trap vector number



Description: If the value of the SRCA operand is less than the value of the SRCB operand, instruction execution continues; otherwise, a trap with the specified vector number occurs.

For programs in the User mode, a Protection Violation trap occurs—instead of the assert trap—if a vector number between 0 and 63 is specified and the assert condition is not met.

ASLTU

ASLTU

Assert Less Than, Unsigned

Operation: IF SRCA < SRCB (unsigned) THEN Continue
ELSE Trap (VN)

Assembler


Syntax: ASLTU vn, ra, rb
 or
 ASLTU vn, ra, const8

Status: Not affected

Operands: SRCA content of register RA

 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)

 VN Trap vector number



OP = 52, 53

ASLTU

Description: If the value of the SRCA operand is less than the value of the SRCB operand, instruction execution continues; otherwise, a trap with the specified vector number occurs. For the comparison, both operands are treated as unsigned integers.

For programs in the User mode, a Protection Violation trap occurs—instead of the assert trap—if a vector number between 0 and 63 is specified and the assert condition is not met.

ASNEQ

ASNEQ

Assert Not Equal To


Operation: IF SRCA <> SRCB THEN Continue
ELSE Trap (VN)

Assembler

Syntax: ASNEQ vn, ra, rb
or
ASNEQ vn, ra, const8

Status: Not affected

Operands: SRCA content of register RA
SRCB M = 0 : content of register RB
M = 1 : I (zero-extended to 32 bits)
VN Trap vector number



Description: If the SRCA operand is not equal to the SRCB operand, instruction execution continues; otherwise, a trap with the specified vector number occurs.

For programs in the User mode, a Protection Violation trap occurs—instead of the assert trap—if a vector number between 0 and 63 is specified and the assert condition is not met.

CALL

CALL

Call Subroutine

Operation: DEST ← PC//00 + 8
 PC ← TARGET
 Execute delay instruction


Assembler

Syntax: CALL ra, target

Status: Not affected

Operands: TARGET A = 0 : I17..I10//I9..I2 (sign-extended to 30 bits) +PC
 A = 1 : I17..I10//I9..I2 (zero-extended to 30 bits)

DEST register RA



OP = A8, A9

CALL

Description: The address of the second following instruction is placed into the DEST location, and a non-sequential instruction fetch occurs to the instruction address given by the TARGET operand. The instruction following the CALL is executed before the non-sequential fetch occurs.

CALLI

CALLI

Call Subroutine, Indirect


Operation: DEST ← PC//00 + 8
PC ← SRCB
Execute delay instruction

Assembler

Syntax: CALLI ra, rb

Status: Not affected

Operands: SRCB content of register RB
DEST register RA



Description: The address of the second following instruction is placed into the DEST location, and a non-sequential instruction fetch occurs to the instruction address given by the SRCB operand. The instruction following the CALLI is executed before the non-sequential fetch occurs.

CONST

CONST

Constant

Operation: DEST ← 0I16


Assembler

Syntax: CONST ra, const16

Status: Not affected

Operands: 0I16 I15..I8//I7..I0, (zero-extended to 32 bits)

DEST register RA



Description: The 0I16 operand is placed into the DEST location.

CONSTH

CONSTH

Constant, High


Operation: Replace high-order half-word of SRCA by I16

Assembler

Syntax: CONSTH ra, const16

Status: Not affected

Operands: SRCA content of register RA
I16 I15..I8//I7..I0
DEST register RA



Description: The low-order half-word of the SRCA operand is appended to the I16 operand, and the result is placed into the DEST operand. Note that the destination register for this instruction is the same as the source register.

CONSTN

CONSTN

Constant, Negative

Operation: DEST ← 1116


Assembler

Syntax: CONSTN ra, const 16

Status: Not affected

Operands: 1116 115..18//17..10, (ones-extended to 32 bits)

DEST register RA



Description: The 1116 operand is placed into the DEST location.

CPBYTE

CPBYTE

Compare Bytes


Operation: IF (SRCA.BYTE0 = SRCB.BYTE0) OR
 (SRCA.BYTE1 = SRCB.BYTE1) OR
 (SRCA.BYTE2 = SRCB.BYTE2) OR
 (SRCA.BYTE3 = SRCB.BYTE3) THEN
 DEST ← TRUE ELSE DEST ← FALSE

Assembler

Syntax: CPBYTE rc, ra, rb
 or
 CPBYTE rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA
 SRCB M = 0 : content of register RB
 M = 1 : 1 (zero-extended to 32 bits)
 DEST register RC



OP = 2E, 2F

CPBYTE

Description: Each byte of the SRCA operand is compared to the corresponding byte of the SRCB operand. If any corresponding bytes are equal, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location.

CPEQ

CPEQ

Compare Equal To


Operation: IF SRCA = SRCB THEN DEST ← TRUE
 ELSE DEST ← FALSE

Assembler

Syntax: CPEQ rc, ra, rb
 or
 CPEQ rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA
 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)
 DEST register RC



Description: If the SRCA operand is equal to the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location.

CPGEU

CPGEU

Compare Greater Than or Equal To, Unsigned


Operation: IF SRCA >= SRCB (unsigned) THEN DEST ← TRUE
ELSE DEST ← FALSE

Assembler

Syntax: CPGEU rc, ra, rb
or
CPGEU rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA
SRCB M = 0 : content of register RB
M = 1 : I (zero-extended to 32 bits)
DEST register RC



Description: If the value of the SRCA operand is greater than or equal to the value of the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location. For the comparison, both operands are treated as unsigned integers.

CPGT

CPGT

Compare Greater Than

Operation: IF SRCA > SRCB THEN DEST ← TRUE
ELSE DEST ← FALSE

Assembler


Syntax: CPGT rc, ra, rb
 or
 CPGT rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA

 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)

 DEST register RC



OP = 48, 49

CPGT

Description: If the value of the SRCA operand is greater than the value of the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location.

CPGTU

CPGTU

Compare Greater Than, Unsigned


Operation: IF SRCA > SRCB (unsigned) THEN DEST ← TRUE
ELSE DEST ← FALSE

Assembler

Syntax: CPGTU rc, ra, rb
or
CPGTU rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA
SRCB M = 0 : content of register RB
M = 1 : 1 (zero-extended to 32 bits)
DEST register RC



Description: If the value of the SRCA operand is greater than the value of the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location. For the comparison, both operands are treated as unsigned integers.

CPLEU

CPLEU

Compare Less Than or Equal To, Unsigned

Operation: IF SRCA <= SRCB (unsigned) THEN DEST ← TRUE
ELSE DEST ← FALSE

Assembler


Syntax: CPLEU rc, ra, rb
 or
 CPLEU rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA

 SRCB M = 0 : content of register RB
 M = 1 : 1 (zero-extended to 32 bits)

 DEST register RC



Description: If the value of the SRCA operand is less than or equal to the value of the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location. For the comparison, both operands are treated as unsigned integers.

CPLTU

CPLTU

Compare Less Than, Unsigned


Operation: IF SRC_A < SRC_B (unsigned) THEN DEST ← TRUE
ELSE DEST ← FALSE

Assembler

Syntax: CPLTU rc, ra, rb
 or
 CPLTU rc, ra, const8

Status: Not affected

Operands: SRC_A content of register RA
 SRC_B M = 0 : content of register RB
 M = 1 : 1 (zero-extended to 32 bits)
 DEST register RC



OP = 42, 43

CPLTU

Description: If the value of the SRC_A operand is less than the value of the SRC_B operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location. For the comparison, both operands are treated as unsigned integers.

CPNEQ

CPNEQ

Compare Not Equal To


Operation: IF SRCA <> SRCB THEN DEST ← TRUE
 ELSE DEST ← FALSE

Assembler

Syntax: CPNEQ rc, ra, rb
 or
 CPNEQ rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA
 SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)
 DEST register RC



OP = 62, 63

CPNEQ

Description: If the SRCA operand is not equal to the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location.

CVDF

CVDF

Convert Floating-Point Double-Precision to Single-Precision


Operation: DEST (single-precision) ← SRCA (double-precision)

Assembler

Syntax: CVDF rc,ra

Status: Not affected

Operands: SRCA contents of register RA and the twin of register RA
DEST register RC



Description: The SRCA operand, treated as a double-precision, floating-point number, is converted to a single-precision, floating-point number, and the result is placed into the DEST location.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes a CVDF trap. When the trap occurs, the IPA and IPC registers are set to reference SRCA and DEST. The IPB register is also affected, but its value has no interpretation.

CVDINT

CVDINT

Convert Floating-Point Double-Precision to Integer


Operation: DEST (integer) ← SRCA (double-precision)

Assembler

Syntax: CVDINT rc,ra

Status: Not affected

Operands: SRCA contents of register RA and the twin of register RA
DEST register RC



OP = E7

CVDINT

Description: The SRCA operand, treated as a double-precision, floating-point number, is converted to an integer, and the result is placed into the DEST location.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes a CVDINT trap. When the trap occurs, the IPA and IPC registers are set to reference SRCA and DEST. The IPB register is also affected, but its value has no interpretation.

CVFD

CVFD

Convert Floating-Point Single-Precision to Double-Precision

Operation: DEST (double-precision) ← SRCA (single-precision)


Assembler

Syntax: CVFD rc,ra

Status: Not affected

Operands: SRCA content of register RA

DEST register RC and the twin of register RC



OP = E8

CVFD

Description: The SRCA operand, treated as a single-precision, floating-point number, is converted to a double-precision, floating-point number, and the result is placed into the DEST location.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes an CVFD trap. When the trap occurs, the IPA and IPC registers are set to reference SRCA and DEST. The IPB register is also affected, but its value has no interpretation.

CVFINT

CVFINT

Convert Floating-Point Single-Precision to Integer


Operation: DEST (integer) ← SRCA (single-precision)

Assembler

Syntax: CVFINT rc, ra

Status: Not affected

Operands: SRCA content of register RA
DEST register RC



Description: The SRCA operand, treated as a single-precision, floating-point number, is converted to an integer, and the result is placed into the DEST location.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes an CVFINT trap. When the trap occurs, the IPA and IPC registers are set to reference SRCA and DEST. The IPB register is also affected, but its value has no interpretation.

CVINTD

CVINTD

Convert Integer to Floating-Point Double-Precision


Operation: DEST (double-precision) ← SRCA (integer)

Assembler

Syntax: CVINTD rc,ra

Status: Not affected

Operands: SRCA content of register RA
DEST register RC and the twin of register RC



Description: The SRCA operand is converted to a double-precision, floating-point number, and the result is placed into the DEST location.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes an CVINTD trap. When the trap occurs, the IPA and IPC registers are set to reference SRCA and DEST. The IPB register is also affected, but its value has no interpretation.

CVINTF

CVINTF

Convert Integer to Floating-Point Single-Precision


Operation: DEST (single-precision) ← SRCA (integer)

Assembler

Syntax: CVINTF rc, ra

Status: Not affected

Operands: SRCA content of register RA
 DEST register RC



Description: The SRCA operand is converted to a single-precision, floating-point number, and the result is placed into the DEST location.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes an CVINTF trap. When the trap occurs, the IPA and IPC registers are set to reference SRCA and DEST. The IPB register is also affected, but its value has no interpretation.

DADD

DADD

Floating-Point Add, Double-Precision


Operation: DEST (double-precision) ← SRC_A (double-precision) + SRC_B (double-precision)

Assembler

Syntax: DADD rc,ra,rb

Status: Not affected

Operands: SRC_A contents of register RA and the twin of register RA
 SRC_B contents of register RB and the twin of register RB
 DEST register RC and the twin of register RC



OP = F1

DADD

Description: The SRC_A operand is added to the SRC_B operand, and the result is placed into the DEST location. The operands and result of the add are treated as double-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes a DADD trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRC_A, SRC_B, and DEST.

DDIV

DDIV

Floating-Point Divide, Double-Precision


Operation: DEST (double-precision) ← SRCA (double-precision) / SRCB (double-precision)

Assembler

Syntax: DDIV rc,ra,rb

Status: Not affected

- Operands:**
- SRCA contents of register RA and the twin of register RA
 - SRCB contents of register RB and the twin of register RB
 - DEST register RC and the twin of register RC



Description: The SRCA operand is divided by the SRCB operand, and the result is placed into the DEST location. The operands and result of the divide are treated as double-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes a DDIV trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCA, SRCB, and DEST.

DEQ

DEQ

Floating-Point Equal To, Double-Precision


Operation: IF SRCA (double-precision) = SRCB (double-precision)
THEN DEST ← TRUE
ELSE DEST ← FALSE

Assembler

Syntax: DEQ rc, ra, rb

Status: Not affected

Operands: SRCA contents of register RA and the twin of register RA
SRCB contents of register RB and the twin of register RB
DEST register RC



Description: If the SRCA operand is equal to the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location. For the comparison, the SRCA and SRCB operands are treated as double-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes a DEQ trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCA, SRCB, and DEST.

DGT

DGT

Floating-Point Greater Than, Double-Precision


Operation: IF SRCB (double-precision) > SRCB (double-precision)
THEN DEST ← TRUE
ELSE DEST ← FALSE

Assembler

Syntax: DGT rc, ra, rb

Status: Not affected

Operands: SRCB contents of register RA and the twin of register RA
SRCB contents of register RB and the twin of register RB
DEST register RC



Description: If the value of the SRCB operand is greater than the value of the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location. For the comparison, the SRCB and SRCB operands are treated as double-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes a DGT trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCB, SRCB, and DEST.

DIV

DIV

Divide Step

Operation: Perform one-bit step of a divide operation (unsigned)

Assembler


Syntax: DIV rc, ra, rb

Status: V, N, Z, C

Operands: SRCB content of register RA

 SRCB M = 0 : content of register RB
 M = 1 : 1 (zero-extended to 32 bits)

 DEST register RC



Description: If the Divide Flag (DF) bit of the ALU Status Register is 1, the SRCB operand is subtracted from the SRCB operand. If the DF bit is 0, the SRCB operand is added to the SRCB operand.

The carry-out of the add or subtract operation is exclusive-ORed with the value of the DF bit and the value of the Negative (N) bit of the ALU Status Register; the resulting value is complemented and placed into the DF bit. The sign of the result of the add or subtract is placed into the N bit.

The content of the Q register is appended to the result of the add or subtract, and the resulting 64-bit value is shifted left by one bit-position; the value computed for the DF bit above fills the vacated bit position. The high-order 32 bits of the 64-bit shifted value are placed into the DEST location. The low-order 32 bits of the shifted value are placed into the Q Register.

Examples of integer divide operations appear in Section 7.1.7.

DIV0

DIV0

Divide Initialize


Operation: Initialize for a sequence of divide steps (unsigned)

Assembler

Syntax: DIV0 rc, rb
or
DIV0 rc, const8

Status: V, N, Z, C

Operands: SRCB M = 0 : content of register RB
M = 1 : 1 (zero-extended to 32 bits)
DEST register RC



OP = 68, 69

DIV0

Description: The Divide Flag (DF) bit of the ALU Status Register is set. The sign of the SRCB operand is placed into the Negative bit of the ALU Status Register.

The content of the Q register is appended to the SRCB operand, and the resulting 64-bit value is shifted left by one bit-position; a 0 fills the vacated bit position. The high-order 32 bits of the 64-bit shifted value are placed into the DEST location. The low-order 32 bits of the shifted value are placed into the Q Register.

Examples of integer divide operations appear in Section 7.1.7.

DIVIDE

DIVIDE

Integer Divide


Operation: $DEST \leftarrow (SRCA/Q) / SRCB$ (unsigned)

Assembler

Syntax: `DIVIDE rc, ra, rb`

Status: Not affected

- Operands:**
- SRCA content of register RA
 - Q content of the Q Register
 - SRCB content of register RB
 - DEST register RC



Description: The content of the Q register is appended to the SRCA operand. The resulting 64-bit value is divided by the SRCB operand, and the result is placed into the DEST location.

This instruction does not check for a divide overflow condition. Checking for divide overflow must occur before the instruction is executed.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes a DIVIDE trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCA, SRCB, and DEST.

DIVREM

DIVREM

Divide Remainder


Operation: Generate remainder for divide operation (unsigned)

Assembler

Syntax: DIVREM rc, ra, rb

Status: V, N, Z, C

Operands: SRCB content of register RA
SRCB M = 0 : content of register RB
 M = 1 : 1 (zero-extended to 32 bits)
DEST register RC



Description: If the Divide Flag (DF) bit of the ALU Status Register is 1, the SRCB operand is placed into the DEST location.

If the DF bit is 0, the SRCB operand is added to the SRCB operand, and the result is placed into the DEST location.

Examples of integer divide operations appear in Section 7.1.7.

DLT

DLT

Floating-Point Less Than, Double-Precision


Operation: IF SRCA (double-precision) < SRCB (double-precision)
THEN DEST ← TRUE
ELSE DEST ← FALSE

Assembler

Syntax: DLT rc, ra, rb

Status: Not affected

Operands: SRCA contents of register RA and the twin of register RA
SRCB contents of register RB and the twin of register RB
DEST register RC



Description: If the value of the SRCA operand is less than the value of the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location. For the comparison, the SRCA and SRCB operands are treated as double-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes a DLT trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCA, SRCB, and DEST.

DMUL

DMUL


Floating-Point Multiply, Double-Precision

Operation: DEST (double-precision) ← SRCA (double-precision) * SRCB (double-precision)

Assembler Syntax: DMUL rc, ra, rb

Status: Not affected

Operands: SRCA contents of register RA and the twin of register RA
 SRCB contents of register RB and the twin of register RB
 DEST register RC and the twin of register RC



Description: The SRCA operand is multiplied by the SRCB operand, and the result is placed into the DEST location. The operands and result of the multiply are treated as double-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes a DMUL trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCA, SRCB, and DEST.

EMULATE

EMULATE

Trap to Software Emulation Routine

Operation: Load IPA and IPB registers with operand register-numbers and Trap (VN)


Assembler

Syntax: EMULATE vn, ra, rb

Status: Not affected

Operands: Absolute register-numbers for registers RA and RB

VN Trap vector number



OP = F8

EMULATE

Description: The IPA and IPB registers are set to the register-numbers of registers RA and RB, respectively. A trap with the specified vector number occurs.

Note that the IPC register is also affected by this instruction, but that its value has no interpretation.

For programs in the User mode, a Protection Violation trap occurs—instead of the EMULATE trap—if a vector number between 0 and 63 is specified.

EXBYTE

EXBYTE


Extract Byte

Operation: DEST ← SRCB, with low-order byte replaced by byte in SRCB selected by BP

Assembler Syntax: EXBYTE rc, ra,rb
or
EXBYTE rc, ra, const8

Status: Not affected

Operands: SRCB content of register RA
SRCB M = 0 : content of register RB
M = 1 : I (zero-extended to 32 bits)
DEST register RC



OP = 0A, 0B

EXBYTE

Description: A byte in the SRCB operand is selected by the Byte Position field of the ALU Status Register and the Byte Order bit of the Configuration Register. The selected byte replaces the low-order byte of the SRCB operand, and the resulting word is placed into the DEST location.

Note: The selection of bytes within words is specified in Section 3.4.3.

EXHWS

EXHWS

Extract Half-Word, Sign-Extended


Operation: DEST ← half-word in SRCA selected by BP,
sign-extended to 32 bits

Assembler

Syntax: EXHWS rc, ra

Status: Not affected

Operands: SRCA content of register RA
DEST register RC



Description: A half-word in the SRCA operand is selected by the Byte Position field of the ALU Status Register and the Byte Order bit of the Configuration Register. The selected half-word is sign-extended to 32 bits, and the resulting word is placed into the DEST location.

Note: The selection of half-words within words is specified in Section 3.4.3.

EXTRACT

EXTRACT

Extract Word, Bit-Aligned


Operation: DEST ← high-order word of (SRCA/SRCB << FC)

Assembler

Syntax: EXTRACT rc, ra ,rb
 or
 EXTRACT rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA
 SRCB M = 0 : content of register RB
 M = 1 : 1 (zero-extended to 32 bits)
 DEST register RC



OP = 7A, 7B

EXTRACT

Description: The SRCB operand is appended to the SRCA operand, and the resulting 64-bit value is shifted left by the number of bit-positions specified by the Funnel Shift Count (FC) field of the ALU Status register. The high-order 32 bits of the 64-bit shifted value are placed into the DEST location.

If the SRCB operand is the same as the SRCA operand, the EXTRACT instruction performs a rotate operation.

FADD

FADD

Floating-Point Add, Single-Precision


Operation: DEST (single-precision) ← SRCA (single-precision) + SRCB (single-precision)

Assembler

Syntax: FADD rc, ra, rb

Status: Not affected

Operands: SRCA content of register RA
 SRCB content of register RB
 DEST register RC



Description: The SRCA operand is added to the SRCB operand, and the result is placed into the DEST location. The operands and result of the add are treated as single-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes an FADD trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCA, SRCB, and DEST.

FDIV

FDIV

Floating-Point Divide, Single-Precision


Operation: DEST(single-precision) \leftarrow SRCA (single-precision) / SRCB (single-precision)

Assembler

Syntax: FDIV rc, ra, rb

Status: Not affected

Operands: SRCA content of register RA
SRCB content of register RB
DEST register RC



Description: The SRCA operand is divided by the SRCB operand, and the result is placed into the DEST location. The operands and result of the divide are treated as single-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes an FDIV trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCA, SRCB, and DEST.

FEQ

FEQ

Floating-Point Equal To, Single-Precision


Operation: IF SRCA (single-precision) = SRCB (single-precision)
THEN DEST ← TRUE
ELSE DEST ← FALSE

Assembler

Syntax: FEQ rc, ra, rb

Status: Not affected

Operands: SRCA content of register RA
SRCB content of register RB
DEST register RC



OP = EA

FEQ

Description: If the SRCA operand is equal to the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location. For the comparison, the SRCA and SRCB operands are treated as single-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes an FEQ trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCA, SRCB, and DEST.

Floating-Point Greater Than, Single-Precision


Operation: IF SRCA (single-precision) > SRCB(single-precision)
 THEN DEST ← TRUE
 ELSE DEST ← FALSE

Assembler

Syntax: FGT rc, ra, rb

Status: Not affected

Operands: SRCA content of register RA
 SRCB content of register RB
 DEST register RC



Description: If the value of the SRCA operand is greater than the value of the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location. For the comparison, the SRCA and SRCB operands are treated as single-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes an FGT trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCA, SRCB, and DEST.

FLT

FLT

Floating-Point Less Than, Single-Precision


Operation: IF SRCA (single-precision) < SRCB (single-precision)
THEN DEST ← TRUE
ELSE DEST ← FALSE

Assembler

Syntax: FLT rc, ra, rb

Status: Not affected

Operands: SRCA content of register RA
SRCB content of register RB
DEST register RC



Description: If the value of the SRCA operand is less than the value of the SRCB operand, a Boolean TRUE is placed into the DEST location; otherwise, a Boolean FALSE is placed into the DEST location. For the comparison, the SRCA and SRCB operands are treated as single-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes an FLT trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCA, SRCB, and DEST.

FMUL

FMUL

Floating-Point Multiply, Single-Precision


Operation: DEST (single-precision) \leftarrow SRC A (single-precision) * SRC B (single-precision)

Assembler

Syntax: FMUL rc, ra, rb

Status: Not affected

Operands: SRC A content of register RA
SRC B content of register RB
DEST register RC



Description: The SRC A operand is multiplied by the SRC B operand, and the result is placed into the DEST location. The operands and result of the multiply are treated as single-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes an FMUL trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRC A, SRC B, and DEST.

FSUB

FSUB

Floating-Point Subtract, Single-Precision


Operation: DEST (single-precision) ← SRCA (single-precision) – SRCB (single-precision)

Assembler

Syntax: FSUB rc, ra, rb

Status: Not affected

Operands: SRCA content of register RA
SRCB content of register RB
DEST register RC



Description: The SRCB operand is subtracted from the SRCA operand, and the result is placed into the DEST location. The operands and result of the subtract are treated as single-precision, floating-point numbers.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes an FSUB trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRCA, SRCB, and DEST.

HALT

HALT

Enter Halt Mode


Operation: Enter Halt mode on next cycle

Assembler

Syntax: HALT

Status: Not affected

Operands: not applicable



Description: The processor is placed into the Halt mode on the next cycle, except that any external data accesses are completed.

This instruction may be executed only by Supervisor-mode programs. An attempted execution by a User-mode program causes a Protection Violation trap to occur.

If the instruction following a Halt instruction has an exception (e.g., TLB Miss), the trap associated with this exception is taken before the processor enters the Halt mode.

INBYTE

INBYTE

Insert Byte


Operation: DEST ← SRCA, with byte selected by BP replaced by low-order byte of SRCB

Assembler

Syntax: INBYTE rc, ra, rb
 or
 INBYTE rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA
 SRCB M=0: content of register RB
 M=1: I (zero-extended to 32 bits)
 DEST register RC



OP = 0C, 0D

INBYTE

Description: A byte in the SRCB operand is selected by the Byte Position field of the ALU Status Register and the Byte Order bit of the Configuration Register. The selected byte is replaced by the low-order byte of the SRCB operand, and the resulting word is placed into the DEST location.

Note: The selection of bytes within words is specified in Section 3.4.3.

INHW

INHW

Insert Half-Word


Operation: DEST ← SRCA, with half-word selected by BP replaced by low-order half-word of SRCB

Assembler

Syntax: INHW rc, ra, rb
or
INHW rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA
SRCB M=0: content of register RB
M=1: I (zero-extended to 32 bits)
DEST register RC



Description: A half-word in the SRCA operand is selected by the Byte Position field of the ALU Status Register and the Byte Order bit of the Configuration Register. The selected half-word is replaced by the low-order half-word of the SRCB operand, and the resulting word is placed into the DEST location.

Note: The selection of half-words within words is specified in Section 3.4.3.

INV

INV

Invalidate


Operation: Reset all Valid bits in Branch Target Cache

Assembler

Syntax: INV

Status: Not affected

Operands: not applicable



Description: This instruction causes all Branch Target Cache Valid bits to be reset, on the execution of the next successful branch. This causes all Branch Target Cache locations to become invalid.

This instruction may be executed only by Supervisor-mode programs. An attempted execution by a User-mode program causes a Protection Violation trap to occur.

IRET

IRET

Interrupt Return


Operation: Perform an interrupt return sequence

Assembler

Syntax: IRET

Status: Not affected

Operands: not applicable



Description: This instruction performs the interrupt return sequence described in Section 3.5.5.

This instruction may be executed only by Supervisor-mode programs. An attempted execution by a User-mode program causes a Protection Violation trap to occur.

IRETINV

IRETINV

Interrupt Return and Invalidate


Operation: Perform an interrupt return sequence, and reset all Valid bits in Branch Target Cache

Assembler

Syntax: IRETINV

Status: Not affected

Operands: not applicable



OP = 8C

IRETINV

Description: This instruction performs the interrupt return sequence described in Section 3.5.5. When the sequence begins, all Branch Target Cache Valid bits are reset to zeros. This causes all Branch Target Cache locations to become invalid.

This instruction may be executed only by Supervisor-mode programs. An attempted execution by a User-mode program causes a Protection Violation trap to occur.

JMP

JMP

Jump


Operation: PC ← TARGET
Execute delay instruction

Assembler

Syntax: JMP target

Status: Not affected

Operands: TARGET A=0: I17..I10//I9..I2 (sign-extended to 30 bits) + PC
A=1: I17..I10//I9..I2 (zero-extended to 30 bits)



OP = A0, A1

JMP

Description: A non-sequential instruction fetch occurs to the instruction address given by the TARGET operand. The instruction following the JMP is executed before the non-sequential fetch occurs.

