

NS32132-6/NS32132-8/NS32132-10 High-Performance Microprocessors

General Description

The NS32132 is a 32-bit, virtual memory microprocessor with a 16-MByte linear address space and a 32-bit external data bus. It has a 32-bit ALU, eight 32-bit general purpose registers, an eight-byte prefetch queue, and a slave processor interface. A dual processing capability is also provided that allows two NS32123 to share the same bus and memory with a resulting performance enhancement of up to 80%. The NS32132 is fabricated with National Semiconductor's advanced XMOS process, and is fully object code compatible with other Series 32000 processors. The Series 32000 instructions set is optimized for modular high-level languages (HLL). The set is very symmetric, it has a two address format, and it incorporates HLL oriented addressing modes. The capabilities of the NS32132 can be expanded with the use of the NS32081 floating point unit (FPU), and the NS32032 demand-paged virtual memory management unit (MMU). Both devices interface to the NS32132 as slave processors. The NS32132 is a general purpose microprocessor that is ideal for a wide range of computational intensive applications.

Features

- 32-bit Architecture and Implementation
- Virtual Memory Support
- 16-MByte Linear Addressing Space
- 32-bit Data Bus
- Tightly Coupled Dual Processing Support
- Powerful Instruction Set
 - General 2-Address Capability
 - High Degree of Symmetry
 - Addressing Modes Optimized for High Level Language
- Series 32000® Slave Processor Support
- High-Speed XMOSTM Technology
- 68-pin Leadless Chip Carrier

Block Diagram

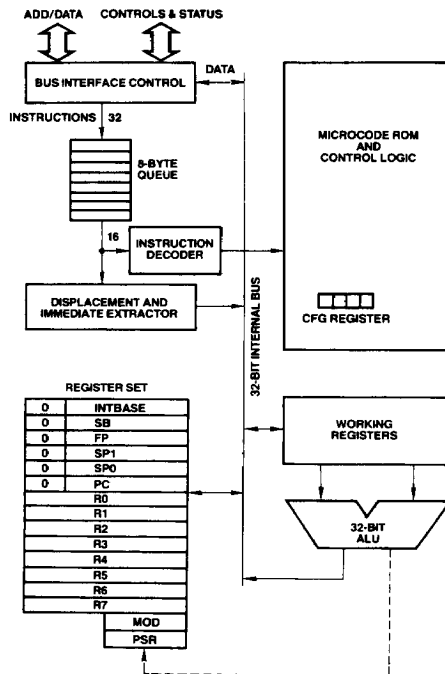


Table of Contents

1.0 PRODUCT INTRODUCTION

2.0 ARCHITECTURAL DESCRIPTION

- 2.1 Programming Model
 - 2.1.1 General Purpose Registers
 - 2.1.2 Dedicated Registers
 - 2.1.3 The Configuration Register (CFG)
 - 2.1.4 Memory Organization
 - 2.1.5 Dedicated Tables
- 2.2 Instruction Set
 - 2.2.1 General Instruction Format
 - 2.2.2 Addressing Modes
 - 2.2.3 Instruction Set Summary

3.0 FUNCTIONAL DESCRIPTION

- 3.1 Power and Grounding
- 3.2 Clocking
- 3.3 Resetting
- 3.4 Bus Cycles
 - 3.4.1 Cycle Extension
 - 3.4.2 Bus Status
 - 3.4.3 Data Access Sequences
 - 3.4.3.1 Bit Accesses
 - 3.4.3.2 Bit Field Accesses
 - 3.4.3.3 Extending Multiply Accesses
 - 3.4.4 Instruction Fetches
 - 3.4.5 Interrupt Control Cycles
 - 3.4.6 Slave Processor Communication
 - 3.4.6.1 Slave Processor Bus Cycles
 - 3.4.6.2 Slave Operand Transfer Sequences
- 3.5 Memory Management Option
 - 3.5.1 Address Translation Strap
 - 3.5.2 Translated Bus Timing
 - 3.5.3 The FLT (Float) Pin
 - 3.5.4 Aborting Bus Cycles
 - 3.5.4.1 The Abort Interrupt
 - 3.5.4.2 Hardware Considerations
- 3.6 Bus Access Control
- 3.7 Dual Processing
 - 3.7.1 Bus Arbitration
 - 3.7.2 Processor Assignment

3.0 FUNCTIONAL DESCRIPTION (Continued)

- 3.8 Instruction Status
 - 3.9 NS32132 Interrupt Structure
 - 3.9.1 General Interrupt/Trap Sequence
 - 3.9.2 Interrupt/Trap Return
 - 3.9.3 Maskable Interrupts (The INT Pin)
 - 3.9.3.1 Non-Vectored Mode
 - 3.9.3.2 Vectored-Mode: Non-Cascaded Case
 - 3.9.3.3 Vectored Mode: Cascaded Case
 - 3.9.4 Non-Maskable Interrupt (The NMI Pin)
 - 3.9.5 Traps
 - 3.9.6 Prioritization
 - 3.9.7 Interrupt/Trap Sequences: Detailed Flow
 - 3.9.7.1 Maskable/Non-Maskable Interrupt Sequences
 - 3.9.7.2 Trap Sequence: Traps Other than Trace
 - 3.9.7.3 Trace Trap Sequence
 - 3.9.7.4 Abort Sequence
 - 3.10 Slave Processor Instructions
 - 3.10.1 Slave Processor Protocol
 - 3.10.2 Floating Point Instructions
 - 3.10.3 Memory Management Instructions
 - 3.10.4 Custom Slave Instructions
- ### 4.0 DEVICE SPECIFICATIONS
- 4.1 Pin Descriptions
 - 4.1.1 Supplies
 - 4.1.2 Input Signals
 - 4.1.3 Output Signals
 - 4.1.4 Input/Output Signals
 - 4.2 Absolute Maximum Ratings
 - 4.3 Electrical Characteristics
 - 4.4 Switching Characteristics
 - 4.4.1 Definitions
 - 4.4.2 Timing Tables
 - 4.4.2.1 Output Signals: Internal Propagation Delays
 - 4.4.2.2 Input Signals Requirements
 - 4.4.2.3 Clocking Requirements
 - 4.4.3 Timing Diagrams
- Appendix A: Instruction Formats
 B: Interfacing Suggestions

List of Illustrations

The General and Dedicated Registers	2-1
Processor Status Register	2-2
CFG Register	2-3
Module Descriptor Format	2-4
A Sample Link Table	2-5
General Instruction Format	2-6
Index Byte Format	2-7
Displacement Encodings	2-8
Recommended Supply Connections	3-1
Clock Timing Relationships	3-2
Power-on Reset Requirements	3-3
General Reset Timing	3-4
Recommended Reset Connections, Non-Memory-Managed System	3-5a
Recommended Reset Connections, Memory-Managed System	3-5b
Bus Connections	3-6
Read Cycle Timing	3-7
Write Cycle Timing	3-8
RDY Pin Timing	3-9
Extended Cycle Example	3-10
Memory Interface	3-11
Slave Processor Connections	3-12
CPU Read from Slave Processor	3-13
CPU Write to Slave Processor	3-14
Read Cycle With Address Translation (CPU Action)	3-15
Write Cycle With Address Translation (CPU Action)	3-16
Memory-Managed Read Cycle	3-17
Memory-Managed Write Cycle	3-18
FLT Timing	3-19
HOLD Timing, Bus Initially Idle	3-20
HOLD Timing, Bus Initially Not Idle	3-21
Dual Processor System Basic Interconnections	3-22
Dual Processor Bus Arbitration Timing	3-23
Instruction Sequence To Assign CPUs to Different Tasks	3-24
Interrupt Dispatch and Cascade Tables	3-25
Interrupt/Trap Service Routine Calling Sequence	3-26
Return from Trap (RETT n) Instruction Flow	3-27
Return from Interrupt (RET) Instruction Flow	3-28
Interrupt Control Connections (16 levels)	3-29
Cascaded Interrupt Control Unit Connections	3-30
Service Sequence	3-31
Slave Processor Protocol	3-32
Slave Processor Status Word Format	3-33
NS32132 Connection Diagram	4-1
Timing Specification Standard (Signal Valid After Clock Edge)	4-2
Timing Specification Standard (Signal Valid Before Clock Edge)	4-3
Write Cycle	4-4
Read Cycle	4-5
Floating by HOLD Timing (CPU Not Idle Initially)	4-6
Floating by HOLD Timing (CPU Initially Idle)	4-7
Release from HOLD	4-8

List of Illustrations (Continued)

FLT Initiated Float Cycle Timing	4-9
Release from FLT Timing	4-10
Ready Sampling (CPU Initially READY)	4-11
Ready Sampling (CPU Initially NOT READY)	4-12
Slave Processor Write Timing	4-13
Slave Processor Read Timing	4-14
SPC Timing	4-15
Reset Configuration Timing	4-16
Clock Waveforms	4-17
Relationship of PFS to Clock Cycles	4-18
Guaranteed Delay, PFS to Non-Sequential Fetch	4-19a
Guaranteed Delay, Non-Sequential Fetch to PFS	4-19b
Relationship of ILO to First Operand of an Interlocked Instruction	4-20a
Relationship of ILO to Last Operand of an Interlocked Instruction	4-20b
Relationship of ILO to Any Clock Cycle	4-21
U/S Relationship to any Bus Cycle - Guaranteed Valid Interval	4-22
Abort Timing, FLT Not Applied	4-23
Abort Timing, FLT Applied	4-24
Power-On Reset	4-25
Non-Power-On Reset	4-26
INT Interrupt Signal Detection	4-27
NMI Interrupt Signal Timing	4-28
Relationship Between Last Data Transfer of an Instruction and PFS Pulse of Next Instruction	4-29
Dual Processor Bus Arbitration Timing	4-30
Single Processor System Connection Diagram	B-1
Dual Processor System Connection Diagram	B-2

List of Tables

NS32132 Addressing Modes	2-1
NS32132 Instruction Set Summary	2-2
Bus Access Types	3-1
Access Sequences	3-2
Interrupt Sequences	3-3
Floating Point Instruction Protocols	3-4
Memory Management Instruction Protocols	3-5
Custom Slave Instruction Protocols	3-6

1.0 Product Introduction

The Series 32000 Microprocessor family is a new generation of devices using National's XMOS and CMOS technologies. By combining state-of-the-art MOS technology with a very advanced architectural design philosophy, this family brings mainframe computer processing power to VLSI processors.

The Series 32000 family supports a variety of system configurations, extending from minimum low-cost system to a powerful 4 gigabyte system. The architecture provides complete upward compatibility from one family member to another. The family consists of a selection of CPUs supported by a set of peripherals and slave processors that provide sophisticated interrupt and memory management facilities as well as high-speed floating-point operation. The architectural features of the Series 32000 family are described briefly below:

Powerful Addressing Modes. Nine addressing modes available to all instructions are included to access data structures efficiently.

Data Types. The architecture provides for numerous data types, such as byte, word, doubleword, and BCD, which may be arranged into a wide variety of data structures.

Symmetric Instruction Set. While avoiding special case instructions that compilers can't use, the Series 32000 family incorporates powerful instructions for control operations, such as array indexing and external procedure calls, which save considerable space and time for compiled code.

Memory-to-Memory Operations. The Series 32000 CPUs represent two-operand machines with each operand addressable by all addressing modes. This powerful memory-to-memory architecture permits memory locations to be treated as registers for all useful operations. This is important for temporary operands as well as for context switching.

Memory Management. Either the NS32382 or the NS32082 Memory Management Unit may be added to the system to provide advanced operating system support functions, including dynamic address translation, virtual memory management, and memory protection.

Large, Uniform Addressing. The NS32132 has 24-bit address pointers that can address up to 16 megabytes without requiring any segmentation; this addressing scheme provides flexible memory management without added-on expense.

Modular Software Support. Any software package for the Series 32000 family can be developed independent of all other packages, without regard to individual addressing. In addition, ROM code is totally relocatable and easy to access, which allows a significant reduction in hardware and software cost.

Software Processor Concept. The Series 32000 architecture allows future expansions of the instruction set that can be executed by special slave processors, acting as extensions to the CPU. This concept of slave processors is unique to the Series 32000 family. It allows software compatibility even for future components because the slave hardware is transparent to the software. With future advances in semiconductor technology, the slaves can be physically integrated on the CPU chip itself.

To summarize, the architectural features cited above provide three primary performance advantages and characteristics:

- High-Level Language Support
- Easy Future Growth Path
- Application Flexibility

2.0 Architectural Description

2.1 PROGRAMMING MODEL

The Series 32000 architecture includes 16 registers on the NS32132 CPU.

2.1.1 General Purpose Registers

There are eight registers for meeting high speed general storage requirements, such as holding temporary variables and addresses. The general purpose registers are free for any use by the programmer. They are thirty-two bits in length. If a general register is specified for an operand that is eight or sixteen bits long, only the low part of the register is used; the high part is not referenced or modified.

2.1.2 Dedicated Registers

The eight dedicated registers of the NS32132 are assigned specific functions.

PC: The PROGRAM COUNTER register is a pointer to the first byte of the instruction currently being executed. The PC is used to reference memory in the program section. (In the NS32132 the upper eight bits of this register are always zero.)

SP0, SP1: The SP0 register points to the lowest address of the last item stored on the INTERRUPT STACK. This stack is normally used only by the operating system. It is used primarily for storing temporary data, and holding return information for operating system subroutines and interrupt and trap service routines. The SP1 register points to the lowest address of the last item stored on the USER STACK. This stack is used by normal user programs to hold temporary data and subroutine return information.

In this document, reference is made to the SP register. The terms "SP register" or "SP" refer to either SP0 or SP1, depending on the setting of the S bit in the PSR register. If the S bit in the PSR is 0 the SP refers to SP0. If the S bit in the PSR is 1 then SP refers to SP1. (In the NS32132 the upper eight bits of these registers are always zero.)

Stacks in the Series 32000 family grow downward in memory. A Push operation pre-decrements the Stack Pointer by the operand length. A Pop operation post-increments the Stack Pointer by the operand length.

FP: The FRAME POINTER register is used by a procedure to access parameters and local variables on the stack. The FP register is set up on procedure entry with the ENTER instruction and restored on procedure termination with the EXIT instruction.

The frame pointer holds the address in memory occupied by the old contents of the frame pointer. (In the NS32132 the upper eight bits of this register are always zero.)

SB: The STATIC BASE register points to the global variables of a software module. This register is used to support relocatable global variables for software modules. The SB register holds the lowest address in memory occupied by the global variables of a module. (In the NS32132 the upper eight bits of this register are always zero.)

INTBASE: The INTERRUPT BASE register holds the address of the dispatch table for interrupts and traps (Sec. 3.8). The INTBASE register holds the lowest address in memory occupied by the dispatch table. (In the NS32132 the upper eight bits of this register are always zero.)

2.0 Architectural Description (Continued)

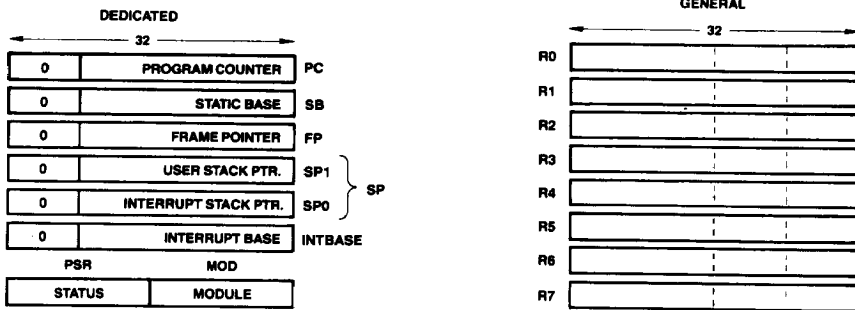



FIGURE 2-1. The General and Dedicated Registers

TL/EE/8583-3

MOD: The MODULE register holds the address of the module descriptor of the currently executing software module. The MOD register is sixteen bits long, therefore the module table must be contained within the first 64K bytes of memory.

PSR: The PROCESSOR STATUS REGISTER (PSR) holds the status codes for the NS32132 microprocessor.

The PSR is sixteen bits long, divided into two eight-bit halves. The low order eight bits are accessible to all programs, but the high order eight bits are accessible only to programs executing in Supervisor Mode.



TL/EE/8583-4

FIGURE 2-2. Processor Status Register

C: The C bit indicates that a carry or borrow occurred after an addition or subtraction instruction. It can be used with the ADDC and SUBC instructions to perform multiple-precision integer arithmetic calculations. It may have a setting of 0 (no carry or borrow) or 1 (carry or borrow).

T: The T bit causes program tracing. If this bit is a 1, a TRC trap is executed after every instruction (Sec. 3.8.5).

L: The L bit is altered by comparison instructions. In a comparison instruction the L bit is set to "1" if the second operand is less than the first operand, when both operands are interpreted as unsigned integers. Otherwise, it is set to "0". In Floating Point comparisons, this bit is always cleared.

F: The F bit is a general condition flag, which is altered by many instructions (e.g., integer arithmetic instructions use it to indicate overflow).

Z: The Z bit is altered by comparison instructions. In a comparison instruction the Z bit is set to "1" if the second operand is equal to the first operand; otherwise it is set to "0".

N: The N bit is altered by comparison instructions. In a comparison instruction the N bit is set to "1" if the second operand is less than the first operand, when both operands are interpreted as signed integers. Otherwise, it is set to "0".

U: If the U bit is "1" no privileged instructions may be executed. If the U bit is "0" then all instructions may be executed. When U = 0 the NS32132 is said to be in Supervisor Mode; when U = 1 the NS32132 is said to be in User Mode. A User Mode program is restricted from executing certain instructions and accessing certain registers which could interfere with the operating system. For example, a User Mode program is prevented from changing the setting of the flag used to indicate its own privilege mode. A Supervisor Mode program is assumed to be a trusted part of the operating system, hence it has no such restrictions.

S: The S bit specifies whether the SP0 register or SP1 register is used as the stack pointer. The bit is automatically cleared on interrupts and traps. It may have a setting of 0 (use the SP0 register) or 1 (use the SP1 register).

P: The P bit prevents a TRC trap from occurring more than once for an instruction (Sec. 3.8.5). It may have a setting of 0 (no trace pending) or 1 (trace pending).

I: If I = 1, then all interrupts will be accepted (Sec. 3.8.). If I = 0, only the NMI interrupt is accepted. Trap enables are not affected by this bit.

2.1.3 The Configuration Register (CFG)

Within the Control section of the NS32132 CPU is the four-bit CFG Register, which declares the presence of certain external devices. It is referenced by only one instruction, SETCFG, which is intended to be executed only as part of system initialization after reset. The format of the CFG Register is shown in Figure 2-3.




FIGURE 2-3. CFG Register

The CFG I bit declares the presence of external interrupt vectoring circuitry (specifically, the NS32202 Interrupt Control Unit). If the CFG I bit is set, interrupts requested through the INT pin are "Vectored." If it is clear, these interrupts are "Non-Vectored." See Sec. 3.8.

The F, M and C bits declare the presence of the FPU, MMU and Custom Slave Processors. If these bits are not set, the corresponding instructions are trapped as being undefined.

2.0 Architectural Description (Continued)


2.1.4 Memory Organization

The main memory of the NS32132 is a uniform linear address space. Memory locations are numbered sequentially starting at zero and ending at $2^{24} - 1$. The number specifying a memory location is called an address. The contents of each memory location is a byte consisting of eight bits. Unless otherwise noted, diagrams in this document show data stored in memory with the lowest address on the right and the highest address on the left. Also, when data is shown vertically, the lowest address is at the top of a diagram and the highest address at the bottom of the diagram. When bits are numbered in a diagram, the least significant bit is given the number zero, and is shown at the right of the diagram. Bits are numbered in increasing significance and toward the left.




Byte at Address A

Two contiguous bytes are called a word. Except where noted (Sec. 2.2.1), the least significant byte of a word is stored at the lower address, and the most significant byte of the word is stored at the next higher address. In memory, the address of a word is the address of its least significant byte, and a word may start at any address.



Word at Address A

Two contiguous words are called a double word. Except where noted (Sec. 2.2.1), the least significant word of a double word is stored at the lowest address and the most significant word of the double word is stored at the address two greater. In memory, the address of a double word is the address of its least significant byte, and a double word may start at any address.



Double Word at Address A

Although memory is addressed as bytes, it is actually organized as double-words. Note that access time to a word or a double-word depends upon its address, e.g. double-words that are aligned to start at addresses that are multiples of four will be accessed more quickly than those not so aligned. This also applies to words that cross a double-word boundary.

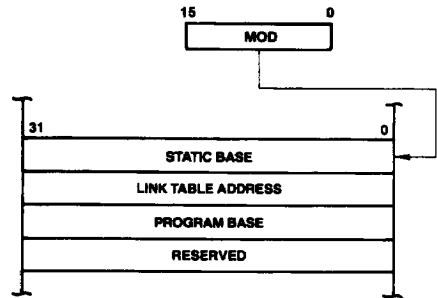
2.1.5 Dedicated Tables

Two of the NS32132 dedicated registers (MOD and INTBASE) serve as pointers to dedicated tables in memory.

The INTBASE register points to the Interrupt Dispatch and Cascade tables. These are described in Sec. 3.8.

The MOD register contains a pointer into the Module Table, whose entries are called Module Descriptors. A Module Descriptor contains four pointers, three of which are used by NS32132. The MOD register contains the address of the Module Descriptor for the currently running module. It is automatically up-dated by the Call External Procedure instructions (CXP and CXPD).

The format of a Module Descriptor is shown in *Figure 2-4*. The Static Base entry contains the address of static data assigned to the running module. It is loaded into the CPU Static Base register by the CXP and CXPD instructions. The Program Base entry contains the address of the first byte of instruction code in the module. Since a module may have multiple entry points, the Program Base pointer serves only as a reference to find them.



TL/EE/8583-5


FIGURE 2-4. Module Descriptor Format

The Link Table Address points to the Link Table for the currently running module. The Link Table provides the information needed for:

- 1) Sharing variables between modules. Such variables are accessed through the Link Table via the External addressing mode.
- 2) Transferring control from one module to another. This is done via the Call External Procedure (CXP) instruction.

The format of a Link Table is given in *Figure 2-5*. A Link Table Entry for an external variable contains the 32-bit address of that variable. An entry for an external procedure contains two 16-bit fields: Module and Offset. The Module field contains the new MOD register contents for the module being entered. The Offset field is an unsigned number giving the position of the entry point relative to the new module's Program Base pointer.

For further details of the functions of these tables, see the Series 32000 Instruction Set Reference Manual.



TL/EE/8583-6

FIGURE 2-5. A Sample Link Table

2.0 Architectural Description (Continued)

2.2 INSTRUCTION SET

2.2.1 General Instruction Format

Figure 2-6 shows the general format of a Series 32000 instruction. The Basic Instruction is one to three bytes long and contains the Opcode and up to two 5-bit General Addressing Mode ("Gen") fields. Following the Basic Instruction field is a set of optional extensions, which may appear depending on the instruction and the addressing modes selected.

Index Bytes appear when either or both Gen fields specify Scaled Index. In this case, the Gen field specifies only the Scale Factor (1, 2, 4 or 8), and the Index Byte specifies which General Purpose Register to use as the index, and which addressing mode calculation to perform before indexing. See Figure 2-7.






TL/EE/8583-8

FIGURE 2-7. Index Byte Format

Following Index Bytes come any displacements (addressing constants) or immediate values associated with the selected address modes. Each Disp/Imm field may contain one or two displacements, or one immediate value. The size of a Displacement field is encoded with the top bits of that field, as shown in Figure 2-8, with the remaining bits interpreted as a signed (two's complement) value. The size of an immediate value is determined from the Opcode field. Both Displacement and Immediate fields are stored most significant byte first. Note that this is different from the memory representation of data (Sec. 2.1.4).

Some instructions require additional, "implied" immediates and/or displacements, apart from those associated with addressing modes. Any such extensions appear at the end of the instruction, in the order that they appear within the list of operands in the instruction definition (Sec. 2.2.3).



TL/EE/8583-9

FIGURE 2-8. Displacement Encodings

2.2.2 Addressing Modes

The NS32132 CPU generally accesses an operand by calculating its Effective Address based on information available when the operand is to be accessed. The method to be used in performing this calculation is specified by the programmer as an "addressing mode."

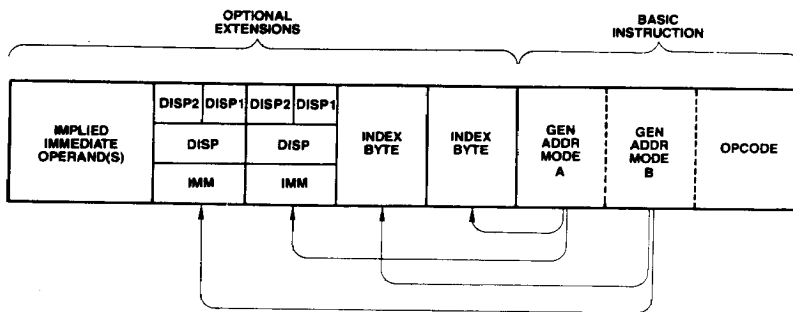


FIGURE 2-6. General Instruction Format

TL/EE/8583-7

2.0 Architectural Description (Continued)

Addressing modes in the NS32132 are designed to optimally support high-level language accesses to variables. In nearly all cases, a variable access requires only one addressing mode, within the instruction that acts upon that variable. Extraneous data movement is therefore minimized. NS32132 Addressing Modes fall into nine basic types:

Register: The operand is available in one of the eight General Purpose Registers. In certain Slave Processor instructions, an auxiliary set of eight registers may be referenced instead.

Register Relative: A General Purpose Register contains an address to which is added a displacement value from the instruction, yielding the Effective Address of the operand in memory.

Memory Space. Identical to Register Relative above, except that the register used is one of the dedicated registers PC, SP, SB or FP. These registers point to data areas generally needed by high-level languages.

Memory Relative: A pointer variable is found within the memory space pointed to by the SP, SB or FP register. A displacement is added to that pointer to generate the Effective Address of the operand.

Immediate: The operand is encoded within the instruction. This addressing mode is not allowed if the operand is to be written.

Absolute: The address of the operand is specified by a displacement field in the instruction.

External: A pointer value is read from a specified entry of the current Link Table. To this pointer value is added a displacement, yielding the Effective Address of the operand.

Top of Stack: The currently-selected Stack Pointer (SP0 or SP1) specifies the location of the operand. The operand is pushed or popped, depending on whether it is written or read.

Scaled Index: Although encoded as an addressing mode. Scaled Indexing is an option on any addressing mode except Immediate or another Scaled Index. It has the effect of calculating an Effective Address, then multiplying any General Purpose Register by 1, 2, 4 or 8 and adding it into the total, yielding the final Effective Address of the operand.

Table 2-1 is a brief summary of the addressing modes. For a complete description of their actions, see the Instruction Set Reference Manual.

2.2.3 Instruction Set Summary

Table 2-2 presents a brief description of the NS32132 instruction set. The Format column refers to the Instruction Format tables (Appendix A). The Instruction column gives the instruction as coded in assembly language, and the Description column provides a short description of the function provided by that instruction. Further details of the exact operations performed by each instruction may be found in the Instruction Set Reference Manual.

Notations:

i = Integer length suffix: B = Byte
W = Word
D = Double Word

f = Floating-Point length suffix: F = Standard Floating
L = Long Floating

gen = General operand. Any addressing mode can be specified.

short = A 4-bit value encoded within the Basic Instruction (see Appendix A for encodings).

imm = Implied immediate operand. An 8-bit value appended after any addressing extensions.

disp = Displacement (addressing constant); 8, 16 or 32 bits. All three lengths legal.

reg = Any General Purpose Register: R0–R7.

areg = Any Dedicated/Address Register: SP, SB, FP, MOD, INTBASE, PSR, US (bottom 8 PSR bits).

mreg = Any Memory Management Status/Control Register.

creg = A Custom Slave Processor Register (Implementation Dependent).

cond = Any condition code, encoded as a 4-bit field within the Basic Instruction (see Appendix A for encodings).

2.0 Architectural Description (Continued)

TABLE 2-1
NS32132 Addressing Modes

ENCODING	MODE	ASSEMBLER SYNTAX	EFFECTIVE ADDRESS
Register			
00000	Register 0	R0 or F0	None: Operand is in the specified register
00001	Register 1	R1 or F1	
00010	Register 2	R2 or F2	
00011	Register 3	R3 or F3	
00100	Register 4	R4 or F4	
00101	Register 5	R5 or F5	
00110	Register 6	R6 or F6	
00111	Register 7	R7 or F7	
Register Relative			
01000	Register 0 relative	disp(R0)	Disp + Register.
01001	Register 1 relative	disp(R1)	
01010	Register 2 relative	disp(R2)	
01011	Register 3 relative	disp(R3)	
01100	Register 4 relative	disp(R4)	
01101	Register 5 relative	disp(R5)	
01110	Register 6 relative	disp(R6)	
01111	Register 7 relative	disp(R7)	
Memory Relative			
10000	Frame memory relative	disp2(disp1(FP))	Disp2 + Pointer; Pointer found at address Disp1 + Register. "SP" is either SP0 or SP1, as selected in PSR.
10001	Stack memory relative	disp2(disp1(SP))	
10010	Static memory relative	disp2(disp1(SB))	
Reserved			
10011	(Reserved for Future Use)		
Immediate			
10100	Immediate	value	None: Operand is input from instruction queue.
Absolute			
10101	Absolute	@disp	Disp.
External			
10110	External	EXT (disp1) + disp2	Disp2 + Pointer; Pointer is found at Link Table Entry number Disp1.
Top of Stack			
10111	Top of stack	TOS	Top of current stack, using either User or Interrupt Stack Pointer, as selected in PSR. Automatic Push/Pop included.
Memory Space			
11000	Frame memory	disp(FP)	Disp + Register; "SP" is either SP0 or SP1, as selected in PSR.
11001	Stack memory	disp(SP)	
11010	Static memory	disp(SB)	
11011	Program memory	* + disp	
Scaled Index			
11100	Index, bytes	mode[Rn:B]	EA (mode) + Rn.
11101	Index, words	mode[Rn:W]	EA (mode) + 2 × Rn.
11110	Index, double words	mode[Rn:D]	EA (mode) + 4 × Rn.
11111	Index, quad words	mode[Rn:Q]	EA (mode) + 8 × Rn.
			'Mode' and 'n' are contained within the Index Byte.
			EA (mode) denotes the effective address generated using mode.

2.0 Architectural Description (Continued)

TABLE 2-2
NS32132 Instruction Set Summary

MOVES

Format	Operation	Operands	Description
4	MOVi	gen,gen	Move a value.
2	MOVQi	short,gen	Extend and move a signed 4-bit constant.
7	MOVMI	gen,gen,disp	Move Multiple: disp bytes (1 to 16).
7	MOVZBW	gen,gen	Move with zero extension.
7	MOVZID	gen,gen	Move with zero extension.
7	MOVXBW	gen,gen	Move with sign extension.
7	MOVXID	gen,gen	Move with sign extension.
4	ADDR	gen,gen	Move Effective Address.

INTEGER ARITHMETIC

Format	Operation	Operands	Description
4	ADDI	gen,gen	Add.
2	ADDQi	short,gen	Add signed 4-bit constant.
4	ADDci	gen,gen	Add with carry.
4	SUBi	gen,gen	Subtract.
4	SUBci	gen,gen	Subtract with carry (borrow).
6	NEGi	gen,gen	Negate (2's complement).
6	ABSi	gen,gen	Take absolute value.
7	MULi	gen,gen	Multiply
7	QUOi	gen,gen	Divide, rounding toward zero.
7	REMi	gen,gen	Remainder from QUO.
7	DIVi	gen,gen	Divide, rounding down.
7	MODi	gen,gen	Remainder from DIV (Modulus).
7	MEIi	gen,gen	Multiply to Extended Integer.
7	DEIi	gen,gen	Divide Extended Integer.

PACKED DECIMAL (BCD) ARITHMETIC

Format	Operation	Operands	Description
6	ADDPi	gen,gen	Add Packed.
6	SUBPi	gen,gen	Subtract Packed.

INTEGER COMPARISON

Format	Operation	Operands	Description
4	CMPI	gen,gen	Compare.
2	CMQi	short,gen	Compare to signed 4-bit constant.
7	CMPMI	gen,gen,disp	Compare Multiple: disp bytes (1 to 16).

LOGICAL AND BOOLEAN

Format	Operation	Operands	Description
4	ANDi	gen,gen	Logical AND.
4	ORi	gen,gen	Logical OR.
4	BICi	gen,gen	Clear selected bits.
4	XORi	gen,gen	Logical Exclusive OR.
6	COMi	gen,gen	Complement all bits.
6	NOTi	gen,gen	Boolean complement: LSB only.
2	Scondi	gen	Save condition code (cond) as a Boolean variable of size i.

2.0 Architectural Description (Continued)

SHIFTS

Format	Operation	Operands	Description
6	LSHi	gen,gen	Logical Shift, left or right.
6	ASHi	gen,gen	Arithmetic Shift, left or right.
6	ROTi	gen,gen	Rotate, left or right.

BITS

Format	Operation	Operands	Description
4	TBITi	gen,gen	Test bit.
6	SBITi	gen,gen	Test and set bit.
6	SBITii	gen,gen	Test and set bit, interlocked
6	CBITi	gen,gen	Test and clear bit.
6	CBITii	gen,gen	Test and clear bit, interlocked.
6	IBITi	gen,gen	Test and invert bit.
8	FFSi	gen,gen	Find first set bit

BIT FIELDS

Bit fields are values in memory that are not aligned to byte boundaries. Examples are PACKED arrays and records used in Pascal. "Extract" instructions read and align a bit field. "Insert" instructions write a bit field from an aligned source.

Format	Operation	Operands	Description
8	EXTi	reg,gen,gen,disp	Extract bit field (array oriented).
8	INSi	reg,gen,gen,disp	Insert bit field (array oriented).
7	EXTSi	gen,gen,imm,imm	Extract bit field (short form).
7	INSSi	gen,gen,imm,imm	Insert bit field (short form).
8	CVTP	reg,gen,gen	Convert to Bit Field Pointer.

ARRAYS

Format	Operation	Operands	Description
8	CHECKi	reg,gen,gen	Index bounds check.
8	INDEXi	reg,gen,gen	Recursive indexing step for multiple-dimensional arrays.

STRINGS

String instructions assign specific functions to the General Purpose Registers:

- R4 - Comparison Value
- R3 - Translation Table Pointer
- R2 - String 2 Pointer
- R1 - String 1 Pointer
- R0 - Limit Count

Options on all string instructions are:

- B** (Backward): Decrement string pointers after each step rather than incrementing.
- U** (Until match): End instruction if String 1 entry matches R4.
- W** (While match): End instruction if String 1 entry does not match R4.

All string instructions end when R0 decrements to zero.

Format	Operation	Operands	Descriptions
5	MOVSi	options	Move String 1 to String 2.
	MOVST	options	Move string, translating bytes.
5	CMPSi	options	Compare String 1 to String 2.
	CMPST	options	Compare translating, String 1 bytes.
5	SKPSi	options	Skip over String 1 entries
	SKPST	options	Skip, translating bytes for Until/While.

2.0 Architectural Description (Continued)

JUMPS AND LINKAGE

Format	Operation	Operands	Description
3	JUMP	gen	Jump.
0	BR	disp	Branch (PC Relative).
0	Bcond	disp	Conditional branch.
3	CASEi	gen	Multway branch.
2	ACBi	short,gen,disp	Add 4-bit constant and branch if non-zero.
3	JSR	gen	Jump to subroutine.
1	BSR	disp	Branch to subroutine.
1	CXP	disp	Call external procedure.
3	CXPD	gen	Call external procedure using descriptor.
1	SVC		Supervisor Call.
1	FLAG		Flag Trap.
1	BPT		Breakpoint Trap.
1	ENTER	[reg list],disp	Save registers and allocate stack frame (Enter Procedure).
1	EXIT	[reg list]	Restore registers and reclaim stack frame (Exit Procedure).
1	RET	disp	Return from subroutine.
1	RXP	disp	Return from external procedure call.
1	RETT	disp	Return from trap. (Privileged)
1	RETI		Return from interrupt. (Privileged)

CPU REGISTER MANIPULATION

Format	Operation	Operands	Description
1	SAVE	[reg list]	Save General Purpose Registers.
1	RESTORE	[reg list]	Restore General Purpose Registers.
2	LPRi	areg,gen	Load Dedicated Register. (Privileged if PSR or INTBASE)
2	SPRi	areg,gen	Store Dedicated Register. (Privileged if PSR or INTBASE)
3	ADJSPI	gen	Adjust Stack Pointer.
3	BISPSRi	gen	Set selected bits in PSR. (Privileged if not Byte length)
3	BICPSRi	gen	Clear selected bits in PSR. (Privileged if not Byte length)
5	SETCFG	[option list]	Set Configuration Register. (Privileged)

FLOATING POINT

Format	Operation	Operands	Description
11	MOVf	gen,gen	Move a Floating Point value.
9	MOVLF	gen,gen	Move and shorten a Long value to Standard.
9	MOVFL	gen,gen	Move and lengthen a Standard value to Long.
9	MOVif	gen,gen	Convert any integer to Standard or Long Floating.
9	ROUNDfi	gen,gen	Convert to integer by rounding.
9	TRUNCfi	gen,gen	Convert to integer by truncating, toward zero.
9	FLOORfi	gen,gen	Convert to largest integer less than or equal to value.
11	ADDf	gen,gen	Add.
11	SUBf	gen,gen	Subtract.
11	MULf	gen,gen	Multiply.
11	DIVf	gen,gen	Divide.
11	CMPf	gen,gen	Compare.
11	NEGf	gen,gen	Negate.
11	ABSf	gen,gen	Take absolute value.
9	LFSR	gen	Load FSR.
9	SFSR	gen	Store FSR.

MEMORY MANAGEMENT

Format	Operation	Operands	Description
14	LMR	mreg,gen	Load Memory Management Register. (Privileged)
14	SMR	mreg,gen	Store Memory Management Register. (Privileged)
14	RDVAL	gen	Validate address for reading. (Privileged)
14	WRVAL	gen	Validate address for writing. (Privileged)
8	MOVSi	gen,gen	Move a value from Supervisor Space to User Space. (Privileged)
8	MOVUSi	gen,gen	Move a value from User Space to Supervisor Space. (Privileged)

2.0 Architectural Description (Continued)

MISCELLANEOUS

Format	Operation	Operands	Description
1	NOP		No Operation.
1	WAIT		Wait for interrupt.
1	DIA		Diagnose. Single-byte "Branch to Self" for hardware breakpointing. Not for use in programming.

CUSTOM SLAVE

Format	Operation	Operands	Description
15.5	CCAL0c	gen,gen	Custom Calculate.
15.5	CCAL1c	gen,gen	
15.5	CCAL2c	gen,gen	
15.5	CCAL3c	gen,gen	
15.5	CMOV0c	gen,gen	Custom Move.
15.5	CMOV1c	gen,gen	
15.5	CMOV2c	gen,gen	
	CMOV3c	gen,gen	
15.5	CCMP0c	gen,gen	Custom Compare.
	CCMP1c	gen,gen	
	CCV0ci	gen,gen	Custom Convert.
15.1	CCV1ci	gen,gen	
15.1	CCV2ci	gen,gen	
15.1	CCV3ic	gen,gen	
15.1	CCV4DQ	gen,gen	
15.1	CCV5QD	gen,gen	
15.1	LCSR	gen	
15.1	SCSR	gen	Store Custom Status Register.
15.0	CATST0	gen	Custom Address/Test. (Privileged)
15.0	CATST1	gen	
15.0	LCR	creg,gen	Load Custom Register. (Privileged)
15.0	SCR	creg,gen	Store Custom Register. (Privileged)

3.0 Functional Description

3.1 POWER AND GROUNDING

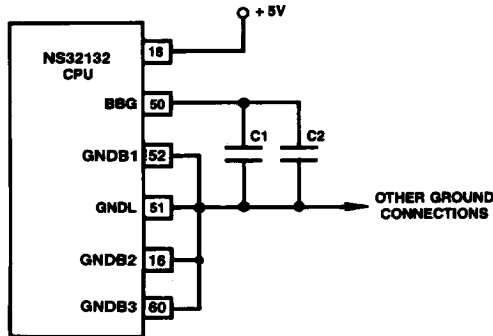
The NS32132 requires a single 5-volt power supply, applied on pin 18 (V_{CC}).

Grounding connections are made on four pins. Logic Ground (GNDL, pin 51) is the common pin for on-chip logic, and Buffer Grounds (GNDB1, pin 52, GNDB2, pin 16 and GNDB3, Pin 60) are the common pins for the output drivers. For optimal noise immunity it is recommended that GNDB1 and GNDB2 be connected together through a single conductor, and GNDL be directly connected to the middle point of this conductor. All other ground connections should be made to the common line as shown in *Figure 3-1*.

In addition to V_{CC} and Ground, the NS32132 CPU uses an internally-generated negative voltage. It is necessary to filter this voltage externally by attaching a pair of capacitors (*Fig. 3-7*) from the BBG pin to ground. Recommended values for these are:

C_1 : 1 μ F, Tantalum.

C_2 : 1000 pF, low inductance. This should be either a disc or monolithic ceramic capacitor.



TL/EE/8583-10


FIGURE 3-1. Recommended Supply Connections

3.2 CLOCKING

The NS32132 inputs clocking signals from the NS32201 Timing Control Unit (TCU), which presents two non-overlapping phases of a single clock frequency. These phases

are called PHI1 (pin 26) and PHI2 (pin 27). Their relationship to each other is shown in *Figure 3-2*.

Each rising edge of PHI1 defines a transition in the timing state ("T-State") of the CPU. One T-State represents the execution of one microinstruction within the CPU, and/or one step of an external bus transfer. See Sec. 4 for complete specifications of PHI1 and PHI2.



TL/EE/8583-11

FIGURE 3-2. Clock Timing Relationships


As the TCU presents signals with very fast transitions, it is recommended that the conductors carrying PHI1 and PHI2 be kept as short as possible, and that they not be connected anywhere except from the TCU to the CPU and, if present, the MMU. A TTL Clock signal (CTTL) is provided by the TCU for all other clocking.

3.3 RESETTING

The $\overline{RST}/\overline{ABT}$ pin serves both as a Reset for on-chip logic and as the Abort input for Memory-Managed systems. For its use as the Abort Command, see Sec. 3.5.4.

The CPU may be reset at any time by pulling the $\overline{RST}/\overline{ABT}$ pin low for at least 64 clock cycles. Upon detecting a reset, the CPU terminates instruction processing, resets its internal logic, and clears the Program Counter (PC) and Processor Status Register (PSR) to all zeroes.

On application of power, $\overline{RST}/\overline{ABT}$ must be held low for at least 50 μ sec after V_{CC} is stable. This is to ensure that all on-chip voltages are completely stable before operation. Whenever a Reset is applied, it must also remain



TL/EE/8583-12

FIGURE 3-3. Power-on Reset Requirements

3.0 Functional Description (Continued)

active for not less than 64 clock cycles. The rising edge must occur while PHI1 is high. See *Figures 3-3 and 3-4*.

The NS32201 Timing Control Unit (TCU) provides circuitry to meet the Reset requirements of the NS32132 CPU. *Figure 3-5a* shows the recommended connections for a non-Memory-Managed system. *Figure 3-5b* shows the connections for a Memory-Managed system.




FIGURE 3-4. General Reset Timing




FIGURE 3-5a. Recommended Reset Connections, Non-Memory-Managed System

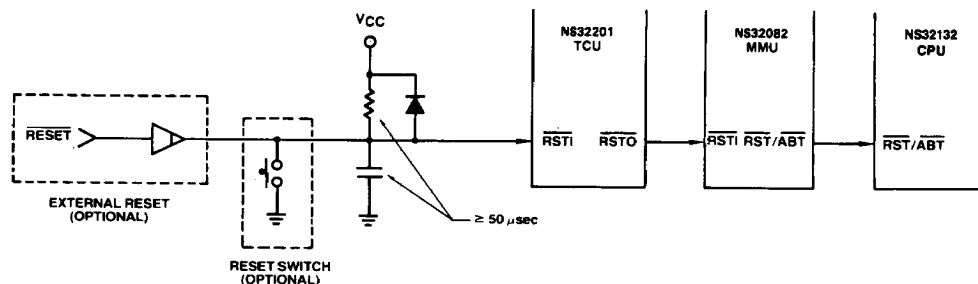


FIGURE 3-5b. Recommended Reset Connections, Memory-Managed System

3.4 BUS CYCLES

The NS32132 CPU has a strap option which defines the Bus Timing Mode as either With or Without Address Translation. This section describes only bus cycles under the No Address Translation option. For details of the use of the strap and of bus cycles with address translation, see Sec. 3.5.

The CPU will perform a bus cycle for one of the following reasons:

- 1) To write or read data, to or from memory or a peripheral interface device. Peripheral input and output are memory-mapped in the Series 32000 family.
- 2) To fetch instructions into the eight-byte instruction queue. This happens whenever the bus would otherwise be idle and the queue is not already full.

- 3) To acknowledge an interrupt and allow external circuitry to provide a vector number, or to acknowledge completion of an interrupt service routine.

- 4) To transfer information to or from a Slave Processor.

In terms of bus timing, cases 1 through 3 above are identical. For timing specifications, see Sec. 4. The only external difference between them is the four-bit code placed on the Bus Status pins (ST0-ST3). Slave Processor cycles differ in that separate control signals are applied (Sec. 3.4.6).

The sequence of events in a non-Slave bus cycle is shown below in *Figure 3-7* for a Read cycle and *Figure 3-8* for a Write cycle. The cases shown assume that the selected memory or interface device is capable of communicating with the CPU at full speed. If it is not, then cycle extension may be requested through the RDY line (Sec. 3.4.1).

3.0 Functional Description (Continued)

A full-speed bus cycle is performed in four cycles of the PHI1 clock signal, labeled T1 through T4. Clock cycles not associated with a bus cycle are designated T_i (for "Idle").

During T1, the CPU applies an address on pins AD0-AD23. It also provides a low-going pulse on the ADS pin, which serves the dual purpose of informing external circuitry that a bus cycle is starting and of providing control to an external latch for demultiplexing Address bits 0-23 from the AD0-AD23 pins. See *Figure 3-6*. During this time also the status signals \overline{DDIN} , indicating the direction of the transfer, and $\overline{BE0}-\overline{BE3}$, indicating which of the four bus bytes are to be referenced, become valid.

During T2 the CPU switches the Data Bus, AD0-AD31 to either accept or present data. It also starts the data strobe (\overline{DS}), signalling the beginning of the data transfer. Associated signals from the NS32201 Timing Control Unit are also activated at this time: \overline{RD} (Read Strobe) or \overline{WR} (Write Strobe), \overline{TSO} (Timing State Output, indicating that T2 has been reached) and DBE (Data Buffer Enable).