JMPF

JMPF

Jump False

Operation: IF SRCA = FALSE THEN PC ← TARGET
Execute delay instruction


Assembler

Syntax: JMPF ra, target

Status: Not affected

Operands: SRCA content of register RA

TARGET A=0: I17..I10//I9..I2 (sign-extended to 30 bits) + PC
 A=1: I17..I10//I9..I2 (zero-extended to 30 bits)



OP = A4, A5

JMPF

Description: If SRCA is a Boolean FALSE, a non-sequential instruction fetch occurs to the instruction address given by the TARGET operand.

If SRCA is a Boolean TRUE, this instruction has no effect.

The instruction following the JMPF is executed regardless of the value of SRCA.

JMPFDEC

JMPFDEC

Jump False and Decrement


Operation: IF SRCA = FALSE THEN
 SRCA ← SRCA - 1
 PC ← TARGET
 ELSE
 SRCA ← SRCA - 1
 Execute delay instruction

Assembler

Syntax: JMPFDEC ra, target

Status: Not affected

Operands: SRCA content of register RA
 TARGET A=0: I17..I10//I9..I2 (sign-extended to 30 bits) +PC
 A=1: I17..I10//I9..I2 (zero-extended to 30 bits)



OP = B4, B5

JMPFDEC

Description: If SRCA is a Boolean FALSE, a non-sequential instruction fetch occurs to the instruction address given by the TARGET operand.

If SRCA is a Boolean TRUE, this instruction has no effect on the instruction-execution sequence.

The SRCA operand is decremented by one, regardless of whether or not the non-sequential instruction fetch occurs. Note that a negative number for the SRCA operand is a Boolean TRUE.

The instruction following the JMPFDEC is executed regardless of the value of SRCA.

JMPFI

JMPFI

Jump False Indirect


Operation: IF SRCA = FALSE THEN PC ← SRCB
Execute delay instruction

Assembler

Syntax: JMPFI ra, rb

Status: Not affected

Operands: SRCA content of register RA
SRCB content of register RB



Description: The SRCA is a Boolean FALSE, a non-sequential instruction fetch occurs to the instruction address given by the SRCB operand.

If SRCA is a Boolean TRUE, this instruction has no effect.

The instruction following the JMPFI is executed regardless of the value of SRCA.

JMPI

JMPI

Jump Indirect


Operation: PC ← SRCB
Execute delay instruction

Assembler

Syntax: JMPI rb

Status: Not affected

Operands: SRCB content of register RB



OP = C0

JMPI

Description: A non-sequential instruction fetch occurs to the instruction address given by the SRCB operand. The instruction following the JMPI is executed before the non-sequential fetch occurs.

JMPT

JMPT

Jump True

Operation: IF SRCA = TRUE THEN PC ← TARGET
Execute delay instruction


Assembler

Syntax: JMPT ra, target

Status: Not affected

Operands: SRCA content of register RA

TARGET A=0: I17..I10//I9..I2 (sign-extended to 30 bits) + PC
A=1: I17..I10//I9..I2 (zero-extended to 30 bits)



Description: If SRCA is a Boolean TRUE, a non-sequential instruction fetch occurs to the instruction address given by the TARGET operand.

If SRCA is a Boolean FALSE, this instruction has no effect.

The instruction following the JMPT is executed regardless of the value of SRCA.

JMPTI

JMPTI

Jump True Indirect


Operation: IF SRCA = TRUE THEN PC ← SRCB
Execute delay instruction

Assembler

Syntax: JMPTI ra, rb

Status: Not affected

Operands: SRCA content of register RA
SRCB content of register RB



Description: The SRCA is a Boolean TRUE, a non-sequential instruction fetch occurs to the instruction address given by the SRCB operand.

If SRCA is a Boolean FALSE, this instruction has no effect.

The instruction following the JMPTI is executed regardless of the value of SRCA.

LOAD

LOAD

Load

Operation: DEST ← EXTERNAL WORD [SRCB]


Assembler

Syntax: LOAD ce, cntl, ra, rb
or
LOAD ce, cntl, ra, const8

Status: Not affected

Operands: SRCB M=0: content of register RB
M=1: 1 (zero-extended to 32 bits)

DEST register RA



Description: If the CE bit is 0, the external word addressed by the SRCB operand is placed into the DEST location.

If the CE bit is 1, a word is transferred from the coprocessor into the DEST location. The SRCB operand has no pre-defined interpretation in this case, though it appears on the Address Bus.

The CNTL field of the LOAD instruction affects the access or transfer as described in Sections 3.4.2 and 6.1.2.

LOADL

LOADL

Load and Lock


Operation: DEST ← EXTERNAL WORD [SRCB],
assert *LOCK output during access

Assembler

Syntax: LOADL ce, cntl, ra, rb
or
LOADL ce, cntl, ra, const8

Status: Not affected

Operands: SRCB M=0: content of register RB
M=1: 1 (zero-extended to 32 bits)
DEST register RA



Description: If the CE bit is 0, the external word addressed by the SRCB operand is placed into the DEST location.

If the CE bit is 1, a word is transferred from the coprocessor into the DEST location. The SRCB operand has no pre-defined interpretation in this case, though it appears on the Address Bus.

The CNTL field of the LOADL instruction affects the access or transfer as described in Sections 3.4.2 and 6.1.2.

The *LOCK output is asserted during the access or transfer.

LOADM

LOADM

Load Multiple


Operation: DEST .. DEST + COUNT ←
EXTERNAL WORD [SRCB]...EXTERNAL WORD
[SRCB+COUNT* 4]

Assembler

Syntax: LOADM ce, cntl, ra, rb
or
LOADM ce, cntl, ra, const8

Status: Not affected

Operands: SRCB M=0: content of register RB
M=1: 1 (zero-extended to 32 bits)
DEST register RA



Description: If the CE bit is 0, external words at consecutive word-addresses, beginning with the word addressed by the SRCB operand, are placed into consecutive registers, beginning with the DEST location.

If the CE bit is 1, multiple words are transferred from the coprocessor into consecutive registers, beginning with the DEST location. The SRCB operand has no pre-defined interpretation in this case.

The total number of words accessed or transferred in the sequence is specified by the Count Remaining field of the Channel Control Register (which also appears in the Load/Store Count Remaining Register) at the beginning of the access. The CNTL field of the LOADM instruction affects the access or transfer as described in Sections 3.4.2 and 6.1.2.

LOADSET

LOADSET

Load and Set


Operation: DEST ← EXTERNAL WORD [SRCB]
 EXTERNAL WORD [SRCB] ← h'FFFFFFFF',
 assert *LOCK output during access

Assembler

Syntax: LOADSET ce, cntl, ra, rb
 or
 LOADSET ce, cntl, ra, const8

Status: Not affected

Operands: SRCB M=0: content of register RB
 M=1: 1 (zero-extended to 32 bits)
 DEST register RA



Description: If the CE bit is 0, the external word addressed by the SRCB operand is placed into the DEST location. After the DEST location is altered, the external word addressed by the SRCB operand is written, atomically, with a word consisting of a 1 in every bit position.

If the CE bit is 1, a word is transferred from the coprocessor into the DEST location. The SRCB operand has no pre-defined interpretation in this case, though it appears on the Address Bus. After the DEST location is altered, a word consisting of a 1 in every bit position is transferred, atomically, to the coprocessor.

The CNTL field of the LOADSET instruction affects the access or transfer as described in Sections 3.4.2 and 6.1.2.

The *LOCK output is asserted throughout the LOADSET operation.

MFSR

MFSR

Move from Special Register

Operation: DEST ← SPECIAL


Assembler

Syntax: MFSR rc, spid

Status: Not affected

Operands: SPECIAL content of special-purpose register SA

DEST register RC



OP = C6

MFSR

Description: The SPECIAL operand is placed into the DEST location.

For programs in the User mode, a Protection Violation trap occurs if SA specifies a protected special-purpose register. If a trap occurs, the DEST location is not altered.

MFTLB

MFTLB

Move from Translation Look-Aside Buffer Register

Operation: DEST ← TLB [SRCA]


Assembler

Syntax: MFTLB rc, ra

Status: Not affected

Operands: SRCA content of register RA, bits 6..0

DEST register RC



Description: The Translation Look-Aside Buffer (TLB) register whose register-number is specified by the SCRA operand is placed into the DEST location.

This instruction may be executed only by Supervisor-mode programs. An attempted execution by a User-mode program causes a Protection Violation trap to occur. If a trap occurs, the DEST location is not altered.

MTSR

MTSR

Move to Special Register

Operation: SPDEST ← SRCB


Assembler

Syntax: MTSR spid, rb

Status: Not affected, unless the destination is the ALU Status Register

Operands: SRCB content of register RB

 SPDEST special-purpose register SA



Description: The SRCB operand is placed into the SPECIAL location.

For programs in the User mode, a Protection Violation trap occurs if SA specifies a protected special-purpose register. If a trap occurs, the SPDEST location is not altered.

MTSRIM

MTSRIM

Move to Special Register Immediate

Operation: SPDEST ← 0I16


Assembler

Syntax: MTSRIM spid, const16

Status: Not affected, unless the destination is the ALU Status Register

Operands: 0I16 I15.. I8//I7.. I0 (zero-extended to 32 bits)

SPDEST special-purpose register SA



OP = 04

MTSRIM

Description: The 0I16 operand is placed into the SPECIAL location.

For programs in the User mode, a Protection Violation trap occurs if SA specifies a protected special-purpose register. If a trap occurs, the SPDEST location is not altered.

MTTLB

MTTLB

Move to Translation Look-Aside Buffer Register

Operation: TLB [SRCA] ← SRCB


Assembler

Syntax: MTTLB ra, rb

Status: Not affected

Operands: SRCA content of register RA, bits 6..0

SRCB content of register RB



Description: The SRCB operand is placed into the Translation Look-Aside Buffer (TLB) register whose register-number is specified by the SCRA operand.

This instruction may be executed only by Supervisor-mode programs. An attempted execution by a User-mode program causes a Protection Violation trap to occur. If a trap occurs, the TLB register is not altered.

MUL

MUL

Multiply Step


Operation: Perform one-bit step of a multiply operation

Assembler

Syntax: MUL rc, ra, rb

Status: V, N, Z, C

Operands: SRCA content of register RA
SRCB M=0: content of register RB
M=1: 1 (zero-extended to 32 bits)
DEST register RC



OP = 64, 65

MUL

Description: If the least-significant bit of the Q Register is 1, the SRCA operand is added to the SRCB operand. If the least-significant bit of the Q register is 0, a zero word is added to the SRCB operand.

The content of the Q register is appended to the result of the add, and the resulting 64-bit value is shifted right by one bit-position; the true sign of the result of the add fills the vacated bit position (i.e. the sign of the result is complemented if an overflow occurred during the add operation). The high-order 32 bits of the 64-bit shifted value are placed into the DEST location. The low-order 32 bits of the shifted value are placed into the Q Register.

Examples of integer multiply operations appear in Section 7.1.6.

MULL

MULL

Multiply Last Step


Operation: Complete a sequence of multiply steps (for signed multiply)

Assembler

Syntax: MULL rc, ra, rb

Status: V, N, Z, C

Operands: SRCA content of register RA
SRCB M=0: content of register RB
 M=1: 1 (zero-extended to 32 bits)
DEST register RC



OP = 66, 67

MULL

Description: If the least-significant bit of the Q Register is 1, the SRCA operand is subtracted from the SRCB operand. If the least-significant bit of the Q register is 0, a zero word is subtracted from the SRCB operand.

The content of the Q register is appended to the result of the subtract, and the resulting 64-bit value is shifted right by one bit-position; the true sign of the result of the subtract fills the vacated bit position (i.e. the sign of the result is complemented if an overflow occurred during the subtract operation). The high-order 32 bits of the 64-bit shifted value are placed into the DEST location. The low-order 32 bits of the shifted value are placed into the Q Register.

Examples of integer multiply operations appear in Section 7.1.6.

MULTIPLY

MULTIPLY

Integer Multiply


Operation: $DEST//Q \leftarrow SRC_A * SRC_B$

Assembler

Syntax: MULTIPLY rc, ra,rb

Status: Not affected

Operands: SRC_A content of register RA
SRC_B content of register RB
DEST register RC
Q Q Register



Description: The SRC_A operand is multiplied by the SRC_B operand. The high-order 32 bits of the 64-bit result are placed into the DEST location. The low-order 32 bits of the result are placed into the Q Register.

Note: This instruction is not directly supported in processor hardware. In the current implementation, this instruction causes a MULTIPLY trap. When the trap occurs, the IPA, IPB, and IPC registers are set to reference SRC_A, SRC_B, and DEST.

MULU

MULU

Multiply Step, Unsigned

Operation: Perform one-bit step of a multiply operation (unsigned)

Assembler


Syntax: MULU rc, ra, rb

Status: V, N, Z, C

Operands: SRCB content of register RA

 SRCB M=0: content of register RB
 M=1: 1 (zero-extended to 32 bits)

 DEST register RC



OP = 74, 75

MULU

Description: If the least-significant bit of the Q Register is 1, the SRCB operand is added to the SRCB operand. If the least-significant bit of the Q register is 0, a zero word is added to the SRCB operand.

The content of the Q register is appended to the result of the add, and the resulting 64-bit value is shifted right by one bit-position; the carry-out of the add fills the vacated bit position. The high-order 32 bits of the 64-bit shifted value are placed into the DEST location. The low-order 32 bits of the shifted value are placed into the Q Register.

Examples of integer multiply operations appear in Section 7.1.6.

NAND

NAND

NAND Logical


Operation: $DEST \leftarrow \sim (SRCA \& SRCB)$

Assembler

Syntax: NAND rc, ra, rb
or
NAND rc, ra, const8

Status: N, Z

Operands: SRCA content of register RA
SRCB M=0: content of register RB
M=1: I (zero-extended to 32 bits)
DEST register RC



Description: The SRCA operand is logically ANDed, bit-by-bit, with the SRCB operand. The one's-complement of the result is placed into the DEST location.

NOR

NOR

NOR Logical


Operation: $DEST \leftarrow \sim (SRCA \mid SRCB)$

Assembler

Syntax: NOR rc, ra, rb
 or
 NOR rc, ra, const8

Status: N, Z

Operands: SRCA content of register RA
 SRCB M=0: content of register RB
 M=1: 1 (zero-extended to 32 bits)
 DEST register RC



Description: The SRCA operand is logically ORed, bit-by-bit, with the SRCB operand. The one's-complement of the result is placed into the DEST location.

OR

OR

OR Logical


Operation: DEST ← SRCA | SRCB

Assembler

Syntax: OR rc, ra, rb
or
OR rc, ra, const8

Status: N, Z

Operands: SRCA content of register RA
SRCB M = 0 : content of register RB
 M = 1 : I (zero-extended to 32 bits)
DEST register RC



OP = 92, 93

OR

Description: The SRCA operand is logically ORed, bit-by-bit, with the SRCB operand, and the result is placed into the DEST location.

SETIP

SETIP

Set Indirect Pointers


Operation: Load IPA, IPB, and IPC registers with operand register-numbers

Assembler

Syntax: SETIP rc, ra, rb

Status: Not affected

Operands: Absolute register-numbers for registers RA, RB, and RC



OP = 9E

SETIP

Description: The IPA, IPB, and IPC registers are set to the register-numbers of registers RA, RB, and RC, respectively.

For programs in the User mode, a Protection Violation trap occurs if RA, RB, or RC specifies a register which is protected by the Register Bank Protect Register.

SLL

SLL

Shift Left Logical

Operation: DEST ← SRCA << SRCB (zero fill)

Assembler


Syntax: SLL rc, ra, rb
 or
 SLL rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA

 SRCB M=0: content of register RB, bits 4..0
 M=1: I, bits 4..0

 DEST register RC



Description: The SRCA operand is shifted left by the number of bit positions specified by the SRCB operand; zeros fill vacated bit positions. The result is placed into the DEST location.

SRA

SRA

Shift Right Arithmetic


Operation: DEST ← SRC A >> SRC B (sign fill)

Assembler

Syntax: SRA rc, ra, rb
or
SRA rc, ra, const8

Status: Not affected

Operands: SRC A content of register RA
SRC B M=0: content of register RB, bits 4..0
M=1: I, bits 4..0
DEST register RC



OP = 86, 87

SRA

Description:

The SRC A operand is shifted right by the number of bit positions specified by the SRC B operand; the sign of the SRC A operand fills vacated bit positions. The result is placed into the DEST location.

SRL

SRL

Shift Right Logical

Operation: DEST ← SRCA >> SRCB (zero fill)

Assembler


Syntax: SRL rc,ra,rb
or
SRL rc, ra, const8

Status: Not affected

Operands: SRCA content of register RA

SRCB M=0: content of register RB, bits 4..0
M=1: I, bits 4..0

DEST register RC



OP = 82, 83

SRL

Description: The SRCA operand is shifted right by the number of bit positions specified by the SRCB operand; zeros fill vacated bit positions. The result is placed into the DEST location.

STORE

STORE

Store


Operation: EXTERNAL WORD [SRCB] ← SRCA

Assembler

Syntax: STORE ce, cntl , ra, rb
or
STORE ce, cntl, ra, const8

Status: Not affected

Operands: SRCA content of register RA
 SRCB M=0: content of register RB
 M=1: I (zero-extended to 32 bits)



Description: If the CE bit is 0, the SRCA operand is placed into the external word addressed by the SRCB operand.

If the CE bit is 1, the SRCA and SRCB operands are transferred to the coprocessor.

The CNTL field of the STORE instruction affects the access or transfer as described in Sections 3.4.2 and 6.1.2.

STOREL

STOREL

Store and Lock


Operation: EXTERNAL WORD [SRCB] ← SRCA,
assert *LOCK output during access

Assembler

Syntax: STOREL ce, cntl, ra, rb
or
STOREL ce, cntl, ra, const8

Status: Not affected

Operands: SRCA content of register RA
SRCB M=0: content of register RB
M=1: I (zero-extended to 32 bits)



Description: If the CE bit is 0, the SRCA operand is placed into the external word addressed by the SRCB operand.

If the CE bit is 1, the SRCA and SRCB operands are transferred to the coprocessor.

The CNTL field of the STOREL instruction affects the access or transfer as described in Sections 3.4.2 and 6.1.2.

The *LOCK output is asserted during the access or transfer.

STOREM

STOREM

Store Multiple


Operation: EXTERNAL WORD [SRCB]...EXTERNAL WORD [SRCB+COUNT * 4]
DEST... DEST + COUNT

Assembler

Syntax: STOREM ce, cntl, ra, rb
or
STOREM ce, cntl, ra, const8

Status: Not affected

Operands: SRCA content of register RA
 SRCB M=0: content of register RB
 M=1: 1 (zero-extended to 32 bits)



Description: If the CE bit is 0, the contents of consecutive registers, beginning with the SRCA operand, are placed into external words a consecutive word-addresses, beginning with the word addressed by the SRCB operand.

If the CE bit is 1, the contents of consecutive registers, beginning with the SRCA operand, are transferred to the coprocessor. The SRCB operand has no pre-defined interpretation in this case.

The total number of words accessed or transferred in the sequence is specified by the Count Remaining field of the Channel Control Register (which also appears in the Load/Store Count Remaining Register) at the beginning of the access. The CNTL field of the STOREM instruction affects the access or transfer as described in Sections 3.4.2 and 6.1.2.

Note: The address and register-number sequences for the STOREM instruction are specified in Section 3.4.2

SUB

Subtract

SUB


Operation: $DEST \leftarrow SRCA - SRCB$

Assembler

Syntax: SUB rc, ra, rb
or
SUB rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA
SRCB M=0: content of register RB
M=1: I (zero-extended to 32 bits)
DEST register RC



Description: The SRCA operand is added to the two's-complement of the SRCB operand, and the result is placed into the DEST location.

SUBC

SUBC

Subtract with Carry

Operation: $DEST \leftarrow SRCA - SRCB - 1 + C$

Assembler


Syntax: SUBC rc, ra, rb
 or
 SUBC rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA

 SRCB M=0: content of register RB
 M=1: I (zero-extended to 32 bits)

 DEST register RC



Description: The SRCA operand is added to the one's-complement of the SRCB operand and the value of the ALU Status Carry bit, and the result is placed into the DEST location.

SUBCS

SUBCS

Subtract with Carry, Signed

Operation: $DEST \leftarrow SRCA - SRCB - 1 + C$
IF signed overflow THEN Trap (Out of Range)

Assembler


Syntax: SUBCS rc, ra, rb
 or
 SUBCS rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA

 SRCB M=0: content of register RB
 M=1: I (zero-extended to 32 bits)

 DEST register RC



OP = 28, 29

SUBCS

Description:

The SRCA operand is added to the one's-complement of the SRCB operand and the value of the ALU Status Carry bit, and the result is placed into the DEST location. If the add operation causes a two's-complement signed overflow, an Out of Range trap occurs.

Note that the DEST location is altered whether or not an overflow occurs.

SUBR

SUBR

Subtract Reverse


Operation: $DEST \leftarrow SRCB - SRCA$

Assembler

Syntax: SUBR rc, ra, rb
or
SUBR rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA
SRCB M=0: content of register RB
M=1: I (zero-extended to 32 bits)
DEST register RC



Description: The SRCB operand is added to the two's-complement of the SRCA operand, and the result is placed into the DEST location.

SUBRC

SUBRC

Subtract Reverse with Carry

Operation: $DEST \leftarrow SRCB - SRCA - 1 + C$

Assembler


Syntax: SUBRC rc, ra, rb
 or
 SUBRC rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA

 SRCB M=0: content of register RB
 M=1: 1 (zero-extended to 32 bits)

 DEST register RC



OP = 3C, 3D

SUBRC

Description: The SRCB operand is added to the one's-complement of the SRCA operand and the value of the ALU Status Carry bit, and the result is placed into the DEST location.

SUBRCS

SUBRCS

Subtract Reverse with Carry, Signed


Operation: $DEST \leftarrow SRCB - SRCA - 1 + C$
 IF signed overflow THEN Trap (Out of Range)

Assembler

Syntax: SUBRCS rc, ra, rb
 or
 SUBRCS rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA
 SRCB M=0: content of register RB
 M=1: 1 (zero-extended to 32 bits)
 DEST register RC



Description: The SRCB operand is added to the one's-complement of the SRCA operand and the value of the ALU Status Carry bit, and the result is placed into the DEST location. If the add operation causes a two's-complement signed overflow, an Out of Range trap occurs.

Note that the DEST location is altered whether or not an overflow occurs.

SUBRCU

SUBRCU

Subtract Reverse with Carry, Unsigned


Operation: $DEST \leftarrow SRCB - SRCA - 1 + C$
 IF unsigned underflow THEN Trap (Out of Range)

Assembler

Syntax: SUBRCU rc, ra, rb
 or
 SUBRCU rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA
 SRCB M=0: content of register RB
 M=1: 1 (zero-extended to 32 bits)
 DEST register RC



Description: The SRCB operand is added to the one's-complement of the SRCA operand and the value of the ALU Status Carry bit, and the result is placed into the DEST location. If the add operation causes an unsigned underflow, an Out of Range trap occurs.

Note that the DEST location is altered whether or not an underflow occurs.

SUBRS

SUBRS

Subtract Reverse, Signed


Operation: DEST \leftarrow SRCB - SRCA
IF signed overflow THEN Trap (Out of Range)

Assembler

Syntax: SUBRS rc, ra, rb
 or
 SUBRS rc, ra, const8

Status: V, N, Z, C

Operands: SRCB content of register RA
 SRCB M=0: content of register RB
 M=1: I (zero-extended to 32 bits)
 DEST register RC



OP = 30, 31

SUBRS

Description: The SRCB operand is added to the two's-complement of the SRCA operand, and the result is placed into the DEST location. If the add operation causes a two's-complement signed overflow, an Out of Range trap occurs.

Note that the DEST location is altered whether or not an overflow occurs.

SUBRU

SUBRU


Subtract Reverse, Unsigned

Operation: $DEST \leftarrow SRCB - SRCA$
IF unsigned underflow THEN Trap (Out of Range)

Assembler Syntax: SUBRU rc, ra, rb
or
SUBRU rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA
SRCB M=0: content of register RB
M=1: 1 (zero-extended to 32 bits)
DEST register RC



OP = 32, 33

SUBRU

Description: The SRCB operand is added to the two's-complement of the SRCA operand, and the result is placed into the DEST location. If the add operation causes an unsigned underflow, an Out of Range trap occurs.

Note that the DEST location is altered whether or not an underflow occurs.

SUBS

SUBS

Subtract, Signed

Operation: DEST ← SRCA – SRCB
IF signed overflow THEN Trap (Out of Range)

Assembler


Syntax: SUBS rc, ra ,rb
 or
 SUBS rc, ra, const8

Status: V, N, Z, C

Operands: SRCA content of register RA

 SRCB M=0: content of register RB
 M=1: I (zero-extended to 32 bits)

 DEST register RC



OP = 20, 21

SUBS

Description: The SRCA operand is added to the two's-complement of the SRCB operand, and the result is placed into the DEST location. If the add operation causes a two's-complement signed overflow, an Out of Range trap occurs.

Note that the DEST location is altered whether or not an overflow occurs.

SUBU

SUBU

Subtract, Unsigned


Operation: $DEST \leftarrow SRC_A - SRC_B$
IF unsigned underflow THEN Trap (Out of Range)

Assembler

Syntax: SUBU rc, ra, rb
or
SUBU rc, ra, const8

Status: V, N, Z, C

Operands: SRC_A content of register RA
SRC_B M=0: content of register RB
M=1: 1 (zero-extended to 32 bits)
DEST register RC



OP = 22, 23

SUBU

Description: The SRC_A operand is added to the two's-complement of the SRC_B operand, and the result is placed into the DEST location. If the add operation causes an unsigned underflow, an Out of Range trap occurs.

Note that the DEST location is altered whether or not an underflow occurs.

XNOR

XNOR

Exclusive-NOR Logical


Operation: DEST ← \sim (SRCA ^ SRCB)

Assembler

Syntax: XNOR rc, ra, rb
or
XNOR rc, ra, const8

Status: N, Z

Operands: SRCA content of register RA
SRCB M=0: content of register RB
M=1: I (zero-extended to 32 bits)
DEST register RC



OP = 96, 97

XNOR

Description:

The SRCA operand is logically exclusive-ORed, bit-by-bit, with the SRCB operand. The one's-complement of the result is placed into the DEST location.

XOR

XOR

Exclusive-OR Logical

Operation: DEST < SRCA ^ SRCB

Assembler


Syntax: XOR rc, ra, rb
 or
 XOR rc, ra, const8

Status: N, Z

Operands: SRCA content of register RA

 SRCB M=0: content of register RB
 M=1: I (zero-extended to 32 bits)

 DEST register RC



Description: The SRCA operand is logically exclusive-ORed, bit-by-bit, with the SRCB operand, and the result is placed into the DEST location.