The T3 state provides for access time requirements, and it occurs at least once in a bus cycle. At the end of T2, on the falling edge of the PHI2 clock, the RDY (Ready) line is sampled to determine whether the bus cycle will be extended (Sec. 3.4.1).

If the CPU is performing a Read cycle, the Data Bus (AD0-AD31) is sampled at the falling edge of PHI2 of the last T3 state. See Timing Specification, Sec. 4. Data must, however, be held at least until the beginning of T4. \overline{DS} and \overline{RD} are guaranteed not to go inactive before this point, so the rising edge of either of them may safely be used to disable the device providing the input data.

The T4 state finishes the bus cycle. At the beginning of T4, the \overline{DS} , \overline{RD} or \overline{WR} , and \overline{TSO} signals go inactive, and at the rising edge of PHI2, \overline{DBE} goes inactive, having provided for necessary data hold times. Data during Write cycles remains valid from the CPU throughout T4. Note that the Bus Status lines (ST0-ST3) change at the beginning of T4, anticipating the following bus cycle (if any).

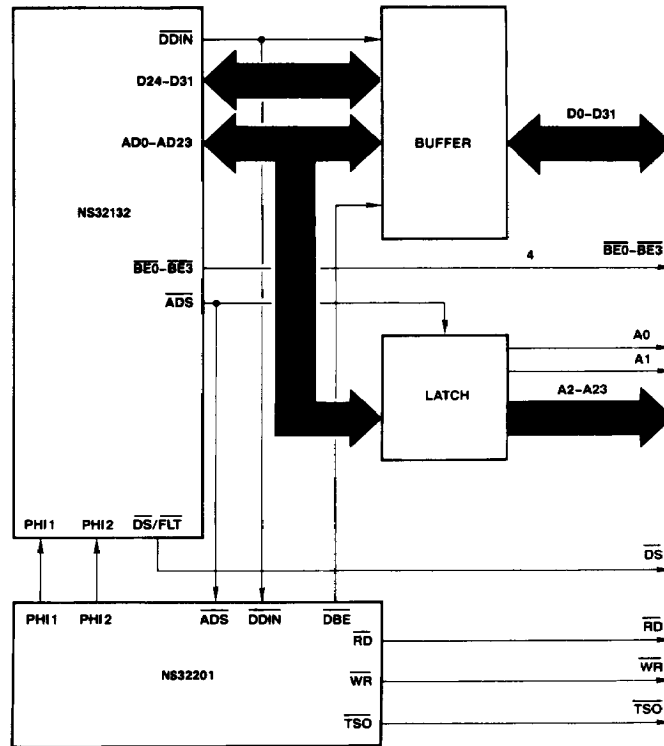


FIGURE 3-6. Bus Connections

TL/EE/8583-16

3.0 Functional Description (Continued)

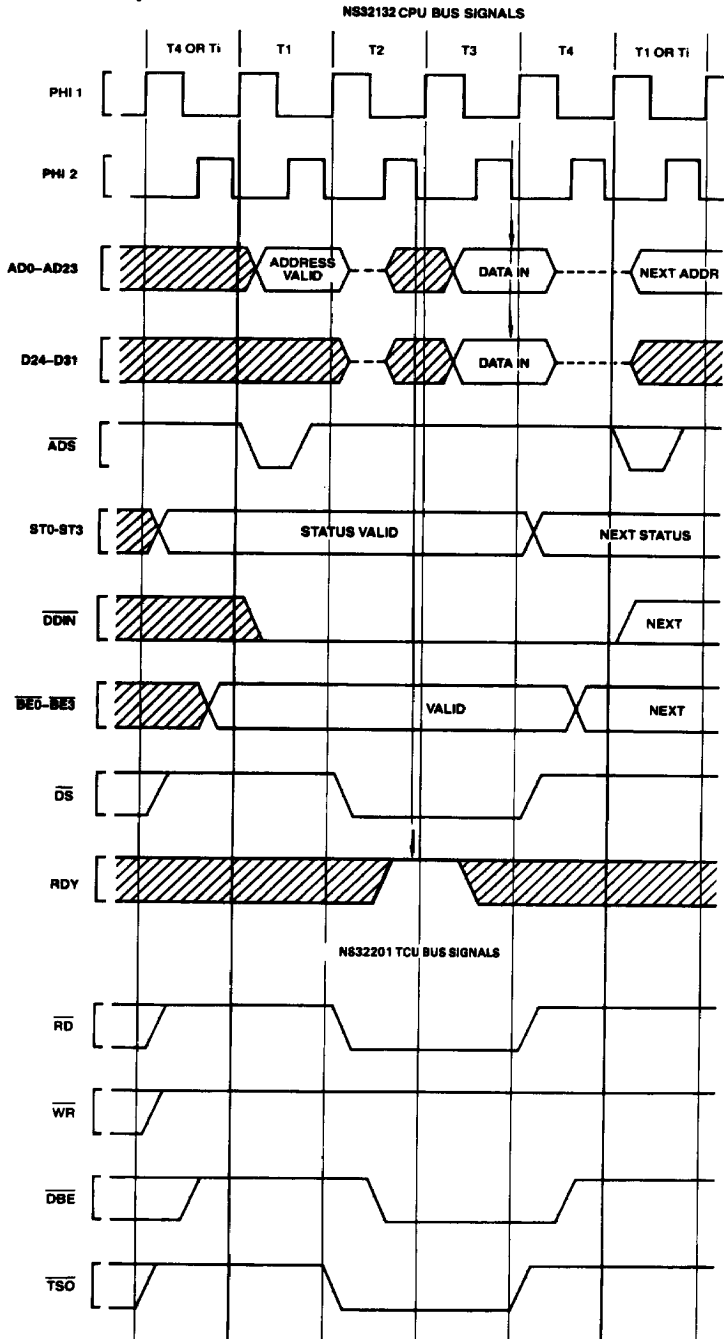


FIGURE 3-7. Read Cycle Timing

TL/EE/8583-17

3.0 Functional Description (Continued)

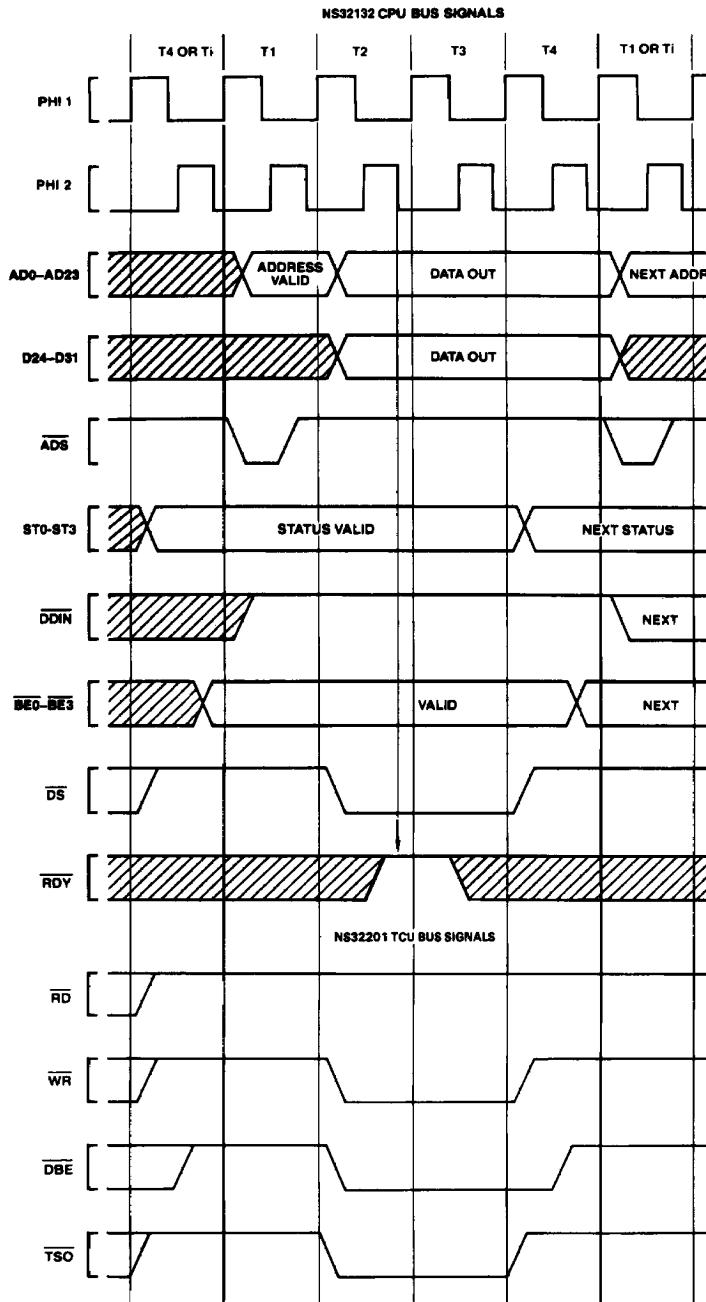


FIGURE 3-8. Write Cycle Timing

TL/EE/8583-18

3.0 Functional Description (Continued)

3.4.1 Cycle Extension

To allow sufficient strobe widths and access times for any speed of memory or peripheral device, the NS32016 provides for extension of a bus cycle. Any type of bus cycle except a Slave Processor cycle can be extended.

In *Figures 3-7* and *3-8*, note that during T3 all bus control signals from the CPU and TCU are flat. Therefore, a bus cycle can be cleanly extended by causing the T3 state to be repeated. This is the purpose of the RDY (Ready) pin.

At the end of T2 on the falling edge of PHI2, the RDY line is sampled by the CPU. If RDY is high, the next T-states will be T3 and then T4, ending the bus cycle. If RDY is low, then another T3 state will be inserted after the next T-state and the RDY line will again be sampled on the falling edge of PHI2. Each additional T3 state after the first is referred to as a "WAIT STATE". See *Figure 3-9*.

The RDY pin is driven by the NS32201 Timing Control Unit, which applies WAIT States to the CPU as requested on three sets of pin:

- 1) $\overline{\text{WAIT}}$ (Continuous WAIT), which holds the CPU in WAIT states until removed.
- 2) $\overline{\text{WAIT1}}$, $\overline{\text{WAIT2}}$, $\overline{\text{WAIT4}}$, $\overline{\text{WAIT8}}$ (Collectively $\overline{\text{WAITn}}$), which may be given a four-bit binary value requesting a specific number of WAIT States from 0 to 15.
- 3) $\overline{\text{PER}}$ (Peripheral), which inserts five additional WAIT states and causes the TCU to reshape the $\overline{\text{RD}}$ and $\overline{\text{WR}}$ strobes. This provides the setup and hold times required by most MOS peripheral interface devices.

Combinations of these various WAIT requests are both legal and useful. For details of their use, see the NS32201 Data Sheet.

Figure 3-10 illustrates a typical Read cycle, with two WAIT states requested through the TCU $\overline{\text{WAITn}}$ pins.




FIGURE 3-9. RDY Pin Timing

TL/EE/8583-19

3.4.2 Bus Status

The NS32132 CPU presents four bits of Bus Status information on pins ST0-ST3. The various combinations on these pins indicate why the CPU is performing a bus cycle, or, if it is idle on the bus, then why it is idle.

Referring to *Figures 3-7* and *3-8*, note that Bus Status leads the corresponding Bus Cycle, going valid one clock cycle before T1, and changing to the next state at T4. This allows the system designer to fully decode the Bus Status and, if desired, latch the decoded signals before $\overline{\text{ADS}}$ initiates the Bus Cycle.

The Bus Status pins are interpreted as a four-bit value, with ST0 the least significant bit. Their values decode as follows:

- 0000 - The bus is idle because the CPU does not need to perform a bus access.
- 0001 - The bus is idle because the CPU is executing the WAIT instruction.
- 0010 - (Reserved for future use.)
- 0011 - The bus is idle because the CPU is waiting for a Slave Processor to complete an instruction.
- 0100 - Interrupt Acknowledge, Master.

The CPU is performing a Read cycle. To acknowledge receipt of a Non-Maskable Interrupt (on NMI) it will read from address FFFF00₁₆, but will ignore any data provided.

To acknowledge receipt of a Maskable Interrupt (on INT) it will read from address FFFE00₁₆, expecting a vector number to be provided from the Master NS32202 Interrupt Control Unit. If the vectoring mode selected by the last SETCFG instruction was Non-Vectored, then the CPU will ignore the value it has read and will use a default vector instead, having assumed that no NS32202 is present. See Sec. 3.4.5.

- 0101 - Interrupt Acknowledge, Cascaded.

The CPU is reading a vector number from a Cascaded NS32202 Interrupt Control Unit. The address provided is the address of the NS32202 Hardware Vector register. See Sec. 3.4.5.

- 0110 - End of Interrupt, Master.

The CPU is performing a Read cycle to indicate that it is executing a Return from Interrupt (RETI) instruction. See Sec. 3.4.5.

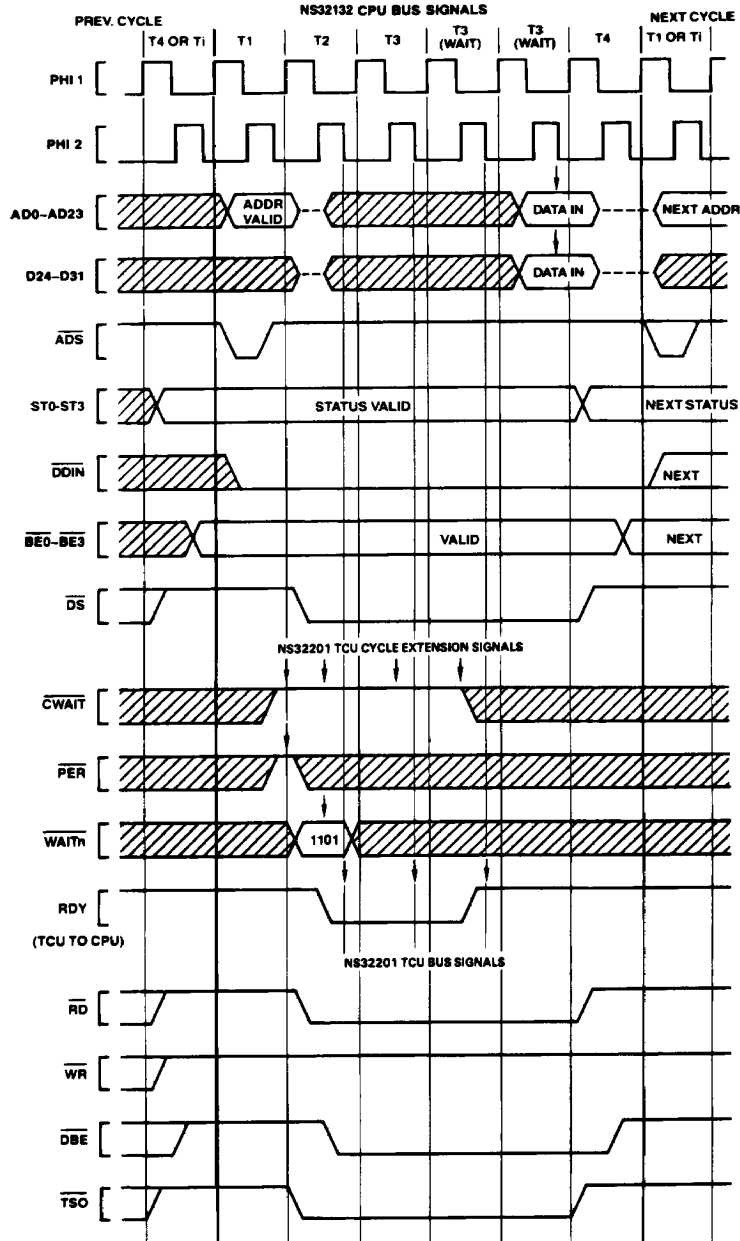
- 0111 - End of Interrupt, Cascaded.

The CPU is reading from a Cascaded Interrupt Control Unit to indicate that it is returning (through RETI) from an interrupt service routine requested by that unit. See Sec. 3.4.5.

- 1000 - Sequential Instruction Fetch.

The CPU is reading the next sequential word from the instruction stream into the instruction

3.0 Functional Description (Continued)



TL/EE/8583-20

FIGURE 3-10. Extended Cycle Example

Note: Arrows on CWAIT, PER, WAITr indicate points at which the TCU samples. Arrows on AD0-AD15 and RDY indicate points at which the CPU samples.

3.0 Functional Description (Continued)

Queue. It will do so whenever the bus would otherwise be idle and the queue is not already full.

1001 – Non-Sequential Instruction Fetch.

The CPU is performing the first fetch of instruction code after the Instruction Queue is purged. This will occur as a result of any jump or branch, or any interrupt or trap, or execution of certain instructions.

1010 – Data Transfer.

The CPU is reading or writing an operand of an instruction.

1011 – Read RMW Operand.

The CPU is reading an operand which will subsequently be modified and rewritten. If memory protection circuitry would not allow the following Write cycle, it must abort this cycle.

1100 – Read for Effective Address Calculation.

The CPU is reading information from memory in order to determine the Effective Address of an operand. This will occur whenever an instruction uses the Memory Relative or External addressing mode.

1101 – Transfer Slave Processor Operand.

The CPU is either transferring an instruction operand to or from a Slave Processor, or it is issuing the Operation Word of a Slave Processor instruction. See Sec. 3.9.1.

1110 – Read Slave Processor Status.

The CPU is reading a Status Word from a Slave Processor. This occurs after the Slave Processor has signalled completion of an instruction. The transferred word tells the CPU whether a trap should be taken, and in some instructions it presents new values for the CPU Processor Status Register bits N, Z, L or F. See Sec. 3.9.1.

1111 – Broadcast Slave ID.


The CPU is initiating the execution of a Slave Processor instruction. The ID Byte (first byte of the instruction) is sent to all Slave Processors, one of which will recognize it. From this point the CPU is communicating with only one Slave Processor. See Sec. 3.9.1.

3.4.3 Data Access Sequences

The 24-bit address provided by the NS32132 is a byte address; that is, it uniquely identifies one of up to 16,777,216 eight-bit memory locations. An important feature of the NS32132 is that the presence of a 32-bit data bus imposes no restrictions on data alignment; any data item, regardless of size, may be placed starting at any memory address. The NS32132 provides special control signals. Byte Enable ($\overline{BE0}$ – $\overline{BE3}$) which facilitate individual byte accessing on a 32-bit bus.

Memory is organized as four eight-bit banks, each bank receiving the double-word address ($A2$ – $A23$) in parallel. One bank, connected to Data Bus pins $AD0$ – $AD7$ is enabled

when $\overline{BE0}$ is low. The second bank, connected to data bus pins $AD8$ – $AD15$ is enabled when $\overline{BE1}$ is low. The third and fourth banks are enabled by $\overline{BE2}$ and $\overline{BE3}$, respectively. See Figure 3-11.



TL/EE/8583-21

FIGURE 3-11. Memory interface

Since operands do not need to be aligned with respect to the double-word bus accessed performed by the CPU, a given double-word access can contain one, two, three, or four bytes of the operand being addressed, and these bytes can begin at various positions, as determined by $A1$, $A0$. Table 3-1 lists the 10 resulting access types.

TABLE 3-1

Bus Access Types

Type	Bytes Accessed	$A1, A0$	$\overline{BE3}$	$\overline{BE2}$	$\overline{BE1}$	$\overline{BE0}$
1	1	00	1	1	1	0
2	1	01	1	1	0	1
3	1	10	1	0	1	1
4	1	11	0	1	1	1
5	2	00	1	1	0	0
6	2	01	1	0	0	1
7	2	10	0	0	1	1
8	3	00	1	0	0	0
9	3	01	0	0	0	1
10	4	00	0	0	0	0

Accesses of operands requiring more than one bus cycle are performed sequentially, with no idle T-States separating them. The number of bus cycles required to transfer an operand depends on its size and its alignment. Table 3-2 lists the bus cycles performed for each situation.

3.0 Functional Description (Continued)

TABLE 3-2
Access Sequences

Cycle	Type	Address	BE3	BE2	BE1	BE0	Data Bus			
							Byte 3	Byte 2	Byte 1	Byte 0
A. Word at address ending with 11							<div style="border: 1px solid black; display: inline-block; padding: 2px;"> BYTE 1 BYTE 0 </div> ← A			
1.	4	A	0	1	1	1	Byte 0	X	X	X
2.	1	A + 1	1	1	1	0	X	X	X	Byte 1
B. Double word at address ending with 01							<div style="border: 1px solid black; display: inline-block; padding: 2px;"> BYTE 3 BYTE 2 BYTE 1 BYTE 0 </div> ← A			
1.	9	A	0	0	0	1	Byte 2	Byte 1	Byte 0	X
2.	1	A + 3	1	1	1	0	X	X	X	Byte 3
C. Double word at address ending with 10							<div style="border: 1px solid black; display: inline-block; padding: 2px;"> BYTE 3 BYTE 2 BYTE 1 BYTE 0 </div> ← A			
1.	7	A	0	0	1	1	Byte 1	Byte 0	X	X
2.	5	A + 2	1	1	0	0	X	X	Byte 3	Byte 2
D. Double word at address ending with 11							<div style="border: 1px solid black; display: inline-block; padding: 2px;"> BYTE 3 BYTE 2 BYTE 1 BYTE 0 </div> ← A			
1.	4	A	0	1	1	1	Byte 0	X	X	X
2.	8	A + 1	1	0	0	0	X	Byte 3	Byte 2	Byte 1
E. Quad word at address ending with 00							<div style="border: 1px solid black; display: inline-block; padding: 2px;"> BYTE 7 BYTE 6 BYTE 5 BYTE 4 BYTE 3 BYTE 2 BYTE 1 BYTE 0 </div> ← A			
1.	10	A	0	0	0	0	Byte 3	Byte 2	Byte 1	Byte 0
Other bus cycles (instruction prefetch or slave) can occur here.										
2.	10	A + 4	0	0	0	0	Byte 7	Byte 6	Byte 5	Byte 4
F. Quad word at address ending with 01							<div style="border: 1px solid black; display: inline-block; padding: 2px;"> BYTE 7 BYTE 6 BYTE 5 BYTE 4 BYTE 3 BYTE 2 BYTE 1 BYTE 0 </div> ← A			
1.	9	A	0	0	0	1	Byte 2	Byte 1	Byte 0	X
2.	1	A + 3	1	1	1	0	X	X	X	Byte 3
Other bus cycles (instruction prefetch or slave) can occur here.										
3.	9	A + 4	0	0	0	1	Byte 6	Byte 5	Byte 4	X
4.	1	A + 7	1	1	1	0	X	X	X	Byte 7
G. Quad word at address ending with 10							<div style="border: 1px solid black; display: inline-block; padding: 2px;"> BYTE 7 BYTE 6 BYTE 5 BYTE 4 BYTE 3 BYTE 2 BYTE 1 BYTE 0 </div> ← A			
1.	7	A	0	0	1	1	Byte 1	Byte 0	X	X
2.	5	A + 2	1	1	0	0	X	X	Byte 3	Byte 2
Other bus cycles (instruction prefetch or slave) can occur here.										
3.	7	A + 4	0	0	1	1	Byte 5	Byte 4	X	X
4.	5	A + 6	1	1	0	0	X	X	Byte 7	Byte 6
H. Quad word at address ending with 11							<div style="border: 1px solid black; display: inline-block; padding: 2px;"> BYTE 7 BYTE 6 BYTE 5 BYTE 4 BYTE 3 BYTE 2 BYTE 1 BYTE 0 </div> ← A			
1.	4	A	0	1	1	1	Byte 0	X	X	X
2.	8	A + 1	1	0	0	0	X	Byte 3	Byte 2	Byte 1
Other bus cycles (instruction prefetch or slave) can occur here.										
1.	4	A + 4	0	1	1	1	Byte 4	X	X	X
2.	8	A + 5	1	0	0	0	X	Byte 7	Byte 6	Byte 5

X = Don't Care

3.0 Functional Description (Continued)

3.4.3.1 Bit Accesses

The Bit Instructions perform byte accesses to the byte containing the designated bit. The Test and Set Bit instruction (SBIT), for example, reads a byte, alters it, and rewrites it, having changed the contents of one bit.