8.5 INSTRUCTION INDEX BY OPERATION CODE

01	CONSTN	Constant, Negative
02	CONSTH	Constant, High
03	CONST	Constant
04	MTSRIM	Move to Special Register Immediate
06,07	LOADL	Load and Lock
08,09	CLZ	Count Leading Zeros
0A,0B	EXBYTE	Extract Byte
0C,0D	INBYTE	Insert Byte
0E,0F	STOREL	Store and Lock
10,11	ADDS	Add, Signed
12,13	ADDU	Add, Unsigned
14,15	ADD	Add
16,17	LOAD	Load
18,19	ADDCS	Add with Carry, Signed
1A,1B	ADDCU	Add with Carry, Unsigned
1C,1D	ADDC	Add with Carry
1E,1F	STORE	Store
20,21	SUBS	Subtract, Signed
22,23	SUBU	Subtract, Unsigned
24,25	SUB	Subtract
26,27	LOADSET	Load and Set
28,29	SUBCS	Subtract with Carry, Signed
2A,2B	SUBCU	Subtract with Carry, Unsigned
2C,2D	SUBC	Subtract with Carry
2E,2F	CPBYTE	Compare Bytes
30,31	SUBRS	Subtract Reverse, Signed
32,33	SUBRU	Subtract Reverse, Unsigned
34,35	SUBR	Subtract Reverse
36,37	LOADM	Load Multiple
38,39	SUBRCS	Subtract Reverse with Carry, Signed
3A,3B	SUBRCU	Subtract Reverse with Carry, Unsigned
3C,3D	SUBRC	Subtract Reverse with Carry
3E,3F	STOREM	Store Multiple
40,41	CPLT	Compare Less Than
42,43	CPLTU	Compare Less Than, Unsigned
44,45	CPLE	Compare Less Than or Equal To
46,47	CPLEU	Compare Less Than or Equal To, Unsigned
48,49	CPGT	Compare Greater Than
4A,4B	CPGTU	Compare Greater Than, Unsigned
4C,4D	CPGE	Compare Greater Than or Equal To

4E,4F	CPGEU	Compare Greater Than or Equal To, Unsigned
50,51	ASLT	Assert Less Than
52,53	ASLTU	Assert Less Than, Unsigned
54,55	ASLE	Assert Less Than or Equal To
56,57	ASLEU	Assert Less Than or Equal To, Unsigned
58,59	ASGT	Assert Greater Than
5A,5B	ASGTU	Assert Greater Than, Unsigned
5C,5D	ASGE	Assert Greater Than or Equal To
5E,5F	ASGEU	Assert Greater Than or Equal To, Unsigned
60,61	CPEQ	Compare Equal To
62,63	CPNEQ	Compare Not Equal To
64,65	MUL	Multiply Step
66,67	MULL	Multiply Last Step
68,69	DIV0	Divide Initialize
6A,6B	DIV	Divide Step
6C,6D	DIVL	Divide Last Step
6E,6F	DIVREM	Divide Remainder
70,71	ASEQ	Assert Equal To
72,73	ASNEQ	Assert Not Equal To
74,75	MULU	Multiply Step, Unsigned
78,79	INHW	Insert Half-Word
7A,7B	EXTRACT	Extract Word, Bit-Aligned
7C,7D	EXHW	Extract Half-Word
7E	EXHWS	Extract Half-Word, Sign-Extended
80,81	SLL	Shift Left Logical
82,83	SRL	Shift Right Logical
86,87	SRA	Shift Right Arithmetic
88	IRET	Interrupt Return
89	HALT	Enter HALT Mode
8C	IRETINV	Interrupt Return and Invalidate
90,91	AND	AND Logical
92,93	OR	OR Logical
94,95	XOR	Exclusive-OR Logical
96,97	XNOR	Exclusive-NOR Logical
98,99	NOR	NOR Logical
9A,9B	NAND	NAND Logical
9C,9D	ANDN	AND-NOT Logical
9E	SETIP	Set Indirect Pointers
9F	INV	Invalidate
A0,A1	JMP	Jump
A4,A5	JMPF	Jump False
A8,A9	CALL	Call Subroutine
AC,AD	JMPT	Jump True

B4,B5	JMPFDEC	Jump False and Decrement
B6	MFTLB	Move from Translation Look-aside Buffer Register
BE	MTTLB	Move to Translation Look-aside Buffer Register
C0	JMPI	Jump Indirect
C4	JMPFI	Jump False Indirect
C6	MFSR	Move from Special Register
C8	CALLI	Call Subroutine, Indirect
CC	JMPTI	Jump True Indirect
CE	MTSR	Move to Special Register
E0	MULTIPLY	Integer Multiply
E1	DIVIDE	Integer Divide
E4	CVINTF	Convert Integer to Floating-Point Single-Precision
E5	CVINTD	Convert Integer to Floating-Point Double-Precision
E6	CVFINT	Convert Floating-Point Single-Precision to Integer
E7	CVDINT	Convert Floating-Point Double-Precision to Integer
E8	CVFD	Convert Floating-Point Single-Precision to Double-Precision
E9	CVDF	Convert Floating-Point Double-Precision to Single-Precision
EA	FEQ	Floating-Point Equal To, Single-Precision
EB	DEQ	Floating-Point Equal To, Double-Precision
EC	FGT	Floating-Point Greater Than, Single-Precision
ED	DGT	Floating-Point Greater Than, Double-Precision
EE	FLT	Floating-Point Less Than, Single-Precision
EF	DLT	Floating-Point Less Than, Double-Precision
F0	FADD	Floating-Point Add, Single-Precision
F1	DADD	Floating-Point Add, Double-Precision
F2	FSUB	Floating-Point Subtract, Single-Precision
F3	DSUB	Floating-Point Subtract, Double-Precision
F4	FMUL	Floating-Point Multiply, Single-Precision
F5	DMUL	Floating-Point Multiply, Double-Precision
F6	FDIV	Floating-Point Divide, Single-Precision
F7	DDIV	Floating-Point Divide, Double-Precision
F8	EMULATE	Trap to Software Emulation Routine

Appendices

APPENDIX A. CHANNEL OPERATION TIMING

Table A-1. Signal Summary

Signal Name	Signal Function	Type †	Synch Async
A0 - A31	Address Bus	3-state output	synch
*BGRT	Bus Grant	output	synch
*BINV	Bus Invalid	output	synch
*BREQ	Bus Request	input	synch
*CDA	Coprocessor Data Accept	input	synch
CNTL0 - CNTL1	CPU Control	input	async
D0 - D31	Data Bus	bi-directional	synch
*DBACK	Data Burst Acknowledge	input	synch
*DBREQ	Data Burst Request	3-state output	synch
*DERR	Data Error	input	synch
*DRDY	Data Ready	input	synch
*DREQ	Data Request	3-state output	synch
DREQT0 - DREQT1	Data Request Type	3-state output	synch
I0 - I31	Instruction Bus	input	synch
*IBACK	Instruction Burst Acknowledge	input	synch
*IBREQ	Instruction Burst Request	3-state output	synch
*IERR	Instruction Error	input	synch
INCLK	Input Clock	input	N/A
*INTR0 - *INTR3	Interrupt Request	input	async
*IRDY	Instruction Ready	input	synch
*IREQ	Instruction Request	3-state output	synch
IREQT	Instruction Request Type	3-state output	synch

08996A-A-T1

(Table Continued)

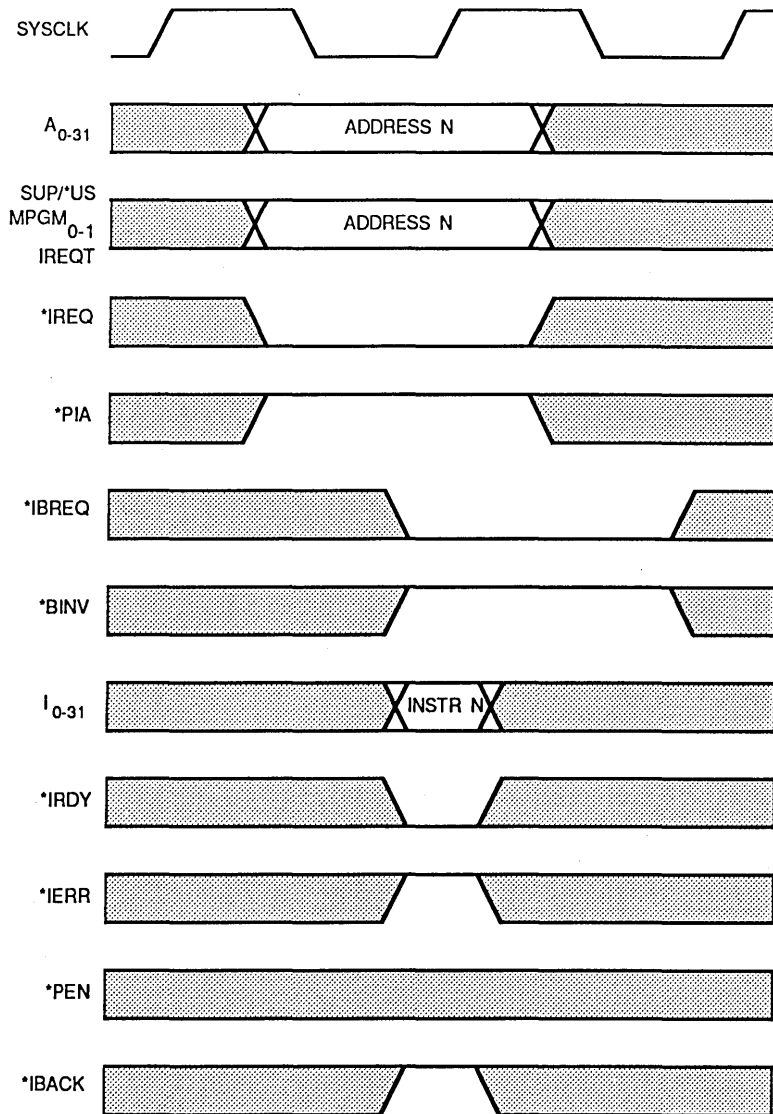
† The signals labeled " 3-state output" and "bi-directional" (except SYSClk) are disabled when the channel is granted to an external master. All outputs (except MSERR) may be disabled by asserting the *TEST input.

Table A-1. Signal Summary (continued)

Signal Name	Signal Function	Type †	Synch Async
*LOCK	Lock	3-state output	synch
MPGM0 - MPGM1	MMU Programmable	3-state output	synch
MSERR	Master/Saive Error	output	synch
OPT0 - OPT2	Option Control	3-state output	synch
*PDA	Pipelined Data Access	3-state output	synch
*PEN	Pipeline Enable	input	synch
*PIA	Pipelined Instruction Access	3-state output	synch
R/*W	Read/Write	3-state output	synch
*RESET	Reset	input	async
STAT0 - STAT2	CPU Status	output	synch
SUP/*US	Supervisor/User Mode	3-state output	synch
SYSCLK	System Clock	bi-directional	N/A
*TEST	Test Mode	input	async
*TRAP0 - *TRAP1	Trap Request	input	async
*WARN	Warn	edge-sensitive input	async


08996A-A-T2

† The signals labeled "3-state output" and "bi-directional" (except SYSCLK) are disabled when the channel is granted to an external master. All outputs (except MSERR) may be disabled by asserting the *TEST input.




INSTRUCTION READ - SIMPLE ACCESS

DIAGRAM 1
SIP01 SF 12-17-86

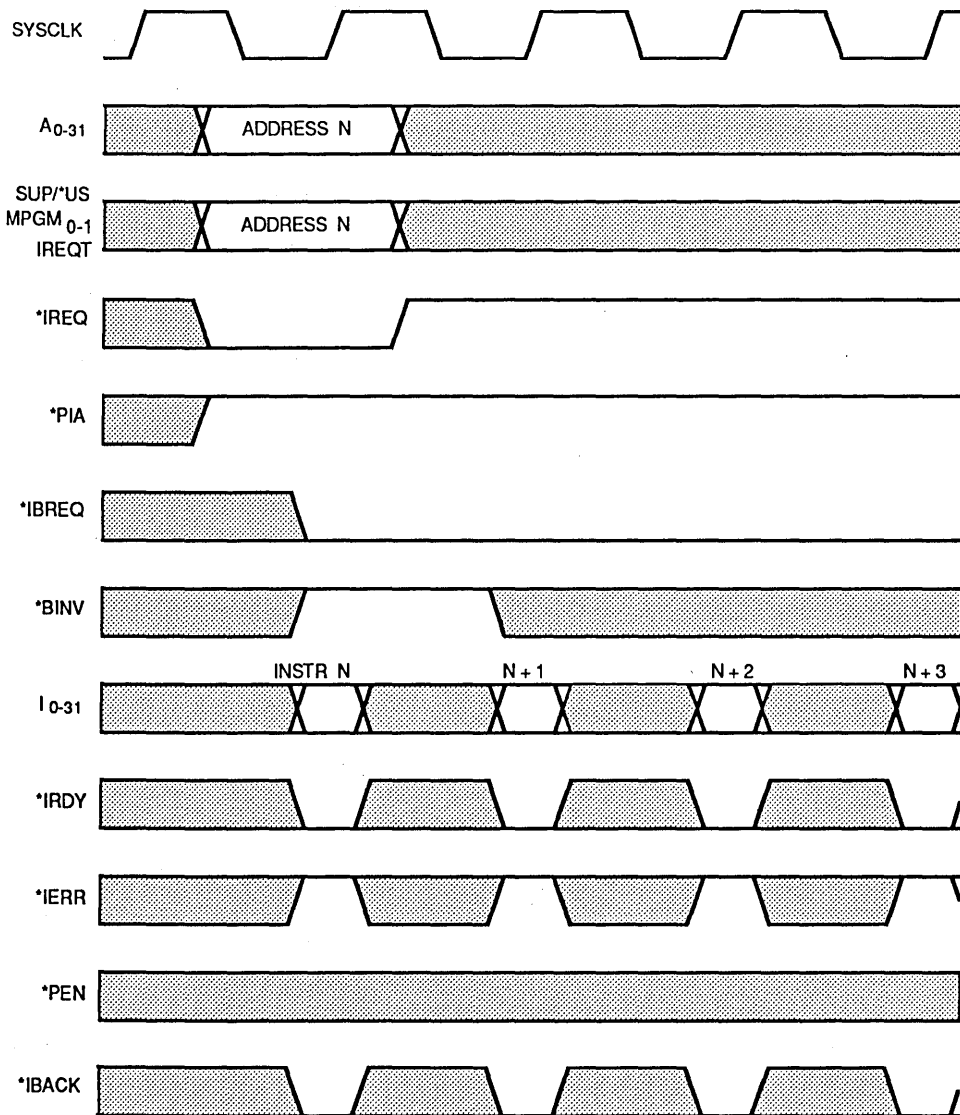


INSTRUCTION READ - SIMPLE ACCESS WITH $*IRDY$ DELAYED

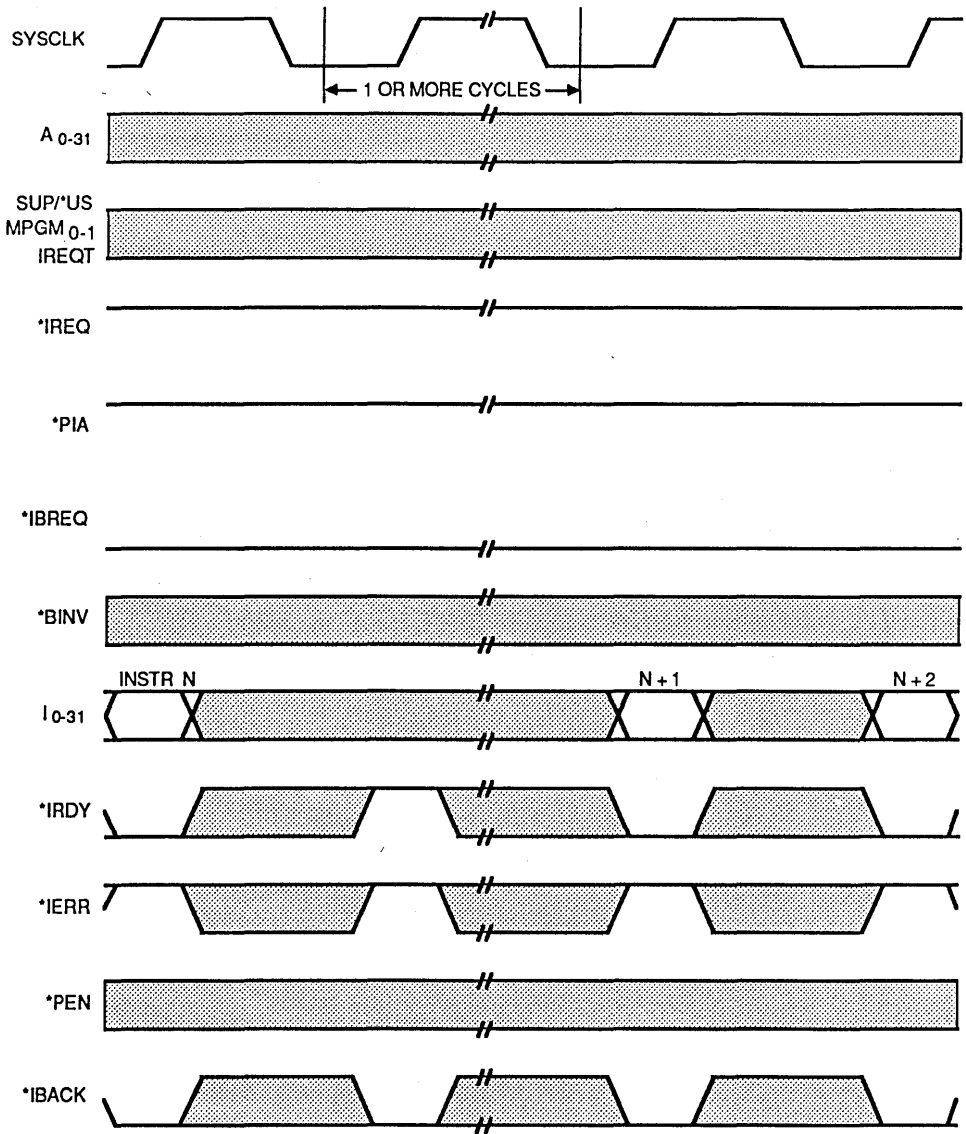


INSTRUCTION READ - PIPELINED ACCESS

DIAGRAM 3
SIP03 SF 12-17-86

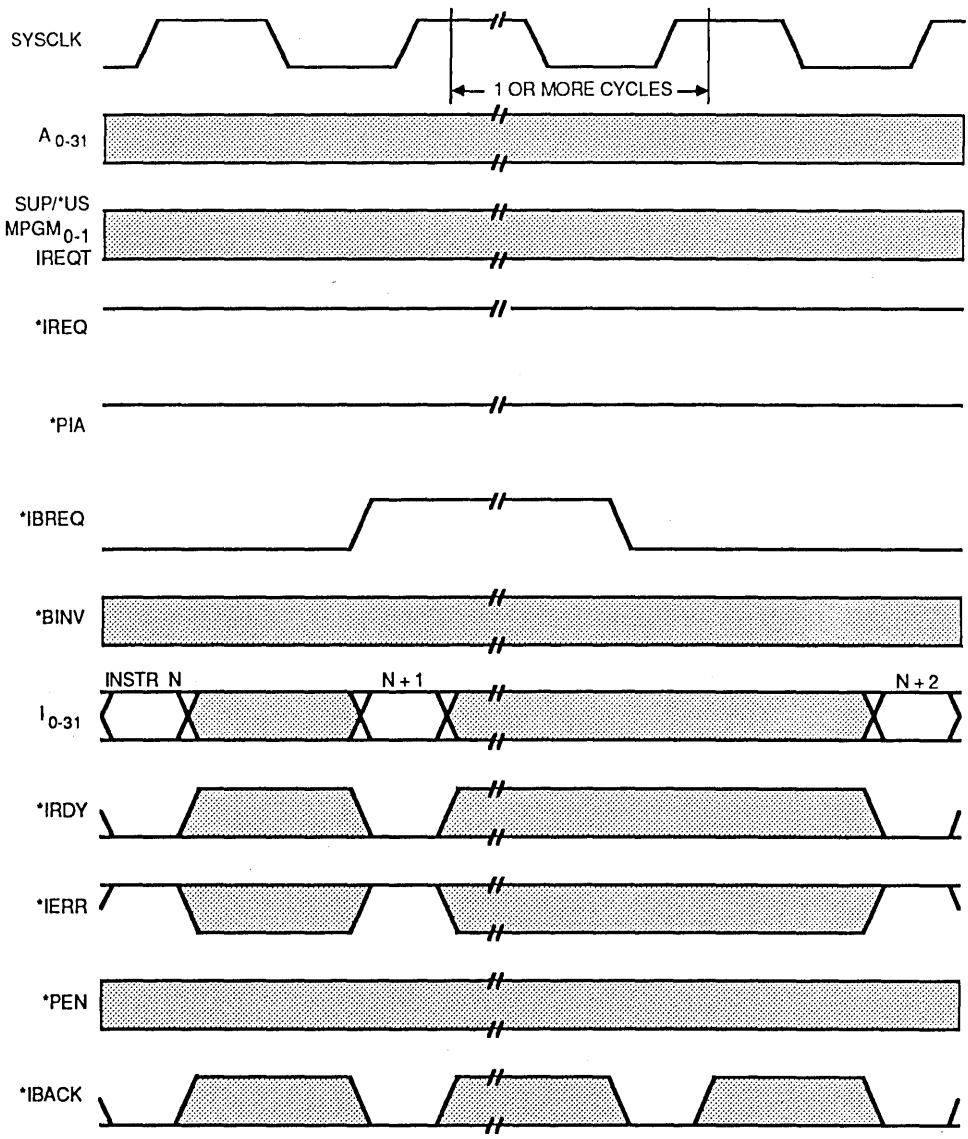


INSTRUCTION READ - ESTABLISHING BURST-MODE ACCESS




INSTRUCTION READ - BURST-MODE ACCESS SUSPENDED BY SLAVE

DIAGRAM 5
SIP05 SF 12-17-86

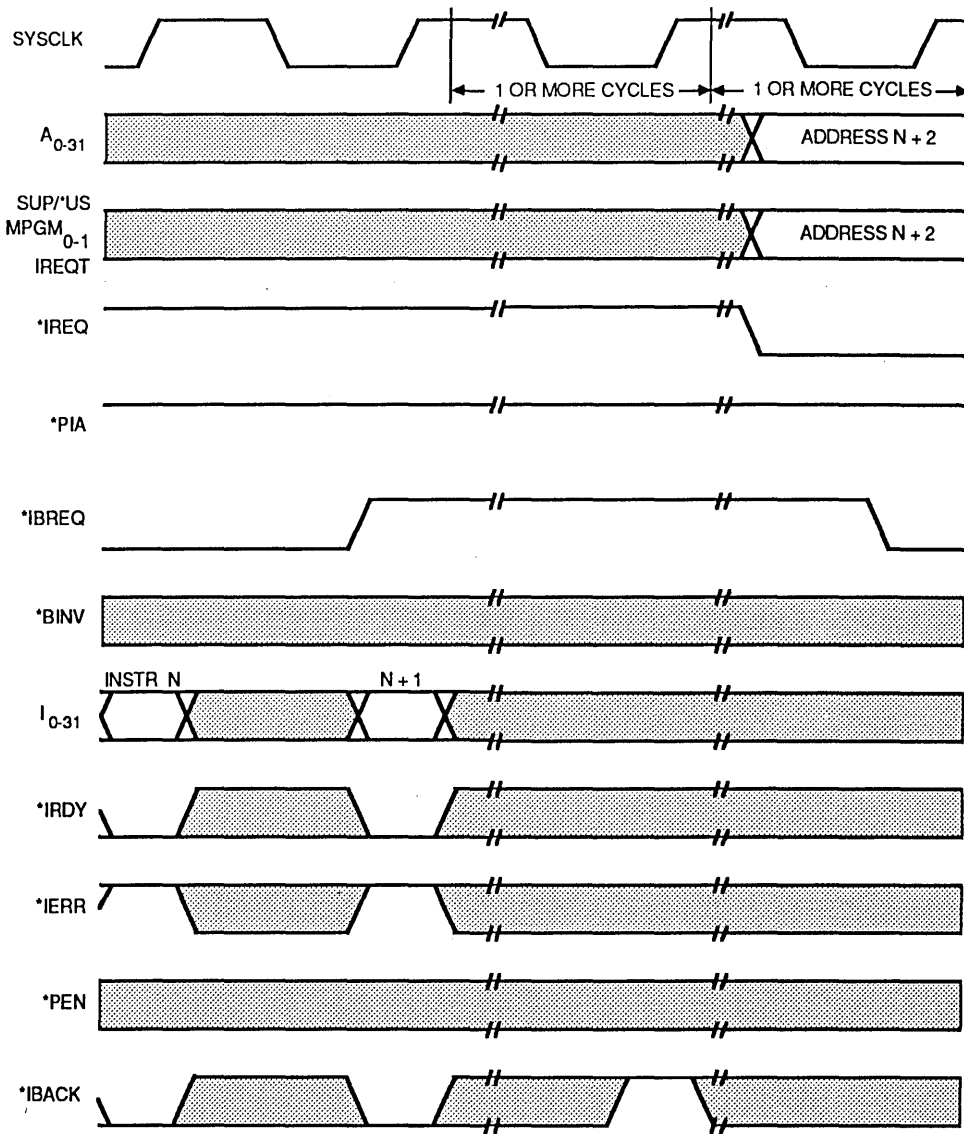


INSTRUCTION READ - BURST-MODE ACCESS SUSPENDED BY MASTER




INSTRUCTION READ - BURST-MODE ACCESS PREEMPTED BY SLAVE

DIAGRAM 7
SIP07 SF 12-17-86




INSTRUCTION READ - BURST-MODE ACCESS SUSPENDED BY MASTER
AND LATER PREEMPTED BY SLAVE




INSTRUCTION READ - BURST-MODE ACCESS CANCELLED BY SLAVE

DIAGRAM 8.5A
SIP36 SF 12-17-86




INSTRUCTION READ - BURST-MODE ACCESS ENDED BY MASTER
(PREEMPTED, TERMINATED, OR CANCELLED)

DIAGRAM 9
SIP08 SF 12-17-86



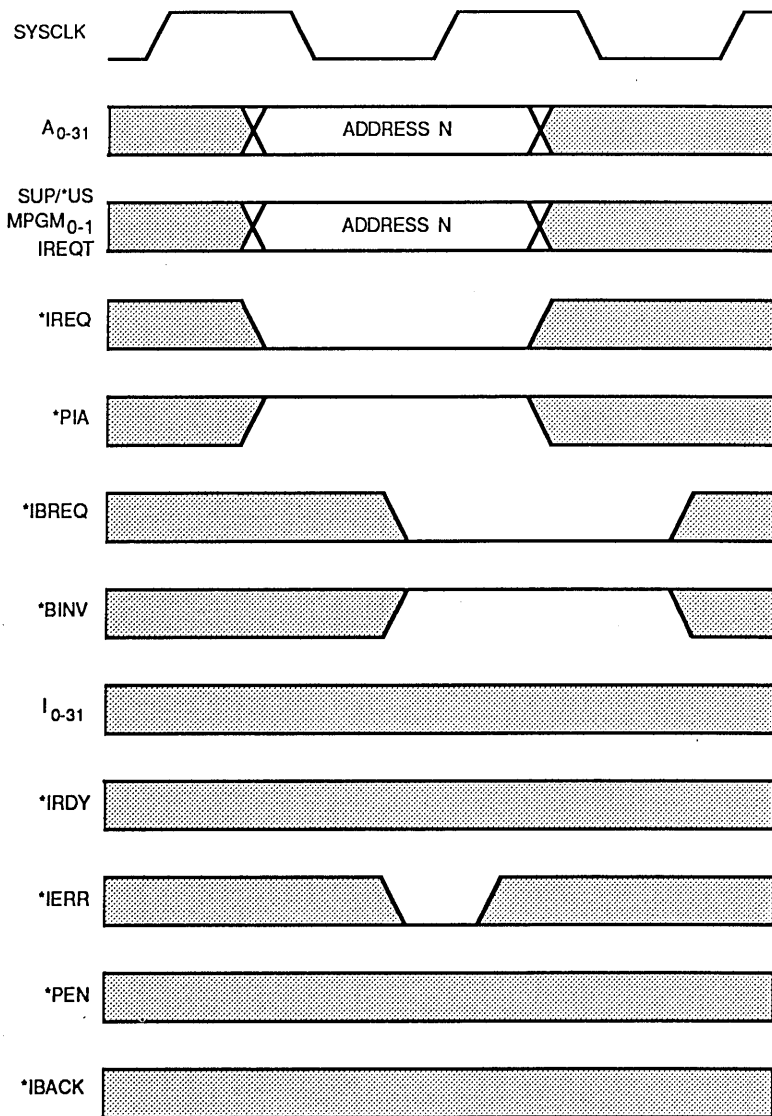
INSTRUCTION READ - TLB MISS OR PROTECTION VIOLATION