3.4.3.2 Bit Field Accesses

An access to a Bit Field in memory always generates a Double-Word transfer at the address containing the least significant bit of the field. The Double Word is read by an Extract instruction; an Insert instruction reads a Double Word, modifies it, and rewrites it.

3.4.3.3 Extending Multiply Accesses

The Extending Multiply Instruction (MEI) will return a result which is twice the size in bytes of the operand it reads. If the multiplicand is in memory, the most-significant half of the result is written first (at the higher address), then the least-significant half. This is done in order to support retry if this instruction is aborted.

3.4.4 Instruction Fetches

Instructions for the NS32132 CPU are "prefetched"; that is, they are input before being needed into the next available entry of the eight-byte Instruction Queue. The CPU performs two types of Instruction Fetch cycles: Sequential and Non-Sequential. These can be distinguished from each other by their differing status combinations on pins ST0-ST3 (Sec. 3.4.2).

A Sequential Fetch will be performed by the CPU whenever the Data Bus would otherwise be idle and the Instruction Queue is not currently full. Sequential Fetches are always type 10 Read cycles (Table 3-1).

A Non-Sequential Fetch occurs as a result of any break in the normally sequential flow of a program. Any jump or branch instruction, a trap or an interrupt will cause the next Instruction Fetch cycle to be Non-Sequential. In addition, certain instructions flush the instruction queue, causing the next instruction fetch to display Non-Sequential status. Only the first bus cycle after a break displays Non-Sequential status, and that cycle depends on the destination address.

Note: During non-sequential fetches BE0-BE3 are all active regardless of the alignment.

3.4.5 Interrupt Control Cycles

Activating the $\overline{\text{INT}}$ or $\overline{\text{NMI}}$ pin on the CPU will initiate one or more bus cycles whose purpose is interrupt control rather than the transfer of instructions or data. Execution of the Return from Interrupt instruction (RETI) will also cause Interrupt Control bus cycles. These differ from instruction or data transfers only in the status presented on pins ST0-ST3. All Interrupt Control cycles are single-byte Read cycles.

This section describes only the Interrupt Control sequences associated with each interrupt and with the return from its service routine. For full details of the NS32132 interrupt structure, see Sec. 3.8.

3.0 Functional Description (Continued)

TABLE 3-3
Interrupt Sequences

Cycle	Status	Address	DDIN	BE3	BE2	BE1	BE0	Data Bus			
								Byte 3	Byte 2	Byte 1	Byte 0
<i>A. Non-Maskable Interrupt Control Sequences</i>											
Interrupt Acknowledge											
1	0100	FFFF00 ₁₆	0	1	1	1	0	X	X	X	X
Interrupt Return											
None: Performed through Return from Trap (RETT) instruction.											
<i>B. Non-Vectored Interrupt Control Sequences</i>											
Interrupt Acknowledge											
1	0100	FFFE00 ₁₆	0	1	1	1	0	X	X	X	X
Interrupt Return											
1	0110	FFFE00 ₁₆	0	1	1	1	0	X	X	X	X
<i>C. Vectored Interrupt Sequences: Non-Cascaded.</i>											
Interrupt Acknowledge											
1	0100	FFFE00 ₁₆	0	1	1	1	0	X	X	X	Vector: Range: 0-127
Interrupt Return											
1	0110	FFFE00 ₁₆	0	1	1	1	0	X	X	X	Vector: Same as in Previous Int. Ack. Cycle
<i>D. Vectored Interrupt Sequences: Cascaded</i>											
Interrupt Acknowledge											
1	0100	FFFE00 ₁₆	0	1	1	1	0	X	X	X	Cascade Index: range -16 to -1
(The CPU here uses the Cascade Index to find the Cascade Address.)											
2	0101	Cascade Address	0			See Note		Vector, range 9-255; on appropriate byte of data bus.			
Interrupt Return											
1	0110	FFFE00 ₁₆	0	1	1	1	0	X	X	X	Cascade Index: Same as in previous Int. Ack. Cycle
(The CPU here uses the Cascade Index to find the Cascade Address)											
2	0111	Cascade Address	0			See Note		X	X	X	X


X = Don't Care

Note: BE0-BE3 signals will be activated according to the cascaded ICU address. The cycle type can be 1, 2, 3 or 4, when reading the interrupt vector. The vector value can be in the range 0-255.

3.0 Functional Description (Continued)

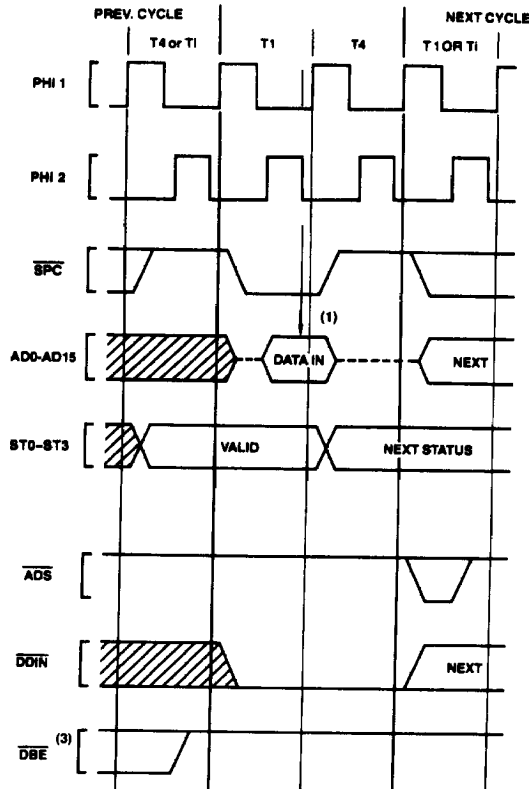
3.4.6 Slave Processor Communication

In addition to its use as the Address Translation strap (Sec. 3.5.1), the $\overline{AT}/\overline{SPC}$ pin is used as the data strobe for Slave Processor transfers. In this role, it is referred to as Slave Processor Control (\overline{SPC}). In a Slave Processor bus cycle, data is transferred on the Data Bus (AD0-AD15), and the status lines ST0-ST3 are monitored by each Slave Processor in order to determine the type of transfer being performed. \overline{SPC} is bidirectional, but is driven by the CPU during all Slave Processor bus cycles. See Sec. 3.9 for full protocol sequences.



TL/EE/8583-22

FIGURE 3-12. Slave Processor Connections



TL/EE/8583-23

Note:

- (1) CPU samples Data Bus here.
- (2) \overline{DBE} and all other NS32201 TCU bus signals remain inactive because no \overline{ADS} pulse is received from the CPU.

FIGURE 3-13. CPU Read from Slave Processor

3.0 Functional Description (Continued)

3.4.6.1 Slave Processor Bus Cycles

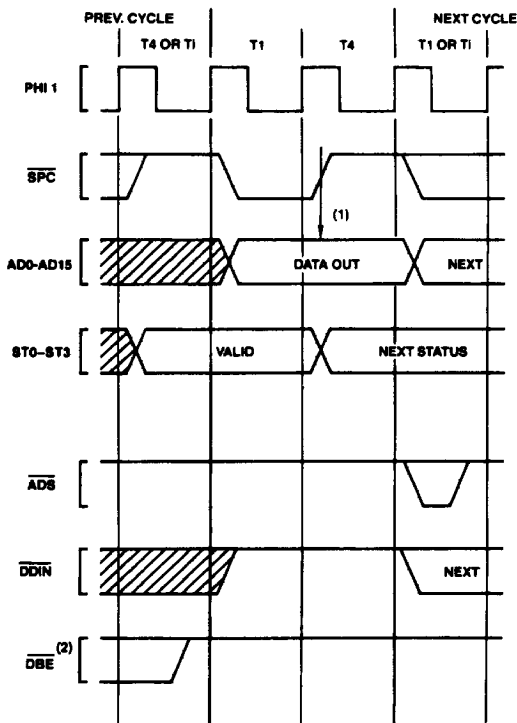
A Slave Processor bus cycle always takes exactly two clock cycles, labeled T1 and T4 (see *Figures 3-13 and 3-14*). During a Read cycle, \overline{SPC} is activated at T1, data is sampled at T4, and \overline{SPC} is removed. The Cycle Status pins lead the cycle by one clock period, and are sampled at the leading edge of \overline{SPC} . During a Write cycle, the CPU applies data and activates \overline{SPC} at T1, removing \overline{SPC} at T4. The Slave Processor latches status on the leading edge of \overline{SPC} and latches data on the trailing edge.

Since the CPU does not pulse the Address Strobe (\overline{ADS}), no bus signals are generated by the NS32201 Timing Control Unit. The direction of a transfer is determined by the sequence ("protocol") established by the instruction under execution; but the CPU indicates the direction on the \overline{DDIN} pin for hardware debugging purposes.

3.4.6.2 Slave Operand Transfer Sequences

A Slave Processor operand is transferred in one or more Slave bus cycles. A Byte operand is transferred on the least-significant byte of the Data Bus (AD0-AD7), and a Word operand is transferred on bits AD0-AD15. A Double Word is transferred in a consecutive pair of bus cycles, least-significant word first. A Quad Word is transferred in two pairs of Slave cycles, with other bus cycles possibly occurring between them. The word order is from least-significant word to most-significant.

Note that the NS32132 uses only the two least significant bytes of the data bus for slave cycles. This is to maintain compatibility with existing slave processors.



TL/EE/8583-24

Note:

(1) Slave Processor samples Data Bus here.

(2) \overline{DBE} , being provided by the NS32201 TCU, remains inactive due to the fact that no pulse is presented on \overline{ADS} . TCU signals \overline{RD} , \overline{WR} and \overline{TSO} also remain inactive.

FIGURE 3-14. CPU Write to Slave Processor

3.0 Functional Description (Continued)

3.5 MEMORY MANAGEMENT OPTION

The NS32132 CPU, in conjunction with the NS32082 Memory Management Unit (MMU), provides full support for address translation, memory protection, and memory allocation techniques up to and including Virtual Memory.

3.5.1 Address Translation Strap

The Bus Interface Control section of the NS32132 CPU has two bus timing modes: With or Without Address Translation. The mode of operation is selected by the CPU by sampling the $\overline{AT}/\overline{SPC}$ (Address Translation/Slave Processor Control) pin on the rising edge of the \overline{RST} (Reset) pulse. If $\overline{AT}/\overline{SPC}$

is sampled as high, the bus timing is as previously described in Sec. 3.4. If it is sampled as low, two changes occur:

- 1) An extra clock cycle, T_{mmu} , is inserted into all bus cycles except Slave Processor transfers.
- 2) The $\overline{DS}/\overline{FLT}$ pin changes in function from a Data Strobe output (\overline{DS}) to a Float Command input (\overline{FLT}).

The NS32082 MMU will itself pull the CPU $\overline{AT}/\overline{SPC}$ pin low when it is reset. In non-Memory-Managed systems this pin should be pulled up to V_{CC} through a 10 k Ω resistor.

Note that the Address Translation strap does not specifically declare the presence of an NS32082 MMU, but only the

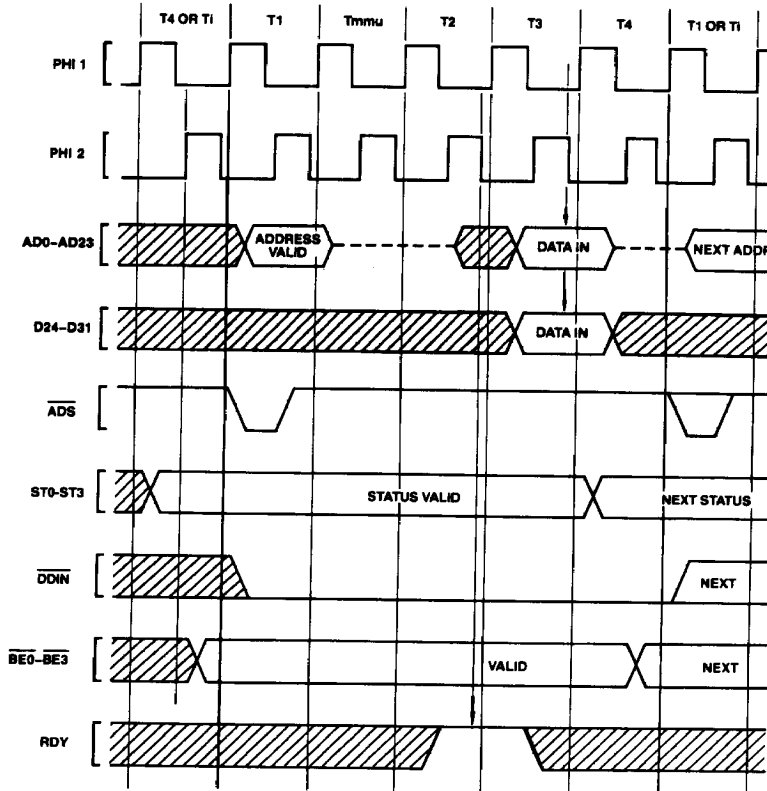


FIGURE 3-15. Read Cycle with Address Translation (CPU Action)

TL/EE/8583-25

3.0 Functional Description (Continued)

presence of external address translation circuitry. MMU instructions will still trap as being undefined unless the SETCFG (Set Configuration) instruction is executed to declare the MMU instruction set valid. See Sec. 2.1.3.

3.5.2 Translated Bus Timing

Figures 3-15 and 3-16 illustrate the CPU activity during a Read cycle and a Write cycle in Address Translation mode. The additional T-State, Tmmu, is inserted between T1 and T2. During this time the CPU places AD0-AD23 into the TRI-STATE® mode, allowing the MMU to assert the translated address and issue the physical address strobe PAV. T2 through T4 of the cycle are identical to their counterparts without Address Translation. Note that in order for the

NS32082 MMU to operate correctly it must be set to the 32132 mode by forcing A24/HBF low during reset. In this mode the bus lines AD16-AD23 are floated after the MMU address has been latched, since they are used by the CPU to transfer data.

Figures 3-17 and 3-18 show a Read cycle and a Write cycle as generated by the 32132/32082/32201 group. Note that with the CPU \overline{ADS} signal going only to the MMU, and with the MMU PAV signal substituting for \overline{ADS} everywhere else, Tmmu through T4 look exactly like T1 through T4 in a non-Memory-Managed system. For the connection diagram, see Appendix B.

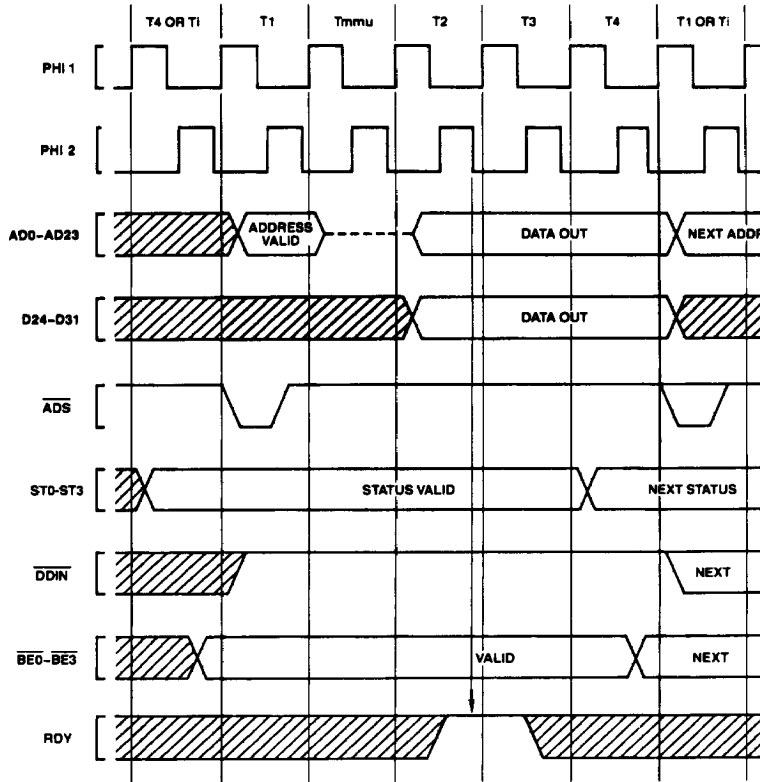


FIGURE 3-16. Write Cycle with Address Translation (CPU Action)

TL/EE/8583-26

3.0 Functional Description (Continued)

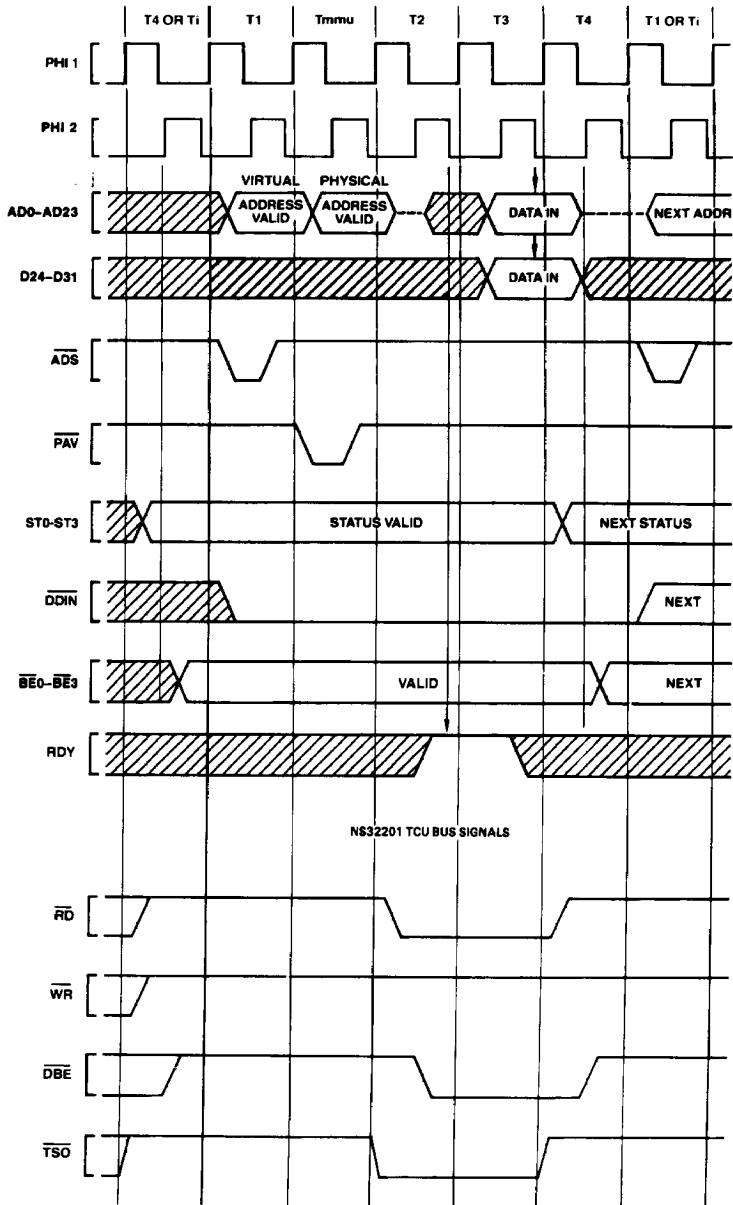


FIGURE 3-17. Memory-Managed Read Cycle

TL/EE/8583-27

3.0 Functional Description (Continued)




FIGURE 3-18. Memory-Managed Write Cycle

TL/EE/8583-28

3.0 Functional Description (Continued)

3.5.3 The FLT (Float) Pin

The FLT pin is used by the CPU for address translation support. Activating FLT during T_{mmu} causes the CPU to wait longer than T_{mmu} for address translation and validation. This feature is used occasionally by the NS32082 MMU in order to update its internal translation look-aside buffer (TLB) from page tables in memory, or to update certain status bits within them.

Figure 3-19 shows the effect of FLT. Upon sampling FLT low, late in T_{mmu}, the CPU enters idle T-States (T_f) during which it:

- 1) Sets AD0-AD23, D24-D31 and DDIN to the TRI-STATE condition ("floating").
- 2) Suspends further internal processing of the current instruction. This ensures that the current instruction remains abortable with retry. (See RST/ABT description, Sec. 3.5.4.)

Note that the AD0-AD23 pins may be briefly asserted during the first idle T-State. The above conditions remain in effect until FLT again goes high. See the Timing Specifications, Sec. 4.

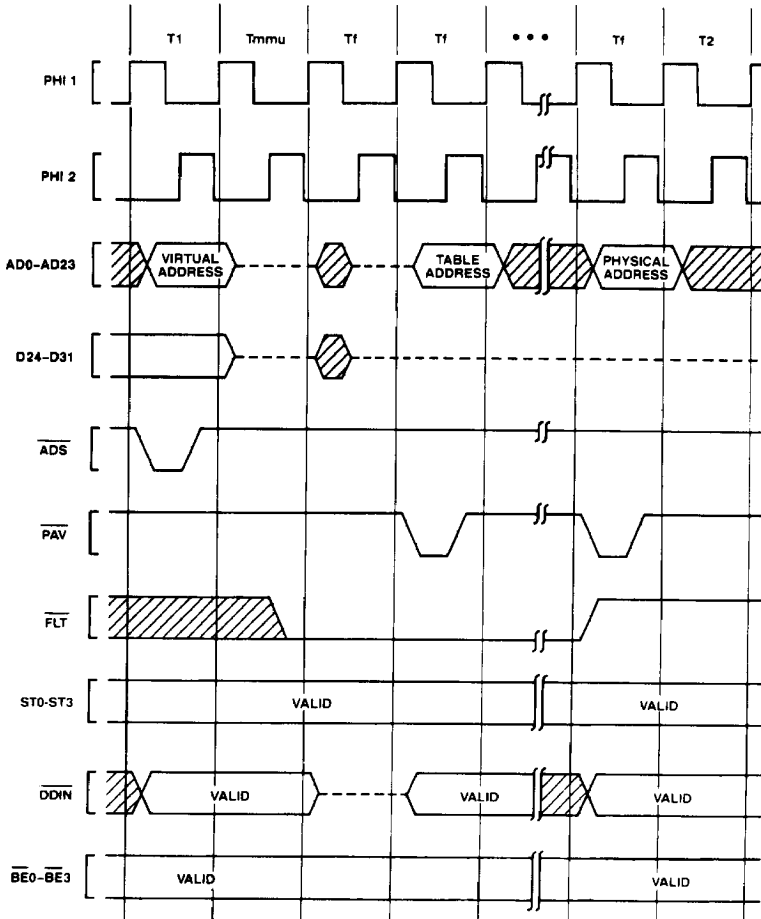


FIGURE 3-19. FLT Timing

TL/EE/8583-29

3.0 Functional Description (Continued)

3.5.4 Aborting Bus Cycles

The $\overline{RST}/\overline{ABT}$ pin, apart from its Reset function (Sec. 3.3), also serves as the means to "abort", or cancel, a bus cycle and the instruction, if any, which initiated it. An Abort request is distinguished from a Reset in that the $\overline{RST}/\overline{ABT}$ pin is held active for only one clock cycle.

If $\overline{RST}/\overline{ABT}$ is pulled low during Tmmu or Tf, this signals that the cycle must be aborted. The CPU itself will enter T2 and then T1, thereby terminating the cycle. Since it is the MMU PAV signal which triggers a physical cycle, the rest of the system remains unaware that a cycle was started.

The NS32082 MMU will abort a bus cycle for either of two reasons:

- 1) The CPU is attempting to access a virtual address which is not currently resident in physical memory. The referenced page must be brought into physical memory from mass storage to make it accessible to the CPU.
- 2) The CPU is attempting to perform an access which is not allowed by the protection level assigned to that page.

When a bus cycle is aborted by the MMU, the instruction that caused it to occur is also aborted in such a manner that it is guaranteed re-executable later. The information that is changed irrecoverably by such a partly-executed instruction does not affect its re-execution.

3.5.4.1 The Abort Interrupt

Upon aborting an instruction, the CPU immediately performs an interrupt through the ABT vector in the Interrupt Table (see Sec. 3.8). The Return Address pushed on the Interrupt Stack is the address of the aborted instruction, so that a Return from Trap (RETT) instruction will automatically retry it.

The one exception to this sequence occurs if the aborted bus cycle was an instruction prefetch. If so, it is not yet certain that the aborted prefetched code is to be executed. Instead of causing an interrupt, the CPU only aborts the bus cycle, and stops prefetching. If the information in the Instruction Queue runs out, meaning that the instruction will actually be executed, the ABT interrupt will occur, in effect aborting the instruction that was being fetched.

3.5.4.2 Hardware Considerations

In order to guarantee instruction retry, certain rules must be followed in applying an Abort to the CPU. These rules are followed by the NS32082 Memory Management Unit.

- 1) If \overline{FLT} has not been applied to the CPU, the Abort pulse must occur during or before Tmmu. See the Timing Specifications, *Figure 4-22*.

- 2) If \overline{FLT} has been applied to the CPU, the Abort pulse must be applied before the T-State in which \overline{FLT} goes inactive. The CPU will not actually respond to the Abort command until \overline{FLT} is removed. See *Figure 4-23*.

- 3) The Write half of a Read-Modify-Write operand access may not be aborted. The CPU guarantees that this will never be necessary for Memory Management functions by applying a special RMW status (Status Code 1011) during the Read half of the access. When the CPU presents RMW status, that cycle must be aborted if it would be illegal to write to any of the accessed addresses.

If $\overline{RST}/\overline{ABT}$ is pulsed at any time other than as indicated above, it will abort either the instruction currently under execution or the next instruction and will act as a very high-priority interrupt. However, the program that was running at the time is not guaranteed recoverable.

3.6 BUS ACCESS CONTROL

The NS32132 CPU has the capability of relinquishing its access to the bus upon request from a DMA device or another CPU. This capability is implemented on the \overline{HOLD} (Hold Request) and $\overline{HLD\overline{A}}$ (Hold Acknowledge) pins. By asserting \overline{HOLD} low, an external device requests access to the bus. On receipt of $\overline{HLD\overline{A}}$ from the CPU, the device may perform bus cycles, as the CPU at this point has set the $\overline{AD0}-\overline{AD23}$, $\overline{D24}-\overline{D31}$, \overline{ADS} , \overline{DDIN} and $\overline{BE0}-\overline{BE3}$ pins to the TRI-STATE condition. To return control of the bus to the CPU, the device sets \overline{HOLD} inactive, and the CPU acknowledges return of the bus by setting $\overline{HLD\overline{A}}$ inactive.

How quickly the CPU releases the bus depends on whether it is idle on the bus at the time the \overline{HOLD} request is made, as the CPU must always complete the current bus cycle. *Figure 3-20* shows the timing sequence when the CPU is idle. In this case, the CPU grants the bus during the immediately following clock cycle. *Figure 3-21* shows the sequence if the CPU is using the bus at the time that the \overline{HOLD} request is made. If the request is made during or before the clock cycle shown (two clock cycles before T4), the CPU will release the bus during the clock cycle following T4. If the request occurs closer to T4, the CPU may already have decided to initiate another bus cycle. In that case it will not grant the bus until after the next T4 state. Note that this situation will also occur if the CPU is idle on the bus but has initiated a bus cycle internally.

In a Memory-Managed system, the $\overline{HLD\overline{A}}$ signal is connected in a daisy-chain through the NS32082, so that the MMU can release the bus if it is using it.

3.0 Functional Description (Continued)

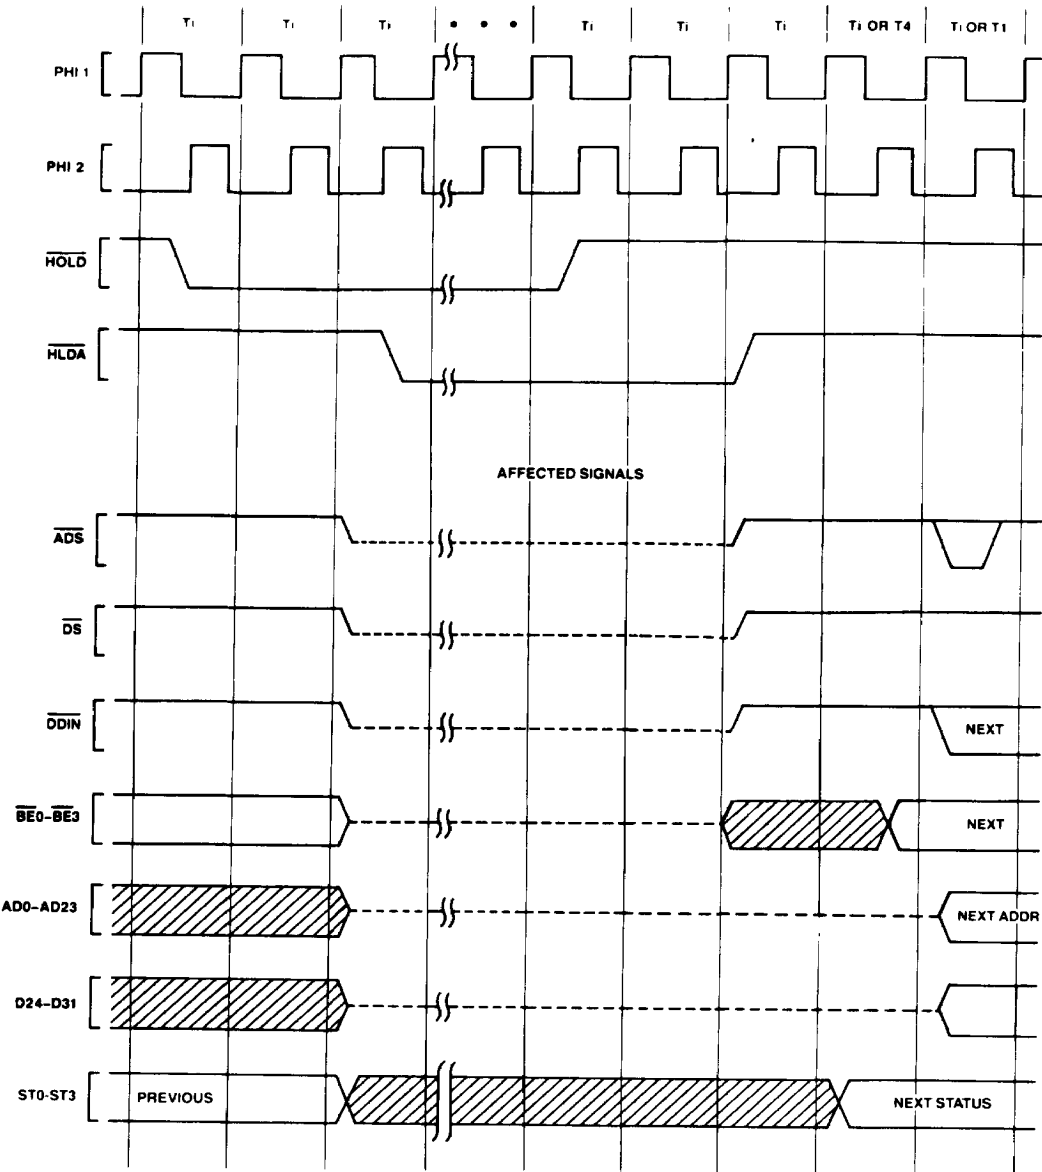
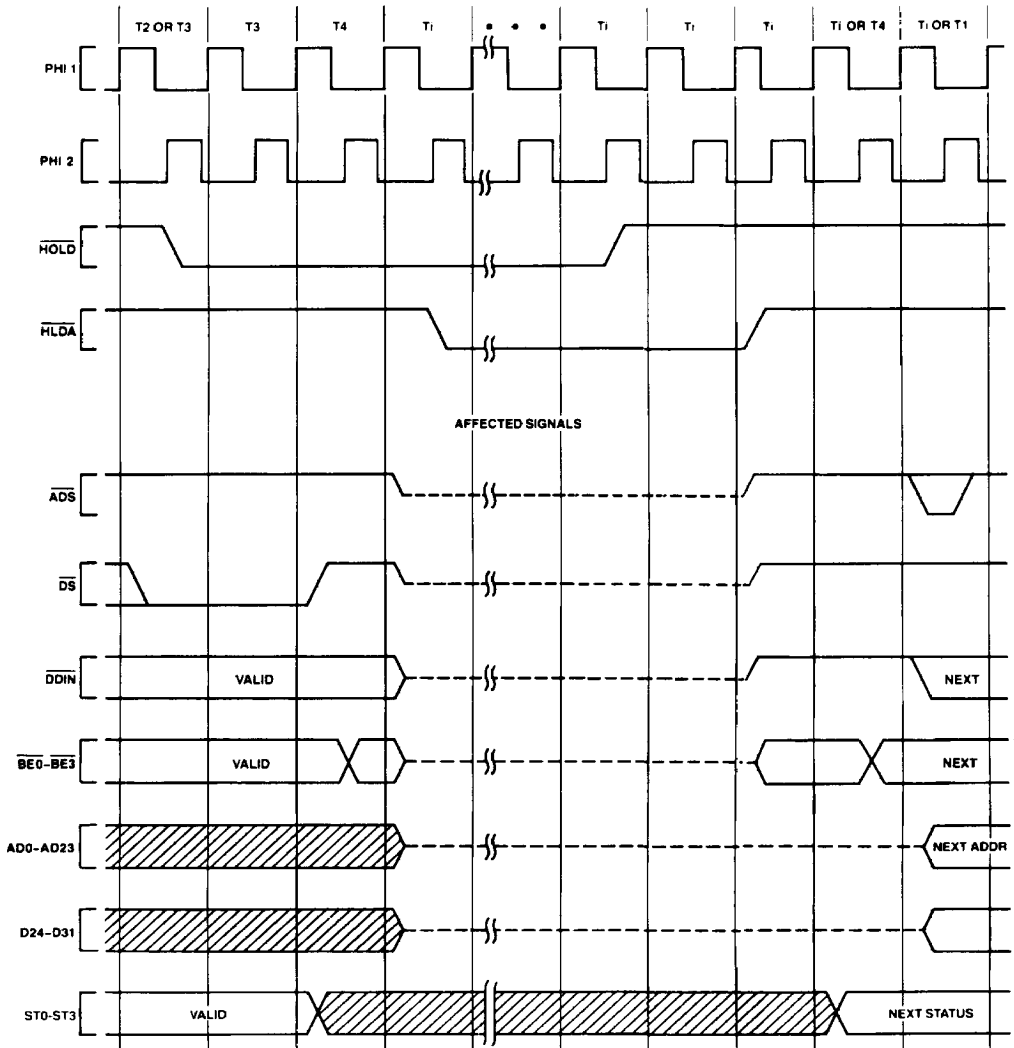


FIGURE 3-20. HOLD Timing, Bus Initially Idle

TL/EE/8583-30

3.0 Functional Description (Continued)



TL/EE/8583-31

FIGURE 3-21. HOLD Timing, Bus Initially Not Idle

3.0 Functional Description (Continued)

3.7 DUAL PROCESSING

An important feature of the NS32132 is its ability to cooperate with another NS32132, sharing the same bus and memory in a mode known as "Tightly-Coupled Dual Processing."

With the addition of one NS32132 and a minimal amount of standard TTL logic, this feature may enhance performance up to 80%. This is possible because the NS32132 average bus usage is in the range of 45% to 55%.

In a memory-managed system, an extra NS32201 TCU is required since a single TCU cannot drive two CPUs and two MMUs. In this case, a delay line may be needed to deskew the clock outputs from the two TCUs to avoid speed degradations.

Figure 3.22 shows a basic dual processing system configuration. A more complete connection diagram of a memory-managed system is given in Appendix B.

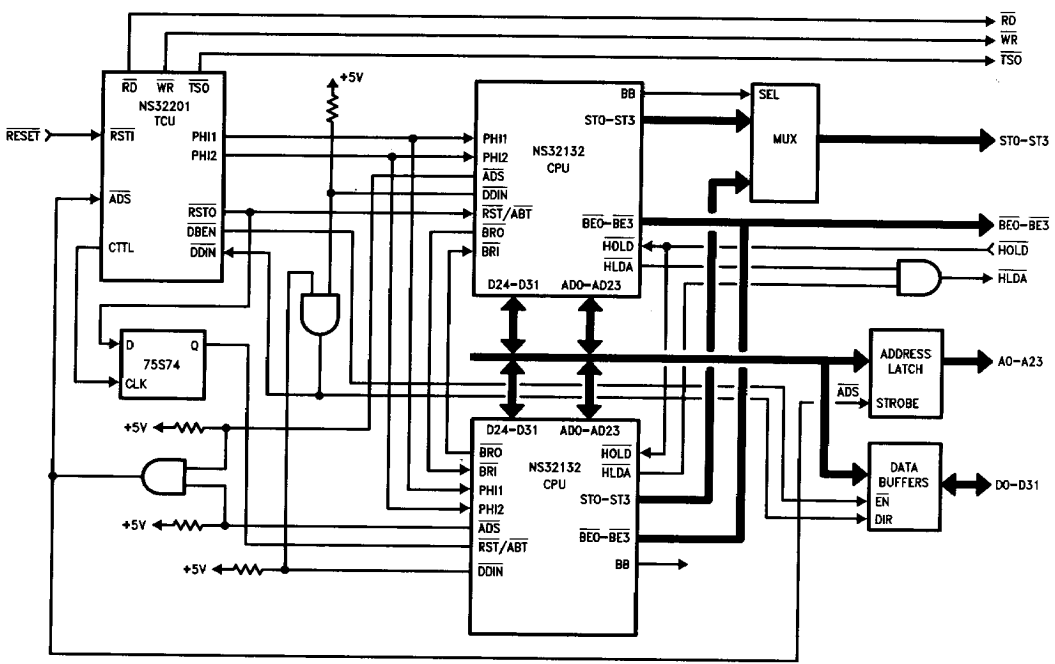


FIGURE 3-22. Dual Processor System Basic Interconnections

TL/EE/8583-32

3.0 Functional Description (Continued)

3.7.1 Bus Arbitration

The NS32132 provides two signals for the arbitration of the bus in a dual processing system:

- bus request out ($\overline{\text{BR}\overline{\text{O}}}$)
- bus request in ($\overline{\text{BR}\overline{\text{I}}}$)

A third signal (BB), activated only when the CPU is in control of the bus, is provided for testing purposes and for controlling external logic.

At any given time only one CPU is in control of the bus, while the other has its address/data lines and other relevant signals floating. Whenever the CPU not in control of the bus needs to perform a bus cycle, it asserts $\overline{\text{BR}\overline{\text{O}}}$ requesting the other CPU to release the bus. The CPU in control of the bus, detecting a low signal on its $\overline{\text{BR}\overline{\text{I}}}$ pin, will keep the bus for no more than two bus cycles and release it. It then notifies the other CPU by deactivating its $\overline{\text{BR}\overline{\text{O}}}$ signal, and bus control switches.

An exception to this rule occurs if the CPU in control of the bus is executing an interlocked instruction. In this case, the bus is released at the end of the instruction rather than at the end of the current bus cycle.

When a bus switch occurs, the address/data lines and the $\text{BE}\overline{0}$ – $\text{BE}\overline{3}$ lines from the CPU in control of the bus are floated a half clock cycle before the other CPU starts driving these lines. This scheme ensures that no contentions occur during the switching when the lines from the two CPUs are hardwired.

If a bus switch follows a read cycle, the T4 cycle of the CPU in control of the BUS is immediately followed by the T1 cycle of the other CPU. If the bus switch follows a write cycle, the CPU in control of the bus adds an extra clock cycle and releases the bus during this cycle. This ensures that the data hold times for single and dual processor configurations are the same. A bus switching timing diagram is shown in Figure 3-23.

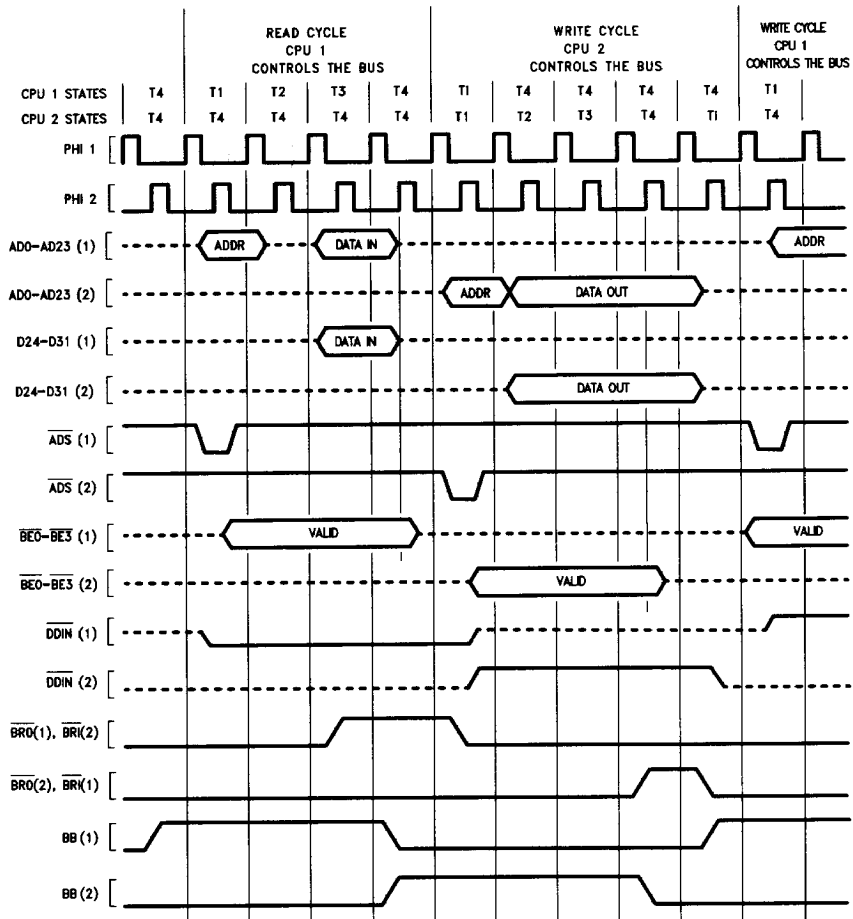


FIGURE 3-23. Dual Processor Bus Arbitration Timing

TL/EE/8589-33

3.0 Functional Description (Continued)

3.7.2 Processor Assignment

One problem with a multiprocessing system is ensuring that after a common initialization sequence, each processor branches to a different location in memory to execute a certain program. In cases where external hardware makes the system asymmetrical, the solution to this problem is simple. In a perfectly symmetrical system it is still possible to identify each processor uniquely so that tasks can be assigned to different processors. It is impossible, however, to know which processor will be assigned to a certain task.

The NS32132 relies on the fact that two tightly-coupled CPUs will switch bus control after at most two bus cycles, except when an interlocked instruction is executed. *Figure 3.24* shows an instruction sequence that can be used to assign processors to different tasks.

Both CPUs clear the same memory location on subsequent bus cycles. This ensures that both CPUs have cleared the memory location before any one of them executes the interlocked instruction. The CPU that exited the reset condition first executes the Set Bit Interlocked Instruction (SBITI), finding the bit clear. The other CPU will then execute the same instruction, but will find the bit set. Both CPUs will execute Branch on Flag Clear (BFC) instruction. The CPU that found the bit clear executes the branch; the other CPU continues to the next instruction. In this case the time constraint which delays one processor's exit from the reset condition makes it possible to know which processor will be assigned to a certain task.

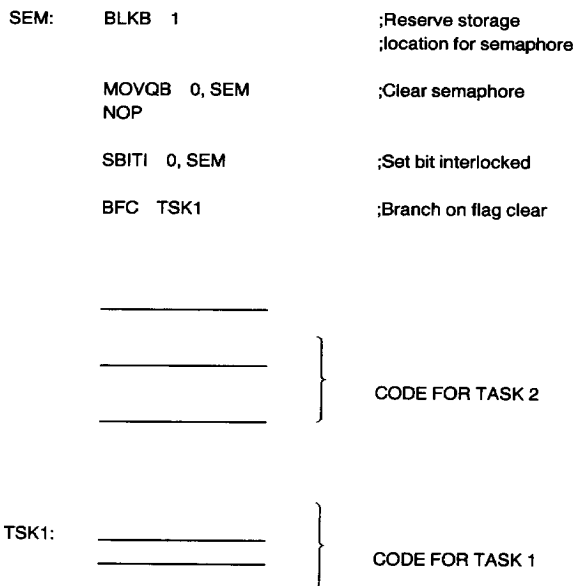


FIGURE 3-24. Instruction Sequence to Assign CPUs to Different Tasks

3.0 Functional Description (Continued)

3.8 INSTRUCTION STATUS

In addition to the four bits of Bus Cycle status (ST0-ST3), the NS32132 CPU also presents Instruction Status information on three separate pins. These pins differ from ST0-ST3 in that they are synchronous to the CPU's internal instruction execution section rather than to its bus interface section.

$\overline{\text{PFS}}$ (Program Flow Status) is pulsed low as each instruction begins execution. It is intended for debugging purposes, and is used that way by the NS32082 Memory Management Unit.

$\text{U}/\overline{\text{S}}$ originates from the U bit of the Processor Status Register, and indicates whether the CPU is currently running in User or Supervisor mode. It is sampled by the MMU for mapping, protection, and debugging purposes. Although it is not synchronous to bus cycles, there are guarantees on its validity during any given bus cycle. See the Timing Specifications, *Figure 4-21*.

$\overline{\text{ILO}}$ (Interlocked Operation) is activated during an SBITI (Set Bit, Interlocked) or CBITI (Clear Bit, Interlocked) instruction. It is made available to external bus arbitration circuitry in order to allow these instructions to implement the semaphore primitive operations for multi-processor communication and resource sharing. As with the $\text{U}/\overline{\text{S}}$ pin, there are guarantees on its validity during the operand accesses performed by the instructions. See the Timing Specification Section, *Figure 4-19*.

3.9 NS32132 INTERRUPT STRUCTURE

$\overline{\text{INT}}$, on which maskable interrupts may be requested,

$\overline{\text{NMI}}$, on which non-maskable interrupts may be requested, and

$\overline{\text{RST}}/\overline{\text{ABT}}$, which may be used to abort a bus cycle and any associated instruction. It generates an interrupt request if an instruction was aborted. See Sec. 3.5.4.