DIAGRAM 10
SIP09 S F 12-17-86




INSTRUCTION READ - PIPELINED ACCESS WITH TLB MISS OR PROTECTION VIOLATION

DIAGRAM 11A
SIP34 SF 12-17-86




INSTRUCTION READ - ERROR DETECTED BY SLAVE



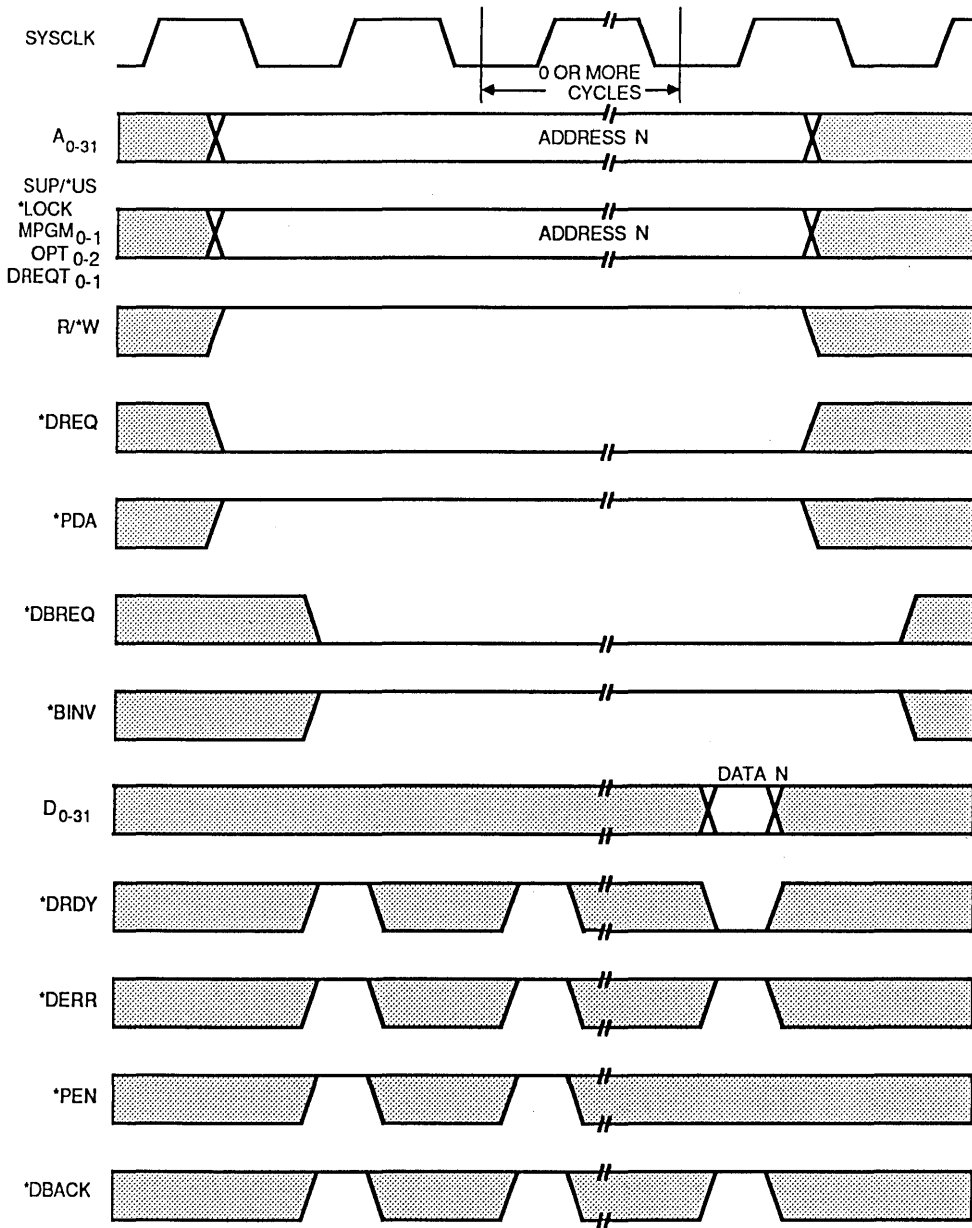
DATA READ - SIMPLE ACCESS

DIAGRAM 13
SIP11 SF 12-17-86




DATA WRITE - SIMPLE ACCESS

DIAGRAM 14
SIP12 SF 12-17-86



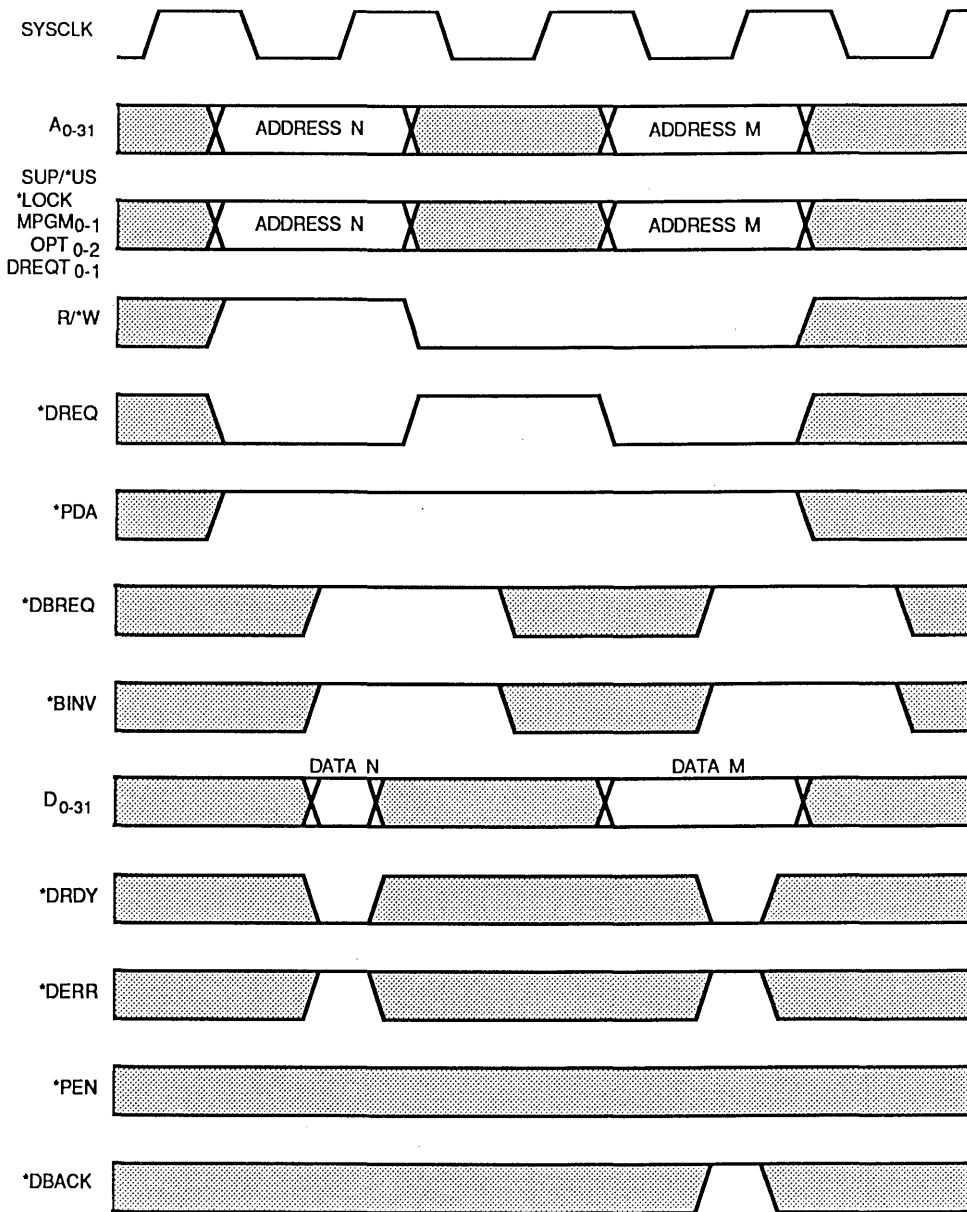
DATA READ - SIMPLE ACCESS WITH $*DRDY$ DELAYED

DIAGRAM 14.1
SIP40 SF 12-17-86




DATA WRITE - SIMPLE ACCESS WITH *DRDY DELAYED

DIAGRAM 14.2
SIP41 SF 12-17-86



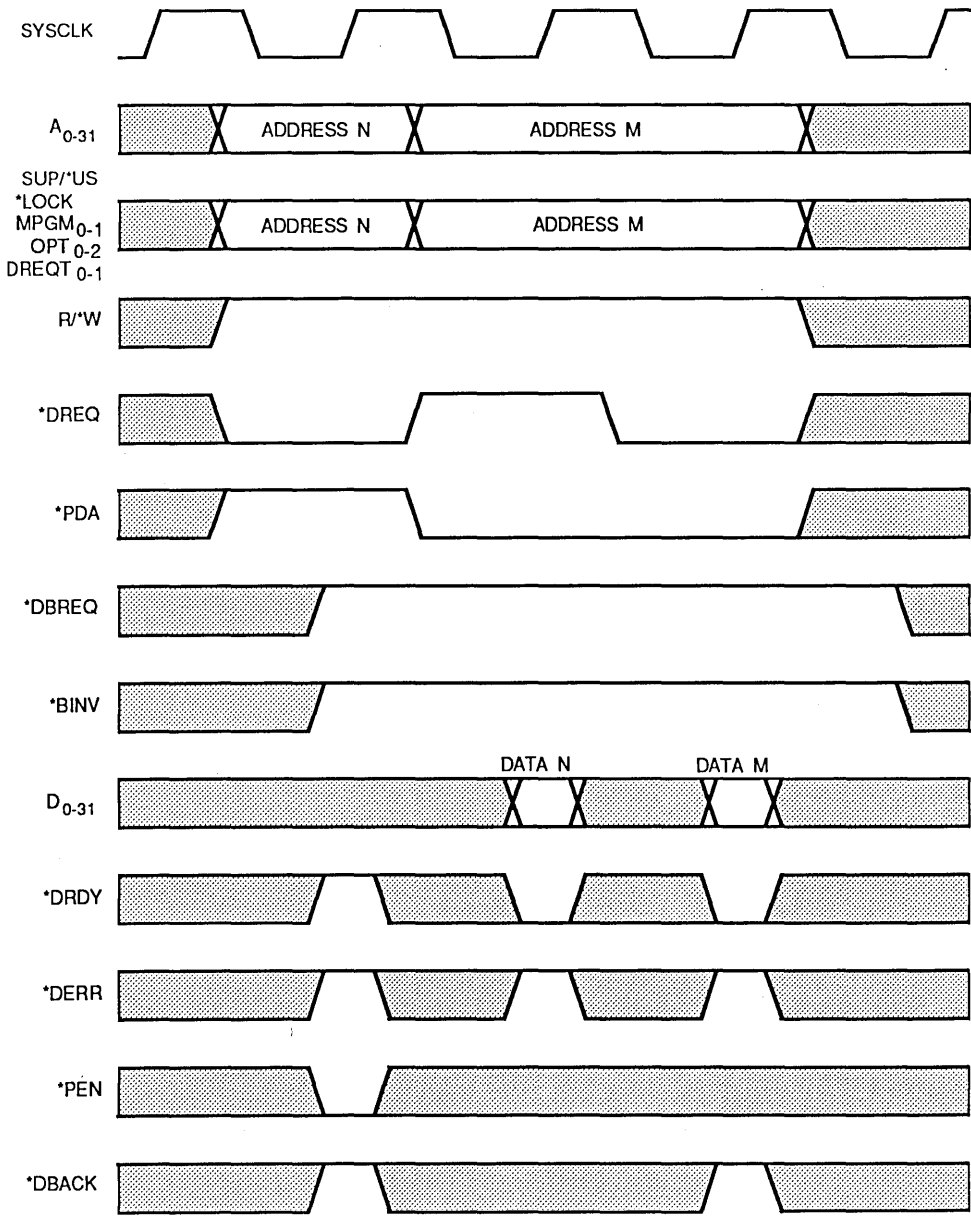
DATA READ FOLLOWED BY DATA WRITE - SIMPLE ACCESS

DIAGRAM 15
SIP13 SF 12-17-86




TEST AND SET INSTRUCTION

DIAGRAM 15.1
SIP46 SF 12-17-86

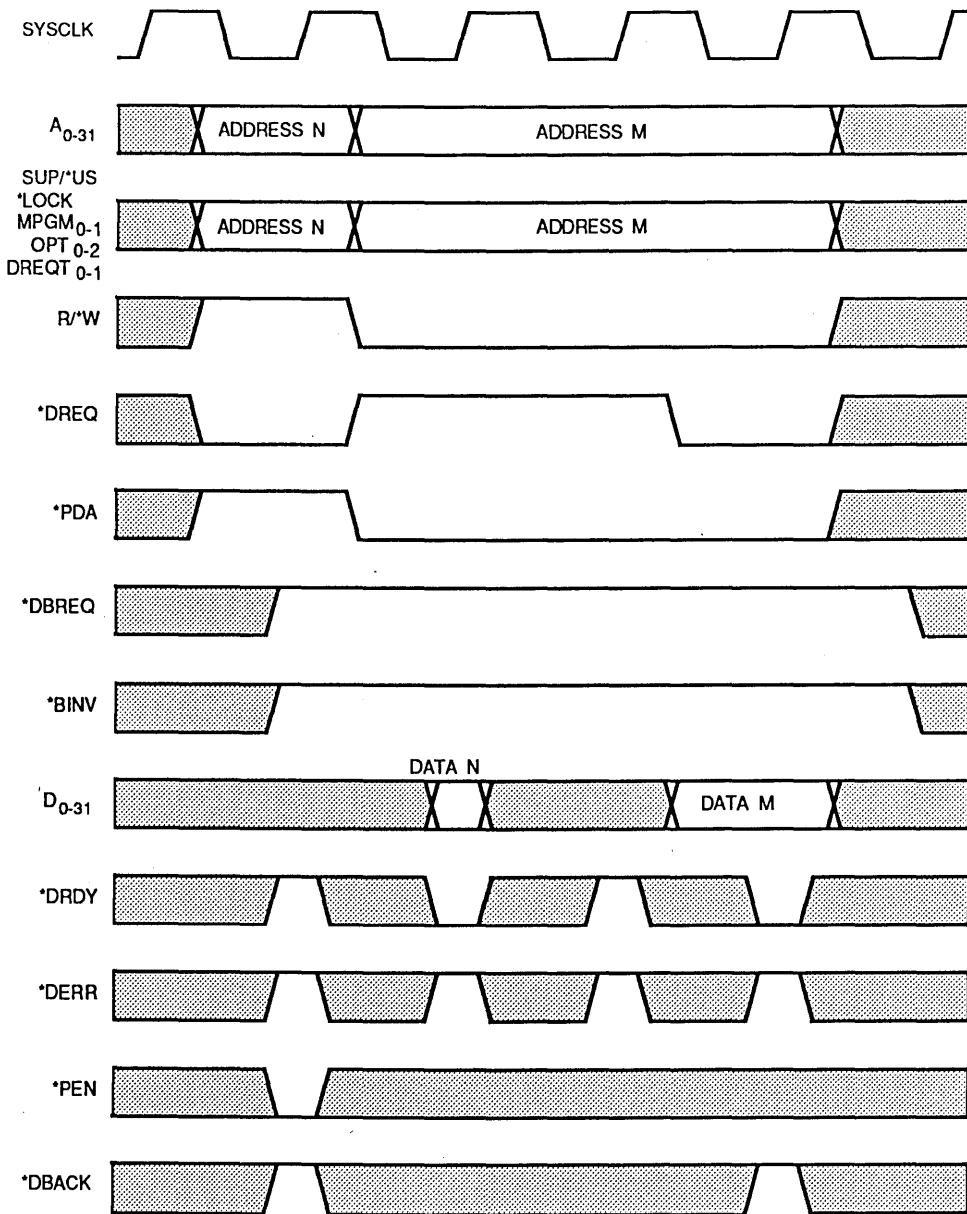


DATA READ - PIPELINED ACCESS



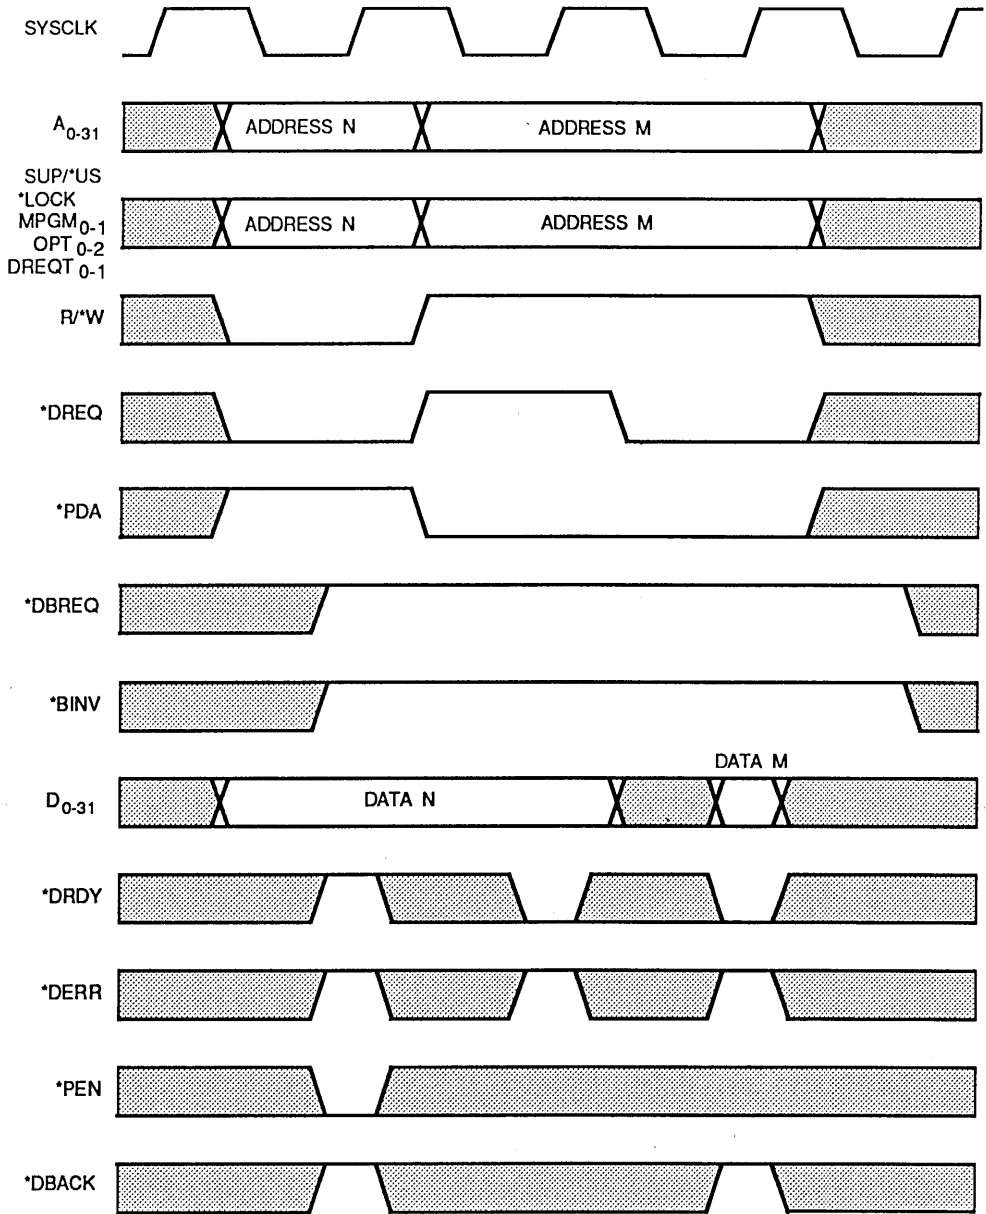
DATA WRITE - PIPELINED ACCESS

DIAGRAM 17
SIP15 SF 12-17-86



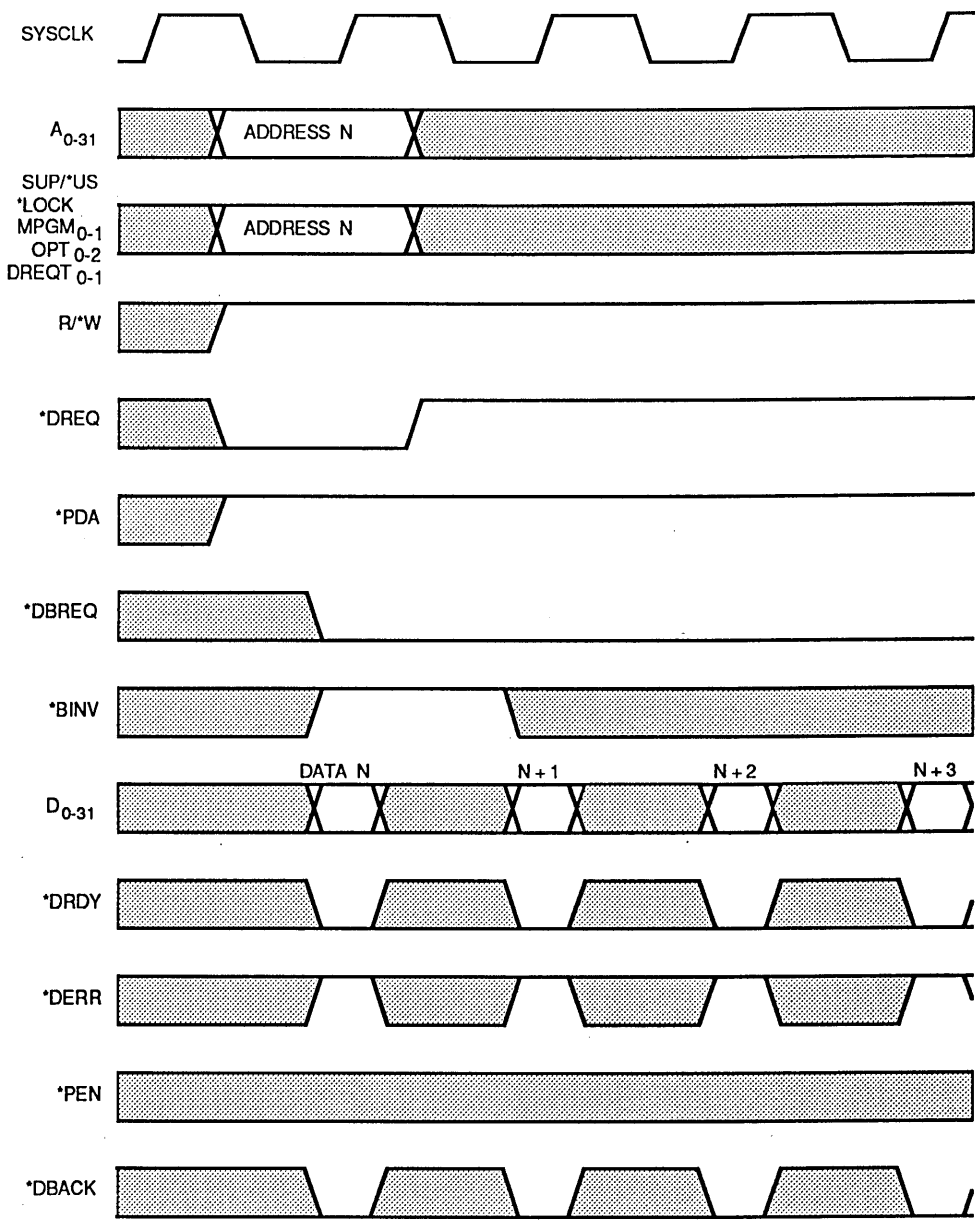
DATA READ FOLLOWED BY DATA WRITE - PIPELINED ACCESS
(NOT USED BY PROCESSOR)