In addition there is a set of internally-generated "traps" which cause interrupt service to be performed as a result either of exceptional conditions (e.g., attempted division by zero) or of specific instructions whose purpose is to cause a trap to occur (e.g., the Supervisor Call instruction).

3.9.1 General Interrupt/Trap Sequence

Upon receipt of an interrupt or trap request, the CPU goes through four major steps:

1) Adjustment of Registers.

Depending on the source of the interrupt or trap, the CPU may restore and/or adjust the contents of the Program Counter (PC), the Processor Status Register (PSR) and the currently-selected Stack Pointer (SP). A copy of the PSR is made, and the PSR is then set to reflect Supervisor Mode and selection of the Interrupt Stack.

2) Saving Processor Status.

The PSR copy is pushed onto the Interrupt Stack as a 16-bit quantity.

3) Vector Acquisition.

A Vector is either obtained from the Data Bus or is supplied by default.

4) Service Call.

The Vector is used as an index into the Interrupt Dispatch Table, whose base address is taken from the CPU Interrupt Base (INTBASE) Register. See *Figure 3-25*. A 32-bit External Procedure Descriptor is read from the table entry, and an External Procedure Call is performed using it. The MOD Register (16 bits) and Program Counter (32 bits) are pushed on the Interrupt Stack.

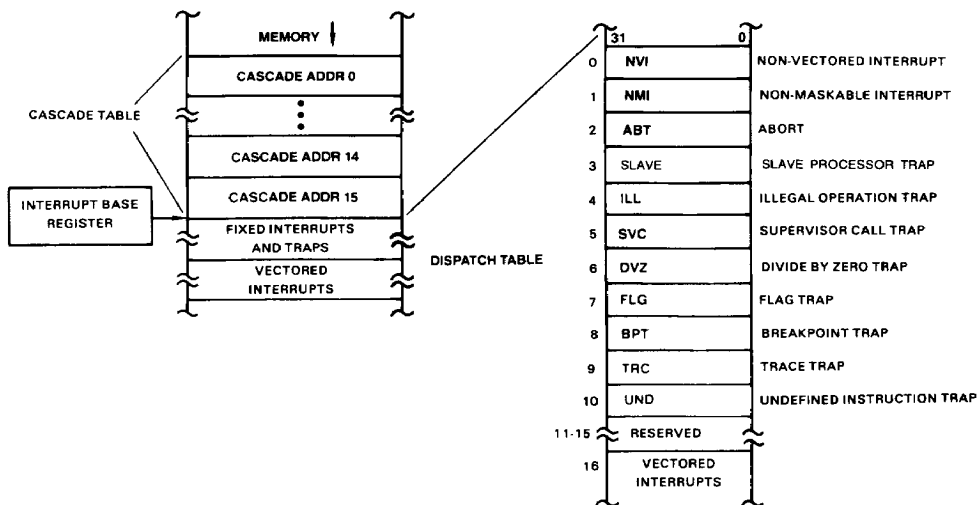


FIGURE 3-25. Interrupt Dispatch and Cascade Tables

TL/EE/B583-34

3.0 Functional Description (Continued)

This process is illustrated in *Figure 3-26*, from the viewpoint of the programmer.

Full sequences of events in processing interrupts and traps may be found as follows:

Interrupt on \overline{INT} or \overline{NMI} pin:

Abort Interrupt:

Traps (except Trace):


Trace Trap:

Sec. 3.9.7.1.

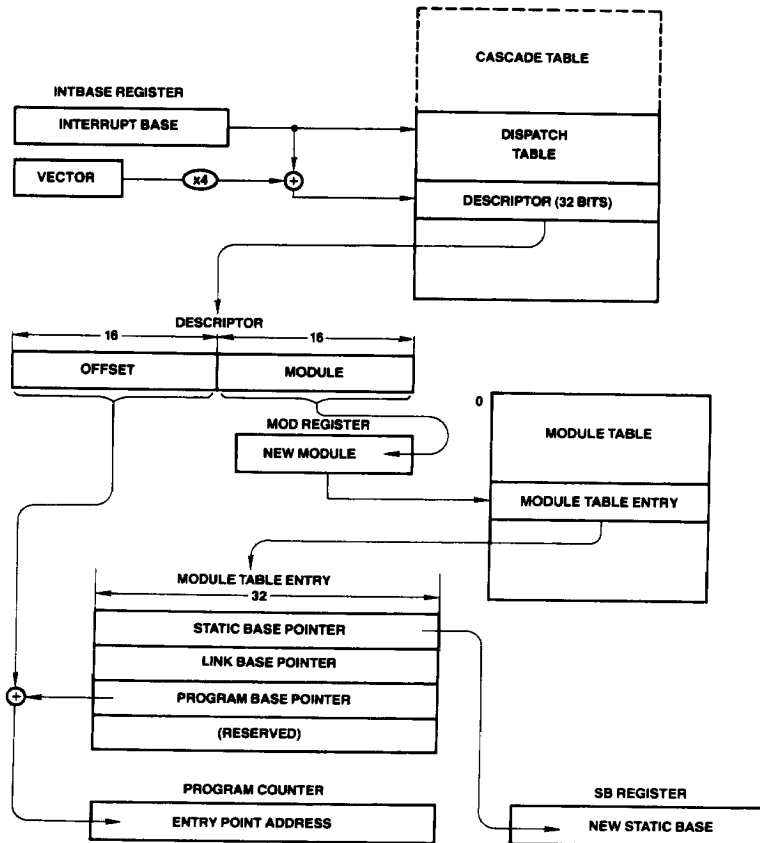
Sec. 3.9.7.4.

Sec. 3.9.7.2.

Sec. 3.9.7.3.



TL/EE/8583-35



TL/EE/8583-36

FIGURE 3-26. Interrupt/Trap Service Routine Calling Sequence

3.0 Functional Description (Continued)

3.9.2 Interrupt/Trap Return

To return control to an interrupted program, one of two instructions is used. The RETT (Return from Trap) instruction (Figure 3-27) restores the PSR, MOD, PC and SB registers to their previous contents and, since traps are often used deliberately as a call mechanism for Supervisor Mode procedures, it also discards a specified number of bytes from the original stack as surplus parameter space. RETT is used to return from any trap or interrupt except the Maskable Interrupt. For this, the RETI (Return from Interrupt) instruction is used, which also informs any external Interrupt Control Units that interrupt service has completed. Since interrupts are generally asynchronous external events, RETI does not pop parameters. See Figure 3-28.

3.9.3 Maskable Interrupts (The INT Pin)

The INT pin is a level-sensitive input. A continuous low level is allowed for generating multiple interrupt requests. The input is maskable, and is therefore enabled to generate interrupt requests only while the Processor Status Register I bit is set. The I bit is automatically cleared during service of an INT, NMI or Abort request, and is restored to its original

setting upon return from the interrupt service routine via the RETT or RETI instruction.

The INT pin may be configured via the SETCFG instruction as either Non-Vectored (CFG Register bit I = C) or Vectored (bit I = 1).

3.9.3.1 Non-Vectored Mode

In the Non-Vectored mode, an interrupt request on the INT pin will cause an Interrupt Acknowledge bus cycle, but the CPU will ignore any value read from the bus and use instead a default vector of zero. This mode is useful for small systems in which hardware interrupt prioritization is unnecessary.

3.9.3.2 Vectored Mode: Non-Cascaded Case

In the Vectored mode, the CPU uses an Interrupt Control Unit (ICU) to prioritize up to 16 interrupt requests. Upon receipt of an interrupt request on the INT pin, the CPU performs an "Interrupt Acknowledge, Master" bus cycle

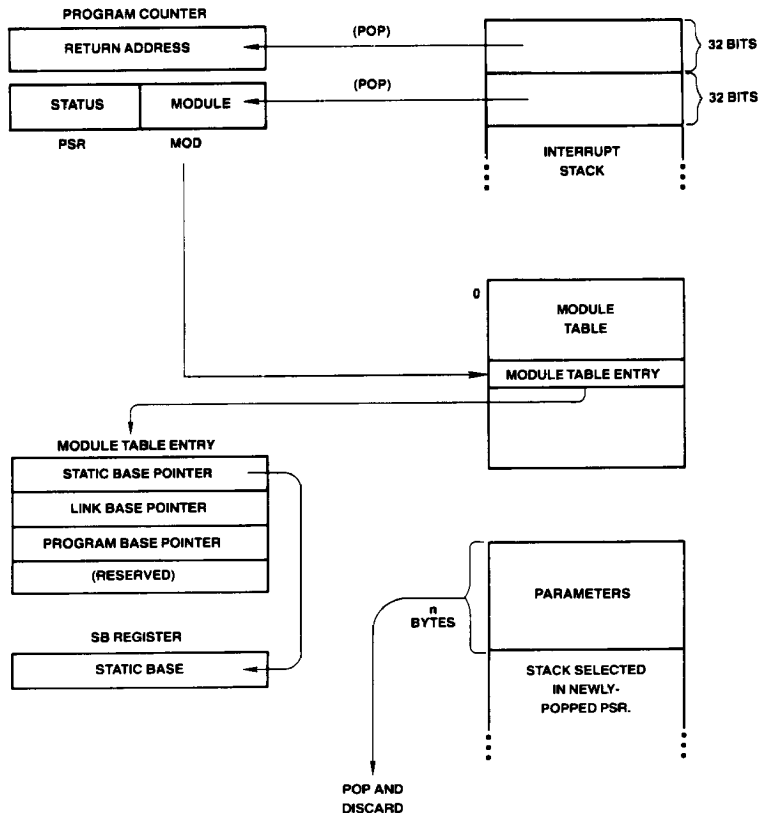


FIGURE 3-27. Return from Trap (RETT n) Instruction Flow

TL/EE/8583-37

3.0 Functional Description (Continued)

NS32132-6/NS32132-8/NS32132-10

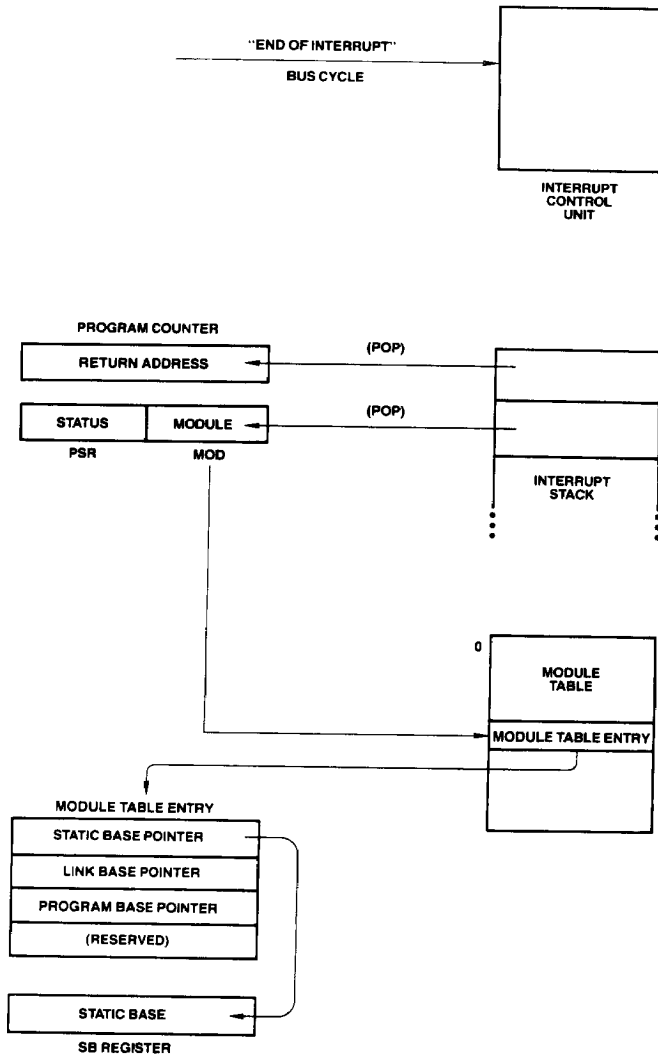


FIGURE 3-28. Return from Interrupt (RETI) Instruction Flow

TL/EE/8583-38

2

3.0 Functional Description (Continued)

(Sec. 3.4.2) reading a vector value from the low-order byte of the Data Bus. This vector is then used as an index into the Dispatch Table in order to find the External Procedure Descriptor for the proper interrupt service procedure. The service procedure eventually returns via the Return from Interrupt (RETI) instruction, which performs an End of Interrupt bus cycle, informing the ICU that it may re-prioritize any interrupt requests still pending. The ICU provides the vector number again, which the CPU uses to determine whether it needs also to inform a Cascaded ICU (see below).

In a system with only one ICU (16 levels of interrupt), the vectors provided must be in the range of 0 through 127; that is, they must be positive numbers in eight bits. By providing a negative vector number, an ICU flags the interrupt source as being a Cascaded ICU (see below).

3.9.3.3 Vectored Mode: Cascaded Case

In order to allow up to 256 levels of interrupt, provision is made both in the CPU and in the NS32202 Interrupt Control Unit (ICU) to transparently support cascading. *Figure 3-30*, shows a typical cascaded configuration. Note that the interrupt output from a Cascaded ICU goes to an Interrupt Request input of the Master ICU, which is the only ICU which drives the CPU INT pin.

In a system which uses cascading, two tasks must be performed upon initialization:

- 1) For each Cascaded ICU in the system, the Master ICU must be informed of the line number (0 to 15) on which it receives the cascaded requests.
- 2) A Cascade Table must be established in memory. The Cascade Table is located in a NEGATIVE direction from the location indicated by the CPU Interrupt Base (INTBASE) Register. Its entries are 32-bit addresses, pointing to the Vector Registers of each of up to 16 Cascaded ICUs.

Figure 3-25 illustrates the position of the Cascade Table. To find the Cascade Table entry for a Cascaded ICU, take its Master ICU line number (0 to 15) and subtract 16 from it,

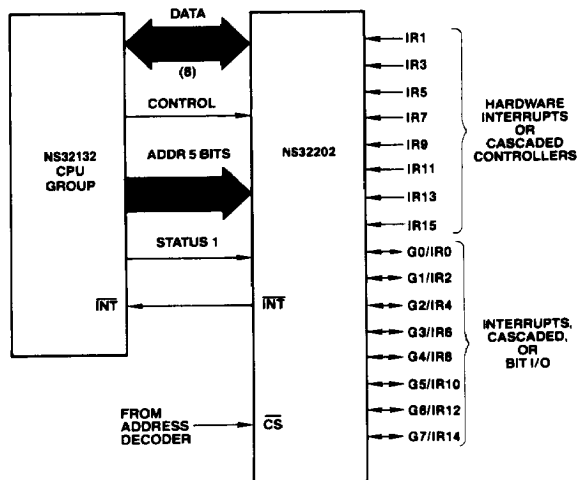
giving an index in the range -16 to -1 . Multiply this value by 4, and add the resulting negative number to the contents of the INTBASE Register. The 32-bit entry at this address must be set to the address of the Hardware Vector Register of the Cascaded ICU. This is referred to as the "Cascade Address."

Upon receipt of an interrupt request from a Cascaded ICU, the Master ICU interrupts the CPU and provides the negative Cascade Table index instead of a (positive) vector number. The CPU, seeing the negative value, uses it as an index into the Cascade Table and reads the Cascade Address from the referenced entry. Applying this address, the CPU performs an "Interrupt Acknowledge, Cascaded" bus cycle (Sec. 3.4.2), reading the final vector value. This vector is interpreted by the CPU as an unsigned byte, and can therefore be in the range of 0 through 255.

In returning from a Cascaded interrupt, the service procedure executes the Return from Interrupt (RETI) instruction, as it would for any Maskable Interrupt. The CPU performs an "End of Interrupt, Master" bus cycle (Sec. 3.4.2), whereupon the Master ICU again provides the negative Cascade Table index. The CPU, seeing a negative value, uses it to find the corresponding Cascade Address from the Cascade Table. Applying this address, it performs an "End of Interrupt, Cascaded" bus cycle (Sec. 3.4.2), informing the Cascaded ICU of the completion of the service routine. The byte read from the Cascaded ICU is discarded.

NOTE: If an interrupt must be masked off, the CPU can do so by setting the corresponding bit in the interrupt mask register of the interrupt controller.

However, if an interrupt is set pending during the CPU instruction that masks off that interrupt, the CPU may still perform an interrupt acknowledge cycle following that instruction since it might have sampled the INT line before the ICU deasserted it. This could cause the ICU to provide an invalid vector. To avoid this problem the above operation should be performed with the CPU interrupt disabled.



TL/EE/8583-39

FIGURE 3-29. Interrupt Control Unit Connections (16 Levels)

3.0 Functional Description (Continued)

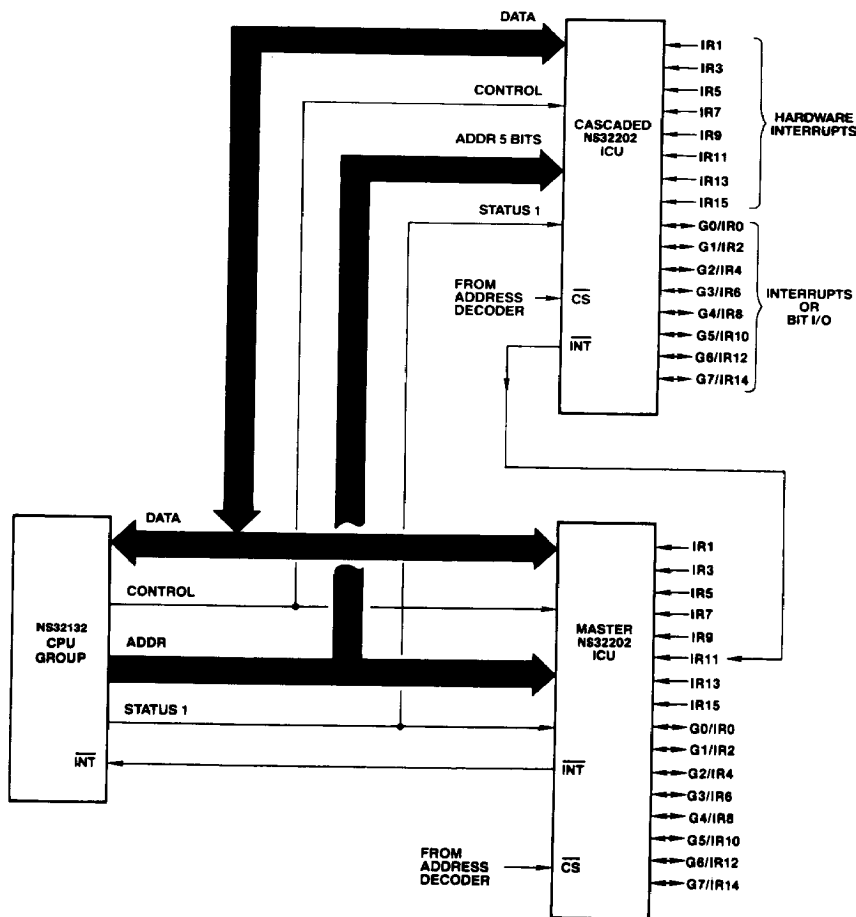


FIGURE 3-30. Cascaded Interrupt Control Unit Connections

TL/EE/8583-40

3.9.4 Non-Maskable Interrupt (The \overline{NMI} Pin)

The Non-Maskable Interrupt is triggered whenever a falling edge is detected on the \overline{NMI} pin. The CPU performs an "Interrupt Acknowledge, Master" bus cycle (Sec. 3.4.2) when processing of this interrupt actually begins. The Interrupt Acknowledge cycle differs from that provided for Maskable Interrupts in that the address presented is $FFFF00_{16}$. The vector value used for the Non-Maskable Interrupt is taken as 1, regardless of the value read from the bus.

The service procedure returns from the Non-Maskable Interrupt using the Return from Trap (RETT) instruction. No special bus cycles occur on return.

For the full sequence of events in processing the Non-Maskable Interrupt, see Sec. 3.9.7.1.

3.9.5 Traps

A trap is an internally-generated interrupt request caused as a direct and immediate result of the execution of an instruction. The Return Address pushed by any trap except Trap (TRC) is the address of the first byte of the instruction during which the trap occurred. Traps do not disable interrupts, as they are not associated with external events. Traps recognized by the NS32132 CPU are:

Trap (SLAVE): An exceptional condition was detected by the Floating Point Unit or another Slave Processor during the execution of a Slave Instruction. This trap is requested via the Status Word returned as part of the Slave Processor Protocol (Sec. 3.10.1).

3.0 Functional Description (Continued)

Trap (ILL): Illegal operation. A privileged operation was attempted while the CPU was in User Mode (PSR bit U = 1).

Trap (SVC): The Supervisor Call (SVC) instruction was executed.

Trap (DVZ): An attempt was made to divide an integer by zero. (The FPU trap is used for Floating Point division by zero.)

Trap (FLG): The FLAG instruction detected a "1" in the CPU PSR F bit.

Trap (BPT): The Breakpoint (BPT) instruction was executed.

Trap (TRC): The instruction just completed is being traced. See below.

Trap (UND): An undefined opcode was encountered by the CPU.

A special case is the Trace Trap (TRC), which is enabled by setting the T bit in the Processor Status Register (PSR). At the beginning of each instruction, the T bit is copied into the PSR P (Trace "Pending") bit. If the P bit is set at the end of an instruction, then the Trace Trap is activated. If any other trap or interrupt request is made during a traced instruction, its entire service procedure is allowed to complete before the Trace Trap occurs. Each interrupt and trap sequence handles the P bit for proper tracing, guaranteeing one and only one Trace Trap per instruction, and guaranteeing that the Return Address pushed during a Trace Trap is always the address of the next instruction to be traced.

3.9.6 Prioritization

The NS32016 CPU internally prioritizes simultaneous interrupt and trap requests as follows:

- 1) Traps other than Trace (Highest priority)
- 2) Abort
- 3) Non-Maskable Interrupt
- 4) Maskable Interrupts
- 5) Trace Trap (Lowest priority)

3.9.7 Interrupt/Trap Sequences: Detailed Flow

For purposes of the following detailed discussion of interrupt and trap service sequences, a single sequence called "Service" is defined in *Figure 3-31*. Upon detecting any interrupt request or trap condition, the CPU first performs a sequence dependent upon the type of interrupt or trap. This sequence will include pushing the Processor Status Register and establishing a Vector and a Return Address. The CPU then performs the Service sequence.

For the sequence followed in processing either Maskable or Non-Maskable interrupts (on the $\overline{\text{INT}}$ or $\overline{\text{NMI}}$ pins, respectively), see Sec. 3.9.7.1 For Abort Interrupts, see Sec. 3.9.7.4. For the Trace Trap, see Sec. 3.9.7.3, and for all other traps see Sec. 3.9.7.2.

3.9.7.1 Maskable/Non-Maskable Interrupt Sequence

This sequence is performed by the CPU when the $\overline{\text{NMI}}$ pin receives a falling edge, or the $\overline{\text{INT}}$ pin becomes active with the PSR I bit set. The interrupt sequence begins either at the next instruction boundary or, in the case of the String instructions, at the next interruptible point during its execution.