DIAGRAM 18
SIP16 SF 12-17-86

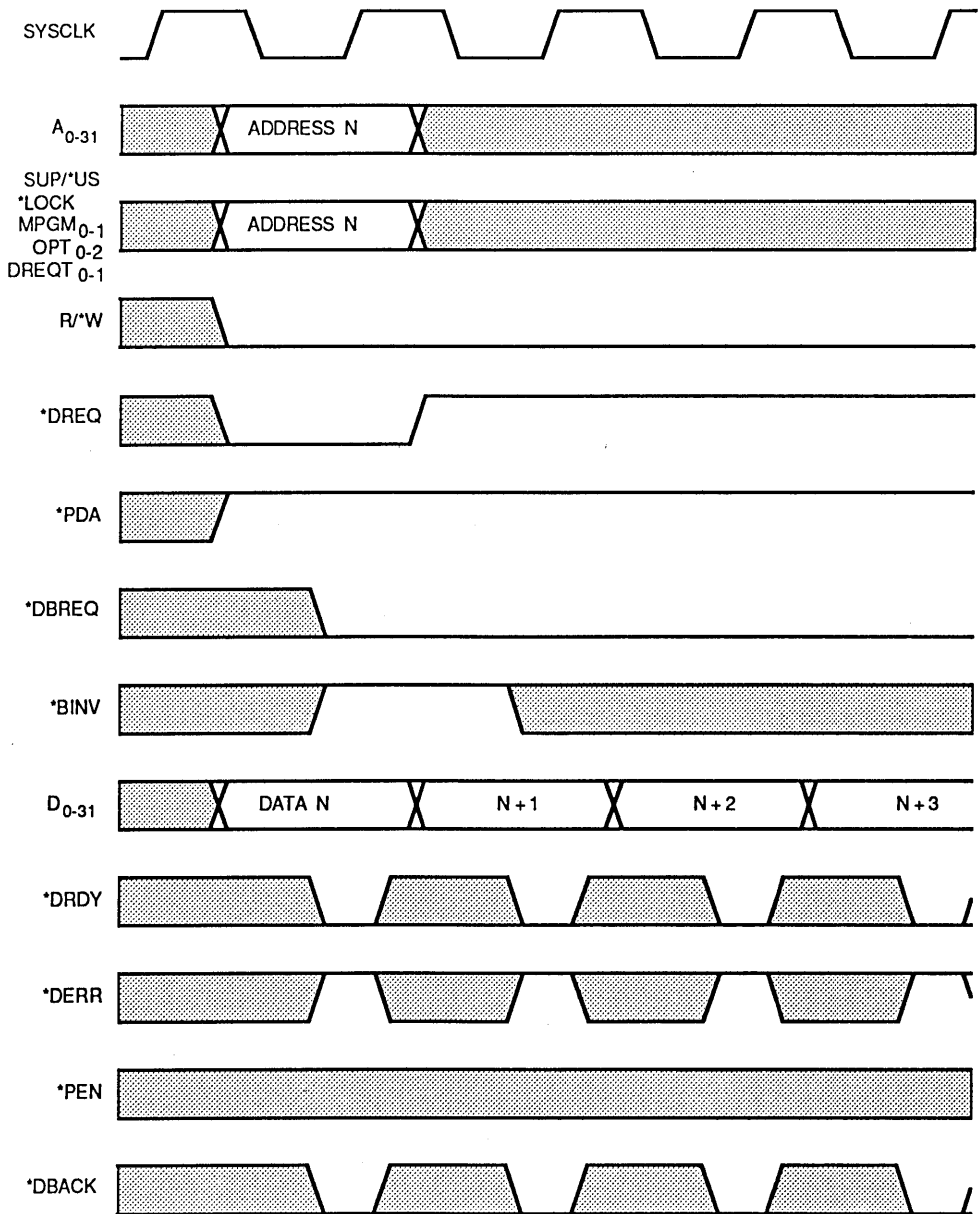


DATA WRITE FOLLOWED BY DATA READ - PIPELINED ACCESS

DIAGRAM 19
SIP17 SF 12-17-86

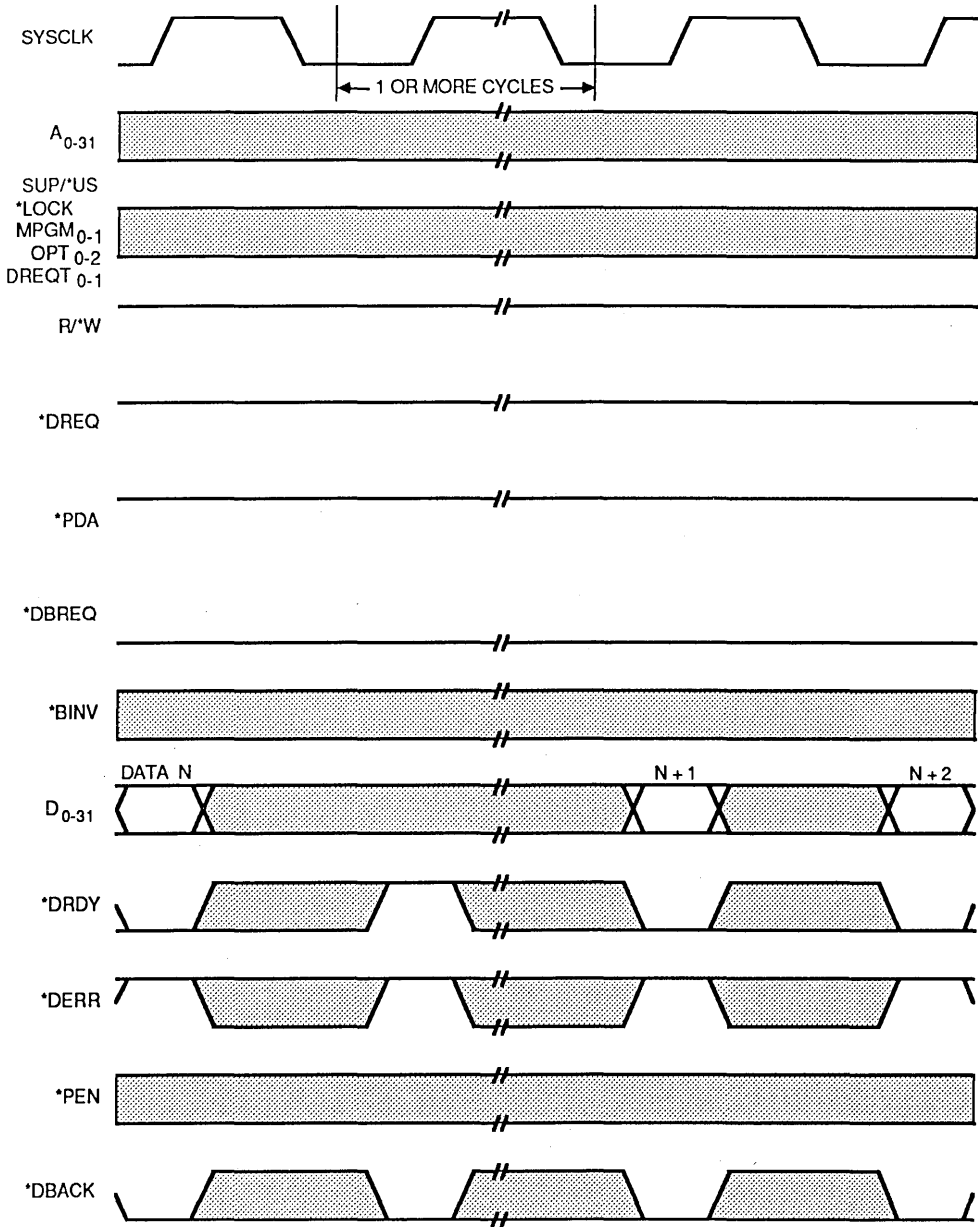


DATA READ - ESTABLISHING BURST-MODE ACCESS



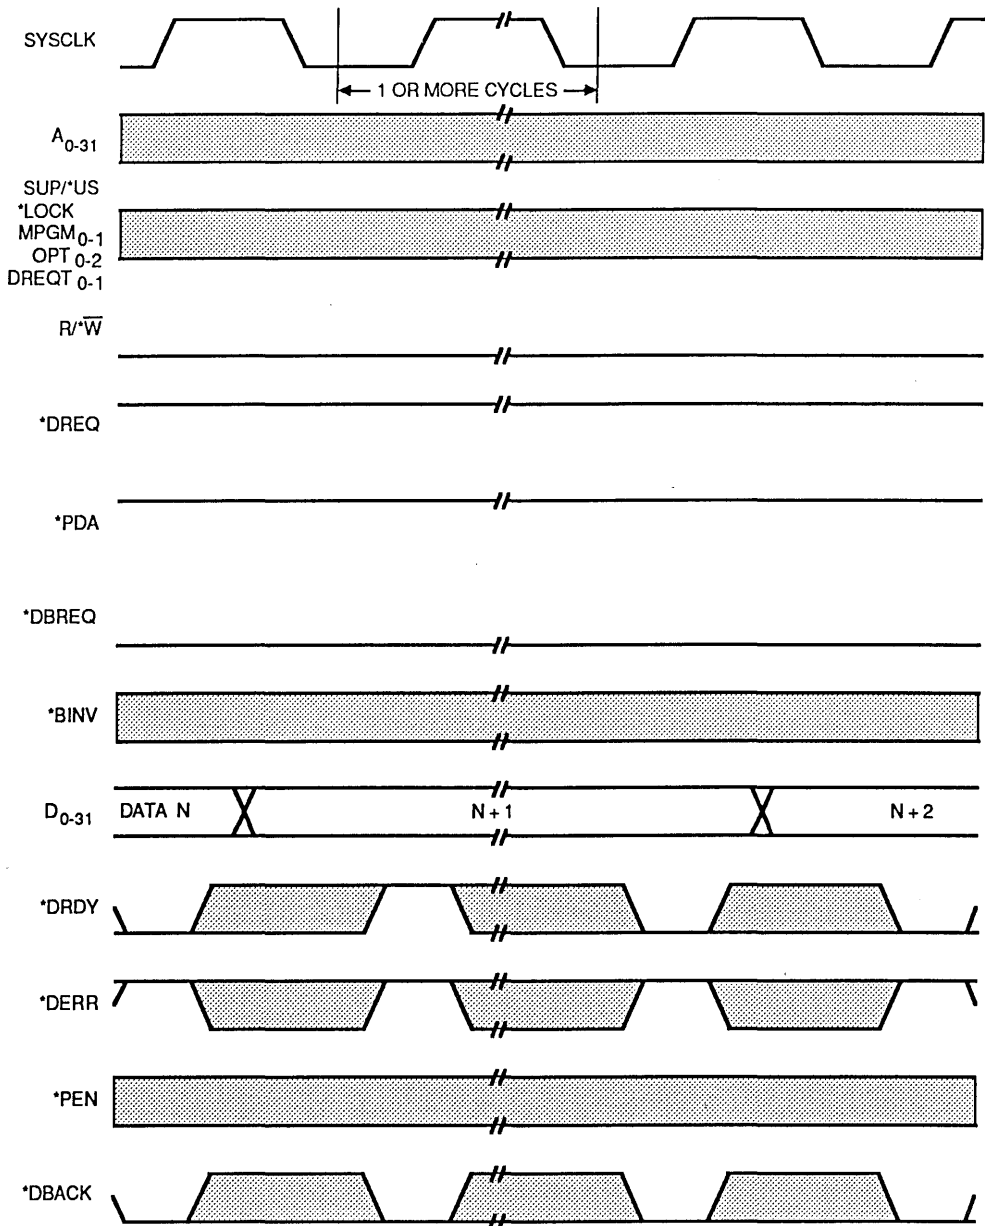
DATA WRITE - ESTABLISHING BURST-MODE ACCESS

DIAGRAM 21
SIP19 SF 12-17-86



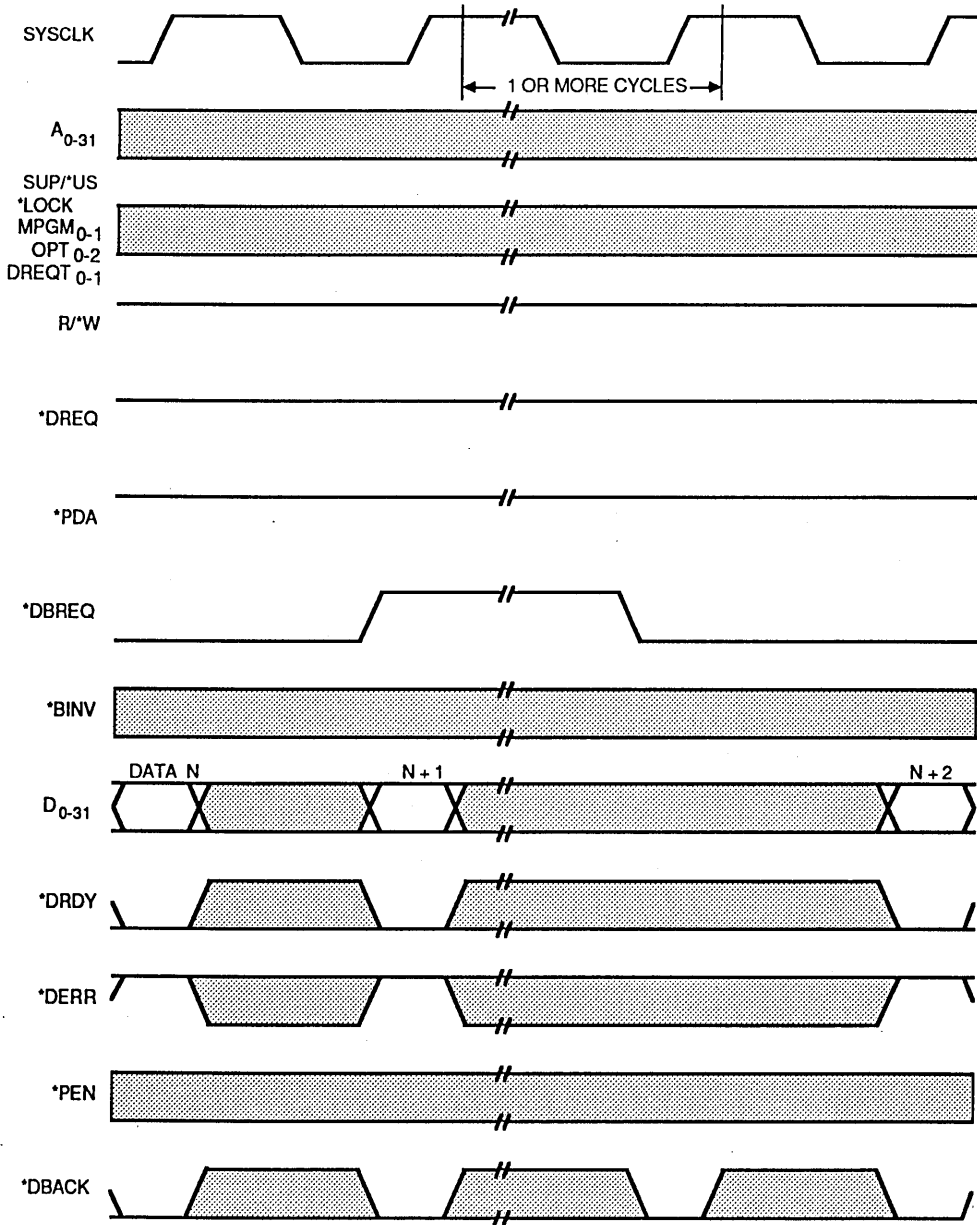
DATA READ - BURST-MODE ACCESS SUSPENDED BY SLAVE

DIAGRAM 22
SIP20 SF 12-17-86



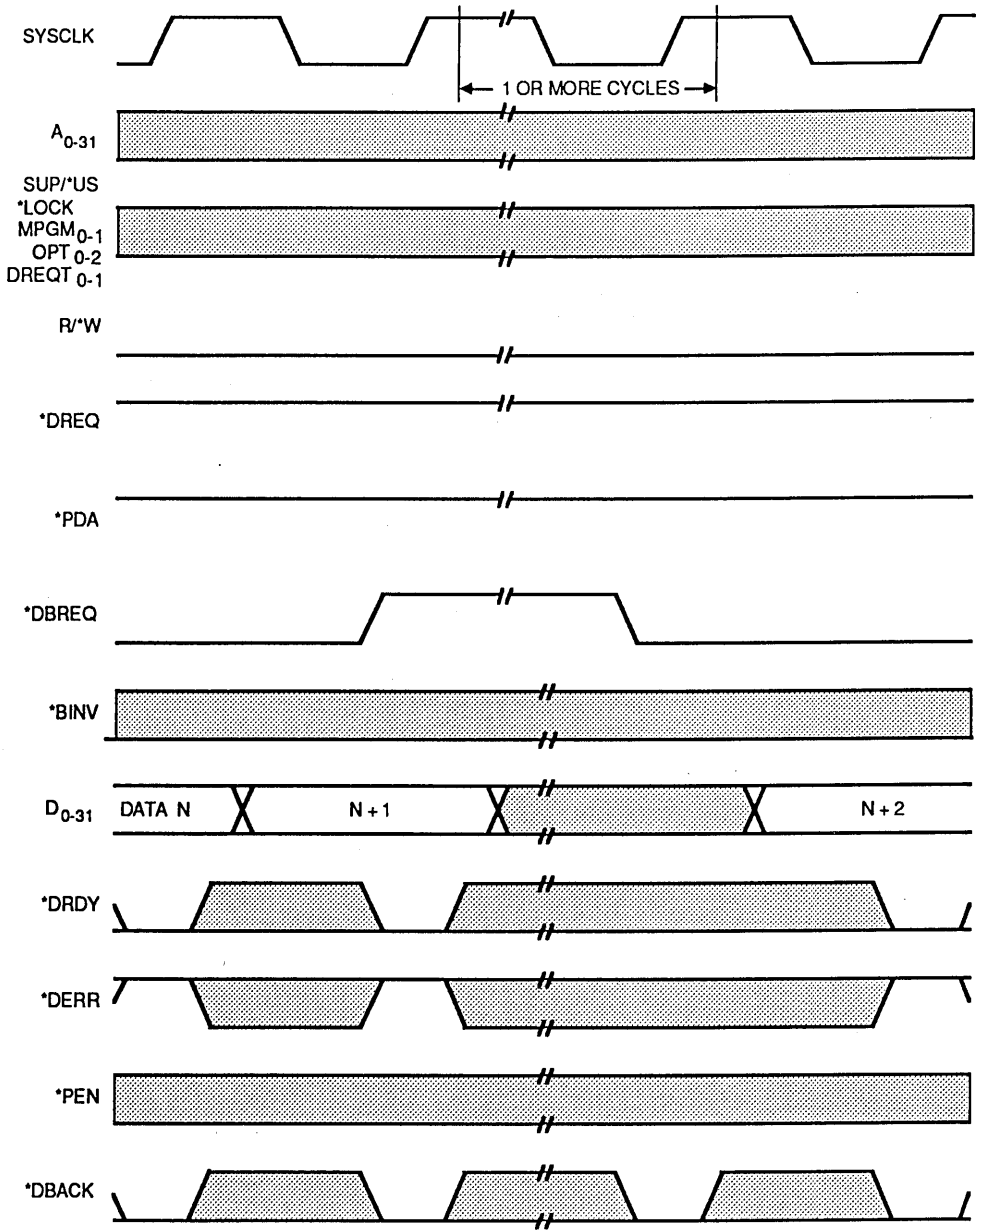
DATA WRITE - BURST-MODE ACCESS SUSPENDED BY SLAVE

DIAGRAM 23
SIP21 SF 12-17-86




DATA READ - BURST-MODE ACCESS SUSPENDED BY MASTER
(NOT USED BY PROCESSOR)

DIAGRAM 24
SIP22 SF 12-17-86




DATA WRITE - BURST-MODE SUSPENDED BY MASTER
(NOT USED BY PROCESSOR)

DIAGRAM 25
SIP23 SF 12-17-86



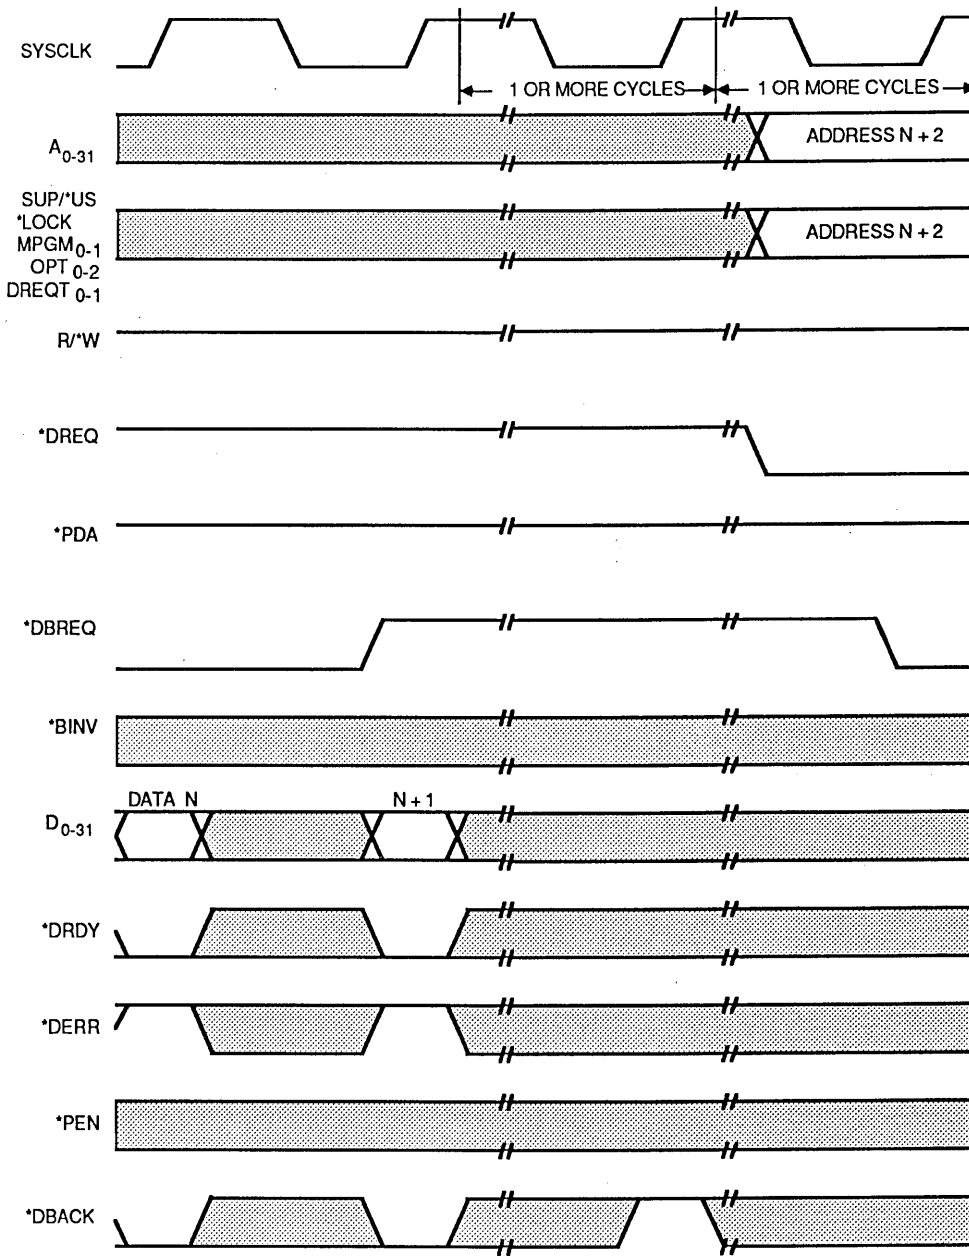
DATA READ - BURST-MODE ACCESS PREEMPTED BY SLAVE

DIAGRAM 26
SIP24 SF 12-17-86

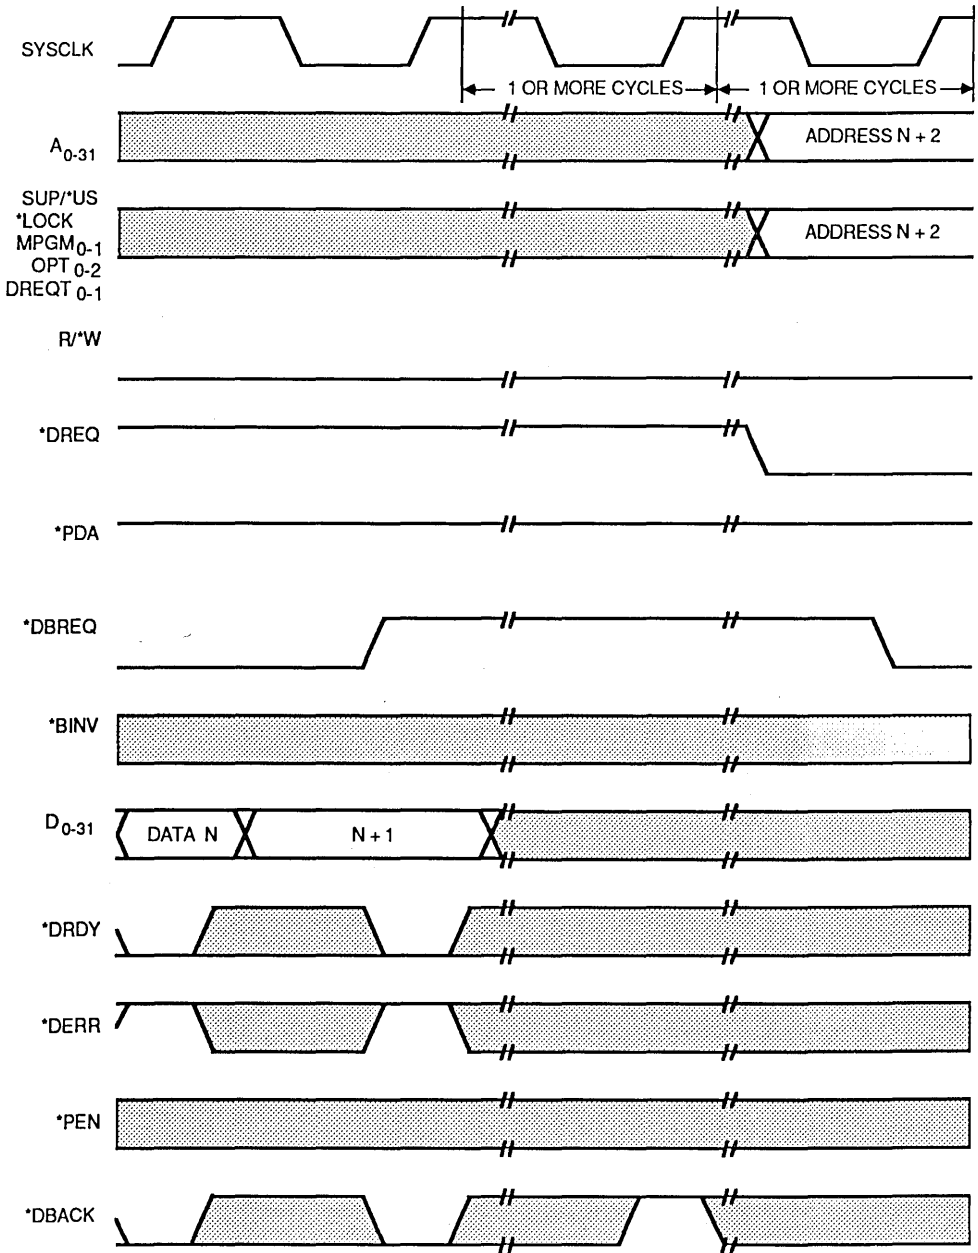


DATA WRITE - BURST-MODE ACCESS PREEMPTED BY SLAVE

DIAGRAM 26.5
SIP25 SF 12-17-86




DATA READ - BURST-MODE ACCESS SUSPENDED BY MASTER AND LATER PREEMPTED BY SLAVE (NOT USED BY PROCESSOR)




DATA WRITE - BURST-MODE ACCESS SUSPENDED BY MASTER AND LATER PREEMPTED BY SLAVE (NOT USED BY PROCESSOR)