1. If a String instruction was interrupted and not yet completed:
 - a. Clear the Processor Status Register P bit.
 - b. Set "Return Address" to the address of the first byte of the interrupted instruction.
 Otherwise, set "Return Address" to the address of the next instruction.
2. Copy the Processor Status Register (PSR) into a temporary register, then clear PSR bits S, U, T, P and I.
3. If the interrupt is Non-Maskable:
 - a. Read a byte from address FFFF0_{16} , applying Status Code 0100 (Interrupt Acknowledge, Master, Sec. 3.4.2). Discard the byte read.
 - b. Set "Vector" to 1.
 - c. Go to Step 8.
4. If the interrupt is Non-Vectored:
 - a. Read a byte from address FFFF0_{16} , applying Status Code 0100 (Interrupt Acknowledge, Master: Sec. 3.4.2). Discard the byte read.
 - b. Set "Vector" to 0.
 - c. Go to Step 8.
5. Here the interrupt is Vectored. Read "Byte" from address FFFE0_{16} , applying Status Code 0100 (Interrupt Acknowledge, Master: Sec. 3.4.2).
6. If "Byte" ≥ 0 , then set "Vector" to "Byte" and go to Step 8.
7. If "Byte" is in the range -16 through -1 , then the interrupt source is Cascaded. (More negative values are reserved for future use.) Perform the following:
 - a. Read the 32-bit Cascade Address from memory. The address is calculated as $\text{INTBASE} + 4 * \text{Byte}$.
 - b. Read "Vector," applying the Cascade Address just read and Status Code 0101 (Interrupt Acknowledge, Cascaded: Sec. 3.4.2).
8. Push the PSR copy (from Step 2) onto the Interrupt Stack as a 16-bit value.
9. Perform Service (Vector, Return Address), *Figure 3-31*.

Service (Vector, Return Address):

- 1) Read the 32-bit External Procedure Descriptor from the Interrupt Dispatch Table: address is $\text{Vector} * 4 + \text{INTBASE}$ Register contents.
- 2) Move the Module field of the Descriptor into the MOD Register.
- 3) Read the new Static Base pointer from the memory address contained in MOD, placing it into the SB Register.
- 4) Read the Program Base pointer from memory address $\text{MOD} + 8$, and add to it the Offset field from the Descriptor, placing the result in the Program Counter.
- 5) Flush queue: Non-sequentially fetch first instruction of interrupt routine.
- 6) Push MOD Register into the Interrupt Stack as a 16-bit value. (The PSR has already been pushed as a 16-bit value.)
- 7) Push the Return Address onto the Interrupt Stack as a 32-bit quantity.

FIGURE 3-31. Service Sequence
Invoked during all interrupt/trap sequences.

3.0 Functional Description (Continued)

3.9.7.2 Trap Sequence: Traps Other Than Trace

- 1) Restore the currently selected Stack Pointer and the Processor Status Register to their original values at the start of the trapped instruction.
- 2) Set "Vector" to the value corresponding to the trap type.
 - SLAVE: Vector = 3.
 - ILL: Vector = 4.
 - SVC: Vector = 5.
 - DVZ: Vector = 6.
 - FLG: Vector = 7.
 - BPT: Vector = 8.
 - UND: Vector = 10.
- 3) Copy the Processor Status Register (PSR) into a temporary register, then clear PSR bits S, U, P and T.
- 4) Push the PSR copy onto the Interrupt Stack as a 16-bit value.
- 5) Set "Return Address" to the address of the first byte of the trapped instruction.
- 6) Perform Service (Vector, Return Address), *Figure 3-31*.

3.9.7.3 Trace Trap Sequence

- 1) In the Processor Status Register (PSR), clear the P bit.
- 2) Copy the PSR into a temporary register, then clear PSR bits S, U and T.
- 3) Push the PSR copy onto the Interrupt Stack as a 16-bit value.
- 4) Set "Vector" to 9.
- 5) Set "Return Address" to the address of the next instruction.
- 6) Perform Service (Vector, Return Address), *Figure 3-31*.

3.9.7.4 Abort Sequence

- 1) Restore the currently selected Stack Pointer to its original contents at the beginning of the aborted instruction.
- 2) Clear the PSR P bit.
- 3) Copy the PSR into a temporary register, then clear PSR bits S, U, T and I.
- 4) Push the PSR copy onto the Interrupt Stack as a 16-bit value.
- 5) Set "Vector" to 2.
- 6) Set "Return Address" to the address of the first byte of the aborted instruction.
- 7) Perform Service (Vector, Return Address), *Figure 3-31*.

3.10 SLAVE PROCESSOR INSTRUCTIONS

The NS32132 CPU recognizes three groups of instructions being executable by external Slave Processor:

- Floating Point Instruction Set
- Memory Management Instruction Set
- Custom Instruction Set

Each Slave Instruction Set is validated by a bit in the Configuration Register (Sec. 2.1.3). Any Slave Instruction which does not have its corresponding Configuration Register bit set will trap as undefined, without any Slave Processor communication attempted by the CPU. This allows software simulation of a non-existent Slave Processor.

3.10.1 Slave Processor Protocol

Slave Processor instructions have a three-byte Basic Instruction field, consisting of an ID Byte followed by an Operation Word. The ID Byte has three functions:

- 1) It identifies the instruction as being a Slave Processor instruction.
- 2) It specifies which Slave Processor will execute it.
- 3) It determines the format of the following Operation Word of the instruction.

Upon receiving a Slave Processor instruction, the CPU initiates the sequence outlined in *Figure 3-32*. While applying Status Code 1111 (Broadcast ID, Sec. 3.4.2), the CPU transfers the ID Byte on the least-significant byte of the Data Bus (AD0-AD7). All Slave Processors input this byte and decode it. The Slave Processor selected by the ID Byte is activated, and from this point the CPU is communicating only with it. If any other slave protocol was in progress (e.g., an aborted Slave instruction), this transfer cancels it.

The CPU next sends the Operation Word while applying Status Code 1101 (Transfer Slave Operand, Sec. 3.4.2). Upon receiving it, the Slave Processor decodes it, and at this point both the CPU and the Slave Processor are aware of the number of operands to be transferred and their sizes. The operation Word is swapped on the Data Bus, that is, bits 0-7 appear on pins AD8-AD15 and bits 8-15 appear on pins AD0-AD7.

Using the Address Mode fields within the Operation Word, the CPU starts fetching operand and issuing them to the Slave Processor. To do so, it references any Addressing Mode extensions which may be appended to the Slave Processor instruction. Since the CPU is solely responsible

Status Combinations:

Send ID (ID): Code 1111

Xfer Operand (OP): Code 1101

Read Status (ST): Code 1110

Step	Status	Action
1	ID	CPU Send ID Byte.
2	OP	CPU Sends Operaton Word.
3	OP	CPY Sends Required Operands
4	—	Slave Starts Execution. CPU Pre-fetches.
5	—	Slave Pulses SPC Low.
6	ST	CPU Reads Status Word. (Trap? Alter Flags?)
7	OP	CPU Reads Results (If Any).

FIGURE 3-32. Slave Processor Protocol

3.0 Functional Description (Continued)

for memory accesses, these extensions are not sent to the Slave processor. The Status Code applied is 1101 (Transfer Slave Processor Operand, Sec. 3.4.2).

After the CPU has issued the last operand, the Slave Processor starts the actual execution of the instruction. Upon completion, it will signal the CPU by pulsing \overline{SPC} low. To allow for this, and for the Address Translation strap function, $\overline{AT}/\overline{SFC}$ is normally held high only by an internal pull-up device of approximately 5 k Ω .

While the Slave Processor is executing the instruction, the CPU is free to prefetch instructions into its queue. If it fills the queue before the Slave Processor finishes, the CPU will wait, applying Status Code 0011 (Waiting for Slave, Sec. 3.4.2).

Upon receiving the pulse on \overline{SPC} , the CPU uses \overline{SPC} to read a Status Word from the Slave Processor, applying Status Code 1110 (Read Slave Status, Sec. 3.4.2). This word has the format shown in *Figure 3-33*. If the Q bit ("Quit", Bit 0) is set, this indicates that an error was detected by the Slave Processor. The CPU will not continue the protocol, but will immediately trap through the SLAVE vector in the Interrupt Table. Certain Slave Processor instructions cause CPU PSR bits to be loaded from the Status Word.

The last step in the protocol is for the CPU to read a result, if any, and transfer it to the destination. The Read cycles from the Slave Processor are performed by the CPU while applying Status Code 1101 (Transfer Slave Operand, Sec. 3.4.2).

An exception to the protocol above is the LMR (Load Memory Management Register) instruction, and a corresponding Custom Slave instruction (LCR: Load Custom Register). In executing these instructions, the protocol ends after the CPU has issued the last operand. The CPU does not wait for an acknowledgement from the Slave Processor, and it does not read status.

3.10.2 Floating Point Instructions

Table 3-4 gives the protocols followed for each Floating Point instruction. The instructions are referenced by their mnemonics. For the bit encodings of each instruction, see Appendix A.

The Operand class columns give the Access Class for each general operand, defining how the addressing modes are interpreted (see Instruction Set Reference Manual).

The Operand Issued columns show the sizes of the operands issued to the Floating Point Unit by the CPU. "D" indicates a 32-bit Double Word. "i" indicates that the instruction specifies an integer size for the operand (B = Byte, W = Word, D = Double Word). "f" indicates that the instruction specifies a Floating Point size for the operand (F = 32-bit Standard Floating, L = 64-bit Long Floating).

The Returned Value Type and Destination column gives the size of any returned value and where the CPU places it. The PSR Bits Affected column indicates which PSR bits, if any, are updated from the Slave Processor Status Word (*Figure 3-33*).

TABLE 3-4

Floating Point Instruction Protocols.

Mnemonic	Operand 1 Class	Operand 2 Class	Operand 1 Issued	Operand 2 Issued	Returned Value Type and Dest.	PSR Bits Affected
ADDf	read.f	rmw.f	f	f	f to Op. 2	none
SUBf	read.f	rmw.f	f	f	f to Op. 2	none
MULf	read.f	rmw.f	f	f	f to Op. 2	none
DIVf	read.f	rmw.f	f	f	f to Op. 2	none
MOVf	read.f	write.f	f	N/A	f to Op. 2	none
ABSf	read.f	write.f	f	N/A	f to Op. 2	none
NEGf	read.f	write.f	f	N/A	f to Op. 2	none
CMPf	read.f	read.f	f	f	N/A	N,Z,L
FLOORfi	read.f	write.i	f	N/A	i to Op. 2	none
TRUNCfi	read.f	write.i	f	N/A	i to Op. 2	none
ROUNDfi	read.f	write.i	f	N/A	i to Op. 2	none
MOVFL	read.F	write.L	F	N/A	L to Op. 2	none
MOVLf	read.L	write.F	L	N/A	F to Op. 2	none
MOVif	read.i	write.f	i	N/A	f to Op. 2	none
LFSR	read.D	N/A	D	N/A	N/A	none
SFSR	N/A	write.D	N/A	N/A	D to Op. 2	none

Note:

D = Double Word

i = Integer size (B,W,D) specified in mnemonic.

f = Floating Point type (F,L) specified in mnemonic.

N/A = Not Applicable to this instruction.

3.0 Functional Description (Continued)




FIGURE 3-33. Slave Processor Status Word Format

Any operand indicated as being of type "F" will not cause a transfer if the Register addressing mode is specified. This is because the Floating Point Registers are physically on the Floating Point Unit and are therefore available without CPU assistance.

3.10.3 Memory Management Instructions

Table 3-5 gives the protocols for Memory Management instructions. Encodings for these instructions may be found in Appendix A.

In executing the RDVAL and WRVAL instructions, the CPU calculates and issues the 32-bit Effective Address of the single operand. The CPU then performs a single-byte Read cycle from that address, allowing the MMU to safely abort the instruction if the necessary information is not currently in physical memory. Upon seeing the memory cycle complete, the MMU continues the protocol, and returns the validation result in the F bit of the Slave Status Word.

The size of a Memory Management operand is always a 32-bit Double Word. For further details of the Memory Management Instruction set, see the Instruction Set Reference Manual and the NS32082 MMU Data Sheet.

TABLE 3-5

Memory Management Instruction Protocols.

Mnemonic	Operand 1	Memory Management Instruction Protocols.		Returned Value		PSR Bits Affected
	Class	Operand 2 Class	Operand 1 Issued	Operand 2 Issued	Type and Dest.	
RDVAL*	addr	N/A	D	N/A	N/A	F
WRVAL*	addr	N/A	D	N/A	N/A	F
LMR*	read.D	N/A	D	N/A	N/A	none
SMR*	write.D	N/A	N/A	N/A	D to Op, 1	none

Note:

In the RDVAL and WRVAL instructions, the CPU issues the address as a Double Word, and performs a single-byte Read cycle from that memory address. For details, see the Instruction Set Reference Manual and the NS32082 Memory Management Unit Data Sheet.

D = Double Word

* = Privileged Instruction: will trap if CPU is in User Mode.

N/A = Not Applicable to this instruction.

3.0 Functional Description (Continued)

3.10.4 Custom Slave Instructions

Provided in the NS32132 is the capability of communicating with a user-defined, "Custom" Slave Processor. The instruction set provided for a Custom Slave Processor defines the instruction formats, the operand classes and the communication protocol. Left to the user are the interpretations of the Op Code fields, the programming model of the Custom Slave and the actual types of data transferred. The protocol specifies only the size of an operand, not its data type.

Table 3-6 lists the relevant information for the Custom Slave instruction set. The designation "c" is used to represent an operand which can be a 32-bit ("D") or 64-bit ("Q") quantity in any format; the size is determined by the suffix on the mnemonic. Similarly, an "i" indicates an integer size (Byte, Word, Double Word) selected by the corresponding mnemonic suffix.

Any operand indicated as being of type "c" will not cause a transfer if the register addressing mode is specified. It is assumed in this case that the slave processor is already holding the operand internally.

For the instruction encodings, see Appendix A.

TABLE 3-6
Custom Slave Instruction Protocols.

Mnemonic	Operand 1 Class	Operand 2 Class	Operand 1 Issued	Operand 2 Issued	Returned Value Type and Dest.	PSR Bits Affected
CCAL0c	read.c	rmw.c	c	c	c to Op. 2	none
CCAL1c	read.c	rmw.c	c	c	c to Op. 2	none
CCAL2c	read.c	rmw.c	c	c	c to Op. 2	none
CCAL3c	read.c	rmw.c	c	c	c to Op. 2	none
CMOV0c	read.c	write.c	c	N/A	c to Op. 2	none
CMOV1c	read.c	write.c	c	N/A	c to Op. 2	none
CMOV2c	read.c	write.c	c	N/A	c to Op. 2	none
CMOV3c	read.c	write.c	c	N/A	c to Op.2	none
CCMPc	read.c	read.c	c	c	N/A	N,Z,L
CCMP1c	read.c	read.c	c	c	N/A	N,Z,L
CCV0ci	read.c	write.i	c	N/A	i to Op. 2	none
CCV1ci	read.c	write.i	c	N/A	i to Op. 2	none
CCV2ci	read.c	write.i	c	N/A	i to Op. 2	none
CCV3ic	read.i	write.c	i	N/A	c to Op. 2	none
CCV4DQ	read.D	write.Q	D	N/A	Q to Op. 2	none
CCV5QD	read.Q	write.D	Q	N/A	D to Op. 2	none
LCSR	read.D	N/A	D	N/A	N/A	none
SCSR	N/A	write.D	N/A	N/A	D to OP. 2	none
CATST0*	addr	N/A	D	N/A	N/A	F
CATST1*	addr	N/A	D	N/A	N/A	F
LCR*	read.D	N/A	D	N/A	N/A	none
SCR*	write.D	N/A	N/A	N/A	D to Op.1	none

Note:

D = Double Word

i = Integer size (B,W,D) specified in mnemonic.

c = Custom size (D:32 bits or Q:64 bits) specified in mnemonic.

* = Privileged instruction: will trap if CPU is in User Mode.

N/A = Not Applicable to this instruction.

4.0 Device Specifications

4.1 NS32132 PIN DESCRIPTIONS

The following is a brief description of all NS32132 pins. The descriptions reference portions of the Functional Description. Sec. 3.

Unless otherwise indicated (see pin 34) reserved pins should be left open.

4.1.1 Supplies

Power (V_{CC}): +5V Positive Supply. Sec. 3.1.

Logic Ground (GN_{DL}): Ground reference for on-chip logic. Sec. 3.1.

Buffer Grounds (GN_{DB1}, GN_{DB2}, GN_{DB3}): Ground references for the on-chip output drivers.

Back-Bias Generator (BBG): Output of on-chip substrate voltage generator. Sec. 3.1.

4.1.2 Input Signals

Clocks (PHI1, PHI2): Two-phase clocking signals. Sec. 3.2.

Ready (RDY): Active high. While RDY is inactive, the CPU extends the current bus cycle to provide for a slower memory or peripheral reference. Upon detecting RDY active, the CPU terminates the bus cycle. Sec. 3.4.1.

Hold Request (HOLD): Active low. Causes the CPU to release the bus for DMA or multiprocessing purposes.

Note 1. HOLD must not be asserted until HLD_A from a previous HOLD/HLD_A sequence is deasserted.

Note 2. If the HOLD signal is generated asynchronously, its set up and hold times may be violated.

In this case it is recommended to synchronize it with CTTL to minimize the possibility of metastable states.

The CPU provides only one synchronization stage to minimize the HLD_A latency. This is to avoid speed degradations in cases of heavy HOLD activity (i.e. DMA controller cycles interleaved with CPU cycles.)

Interrupt (INT): Active low. Maskable Interrupt request. Sec. 3.9.

Non-Maskable Interrupt (NMI): Active low. Non-Maskable Interrupt request. Sec. 3.9.

Reset/Abort (RST/ABT): Active low. If held active for one clock cycle and released, this pin causes an instruction Abort, Sec. 3.5.4. If held longer, it initiates a Reset. Sec. 3.3.

Bus Request IN (BRI): Active Low. Used in Dual-Processing systems to signal one NS32132 that the other NS32132 in the system is requesting the bus. Sec. 3.7.

This pin should be connected to V_{CC} through a 4.7 kΩ resistor if dual processing is not used.

4.1.3 Output Signals

Address Strobe (ADS): Active low. Controls address latches; indicates start of a bus cycle. Sec. 3.4.

Data Direction in (DDIN): Active low. Status signal indicating direction of data transfer during a bus cycle. Sec. 3.4.

Byte Enable (BE₀–BE₃): Active low. Four control signals enabling data transfers on individual bus bytes. Sec. 3.4.3.

Status (ST₀–ST₃): Bus cycle status code, ST₀ least significant. Sec. 3.4.2. Encodings are:

0000 — Idle: CPU Inactive on Bus.
 0001 — Idle: WAIT Instruction.
 0010 — (Reserved).
 0011 — Idle: Waiting for Slave.
 0100 — Interrupt Acknowledge, Master.
 0101 — Interrupt Acknowledge, Cascaded.
 0110 — End of Interrupt, Master.
 0111 — End of Interrupt, Cascaded.
 1000 — Sequential Instruction Fetch.
 1001 — Non-Sequential Instruction Fetch.
 1010 — Data Transfer.
 1011 — Read Read-Modify-Write Operand.
 1100 — Read for Effective Address.
 1101 — Transfer Slave Operand.
 1110 — Read Slave Status Word.
 1111 — Broadcast Slave ID.

Hold Acknowledge (HLD_A): Active low. Applied by the CPU in response to HOLD input, indicating that the bus has been released for DMA or multiprocessing purposes. Sec. 3.6.

User/Supervisor (U/S): User or Supervisor Mode status. Sec. 3.7. High state indicates User Mode, low indicates Supervisor Mode. Sec. 3.8.

Interlocked Operation (ILO): Active low. Indicates that an interlocked instruction is being executed. Sec. 3.8.

Bus Request OUT (BRO): Active Low. Used in Dual-Processing Systems by the NS32132 to signal the other CPU in the system that the bus is needed. Sec. 3.7.

Bus Busy (BB): This signal is used in Dual-Processing Systems and is activated by the CPU in control of the bus. When BB is inactive the Address-data-bus, BE₀–BE₃ and DDIN are floated. Sec. 3.7.

Program Flow Status (PFS): Active low. Pulse indicates beginning of an instruction execution. Sec. 3.8.

4.1.4 Input/Output Signals

Address/Data 0–23 (AD₀–AD₂₃): Multiplexed Address/Data information. Bit 0 is the least significant bit of each. Sec. 3.4.

Data Bits 24–31 (D₂₄–D₃₁): The high order 8 bits of the data bus.

Address Translation/Slave Processor Control (AT/SPC): Active low. Used by the CPU as the data strobe output for Slave Processor transfers; used by Slave Processors to acknowledge completion of a SLAVE instruction. Sec. 3.4.6; Sec. 3.9. Sampled on the rising edge of Reset as Address Translation Strap. Sec. 3.5.1.

In Non-Memory-Managed Systems this PIN should be pulled-up to V_{CC} through a 10 kΩ resistor.

Data Strobe/Float (DS/FLT): Active low. Data Strobe output, Sec. 3.4, or Float Command input, Sec. 3.5.3. Pin function is selected on AT/SPC pin, Sec. 3.5.1.

4.0 Device Specifications (Continued)

4.4 SWITCHING CHARACTERISTICS

4.4.1 Definitions

All the timing specifications given in this section refer to 2.0V on rising or falling edges of the clock phases PHI1 and

PHI2, and 0.8V or 2.0V on all other signals as illustrated in Figures 4-2 and 4-3, unless specifically stated otherwise.

ABBREVIATIONS:

L.E. — leading edge R.E. — rising edge
T.E. — trailing edge F.E. — falling edge




FIGURE 4-2. Timing Specification Standard
(Signal Valid After Clock Edge)

TL/EE/8583-42




FIGURE 4-3. Timing Specification Standard
(Signal Valid Before Clock Edge)

TL/EE/8583-43

4.4.2 Timing Tables

4.4.2.1 Output Signals: Internal Propagation Delays NS32132-6, NS32132-8, NS32132-10

Maximum times assume capacitive loading of 100 pF.