DIAGRAM 26.7
SIP43 SF 12-17-86




DATA READ - BURST-MODE ACCESS CANCELLED BY SLAVE

DIAGRAM 27
SIP37 SF 12-17-86




DATA WRITE - BURST-MODE ACCESS CANCELLED BY SLAVE

DIAGRAM 27.5
SIP38 SF 12-17-86



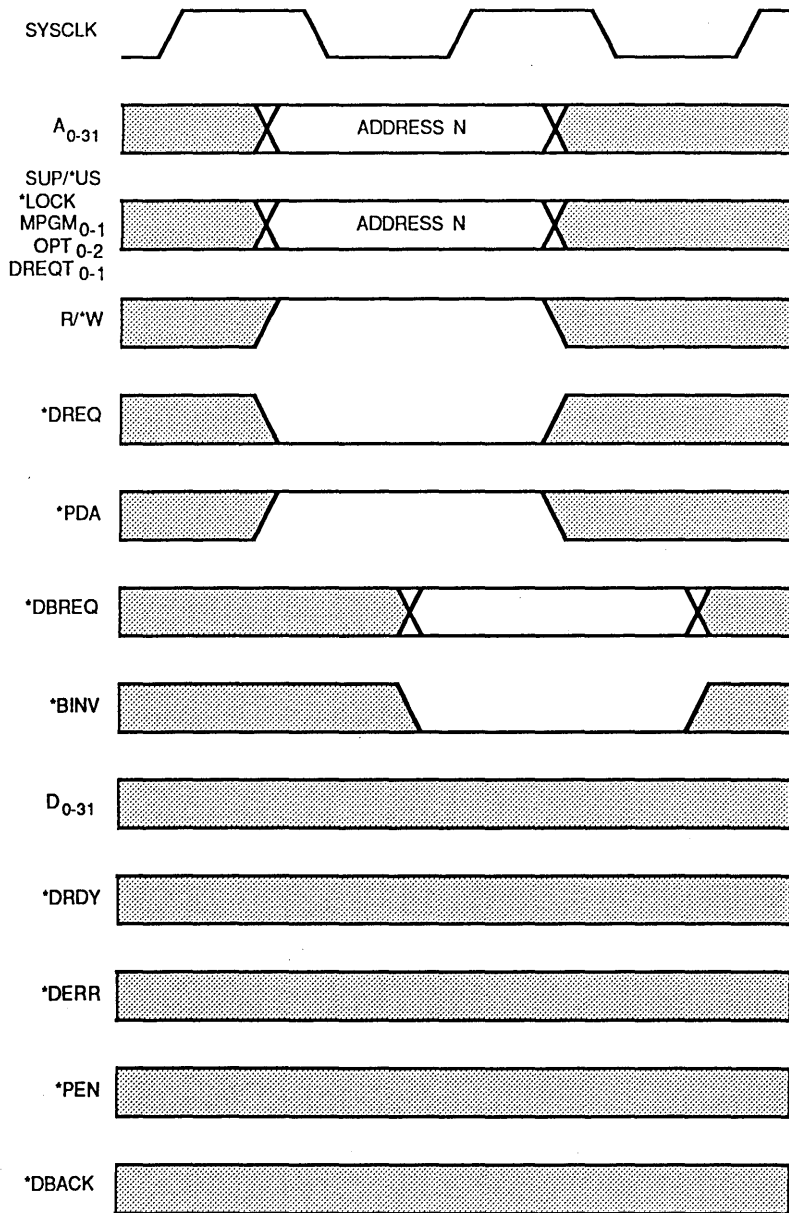
DATA READ - BURST-MODE ACCESS ENDED BY MASTER
(PREEMPTED, TERMINATED, OR CANCELLED)

DIAGRAM 28
SIP26 SF 12-17-86




DATA WRITE - BURST-MODE ACCESS ENDED BY MASTER
(PREEMPTED, TERMINATED, OR CANCELLED)

DIAGRAM 29
SIP27 SF 12-17-86




DATA READ - TLB MISS OR PROTECTION VIOLATION

DIAGRAM 30
SIP28 SF 12-17-86



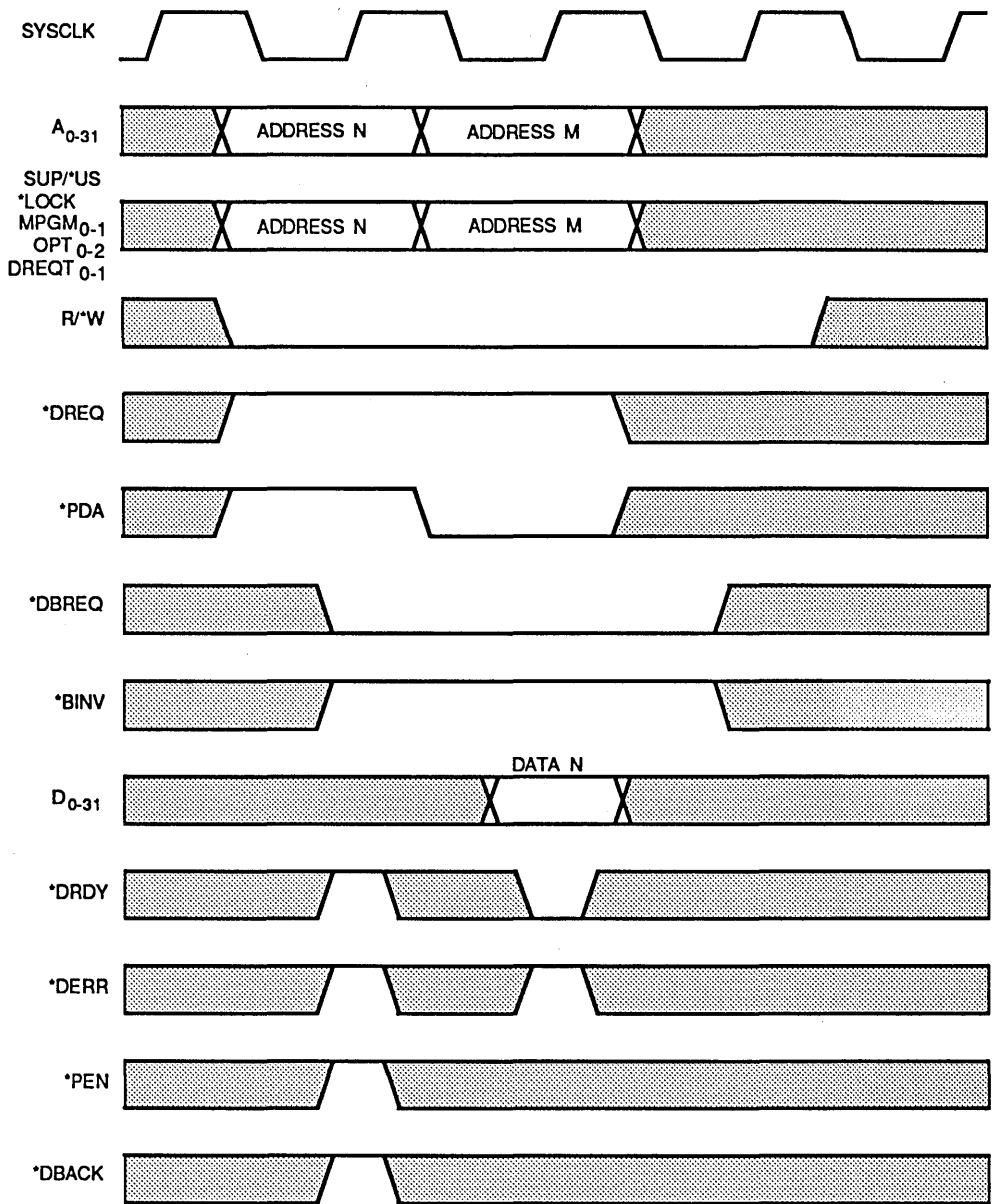
DATA WRITE - TLB MISS OR PROTECTION VIOLATION

DIAGRAM 31
SIP29 SF 12-17-86




DATA READ - PIPELINED ACCESS WITH TLB MISS OR PROTECTION VIOLATION

DIAGRAM 31.1A
SIP44 SF 12-17-86



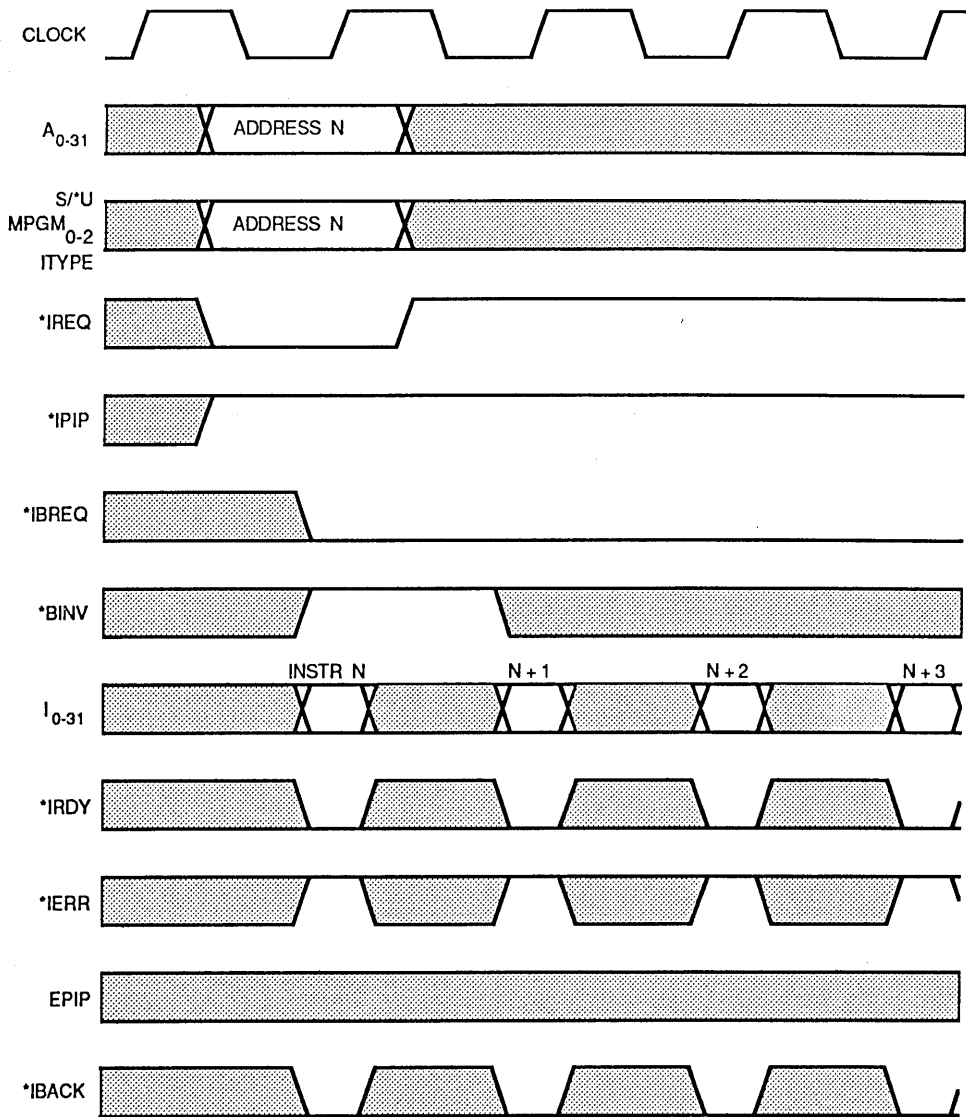
DATA WRITE - PIPELINED ACCESS WITH TLB MISS OR PROTECTION VIOLATION

DIAGRAM 31.2A
SIP45 SF 12-17-86

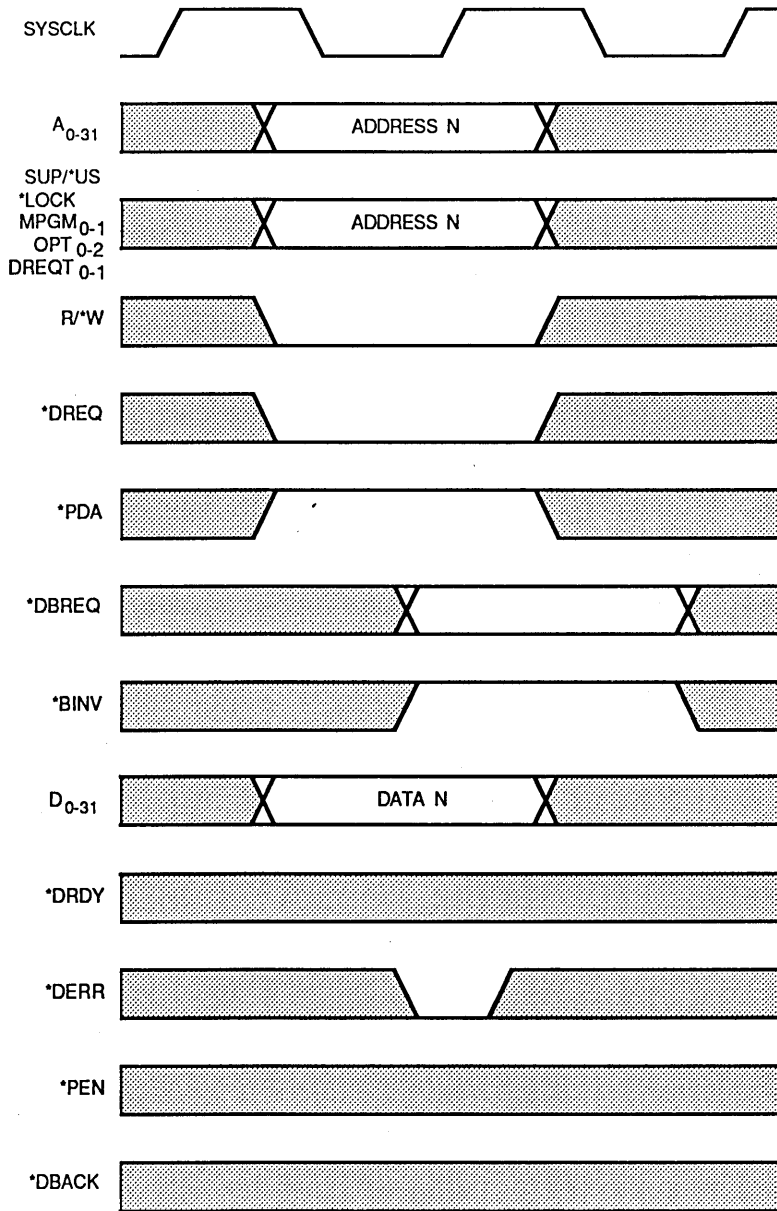


DATA READ - ERROR DETECTED BY SLAVE

DIAGRAM 32
SIP30 SF 12-17-86



INSTRUCTION READ - BURST TRANSFER CYCLE START



DATA WRITE - ERROR DETECTED BY SLAVE

DIAGRAM 33
SIP31 SF 12-17-86

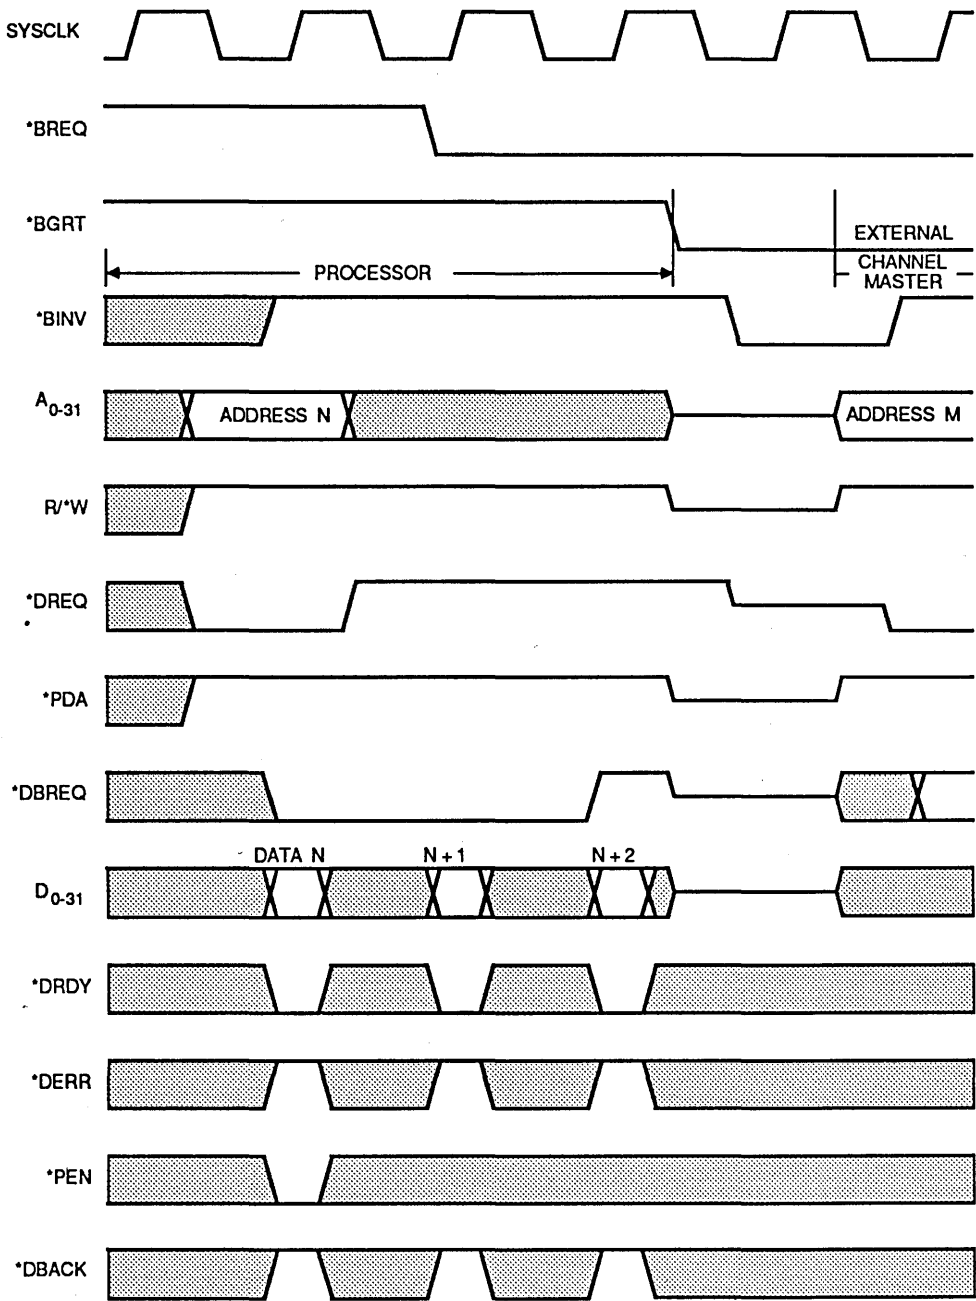


DIAGRAM 34A
SIP35 SF 12-17-86

CHANNEL TRANSFER FROM PROCESSOR TO EXTERNAL MASTER


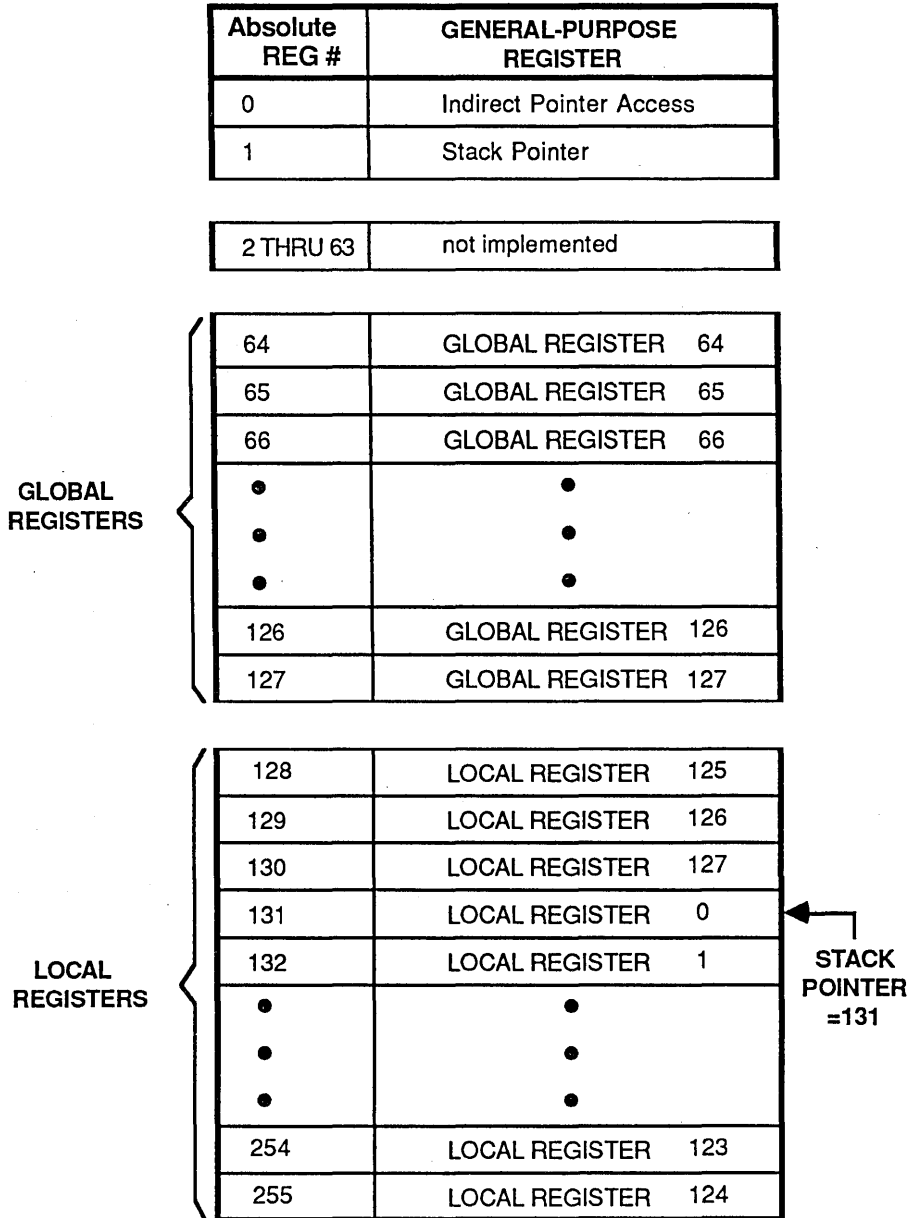


DIAGRAM 35A
SIP39 SF 12-17-86

CHANNEL TRANSFER FROM EXTERNAL MASTER TO PROCESSOR

APPENDIX B. REGISTER SUMMARY



08996A-Fig-B1

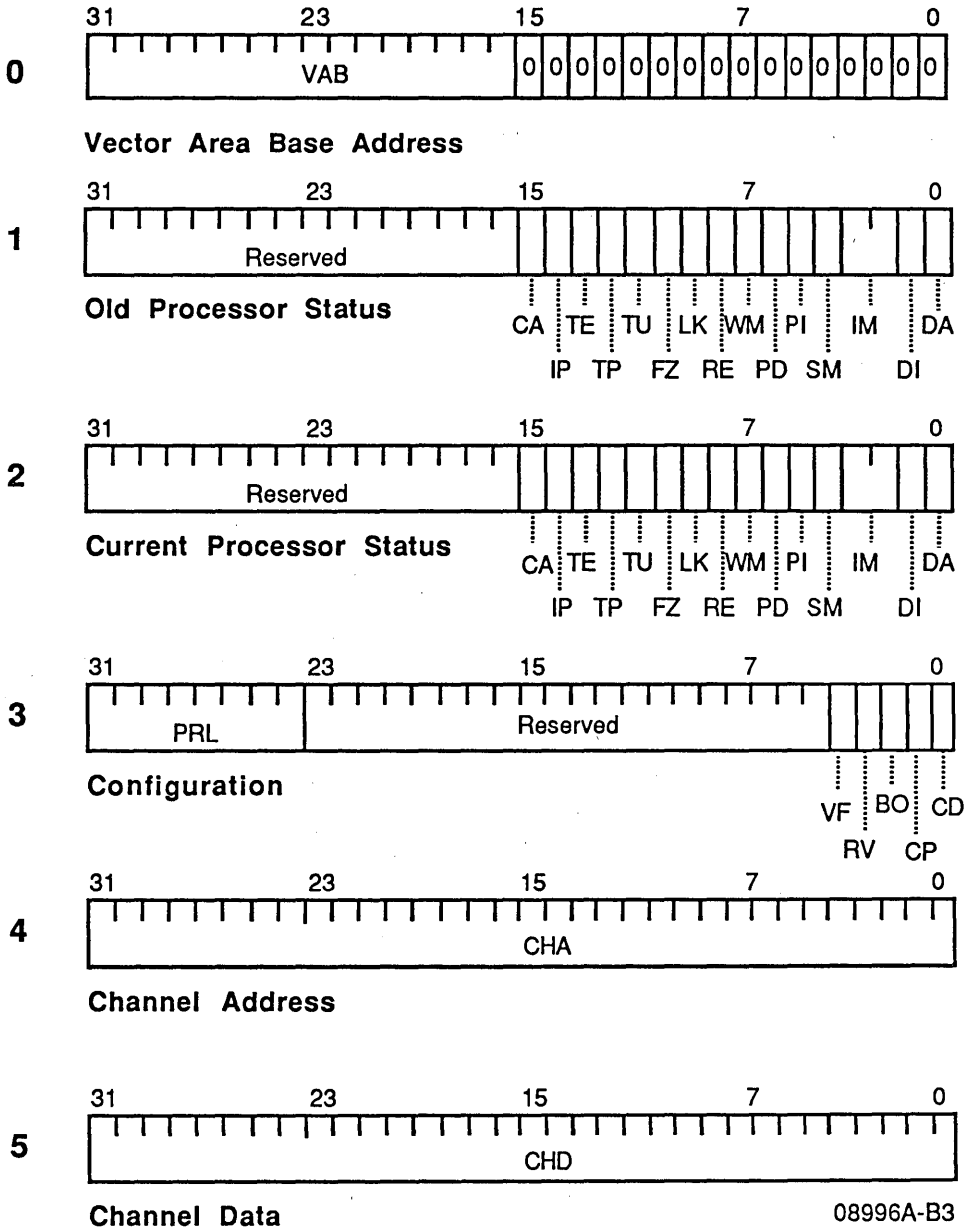
Figure B-1. General-Purpose Register Organization

Register Bank Protect Register Bit	Absolute Register-Numbers	General-Purpose Registers
0	2 thru 15	Bank 0 (unimplemented)
1	16 thru 31	Bank 1 (unimplemented)
2	32 thru 47	Bank 2 (unimplemented)
3	48 thru 63	Bank 3 (unimplemented)
4	64 thru 79	Bank 4
5	80 thru 95	Bank 5
6	96 thru 111	Bank 6
7	112 thru 127	Bank 7
8	128 thru 143	Bank 8
9	144 thru 159	Bank 9
10	160 thru 175	Bank 10
11	176 thru 191	Bank 11
12	192 thru 207	Bank 12
13	208 thru 223	Bank 13
14	224 thru 239	Bank 14
15	240 thru 255	Bank 15

08996A-Fig-B2

Figure B-2. Register Bank Organization

REG #



08996A-B3

Figure B-3. Special Purpose Registers (continued)

REG #




Figure B-3. Special Purpose Registers (continued)

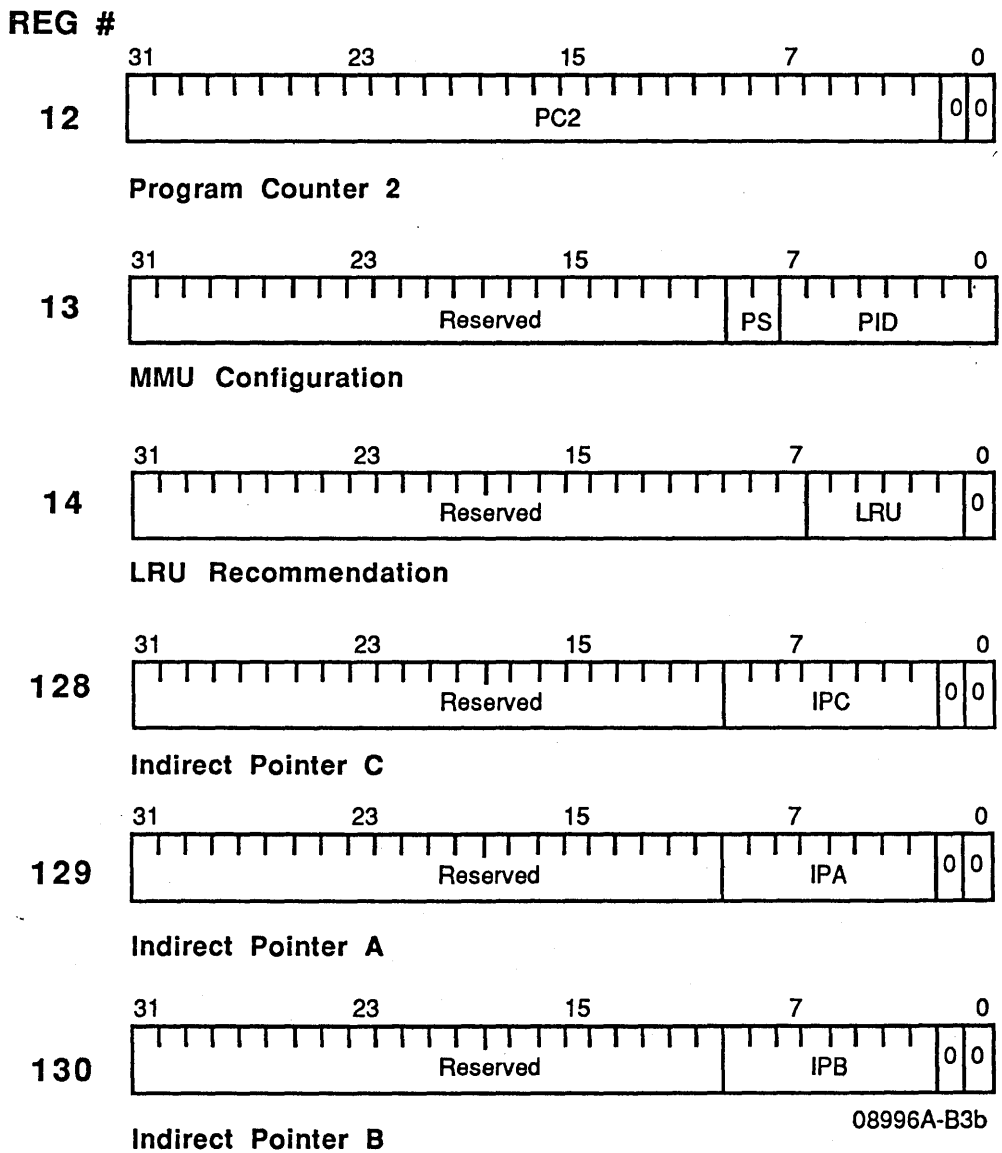






Figure B-3. Special Purpose Registers (continued)


REG #



Byte Pointer



Funnel Shift Count



Load/Store Count Remaining

08996A-Fig-B3c

Figure B-3. Special Purpose Registers

REG #	TLB Set 0
0	TLB Entry Line 0 Word 0
1	TLB Entry Line 0 Word 1
2	TLB Entry Line 1 Word 0
3	TLB Entry Line 1 Word 1
⋮	
60	TLB Entry Line 30 Word 0
61	TLB Entry Line 30 Word 1
62	TLB Entry Line 31 Word 0
63	TLB Entry Line 31 Word 1

TLB Set 1	
64	TLB Entry Line 0 Word 0
65	TLB Entry Line 0 Word 1
66	TLB Entry Line 1 Word 0
67	TLB Entry Line 1 Word 1
⋮	
124	TLB Entry Line 30 Word 0
125	TLB Entry Line 30 Word 1
126	TLB Entry Line 31 Word 0
127	TLB Entry Line 31 Word 1

08996A-Fig-B4

Figure B-4. Translation Look-aside Buffer Registers




Figure B-5. Translation Look-aside Buffer Entries

Table B-6. Register Field Summary

Label	Field Name	Register	Bit
B0	Bank 0 Protection Bit	Register Bank Protect	0
B1	Bank 1 Protection Bit	Register Bank Protect	1
B2	Bank 2 Protection Bit	Register Bank Protect	2
B3	Bank 3 Protection Bit	Register Bank Protect	3
B4	Bank 4 Protection Bit	Register Bank Protect	4
B5	Bank 5 Protection Bit	Register Bank Protect	5
B6	Bank 6 Protection Bit	Register Bank Protect	6
B7	Bank 7 Protection Bit	Register Bank Protect	7
B8	Bank 8 Protection Bit	Register Bank Protect	8
B9	Bank 9 Protection Bit	Register Bank Protect	9
B10	Bank 10 Protection Bit	Register Bank Protect	10
B11	Bank 11 Protection Bit	Register Bank Protect	11
B12	Bank 12 Protection Bit	Register Bank Protect	12
B13	Bank 13 Protection Bit	Register Bank Protect	13
B14	Bank 14 Protection Bit	Register Bank Protect	14
B15	Bank 15 Protection Bit	Register Bank Protect	15
BO	Byte Order	Configuration	2
BP	Byte Pointer	ALU Status Byte Pointer	6-5 1-0
C	Carry	ALU Status	7
CA	Coprocessor Active	Current Processor Status Old Processor Status	15 15
CD	Branch Target Cache Disable	Configuration	0
CE	Coprocessor Enable	Channel Control	31
CHA	Channel Address	Channel Address	31-0
CHD	Channel Data	Channel Data	31-0
CNTL	Control	Channel Control	30-24

08996A-Fig-B6

(Table continued)

Table B-6. Register Field Summary (continued)

Label	Field Name	Register	Bit
CP	Coprocessor Present	Configuration	1
CR	Load/Store Count Remaining	Channel Control	23-16
		Load/Store Count Remaining	7-0
CV	Contents Valid	Channel Control	0
DA	Disable All Interrupts and Traps	Current Processor Status	0
		Old Processor Status	0
DF	Divide Flag	ALU Status	11
DI	Disable Interrupts	Current Processor Status	1
		Old Processor Status	1
F	Flag	TLB Entry Word 1	0
FC	Funnel Shift Count	ALU Status	4-0
		Funnel Shift Count	4-0
FZ	Freeze	Current Processor Status	10
		Old Processor Status	10
IE	Interrupt Enable	Timer Reload	24
IM	Interrupt Mask	Current Processor Status	3-2
		Old Processor Status	3-2
IN	Interrupt	Timer Reload	25
IP	Interrupt Pending	Current Processor Status	14
		Old Processor Status	14
IPA	Indirect Pointer A	Indirect Pointer A	9-2
IPB	Indirect Pointer B	Indirect Pointer B	9-2
IPC	Indirect Pionter C	Indirect Pionter C	9-2
LA	Look Active	Channel Control	12
LK	Lock	Current Processor Status	9
		Old Processor Status	9

08996A-Fig-B6a

(Table continued)

Table B-6. Register Field Summary (continued)

Label	Field Name	Register	Bit
LRU	Least Recently Used Entry	LRU Recommendation	6-1
LS	Load /Store	Channel Control	15
ML	Multiple Operation	Channel Control	14
N	Negative	ALU Status	9
NN	Not Needed	Channel Control	1
OV	Overflow	Timer Reload	26
PC0	Program Counter 0	Program Counter 0	31-2
PC1	Program Counter 1	Program Counter 1	31-2
PC2	Program Counter 2	Program Counter 2	31-2
PD	Physical Addressing/Data	Current Processor Status	6
		Old Processor Status	6
PGM	User Programmable	TLB Entry Word 1	7-6
PI	Physical Addressing/Instructions	Current Processor Status	5
		Old Processor Status	5
PID	Process Identifier	MMU Configuration	7-0
PRL	Processor Release Level	Configuration	31-24
PS	Page Size	MMU Configuration	9-8
Q	Quotient/Multiplier	Q Register	31-0
RE	ROM Enable	Current Processor Status	8
		Old Processor Status	8
RPN	Real Page Number	TLB Entry Word 1	31-10
RV	ROM Vector Area	Configuration	3
SE	Supervisor Execute	TLB Entry Word 0	11
SM	Supervisor Mode	Current Processor Status	4
		Old Processor Status	4

08996A-Fig-B6b

(Table continued)

Table B-6. Register Field Summary (continued)

Label	Field Name	Register	Bit
SR	Supervisor Read	TLB Entry Word 0	13
ST	Set	Channel Control	13
SW	Supervisor Write	TLB Entry Word 0	12
TCV	Timer Count Value	Timer Counter	23-0
TE	Trace Enable	Current Processor Status	13
		Old Processor Status	13
TF	Transaction Faulted	Channel Control	10
TID	Task Identifier	TLB Entry Word 0	7-0
TP	Trace Pending	Current Processor Status	12
		Old Processor Status	12
TR	Traget Register	Channel Control	9-2
TRV	Timer Reload Value	Timer Reload	23-0
TU	Trap Unaligned Access	Current Processor Status	11
		Old Processor Status	11
U	Usage	TLB Entry Word 1	1
UE	User Execute	TLB Entry Word 0	8
UR	User Read	TLB Entry Word 0	10
UW	User Write	TLB Entry Word 0	9
V	Overflow	ALU Status	10
VAB	Vector Area Base	Vector Area Base Address	31-16
VE	Valid Entry	TLB Entry Word 0	14
VF	Vector Fetch	Configuration	4
VTAG	Virtual Tag	TLB Entry Word 0	31-15
WM	WAIT Mode	Current Processor Status	7
		Old Processor Status	7
Z	Zero	ALU Status	8

08996A-Fig-B6c

Index

INDEX

A

- A (Absolute) 8-8
- A0-A31 (Address Bus) 1-4, 5-1
- access privilege 5-23
- access protocol 2-20, 5-9
- access types 3-71, 7-23
- access, burst-mode 1-4
- access, simple 2-21
- access, simultaneous 5-21
- access, triple-port 1-5
- activation record 7-1, 7-2, 7-6, 7-7
- activation record mapping 7-4
- ADD 7-33
- addition, integer 7-9
- Address Bus (A0-A31) 1-4, 5-1
- address bus, coprocessor operations 6-9
- address bus, shared 2-21
- address space, Coprocessor 2-11
- address space, Input/Output 2-11
- address space, Instruction ROM 2-11
- address space, Instruction/Data 2-11
- address spaces 3-46
- Address Tag 4-8, 4-9
- address transfer 2-21
- address translation 2-14, 3-67, 3-68, 3-69, 4-18, 7-24
- address translation controls 3-66
- address translation exceptions 1-7
- Address Unit 2-17, 4-12, 4-15, 4-16
- address, absolute 4-16
- address, physical 2-11
- address, PC relative 4-16
- address, virtual 2-11
- addresses, intermediate 3-45
- addresses, pipelined 1-4
- addressing 2-11, 3-46
- addressing, indirect 7-7
- addressing, register 4-13
- ADRF Latch 4-15, 4-16
- alignment 2-11, 3-46
- alignment, Branch Target Cache 4-10

- alignment, bytes 7-17
- alignment, half-words 3-50
- alignment, instructions 3-51
- alignment, words 3-50
- ALU (Arithmetic/Logic Unit) 2-17, 4-12, 4-17, 8-6
- ALU Status Register 3-8, 3-24, 3-57, 5-28, 7-19, 7-28
- Am29000 1-3
- Am29000 features 1-1
- Am29000 special features 1-13
- applications 7-1
- arbitration 2-21, 5-7, 5-20
- arguments, incoming 7-6
- arguments, outgoing 7-6
- arithmetic instruction traps 7-14
- arithmetic operation 8-6
- Arithmetic/Logic Unit (ALU) 2-17, 4-12, 4-17, 8-6
- AS 3-41
- ASEQ 7-16, 7-34
- ASLEU instruction 7-7
- ASNE 7-13, 7-9
- assembler syntax 8-5
- assert compare 3-30
- Assert instructions 3-30, 7-8, 7-9

B

- B-Bus 4-15
- *BGRT (Bus Grant) 5-1, 5-20
- BINV 3-69
- *BINV (Bus Invalid) 5-1, 5-20
- B0-B15 Register Bank Protect bits (Banks 0-15) 3-16, 3-17
- bank protect 3-6
- BO (Byte Order) Configuration Reg. 3-13, 3-47, 3-48, 3-49
- Boolean 3-30, 7-14
- Boolean Compare 3-30
- Boolean data 3-40
- Boolean FALSE 7-15
- Boolean TRUE 7-15
- boundary crossings 4-10
- BP (Byte Pointer) ALU Status Reg. 3-24
- BP (Byte Pointer) Byte Pointer Reg. 3-25
- BP 3-49, 8-1
- branch displacement, relative 7-15
- Branch instructions 3-36, 3-38
- branch target 4-15, 4-16
- Branch Target Cache 1-5, 2-16, 3-56, 3-58, 4-3, 4-6, 4-16, 7-26, 7-27

Branch Target Cache lookup process 4-8
 branch, relative 1-6, 2-5, 4-9, 7-32, 7-33
 Branches, immediately adjacent 7-33
 Branch Target Cache Disable (CD), Configuration Reg. 3-13, 4-7, 7-26
 *BREQ (Bus Request) 5-1, 5-20
 Burst 5-10, 5-13
 Burst mode 4-16, 5-13, 5-15, 5-16, 5-25
 Burst mode access 1-4, 5-9, 5-13, 5-17, 5-18
 burst mode access protocol 2-20
 burst mode cancellation 5-19
 burst mode preemption 5-19
 burst mode termination 5-19
 Bus Grant (*BGRT) 5-1, 5-20
 Bus Invalid (*BINV) 5-1, 5-20
 Bus Request (*BREQ) 5-1, 5-20
 bus sharing 5-21
 byte addressing 3-47, 3-48, 3-49
 byte alignment 7-17
 byte operations 3-39
 Byte Pointer Register 3-8, 3-25
 Byte Order (BO), Configuration Reg. 3-13
 Byte Pointer (BP) ALU Status Reg. 3-24
 Byte Pointer (BP), Byte Pointer Reg. 3-25

C

C (Carry) ALU Status Reg. 3-24, 8-1, 8-6
 CA (Coprocessor Active) 3-73, 6-5
 CA (Coprocessor Active) Current Processor Status Reg. 3-10
 CA (Coprocessor Active) Old Processor Status Reg. 3-10
 Cache Block 4-9
 Cache Disable (CD) 3-13, 4-7, 7-26
 Cache replacement, random 4-9
 Cache tag 4-7
 Cache-block boundary 4-7
 CALL 7-32
 call, large range 7-15
 CALLI 7-15
 calls, operating system 7-9
 Carry (C), ALU Status Reg. 3-24, 7-9, 8-6
 CD (Cache Disable) 3-13, 4-7, 7-26
 *CDA (Coprocessor Data Accept) 5-4, 6-8
 *CDA sequencing 6-9
 CE (Coprocessor Enable) Channel Control Reg. 3-15, 3-41, 6-2
 CE/CNTL 8-9
 CHA (see Channel Address)

Channel 2-20, 5-6
Channel Address (CHA), Channel Addr. Reg. 3-8, 3-14, 3-45, 3-57, 3-62, 5-28, 7-19, 7-22
channel arbiter, external 2-21
Channel Control 3-57, 3-58, 3-62, 5-28, 7-19, 7-22
Channel Control Register 3-8, 3-15, 3-45
Channel Data (CHD), Channel Data Reg. 3-14, 3-57, 3-62, 5-28, 7-19, 7-22
Channel Data Register 3-8, 3-14
Channel Registers 3-62
character detection 7-17
character-string 7-16, 7-17
CHD (see Channel Data)
clock synchronization 5-33
clock, processor-generated 5-32
clock, system-generated 5-32
clocks 2-22
CNTL (Control) Channel Control Reg. 3-15, 3-41
CNTL0-CNTL1 (CPU Control) 3-15, 5-5, 5-24, 5-25, 5-26, 5-27, 5-28, 5-31, 5-32
Compare 3-30
Compare Bytes (CPBYTE) 7-17
Compare instructions 3-32
compiler's run-time stack 1-5
compiler, optimizing 1-10
compilers 1-9
complementing a Boolean 7-14
Configuration Register 2-3, 3-8, 3-13, 4-7, 7-26
CONST 7-13, 7-15, 7-32, 7-33, 7-34
Constant 3-33
Constant instructions 3-36
constant, 32-bit 7-15
constant, 8-bit 2-5
CONSTH 7-15
CONSTN 7-15
Contents Valid (CV), Channel Control Reg. 3-15, 3-44, 3-57, 3-58, 7-27
context switching 2-12, 2-13, 7-21, 7-22
context switching, temporary 2-13
contexts, saving and restoring 2-13
Control (CNTL), Channel Control Reg. 3-15
Coprocessor 6-1
Coprocessor Active (CA) 3-10, 6-5
coprocessor attachment 2-22
coprocessor communication 6-8
Coprocessor Data Accept (*CDA) 5-4, 6-8
Coprocessor Enable (CE), Channel Control Reg. 3-15, 6-2
coprocessor exception 3-53, 6-4, 6-9
Coprocessor exception trap 5-9
coprocessor interrupts 6-5

- coprocessor Load/Store 6-2
- coprocessor not present trap 3-53
- coprocessor operations 6-1
- Coprocessor transfer 5-3, 6-1, 6-2, 6-4, 6-6, 6-8
- Coprocessor Present (CP), Configuration Reg. 3-13, 6-5
- COUNT 8-1
- CP (Coprocessor Present) Configuration Reg. 3-13, 6-5
- CPBYTE (Compare Bytes) 3-30, 3-39, 7-17
- CPEQ 7-13, 7-14
- CPLT 7-14
- CPNEQ 7-14, 7-32
- CPU Control (CNTL0-CNTL1) 5-5, 5-24, 5-25, 5-26, 5-27, 5-28, 5-31, 5-32
- CPU Status (STAT0-STAT2) 5-5, 5-23, 5-25, 5-26, 5-27, 5-28, 5-31
- CR (Load/Store Count Remaining) Channel Control Reg. 3-15
- CR (Load/Store Count Remaining) Load/Store Count Remaining Reg. 3-26
- CR 3-44
- Current Processor Status 2-3
- Current Processor Status Register 3-8, 3-10, 7-31
- CV (Contents Valid) Channel Control Reg. 3-15, 3-44, 3-57, 3-58, 7-27
- cycle time 1-2
- CD (Branch Target Cache Disable) Configuration Reg. 3-13

D

- D0-D31 (Data Bus) 1-4, 5-3
- DA (Disable All Interrupts) 3-10, 3-52, 3-57, 3-73, 5-30
- daisy-chain 2-21
- data access 5-8
- Data access exception trap 5-9
- Data Access request 5-7
- data accesses, external 2-10, 2-11, 3-41
- Data Address Transfer 5-7
- data blocks, movement of large 7-17
- Data Burst Acknowledge (*DBACK) 5-4, 5-10, 5-11, 5-16, 5-176, 5-18
- Data Burst Request (*DBREQ) 5-4, 5-11, 5-16, 5-17, 5-18
- Data Bus (D0-D31) 1-4, 5-3
- data dependencies, pipeline 4-14
- Data Error (*DERR) 3-63, 5-3, 5-9, 5-16
- data exceptions 3-62
- data formats 2-6, 3-39
- data forwarding 4-14
- data handling 3-39
- Data Memory 4-7
- Data movement instructions 3-34, 3-35
- Data Ready (*DRDY) 5-3, 5-10, 5-12, 5-16
- Data Request (*DREQ) 5-3, 5-10, 5-12, 5-19

Data Request Type (DREQT0-DREQT1) 5-3, 5-8, 6-6
Data transfer 5-7
data types 2-6, 3-39
data-flow organization 1-7
data-unit numbering conventions 2-10
*DBACK (Data Burst Acknowledge) 5-4, 5-10, 5-11, 5-16, 5-17, 5-18
*DBREQ (Data Burst Request) 5-4, 5-11, 5-16, 5-17, 5-18
de-reference 7-8
Decode PC Register 4-15, 4-16
decode stage 4-1
delay cycle, indirect addressing 7-8
delayed branch 7-32, 7-33
Delayed effects, registers 7-35
demand paging 7-23, 7-25
*DERR (Data Error) 3-63, 5-3, 5-9, 5-16
DEST 8-2
DF (Divide Flag) ALU Status Reg. 3-24
DI (Disable Interrupts) 3-52, 3-57, 3-73, 5-30, 6-6
DI (Disable Interrupts) Current Processor Status Reg. 3-10
DI (Disable Interrupts) Old Processor Status Reg. 3-10
Disable All Interrupts and Traps (DA), Curr. Proc. Status Reg. 3-10
Disable All Interrupts and Traps (DA), Old Proc. Status Reg. 3-10
Disable Interrupts (DI), Current Processor Status Reg. 3-10
Disable Interrupts (DI), Old Processor Status Reg. 3-10
DIV 7-12, 7-13
DIV0 7-12, 7-13
divide instructions 3-30, 3-59, 7-12
Divide Flag (DF), ALU Status Reg. 3-24
DIVL 7-12, 7-13
DIVREM 7-12, 7-13, 7-14
double-precision 3-40
*DRDY (Data Ready) 5-3, 5-10, 5-12, 5-16
*DREQ (Data Request) 5-3, 5-10, 5-12, 5-19
DREQT0-DREQT1 (Data Request Type) 5-3, 5-8, 6-6
DTR 4-13
dynamic-nesting-depth 7-4

E

EMULATE 7-8
Entry Word 3-28
ETR 4-13, 4-14
exception reporting 3-62
exceptions, address translation 1-7
execute stage 4-1

Executing mode 2-18, 5-5
Execution Unit 2-15, 2-17, 4-1, 4-12
execution, single cycle 1-13
EXBYTE 3-39
EXHW (Extract Half-Word) 3-40
EXHWS (Extract Half-Word, Sign-extended) 3-40
extended arithmetic 3-59
external access 7-27
external access protection 7-19
external data access 3-40
external hardware support 3-49
external interrupts 5-29
external traps 5-29
EXTERNAL WORD[n] 8-2
EXTRACT 3-32, 7-18
Extract Half-Word (EXHW) 3-40
Extract Half-Word, Sign-extended (EXHWS) 3-40

F

F (Flag) TLB Entry Word 1 3-29, 3-68
FALSE 8-2
fast context switching 7-21, 7-22
FC (Funnel Shift Count) Funnel Shift Count Reg. 3-26, 7-17, 8-2
FC (Funnel Shift Count) ALU Status Reg. 3-24
fetch special instruction 4-16
fetch stage 4-1
Fetch-Ahead Adder 4-10, 4-11
Fetch-Ahead Adder overflow 4-11
fetch-ahead disabling 4-10
Field Shift Unit 2-18, 4-12, 4-17, 4-18
FIFO 4-3
first unmapped location 7-7
Flag (F), TLB Entry Word 1 3-29
floating-point 1-9, 7-14
Floating-Point data 3-40
Floating-Point instructions 3-36, 3-37, 7-8
Floating-Point Standard P754 3-40
Freeze (FZ) 3-56, 3-57, 3-58, 3-73, 4-11, 7-19, 8-6
Freeze (FZ), Current Processor Status Reg. 3-10
Freeze (FZ), Old Processor Status Reg. 3-10
Funnel Shift Count (FC), Funnel Shift Count Reg. 3-26
Funnel Shift Count Register 3-8, 3-26

Funnel-Shift Unit 4-17
Funnel Shift Count (FC), ALU Status Reg. 3-24
FZ (Freeze bit) 3-56, 3-57, 3-58, 3-73, 4-11, 7-19, 8-6
FZ (Freeze) Current Processor Status Reg. 3-10
FZ (Freeze) Old Processor Status Reg. 3-10

G

general-purpose registers 1-4, 2-1, 2-2, 3-2
generator, register address 4-13
global registers 3-3, 3-4, 4-13

H

Half-Word addressing 3-47, 3-48, 3-49
half-word constant 3-33
half-word operations 3-39, 3-40
Halt 5-5
Halt mode 2-19, 5-24, 5-25, 5-28
handler starting address 3-53
hardware development system 2-21, 5-27
hardware testing 5-29

I

OI16 (16-bit immediate data zero-extended to 32 bits) 8-1
II16 (16-bit immediate data, ones-extended to 32 bits) 8-1
*IBACK (Instruction Burst Acknowledge) 5-3, 5-10, 5-11, 5-14, 5-15, 5-18
*IBREQ (Instruction Burst Request) 5-3, 5-11, 5-14, 5-15, 5-18
I-Bus 4-15
IO-I31 (Instruction Bus) 5-2, 8-9,
I16 (16-bit immediate data) 8-2
IE (Interrupt Enable) Timer Reload Reg. 3-18, 7-29
*IERR (Instruction Error) 3-63, 4-4, 5-2, 5-9, 5-14, 5-27
IFP (Instruction Fetch Pointer) 4-3
IFU 4-18
illegal opcode trap 3-53
IM (Interrupt Mask) Current Processor Status Reg. 3-10
IM (Interrupt Mask) Old Processor Status Reg. 3-10
IM 3-52, 3-57, 3-73
immediate move to TLB 3-72

IN (Interrupt) Timer Reload Reg. 3-18, 7-29
 in args 7-2
 INBYTE 3-39, 3-49
 INCLK (Input Clock) 5-6, 5-32, 5-33, 5-35
 indirect access 3-6
 indirect addressing 7-7, 7-8
 indirect addressing delay cycle 7-8
 indirect pointers 3-3, 7-7, 7-8, 7-35
 Indirect Pointer A (IPA), Indirect Pointer A Reg. 3-8, 3-22
 Indirect Pointer B (IPB), Indirect Pointer B Reg. 3-8, 3-23
 Indirect Pointer C (IPC), Indirect Pointer C Reg. 3-8, 3-22
 INHW (Insert Half-Word) 3-40, 3-49
 initialization 3-72
 initialization, timer facility 7-29
 Input Clock (INCLK) 5-6
 Input/Output access 5-3
 Insert Half-Word (INHW) 3-40
 instruction access 3-66, 5-7, 5-8
 instruction access as data 3-51
 Instruction Access Exception 4-4
 Instruction Address Transfer 5-7
 instruction boundary 2-15
 Instruction Burst Acknowledge (*IBACK) 5-3, 5-10, 5-11, 5-14, 5-15, 5-18
 Instruction Burst Request (*IBREQ) 5-3, 5-11, 5-14, 5-15, 5-18
 Instruction Bus (I0-I31) 1-4, 5-2
 instruction description format 8-11
 Instruction Error (*IERR) 3-63, 4-4, 5-2, 5-9, 5-14, 5-27
 instruction exceptions 3-62
 Instruction Fetch Pointer (IFP) 4-3, 4-18
 Instruction Fetch Unit 2-15, 2-16, 4-1, 4-2
 instruction fetch, external 4-10
 instruction fetch-ahead 4-10
 instruction formats 8-8
 instruction overview 2-5
 Instruction Prefetch Buffer (IPB) 4-34-16
 instruction prefetch stream 4-3
 Instruction Ready (*IRDY) 4-4, 5-2, 5-9, 5-10, 5-12, 5-14, 5-15, 5-16, 5-27
 Instruction Register (IR) 5-27
 Instruction Request (*IREQ) 5-2, 5-10, 5-12, 5-15, 5-19
 Instruction Request Type (IREQT) 5-2
 instruction ROM 3-47, 5-8
 instruction set 2-7, 8-1
 Instruction set 3-30
 Instruction Transfer 5-7
 instruction, listing by operation code 8-127
 instruction, special 4-16

instruction-field uses 8-10
 Instruction/Data memory 5-8
 Instruction/Data memory access 5-3
 instructions, Branch 3-36, 3-38
 instructions, Compare 3-30, 3-32
 instructions, Constant 3-36
 instructions, Data Movement 3-34, 3-35
 instructions, divide 3-59
 instructions, Floating-Point 3-36, 3-37
 instructions, integer arithmetic 3-31
 instructions, logical 3-33, 3-34
 instructions, miscellaneous 3-36, 3-38
 instructions, shift 3-34
 instructions, three address 2-2
 integer addition 7-9
 integer arithmetic 3-30, 3-31
 integer division 7-12
 integer multiplication 7-10
 integer subtraction 7-9
 interrupt and trap priority, 3-60, 3-61
 Interrupt Enable (IE), Timer Reload Reg. 3-18
 interrupt exceptions 3-63
 interrupt handling 3-55, 7-20
 Interrupt Mask (IM), Current Processor Status Reg. 3-10
 Interrupt Mask (IM), Old Processor Status Reg. 3-10
 interrupt occurrence 3-51
 Interrupt Pending (IP), Current Processor Status Reg. 3-10
 Interrupt Pending (IP), Old Processor Status Reg. 3-10
 interrupt processing, user-defined 2-13
 Interrupt Request (*INTR0-INTR3) 3-52, 3-54, 5-5, 5-29, 5-30
 interrupt return 5-5, 3-72, 7-21
 interrupt sequencing 3-60
 interrupt simulation 7-21
 interrupt taking 3-51, 3-55
 Interrupt (IN), Timer Reload Reg. 3-18, 7-29
 interrupt, fast processing 3-58
 interrupts 1-8, 2-12, 2-15, 3-51, 3-52, 3-62, 3-63, 5-22, 7-19, 7-28
 interrupts, coprocessor 6-5
 interrupts, dynamically nested 2-13, 7-20
 interrupts, external 5-29
 *INTR0-INTR3 (Interrupt Request) 3-52, 3-54, 5-5, 5-29, 5-30
 INV 4-6, 7-26
 Invalidation 3-70
 IP (Interrupt Pending) Current Processor Status Reg. 3-10
 IP (Interrupt Pending) Old Processor Status Reg. 3-10
 IP 3-73

IPA (Indirect Pointer A) Indirect Pointer A Reg. 3-22, 4-13, 8-2
IPB (Indirect Pointer B) Indirect Pointer B Reg. 3-23, 8-2
IPB (Instruction Prefetch Buffer) 4-3
IPB 4-13, 5-14
IPB allocated state 4-4
IPB available state 4-4
IPB error state 4-4, 4-5
IPB state transitions 4-4
IPB valid state 4-4
IPC (Indirect Pointer C) Indirect Pointer C Reg. 3-22, 4-13, 8-2
IR (Instruction Register) 5-27, 5-28
*IRDY (Instr. Ready) 4-4, 5-2, 5-9, 5-10, 5-12, 5-14, 5-15, 5-16, 5-27
*IREQ (Instruction Request) 5-2, 5-10, 5-12, 5-15, 5-19
IREQT (Instruction Request Type) 5-2
IRET 3-56, 7-28
IRETINV 3-56, 4-6, 7-26, 7-28

J

JMP 7-33
JMPF 7-13, 7-14, 7-32, 7-34
jump, large range 7-15

L

LA (Lock Active) Channel Control Reg. 3-15
large call range 7-15
large constants 7-15
large data blocks, movement 7-17
large jump range 7-15
Least Recently Used Entry (LRU), LRU Rec. Reg. 3-8, 3-21, 3-70, 7-24, 7-26
LK (Lock) 3-73, 7-29
LK (Lock) Current Processor Status Reg. 3-10
LK (Lock) Old Processor Status Reg. 3-10
LOAD (Load) 3-43, 7-34
Load and Lock (LOADL) 3-43, 5-22, 7-28
Load and Set (LOADSET) 3-43, 3-71, 7-28
load data, forwarding 1-8
Load Multiple (LOADM) 1-6, 3-43, 3-59, 4-15, 4-16, 5-23, 7-17
load operations 3-43
Load Test Instruction 5-5, 5-24, 5-25, 5-27, 5-28
Load Test Instruction mode 2-19
Load/ Store Count Remaining Register 3-8, 3-26
Load/Store instruction format 3-41
Load/Store instruction format, non-coprocessor 3-41

- Load/Store instruction format 3-41
- Load/Store instruction format, non-coprocessor 3-41
- Load/Store (LS), Channel Control Reg. 3-15
- Load/Store Count Remaining (CR), Channel Cont. Reg. 3-15
- Load/Store Count Remaining (CR), Load/Store Cnt Rem. Reg. 3-26
- LOADL (Load and Lock) 3-43, 5-22, 7-28
- LOADM 1-6, 3-43, 3-59, 4-15, 4-16, 5-23, 7-17
- loads and stores 1-6
- Loads and Stores, overlapped 7-34
- LOADSET (Load and Set) 3-43, 3-71, 7-28
- local registers 3-3, 3-5, 4-13, 7-4
- local registers, stack pointer 2-2
- Lock (*LOCK) 5-1, 5-22
- Lock (LK) 7-29
- lock output 5-22
- Lock (LK), Current Processor Status Reg. 3-10
- Lock (LK), Old Processor Status Reg. 3-10
- Lock Active (LA), Channel Control Reg. 3-15
- logical instructions 3-33, 3-34
- logical operation 8-7
- lower bound, Stack Cache 7-7
- LRU (Least Recently Used Entry) LRU Rec. Reg. 3-8, 3-21, 3-70, 7-25, 7-26
- LS (Load/Store) Channel Control Reg. 3-15

M

- M (IMmediate) 8-8
- mapping activation record 7-4
- master and slave switching 5-34
- master/slave checking 5-33
- Master/Slave Error (MSERR) 5-6, 5-33
- master/slave operation 2-22, 5-34
- memory management 1-8, 2-13, 3-64, 7-23
- Memory Management Unit 2-15, 2-18
- memory protection 7-18, 7-23
- memory, critical areas 7-24
- merge, byte-aligned 7-17
- MFSR 7-11, 7-12, 7-13, 7-14
- MFTLB 3-70
- MIPs 1-2
- miscellaneous instructions 3-36, 3-38
- ML (Multiple Operation) Channel Control Reg. 3-15, 3-45, 3-57
- MMU 3-66, 3-67, 4-9, 4-16, 4-18, 4-19, 7-19, 7-23
- MMU Configuration Register 3-8, 3-20, 7-35
- MMU Programmable (MPGM0-MPGM1) 3-69, 5-2, 5-7
- Mode, Executing 5-5

Mode, Halt 5-5
Mode, Pipeline Hold 5-5
Mode, Step 5-5
Mode, Wait 5-5
Move To Special Register (MTSR) 3-72, 5-28, 7-8, 7-10, 7-11, 7-13, 8-6
monitoring critical areas 7-24
move immediate to TLB 3-72
MPGM0-MPGM1 (MMU Programmable) 3-69, 5-2, 5-7
MSERR (Master/Slave Error) 5-6, 5-33
MTSP 3-7, 7-19
MTSR (Move To Special Register) 3-72, 5-28, 7-8, 7-10, 7-11, 7-13, 8-6
MTTLB 3-70
MUL 7-10, 7-11
MULL 7-11
multi-precision 7-9
multi-processing 7-28
Multiple Access 3-44
multiple masters 2-21, 5-21
Multiple Operation (ML), Channel Control Reg. 3-15
multiple slaves 5-21
multiplication, integer 7-10
MULTIPLY 3-30, 7-8
MULU 7-11, 7-12

N

N (Negative) ALU Status Reg. 3-24, 8-6
NN (Not Needed) Channel Control Reg. 3-15, 3-57, 4-14, 7-27
NO-OP 7-6, 7-16, 7-33
nomenclature 8-1
non-sequential fetch 4-9
non-sequential instruction fetch 4-10, 5-5
non-sequential reference 4-3
Normal 5-5
normal execution 5-26
Normal mode 5-24
notation 8-1
Not Needed (NN), Channel Control Reg. 3-15, 3-57, 4-14, 7-27
numbering conventions, data-unit 2-10

O

Old Processor Status Register 2-3, 3-8, 7-22, 7-31
OP (operation code) 8-8
operand checking, run-time 3-31

- operating system calls 7-9
- operation code (OP) 8-8
- operator symbols 8-3
- OPT (Option) 3-41, 3-46, 6-3
- OPT0-OPT2 (Option Control) 3-49, 3-50, 5-4, 5-7
- Option (OPT) 3-41, 3-46, 6-3
- Option Control (OPT0-OPT2) 3-49, 3-50, 5-4, 5-7
- OR 7-13, 7-33
- organization, Branch Target Cache 4-6
- organization, data flow 1-7
- out args 7-2
- out of range 3-53, 8-7
- Out of Range trap 8-7
- OV (Overflow) 7-29
- OV (Overflow) Timer Reload Reg. 3-18
- Overflow (OV) 7-29
- overflow, signed 8-7
- overflow, unsigned 8-7
- Overflow (OV), Timer Reload Reg. 3-18
- Overflow (V), ALU Status Reg. 3-24, 7-4, 8-6
- overlapped loads 1-6
- overlapped store 1-6

P

- PA 3-41, 3-66
- page change information 7-24
- page fault 7-27
- Page number, real 3-68
- page offset 3-66, 3-69
- page reference 7-23
- page size 3-67, 3-69
- page size, virtual 7-23
- paging 7-25, 7-27
- Page Size (PS), MMU Configuration Reg. 3-20
- PC (Program Counter) 4-11, 8-2
- PC Buffer 4-12
- PC MUX 4-11, 4-12
- PC-Bus 4-15
- PC0 (Program Counter 0) Program Counter 0 Reg. 3-19
- PC0-PC2 3-62, 4-11, 4-12, 5-28, 7-19, 7-22
- PC1 (Program Counter 1) Program Counter 1 Reg. 3-19, 5-22
- PC2 (Program Counter 2) Program Counter 2 Reg. 3-20
- PD (Physical Addressing/Data) Current Processor Status Reg. 3-10
- PD (Physical Addressing/Data) Old Processor Status Reg. 3-10
- PD 3-66, 3-73

- *PDA (Pipelined Data Access) 5-4, 5-10, 5-12
- *PEN (Pipeline Enable) 5-2, 5-10, 5-11, 5-12
- PGM (User Programmable) TLB Entry Word 1 3-29, 3-68, 3-69
- Physical Addressing/Data (PD), Curr. Proc. Status Reg. 3-10
- Physical Addressing/Data (PD), Old Processor Status Reg. 3-10
- Physical Addressing/Instructions (PI), Curr. Proc. Status Reg. 3-10
- Physical Addressing/Instructions (PI), Old Proc. Status Reg. 3-10
- PI (Physical Addressing/Instructions) Current Proc. Status Reg. 3-10
- PI (Physical Addressing/Instructions) Old Proc. Status Reg. 3-10
- PI 3-66, 3-73
- *PIA (Pipelined Instruction Acknowledge) 5-3, 5-10, 5-12
- PID (Process Identifier) MMU Configuration Reg. 3-20, 3-70, 3-71, 7-26
- pipeline 1-8, 2-16, 4-1
- pipeline data dependencies 4-14
- pipeline dependency 4-13
- Pipeline Enable (*PEN) 5-2
- pipeline features exposed 7-1, 7-32
- Pipeline Hold 3-45, 3-72, 4-1, 4-14
- Pipeline Hold mode 2-19, 4-19, 5-5
- pipeline interlocks 1-13
- pipelined access 5-10, 5-9, 5-11
- pipelined access protocol 2-20
- Pipelined addresses 1-4
- Pipelined Data Access (*PDA) 5-4, 5-10, 5-12
- Pipelined Instruction Acknowledge (*PIA) 5-3, 5-10, 5-12
- Port A 4-13
- Port B 4-13
- Port C 4-13
- prefetching 1-5
- primary access 5-10, 5-12
- Prioritizer 2-18, 4-12, 4-18
- Priority 5-29
- priority, interrupts and traps 3-60, 3-61
- PRL (Processor Release Level) Configuration Reg. 3-13
- procedure calls 7-1
- procedure epilogue 7-5, 7-6
- procedure linkage 7-5
- procedure prologue 7-5
- procedure returns 7-1
- Process Identifier (PID), MMU Configuration Reg. 3-20, 3-70, 3-71 7-26
- processor 5-10
- processor cancellation 5-19
- processor modes 2-18
- processor preemption 5-19
- processor reset 5-30
- processor state 7-22

- processor termination 5-19
- processor-generated clock 5-32
- Processor Release Level (PRL), Configuration Reg. 3-13
- Program Counter (PC) 3-51, 3-58, 3-62, 4-11, 7-22
- Program Counter 0 Register (PC0) 3-8, 3-19
- Program Counter 1 Register (PC1) 3-8, 3-19
- Program Counter 2 Register (PC2) 3-8, 3-20
- Program Counter Unit 2-17, 4-12
- program modes 2-1
- programmer reference 3-1
- programming 7-1
- programming, Coprocessor 2-14
- protected segment 2-2
- Protection 3-71
- protection bits, supervisor mode 7-18
- protection bits, TLB 7-24
- protection bits, user mode 7-18
- protection checking 4-16
- protection violation 3-71
- Protection Violation Trap 3-31, 7-8
- protection violation, TLB 7-18
- protection, external access 7-19
- protection, memory 7-18
- protection, register 7-18
- protection, system 7-18
- PS (Page Size) MMU Configuration Reg. 3-20

Q

- Q (Quotient/Multiplier) Q Register 2-4, 3-8, 3-23, 3-57, 7-22, 8-2

R

- R 3-53
- R/*W (Read/Write) 5-1
- RA Register 3-43, 8-2, 8-8, 8-9
- RB or I 3-43, 8-8, 8-9
- RB register 8-2, 8-8, 8-9
- RC register 8-2, 8-8
- RE (ROM Enable) Current Processor Status Reg. 3-10
- RE (ROM Enable) Old Processor Status Reg. 3-10
- RE 3-47, 3-66, 3-73
- Read-Only Memory 4-7
- Read/Write (R/*W) 5-1
- Real Page Number (RPN), TLB Entry Word 1 3-29

- recursion 7-2
- reference, non-sequential 4-3
- register address generator 4-13
- register addressing 3-4, 4-13
- register bank protect 3-6, 7-8
- Register Bank Protect bits B0-B15 (Banks 0-15) 3-16, 3-17
- Register Bank Protect Register 3-8, 3-16
- register banks 3-5, 3-6, 7-21
- register file 1-4, 2-17, 4-12, 4-13
- register file port 4-14
- register number 3-4
- register protection 7-18
- register RA 8-2
- register RB 8-2
- register RC 8-2
- register read-address comparators 4-14
- Register, ALU Status 2-4, 3-8, 3-24
- Register, Byte Pointer 2-4, 3-8, 3-25
- Register, Channel Address 2-3, 3-8, 3-14
- Register, Channel Control 2-3, 3-8, 3-15
- Register, Channel Data 2-3, 3-8, 3-14
- Register, Configuration 2-3, 3-8, 3-13
- Register, Current Processor Status 3-8, 3-10
- Register, Funnel Shift Count 2-4, 3-8, 3-26
- Register, Indirect Pointer A 2-4, 3-8, 3-22
- Register, Indirect Pointer B 2-4, 3-8, 3-23
- Register, Indirect Pointer C 2-4, 3-8, 3-22
- Register, Load/ Store Count Remaining 2-4, 3-8, 3-26
- Register, LRU Recommendation 2-4, 3-8, 3-21
- Register, MMU Configuration 2-4, 3-8, 3-20
- Register, Old Processor Status 3-8
- Register, Program Counter 0 2-3, 3-8, 3-19
- Register, Program Counter 1 2-4, 3-8, 3-19
- Register, Program Counter 2 2-4, 3-8, 3-20
- Register, Q 2-4, 3-8, 3-23, 3-57, 7-22, 8-2
- Register, Register Bank Protect 2-3, 3-8, 3-16, 3-17
- register, special-purpose 3-7
- Register, Timer Counter 2-3, 3-8, 3-17
- Register, Timer Reload 2-3, 3-8, 3-18
- register, TLB 2-5
- Register, Vector Area Base Address 3-8, 3-9
- register-numbers, absolute 3-45
- registers 3-2
- registers, delayed effects 7-35
- registers, global 2-2, 3-3, 3-4
- registers, local 2-2, 3-3, 3-5

- registers, local, stack pointer 2-2
- registers, protected 3-7, 3-8
- registers, special-purpose, protected 2-2
- registers, unimplemented 3-6
- registers, unprotected 3-8
- relational operators 4-17
- relative branch 2-5
- Reload 3-70
- reserved fields 3-7
- Reset (*RESET) 4-19, 5-6, 5-25
- *RESET (Reset) 4-19, 5-6, 5-25
- Reset mode 2-20, 3-73, 5-30, 5-31
- resident pages 7-25
- restart 7-28
- restarting after faulty external access 7-27
- ROM 4-7
- ROM Enable (RE), Current Processor Status Reg. 3-10
- ROM Enable (RE), Old Processor Status Reg. 3-10
- ROM Vector Area (RV), Configuration Reg. 3-13
- RPN (Real Page Number) TLB Entry Word 1 3-29, 3-67, 3-69
- Run-time checking 7-8
- Run-time Stack 7-1, 7-2, 7-4
- RV (ROM Vector Area) Configuration Reg. 3-13, 3-56

S

- SA (Set Coprocessor Active) 6-3
- SA (Special-Purpose Register number) 8-2
- SB 3-41
- SE (Supervisor Execute) TLB Entry Word 0 3-28, 3-71
- segment, protected 2-2
- serialization 3-72
- Set (ST), Channel Control Reg. 3-15
- Set Coprocessor Active (SA) 6-3
- Set Indirect Pointers (SETIP) 7-8
- SETIP (Set Indirect Pointers) 7-8
- Shift instructions 3-34
- shift, byte-aligned 7-17
- simple access 5-9, 5-11
- simulation, interrupts 7-21
- slave cancellation 5-20
- slave device 5-10
- Slave Mode 5-11
- slave preemption 5-20
- SM (Supervisor Mode) Current Processor Status Reg. 3-10
- SM (Supervisor Mode) Old Processor Status Reg. 3-10

SM 3-71, 3-73
 SORT 7-32
 SP (Stack Pointer) 7-7
 Space ID 4-6
 Space Identification Field 4-7
 SPDEST 8-2
 SPECIAL 8-2
 special-purpose registers 3-7
 spurious errors 5-34
 SR (Supervisor Read) TLB Entry Word 0 3-28, 3-71
 SRCA 8-3
 SRCA.BYTE_n 8-3
 SRCB 8-3
 SRCB.BYTE_n 8-3
 ST (Set) Channel Control Reg. 3-15
 Stack Cache 7-3, 7-7
 Stack Cache implementation 7-5
 Stack Cache, lower bound 7-7
 Stack Pointer (SP) 1-13, 2-2, 3-3, 3-5, 4-13, 7-3, 7-7, 7-16, 7-35
 Stack Pointer adjustment 7-5
 stack, compiler's run-time 1-5
 stack, run-time 7-1, 7-2
 STAT0-STAT2 (CPU Status) 5-5, 5-23, 5-25, 5-26, 5-27, 5-28, 5-31
 status results, arithmetic 8-6
 status results, logic 8-6
 Step 5-5, 5-28
 Step mode 2-19, 5-24, 5-25, 5-26
 STORE 3-44
 store and lock 3-44, 5-22, 7-29
 Store Multiple 1-7, 3-44, 3-59, 4-15, 4-16, 5-23, 7-17
 Supervisor Execute (SE), TLBstore operations 3-44
 STOREL 3-44, 5-22, 7-29
 STOREM 1-7, 3-44, 3-59, 4-15, 4-16
 SUB 7-32, 7-33
 SUBR 7-13, 7-14, 7-33
 SUBRC 7-13
 subtraction, integer 7-9
 SUP/*US (Supervisor/User) 3-1, 5-1,
 Supervisor Instruction 4-7
 Supervisor mode (SM) 2-1, 3-1, 5-9
 Supervisor Read (SR), TLB Entry Word 0 3-28
 Supervisor Write (SW), TLB Entry Word 0 3-28
 Supervisor-mode access 3-66
 Supervisor/User (SUP/*US) 3-1, 5-1
 Supervisor Mode (SM), Current Processor Status Reg. 3-10
 Supervisor Mode (SM), Old Processor Status Reg. 3-10

- switch task 7-22
- symbols 8-1
- synchronization, clock 5-33
- syntax, assembler 8-5
- SYSClk (System Clock) 5-6, 5-29, 5-32, 5-33, 5-34, 5-35
- system diagram 1-3
- system interface 2-20
- system protection 7-18
- system-generated clock 5-32
- systems programming 7-18

T

- Taking Interrupt or Trap 5-5
- target 4-8
- TARGET 8-3
- target instruction 4-6
- Target Register (TR), Channel Control Reg. 3-15
- Task ID 3-68
- task identifiers 1-8
- task switch 7-22
- Task Identifier (TID), TLB Entry Word 0 2-14, 3-28
- TC (Transfer Control) 6-3
- TCV (Timer Count Value) Timer Counter Reg. 3-17, 7-29
- TE (Trace Enable) 3-73, 7-31
- TE (Trace Enable) Current Processor Status Reg. 3-10
- TE (Trace Enable) Old Processor Status Reg. 3-10
- temporary variables 7-3
- terminology 8-4
- *TEST (Test mode) 2-20, 5-6, 5-29, 5-31
- Test/Development interface 2-21, 5-23
- TF (Transaction Faulted) Channel Control Reg. 3-15
- TID (Task Identifier) TLB Entry Word 3-28, 3-66, 3-67, 3-71
- Timer Count Value (TCV) Timer Counter Reg. 3-17, 3-18, 7-29
- Timer Counter Register 3-8, 3-17, 5-26, 7-30
- Timer Facility 2-14, 5-26, 7-29
- timer interrupts 7-30
- Timer Reload Register 3-8, 3-18, 7-30
- Timer Reload Value (TRV) 7-29
- TLB (Translation Look-aside Buffer) 1-8, 2-14, 3-64, 3-66, 3-67, 3-68, 4-19
- TLB Entry Word 3-27, 3-28, 3-29
- TLB line select 3-66
- TLB miss 5-14
- TLB Miss handling 7-24

TLB organization 3-65
 TLB registers 2-5
 TLB reload 7-20, 7-24
 TLB set 3-27, 3-65
 TLB, second-level 7-25
 TLB[N] 8-3
 TP (Trace Pending) 3-73, 7-31
 TP (Trace Pending) Current Processor Status Reg. 3-10
 TP (Trace Pending) Old Processor Status Reg. 3-10
 TR (Target Register) Channel Control Reg. 3-15, 3-45
 Trace Enable (TE) 3-73, 7-31
 Trace Facility 2-15, 7-31
 Trace Pending (TP) 7-31
 Trace Enable (TE), Current Processor Status Reg. 3-10
 Trace Enable (TE), Old Processor Status Reg. 3-10
 Trace Pending (TP), Current Processor Status Reg. 3-10
 Trace Pending (TP), Old Processor Status Reg. 3-10
 Trace Trap 7-31
 Transaction Faulted (TF), Channel Control Reg. 3-15
 Transfer Control (TC) 6-3
 transfer, coprocessor 6-1, 6-6
 Translation Look-aside Buffer (TLB) 1-8, 2-14, 3-64, 3-66, 3-67, 3-68, 4-19
 translation, early address 1-8
 translation, instruction address 4-18
 translation, Load Multiple address 4-18
 translation, Store Multiple address 4-18
 translation, virtual to physical 1-8
 trap handler 3-51, 3-55, 7-24
 Trap Request (*TRAP0-TRAP1) 3-52, 3-54, 5-5, 5-30
 trap sequencing 3-60
 trap taking 3-55
 Trap Unaligned Access 3-46, 3-46, 3-50, 3-73
 *TRAP0-TRAP1 (Trap Request) 3-52, 3-54, 5-5, 5-30
 traps 1-8, 2-12, 2-15, 3-51, 3-52, 3-62, 3-63, 5-22, 7-19, 7-28
 Traps, arithmetic instructions 7-14
 traps, external 5-29
 Trap Unaligned Access (TU), Current Processor Status Reg. 3-10
 Trap Unaligned Access (TU), Old Processor Status Reg. 3-10
 Triple-port access 1-5
 TRUE 8-3
 TRV (Timer Reload Value) Timer Reload Reg. 3-18, 7-29
 TU (Trap Unaligned Access) Current Processor Status Reg. 3-10
 TU (Trap Unaligned Access) Old Processor Status Reg. 3-10
 TU 3-46, 3-50, 3-73
 TWIN 8-3

U

- U (Usage) TLB Entry Word 1 3-29, 3-68
- UA (User Access) 3-41, 6-3
- UE (User Execute) TLB Entry Word 0 3-28, 3-71
- unaligned access 3-50, 3-53
- underflow 7-4
- underflow, signed 8-7
- underflow, unsigned 8-7
- UR (User Read) TLB Entry Word 0 3-28, 3-71
- Usage (U), TLB Entry Word 1 3-29, 3-68
- User Access (UA) 3-41, 6-3
- User Execute (UE), TLB Entry Word 0 3-28
- User Instruction/Data Memory 4-7
- User mode 2-1, 3-1
- User Read (UR), TLB Entry Word 0 3-28, 3-71
- User Write (UW), TLB Entry Word 0 3-28
- User-defined 5-7
- User-mode access 3-66
- User Programmable (PGM), TLB Entry Word 1 3-29, 3-68, 3-69
- UW (User Write) TLB Entry Word 0 3-28, 3-71

V

- V (Overflow) ALU Status Reg. 3-24, 8-6
- V. PROT 3-68
- VAB (Vector Area Base) Vector Area Base Address Reg. 3-9
- valid bits, Branch Target Cache 4-7, 4-8
- valid instructions in Cache 4-9
- valid transitions 5-24
- VE (Valid Entry) TLB Entry Word 0 3-28, 3-67, 4-6
- Vector Area 1-9, 2-13, 3-53, 7-20
- Vector Area Base (VAB), Vector Area Base Address Reg. 3-9
- Vector Area Base address 2-3
- Vector Area Base Address Register 3-8, 3-9, 3-53, 3-55
- Vector Fetch (VF), Configuration Reg. 3-13
- vector number 3-53, 3-55, 7-9
- vector number assignment 3-53, 3-54
- vector table entry 3-53
- vectors, table of 2-13
- VF (Vector Fetch) Configuration Reg. 3-13, 3-53, 3-55
- virtual address 3-68, 3-69, 7-7
- virtual address for page sizes 3-66
- virtual address space 4-18


Virtual page size 7-23
Virtual Tag (VTAG) 3-28, 3-68
virtual-page boundary 5-19
visual to physical address translation 1-8
VN 8-5, 8-9
VTAG (Virtual Tag) TLB Entry Word 0 3-28, 3-67

W

Wait mode 2-19, 3-52, 3-73, 5-5
WAIT Mode (WM), Current Processor Status Reg. 3-10
WAIT Mode (WM), Old Processor Status Reg. 3-10
warm start 7-25
Warn (*WARN) 3-52, 3-59, 4-19, 5-5, 5-31, 5-32
WARN trap differences 3-59
WM (WAIT Mode) Current Processor Status Reg. 3-10
WM (WAIT Mode) Old Processor Status Reg. 3-10
word constant 3-33
write-back 4-1, 4-14

Z

Z (Zero) ALU Status Reg. 3-24, 8-6
zero (Z) 8-6





**ADVANCED
MICRO
DEVICES, INC.**

*901 Thompson Place
P.O. Box 3453
Sunnyvale,
California 94088
(408) 732-2400
TWX: 910-339-9280
TELEX: 34-6306
TOLL FREE
(800) 538-8450*