Name	Figure	Description	Reference/Conditions	NS32132-6		NS32132-8		NS32132-10		Units
				Min	Max	Min	Max	Min	Max	
t _{ALv}	4-4	Address bits 0-23 valid	after R.E., PHI1 T1		80		65		50	ns
t _{ALh}	4-4	Address bits 0-23 hold	after R.E., PHI1 Tmmu or T2	5		5		5		ns
t _{Dv}	4-4	Data valid (write cycle)	after R.E., PHI1 T2		80		65		50	ns
t _{Dh}	4-4	Data hold (write cycle)	after R.E., PHI1 next T1 or Ti	0		0		0		ns
t _{ALADs}	4-4	Address bits 0-23 set up	before \overline{ADS} T. E.	25		25		25		ns
t _{ALADsh}	4-10	Address bits 0-23 hold	after \overline{ADS} T. E.	15		15		15		ns
t _{ALf}	4-5	Address bits 0-23 floating (no MMU)	after R.E., PHI1 T2		25		25		25	ns
t _{ADf}	4-5	Data bits D24-D31 floating (no MMU)	after R.E., PHI1 T2		25		25		25	ns
t _{ALMf}	4-9	Address bits 0-23 floating (with MMU)	after R.E., PHI1 Tmmu		25		25		25	ns
t _{ADMf}	4-9	Data bits 21-31 floating (with MMU)	after R.E., PHI1 Tmmu		25		25		25	ns
t _{BEv}	4-4	$\overline{BE}n$ signals valid	after R.E., PHI2 T4		95		70		45	ns
t _{BEh}	4-4	$\overline{BE}n$ signals hold	after R.E., PHI2 T4 or Ti	0		0		0		ns
t _{STv}	4-4	Status (ST0-ST3) valid	after R.E., PHI1 T4 (before T1, see note)		90		70		45	ns
t _{STh}	4-4	Status (ST0-ST3) hold	after R.E., PHI1 T4 (after T1)	0		0		0		ns
t _{DDInv}	4-4	\overline{DDIN} signal valid	after R.E., PHI1 T1		110		90		65	ns

4.0 Device Specifications (Continued)**4.4.2.1 Output Signals: Internal Propagation Delays NS32132-6, NS32132-8, NS32132-10** (Continued)

Name	Figure	Description	Reference/ Conditions	NS32132-6		NS32132-8		NS32132-10		Units
				Min	Max	Min	Max	Min	Max	
t_{DDINh}	4-4	\overline{DDIN} signal hold	after R.E., PHI1 next T1 or Ti	0		0		0		ns
t_{ADSa}	4-4	\overline{ADS} signal active (low)	after R.E., PHI1 T1		55		45		35	ns
t_{ADSi}	4-4	\overline{ADS} signal inactive	after R.E., PHI2 T1		60		55		45	ns
t_{ADSw}	4-4	\overline{ADS} pulse width	at 0.8V (both edges)	50		40		30		ns
t_{DSa}	4-4	\overline{DS} signal active (low)	after R.E., PHI1 T2		70		60		45	ns
t_{DSi}	4-4	\overline{DS} signal inactive	after R.E., PHI1 T4		50		50		40	ns
t_{ALf}	4-6	AD0-AD23 floating (caused by HOLD)	after R.E., PHI1 T1		100		65		25	ns
t_{ADf}	4-6	D24-D31 floating (caused by HOLD)	after R.E., PHI1 T1		100		65		25	ns
t_{DSf}	4-6	\overline{DS} floating (caused by HOLD)	after R.E., PHI1 T1		100		80		55	ns
t_{ADsf}	4-6	\overline{ADS} floating (caused by HOLD)	after R.E., PHI1 Ti		100		80		55	ns
t_{BEf}	4-6	\overline{BEn} floating (caused by HOLD)	after R.E., PHI1 Ti		100		80		55	ns
t_{DDINf}	4-6	\overline{DDIN} floating (caused by HOLD)	after R.E., PHI1 Ti		100		80		55	ns
t_{HLDAa}	4-6	HLDA signal active (low)	after R.E., PHI1 Ti		100		90		75	ns
t_{HLDAi}	4-8	HLDA signal inactive	after R.E., PHI1 Ti		100		90		75	ns
t_{DSr}	4-8	\overline{DS} signal returns from floating (caused by HOLD)	after R.E., PHI1 T1		100		80		55	ns
t_{ADsr}	4-8	\overline{ADS} signal returns from floating (caused by HOLD)	after R.E., PHI1 Ti		100		80		55	ns
t_{BEr}	4-8	\overline{BEn} signals return from floating (caused by HOLD)	after R.E., PHI1 Ti		100		80		55	ns
t_{DDINr}	4-8	\overline{DDIN} signal returns from floating (caused by HOLD)	after R.E., PHI1 Ti		100		80		55	ns
t_{DDINf}	4-9	\overline{DDIN} signal floating (caused by FLT)	after FLT F.E.		80		65		50	ns
t_{DDINr}	4-10	\overline{DDIN} signal returns from floating (caused by FLT)	after FLT R.E.		75		65		50	ns
t_{SPCa}	4-13	SPC output active (low)	after R.E., PHI1 T1		50		45		35	ns
t_{SPCi}	4-13	SPC output inactive	after R.E., PHI1 T4		50		45		35	ns
t_{SPCnf}	4-15	SPC output nonforcing	after R.E., PHI2 T4		40		25		10	ns
t_{Dv}	4-13	Data valid (slave processor write)	after R.E., PHI1 T1		80		65		50	ns
t_{Dh}	4-13	Data hold (slave processor write)	after R.E., PHI1 next T1 or Ti	0		0		0		ns
t_{PFSw}	4-18	PFS pulse width	at 0.8V (both edges)	70		70		70		ns
t_{PFsa}	4-18	PFS pulse active (low)	after R.E., PHI2		70		60		50	ns
t_{PFsia}	4-18	PFS pulse inactive	after R.E., PHI2		70		60		50	ns
t_{ILOs}	4-20a	\overline{ILO} signal setup	before R.E., PHI1 T1 of first interlocked read cycle	30		30		30		ns

4.0 Device Specifications (Continued)

4.4.2.1 Output Signals: Internal Propagation Delays, NS32132-6, NS32132-8, NS32132-10 (Continued)

Name	Figure	Description	Reference/ Conditions	NS32132-6		NS32132-8		NS32132-10		Units
				Min	Max	Min	Max	Min	Max	
t _{ILOh}	4-20b	\overline{ILO} signal hold	after R.E., PHI1 T3 of last interlocked write cycle	10		10		10		ns
t _{ILOa}	4-21	\overline{ILO} signal active (low)	after R.E., PHI1		70		65		55	ns
t _{ILOia}	4-21	\overline{ILO} signal inactive	after R.E., PHI1		70		65		55	ns
t _{USv}	4-22	$\overline{U/S}$ signal valid	after R.E., PHI1 T4		70		60		45	ns
t _{USh}	4-22	$\overline{U/S}$ signal hold	after R.E., PHI1 T4	10		10		10		ns
t _{NSPF}	4-19b	Nonsequential fetch to next PFS clock cycle	after R.E., PHI1 T1	4		4		4		t _{Cp}
t _{PFNS}	4-19a	PFS clock cycle to next non-sequential fetch	before R.E., PHI1 T1	4		4		4		t _{Cp}
t _{LXPF}	4-29	Last operand transfer of an instruction to next PFS clock cycle	before R.E., PHI1 T1 of first bus cycle of transfer	0		0		0		t _{Cp}
t _{BBa}	4-30	BB signal active	after R.E., PHI1 T4		90		75		60	ns
t _{BBia}	4-20	BB signal inactive	after R.E., PHI1 T4		90		75		60	ns
t _{BROa}	4-20	\overline{BRO} signal active	after R.E., PHI2		90		75		60	ns
t _{BROia}	4-20	\overline{BRO} signal inactive	after R.E., PHI2		90		75		60	ns
t _{BEf}	4-30	\overline{BE} n signals floating (caused by \overline{BRI})	after R.E., PHI2 T4		100		80		55	ns
t _{BEr}	4-30	\overline{BE} n signals return from floating (caused by \overline{BRI})	after R.E., PHI1 T1		100		80		55	ns
t _{DDInf}	4-30	\overline{DDIN} floating (caused by \overline{BRI})	after R.E., PHI1 T1		100		80		55	ns
t _{DDInr}	4-30	\overline{DDIN} return from floating (caused by \overline{BRI})	after R.E., PHI1 T1		100		80		55	ns
t _{ALf}	4-30	A0–A23 floating (caused by \overline{BRI})	after R.E., PHI1 T4		45		35		25	ns
t _{ADf}	4-30	D24–D31 Floating (caused by \overline{BRI})	after R.E., PHI1 T4		45		35		25	ns

Note: Every memory cycle starts with T4, during which Cycle Status is applied. If the CPU was idling, the sequence will be: ". . . T4, T1 . . .". If the CPU was not idling, the sequence will be: ". . . T4, T1 . . .".

4.4.2.2 Input Signal Requirements: NS32132-6, NS32132-8, NS32132-10

Name	Figure	Description	Reference/Conditions	NS32132-6		NS32132-8		NS32132-10		Units
				Min	Max	Min	Max	Min	Max	
t _{PWR}	4-25	Power stable to RST T.E.	after V _{CC} reaches 4.5V	50		50		50		μs
t _{Dis}	4-5	Data in setup (read cycle)	before F.E., PHI2 T3	20		15		10		ns
t _{Dih}	4-5	Data in hold (read cycle)	after R.E., PHI1 T4	10		10		10		ns
t _{HLDa}	4-6	\overline{HOLD} active (low) setup time (see note)	before F.E., PHI2 TX1	25		25		25		ns
t _{HLDia}	4-8	\overline{HOLD} inactive setup time	before F.E., PHI2 T1	25		25		25		ns
t _{HL Dh}	4-6	\overline{HOLD} hold time	after R.E., PHI1 TX2	0		0		0		ns
t _{FLTa}	4-9	\overline{FLT} active (low) setup time	before F.E., PHI2 Tmmu	25		25		25		ns

4.0 Device Specifications (Continued)

4.4.2.2 Input Signal Requirements NS32132-6, NS32132-8, NS32132-10 (Continued)

Name	Figure	Description	Reference/ Conditions	NS32132-6		NS32132-8		NS32132-10		Units
				Min	Max	Min	Max	Min	Max	
t_{FLTia}	4-10	FLT inactive setup time	before F.E., PHI2 T2	25		25		25		ns
t_{RDYs}	4-11, 4-12	RDY setup time	before F.E., PHI2 T2 or T3	25		20		15		ns
t_{RDYh}	4-11, 4-12	RDY hold time	after F.E., PHI1 T3	0		0		0		ns
t_{ABTs}	4-23	ABT setup time (FLT inactive)	before F.E., PHI2 Tmmu	30		25		20		ns
t_{ABTs}	4-24	ABT setup time (FLT active)	before F.E., PHI2 T2	30		25		20		ns
t_{ABTh}	4-23	ABT hold time	after R.E., PHI1	0		0		0		ns
t_{RSTs}	4-25, 4-26	RST setup time	before F.E., PHI1	20		15		10		ns
t_{RSTw}	4-26	RST pulse width	at 0.8V (both edges)	64		64		64		t_{Cp}
t_{INTs}	4-27	INT setup time	before F.E., PHI1	20		20		20		ns
t_{NMlw}	4-28	NMI pulse width	at 0.8V (both edges)	70		70		70		ns
t_{Dis}	4-14	Data setup (slave read cycle)	before F.E., PHI2 T1	20		20		10		ns
t_{Dih}	4-14	Data hold (slave read cycle)	after R.E., PHI1 T4	10		10		10		ns
t_{SPCd}	4-15	SPC pulse delay from slave	after R.E., PHI2 T4	17		13		10		ns
t_{SPCs}	4-15	SPC setup time	before F.E., PHI1	42		32		25		ns
t_{SPCw}	4-15	SPC pulse width (from slave processor)	at 0.8V (both edges)	30		25		20		ns
t_{ATs}	4-16	AT/SPC setup for address translation strap	before R.E., PHI1 of cycle during which RST pulse is removed	1		1		1		t_{Cp}
t_{ATh}	4-16	AT/SPC hold for address translation strap	after F.E., PHI1 of cycle during which RST pulse is removed	2		2		2		t_{Cp}
t_{BRIs}	4-20	BRI Setup Time	before F.E. PHI2	15		15		15		ns
t_{BRIn}	4-20	BRI Hold Time	after R.E. PHI1	10		10		10		ns

Note: This setup time is necessary to ensure prompt acknowledgement via HLDA and the ensuing floating of CPU off the buses. Note that the time from the receipt of the HOLD signal until the CPU floats is a function of the time HOLD signal goes low, the state of the RDY input (in MMU systems), and the length of the current MMU cycle.

4.4.2.3 Clocking Requirements: NS32132-6, NS32132-8, NS32132-10

Name	Figure	Description	Reference/ Conditions	NS32132-6		NS32132-8		NS32132-10		Units
				Min	Max	Min	Max	Min	Max	
t_{CLr}	4-17	PHI1, PHI2 rise time	0.8V to $V_{CC} - 0.9V$ on R.E., PHI1, PHI2		9		8		7	ns
t_{CLf}	4-17	PHI1, PHI2 fall time	$V_{CC} - 0.9V$ to 0.8V on F.E., PHI1, PHI2		9		8		7	ns
t_{Cp}	4-17	Clock period	R.E., PHI1, PHI2 to next R.E., PHI1, PHI2	170	5000	130	5000	100	5000	ns
$t_{CLw(1,2)}$	4-17	PHI1, PHI2 pulse width	At 2.0V on PHI1, PHI2 (both edges)	$0.5t_{Cp} - 14$		$0.5t_{Cp} - 12$		$0.5t_{Cp} - 10$		ns
$t_{CLh(1,2)}$	4-17	PHI1, PHI2 high time	At $V_{CC} - 0.9V$ on PHI1, PHI2 (both edges)	$0.5t_{Cp} - 18$		$0.5t_{Cp} - 17$		$0.5t_{Cp} - 15$		ns
$t_{nOVL(1,2)}$	4-17	Non-overlap time	0.8V on F.E., PHI1, PHI2 to 0.8V on R.E., PHI2, PHI1	0	7	0	7	0	7	ns
t_{nOVLas}		Non-overlap asymmetry ($t_{nOVL(1)} - t_{nOVL(2)}$)	at 0.8V on PHI1, PHI2	-4	4	-4	4	-4	4	ns
t_{CLwas}		PHI1, PHI2 asymmetry ($t_{CLw(1)} - t_{CLw(2)}$)	at 2.0V on PHI1, PHI2	-5	5	-5	5	-5	5	ns

4.0 Device Specifications (Continued)

4.4.3 Timing Diagrams

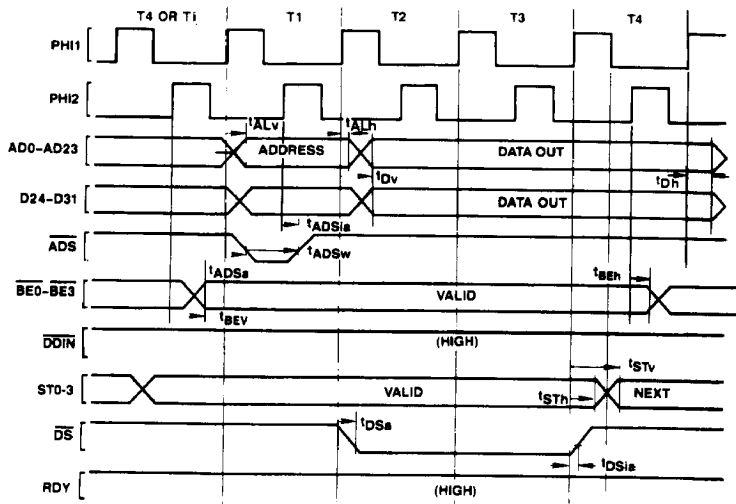


FIGURE 4-4. Write Cycle

TL/EE/8583-44


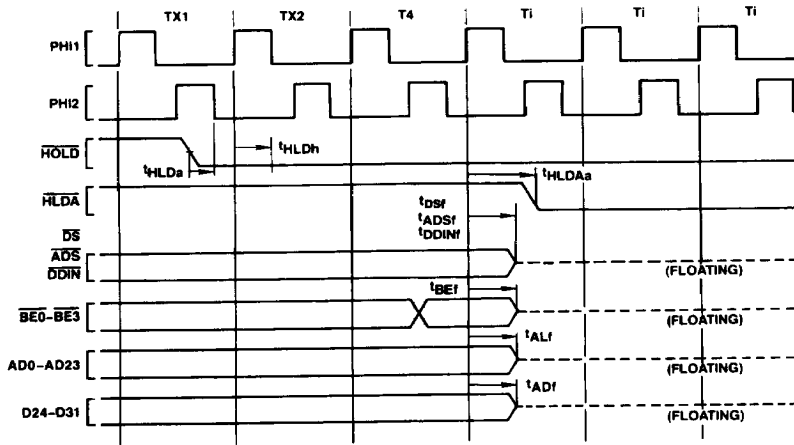


FIGURE 4-5. Read Cycle

TL/EE/8583-45

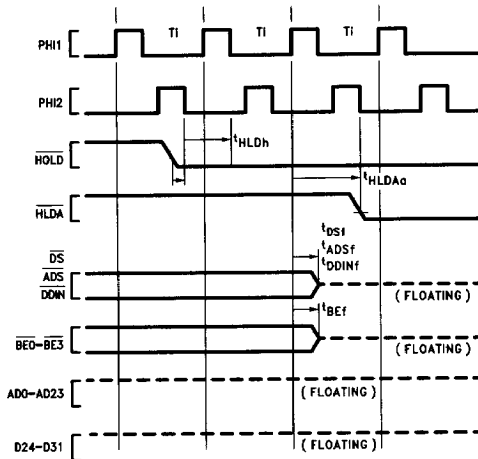
4.0 Device Specifications (Continued)



TL/EE/8583-46

FIGURE 4-6. Floating by $\overline{\text{HOLD}}$ Timing (CPU Not Idle Initially)

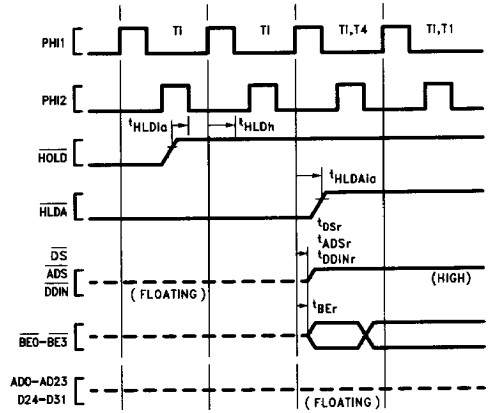
Note that whenever the CPU is not idling (not in TI), the $\overline{\text{HOLD}}$ request ($\overline{\text{HOLD}}$ low) must be active t_{HLDa} before the falling edge of PH12 of the clock cycle that appears two clock cycles before T4 (TX1) and stay low until t_{HLDh} after the rising edge of PH11 of the clock cycle that precedes T4 (TX2) for the request to be acknowledged.



TL/EE/8583-47

FIGURE 4-7. Floating by $\overline{\text{HOLD}}$ Timing (CPU initially idle)

Note that during TI1 the CPU is already idling.



TL/EE/8583-48

FIGURE 4-8. Release from $\overline{\text{HOLD}}$

4.0 Device Specifications (Continued)

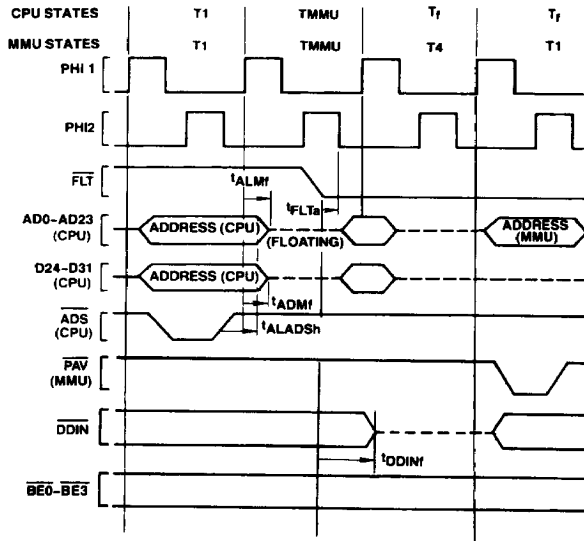


FIGURE 4-9. FLT Initiated Float Cycle Timing

TL/EE/8583-49




FIGURE 4-10. Release from \overline{FLT} Timing

TL/EE/8583-50

Note that when \overline{FLT} is deasserted the CPU restarts driving \overline{DDIN} before the MMU releases it. This, however, does not cause any conflict, since both CPU and MMU force \overline{DDIN} to the same logic level.





FIGURE 4-11. Ready Sampling (CPU Initially READY)


TL/EE/8583-51

4.0 Device Specifications (Continued)



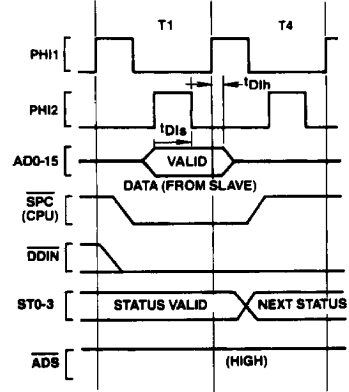
TL/EE/8583-52

FIGURE 4-12. Ready Sampling (CPU Initially NOT READY)



TL/EE/8583-53

FIGURE 4-13. Slave Processor Write Timing



TL/EE/8583-54

FIGURE 4-14. Slave Processor Read Timing




FIGURE 4-15. SPC Timing

TL/EE/8583-83

After transferring last operand to a Slave Processor, CPU turns OFF driver and holds SPC high with internal 5 kΩ pullup.




FIGURE 4-16. Reset Configuration Timing

TL/EE/8583-56

4.0 Device Specifications (Continued)




FIGURE 4-17. Clock Waveforms

TL/EE/8583-57




FIGURE 4-18. Relationship of \overline{PFS} to Clock Cycles

TL/EE/8583-58




FIGURE 4-19a. Guaranteed Delay, \overline{PFS} to Non-Sequential Fetch

TL/EE/8583-59

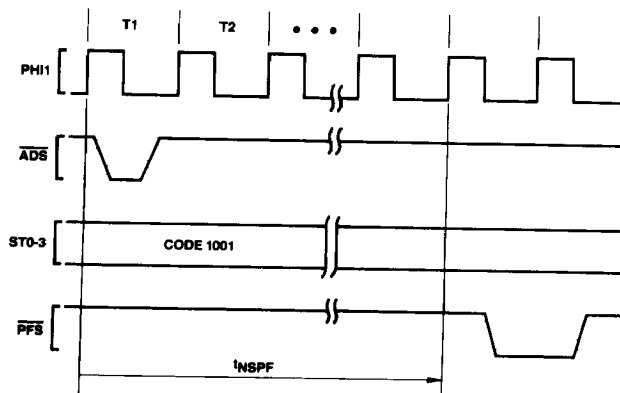



FIGURE 4-19b. Guaranteed Delay, Non-Sequential Fetch to \overline{PFS}


TL/EE/8583-60

4.0 Device Specifications (Continued)




TL/EE/8583-61

FIGURE 4-20a. Relationship of \overline{ILO} to First Operand Cycle of an Interlocked Instruction



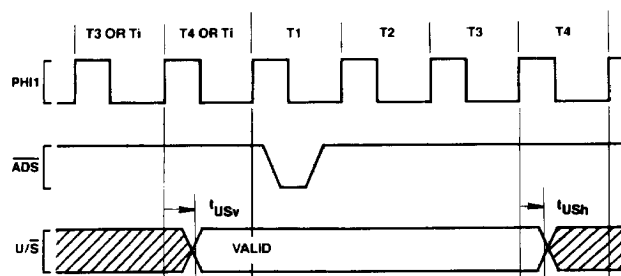
TL/EE/8583-62

FIGURE 4-20b. Relationship of \overline{ILO} to Last Operand Cycle of an Interlocked Instruction



TL/EE/8583-63

FIGURE 4-21. Relationship of \overline{ILO} to Any Clock Cycle



TL/EE/8583-64

FIGURE 4-22. U/\overline{S} Relationship to Any Bus Cycle — Guaranteed Valid Interval

4.0 Device Specifications (Continued)




FIGURE 4-23. Abort Timing, FLT Not Applied

TL/EE/8583-65




FIGURE 4-24. Abort Timing, FLT Applied

TL/EE/8583-66

2




FIGURE 4-25. Power-On Reset

TL/EE/8583-67




FIGURE 4-26. Non-Power-On Reset

TL/EE/8583-68

4.0 Device Specifications (Continued)




FIGURE 4-27. \overline{INT} Interrupt Signal Detection




FIGURE 4-28. \overline{NMI} Interrupt Signal Timing




FIGURE 4-29. Relationship Between Last Data Transfer of an Instruction and \overline{PFS} Pulse of Next Instruction

Note: In a transfer of a Read-Modify-Write type operand, this is the Read transfer, displaying RMW Status (Code 1011).

4.0 Device Specifications (Continued)

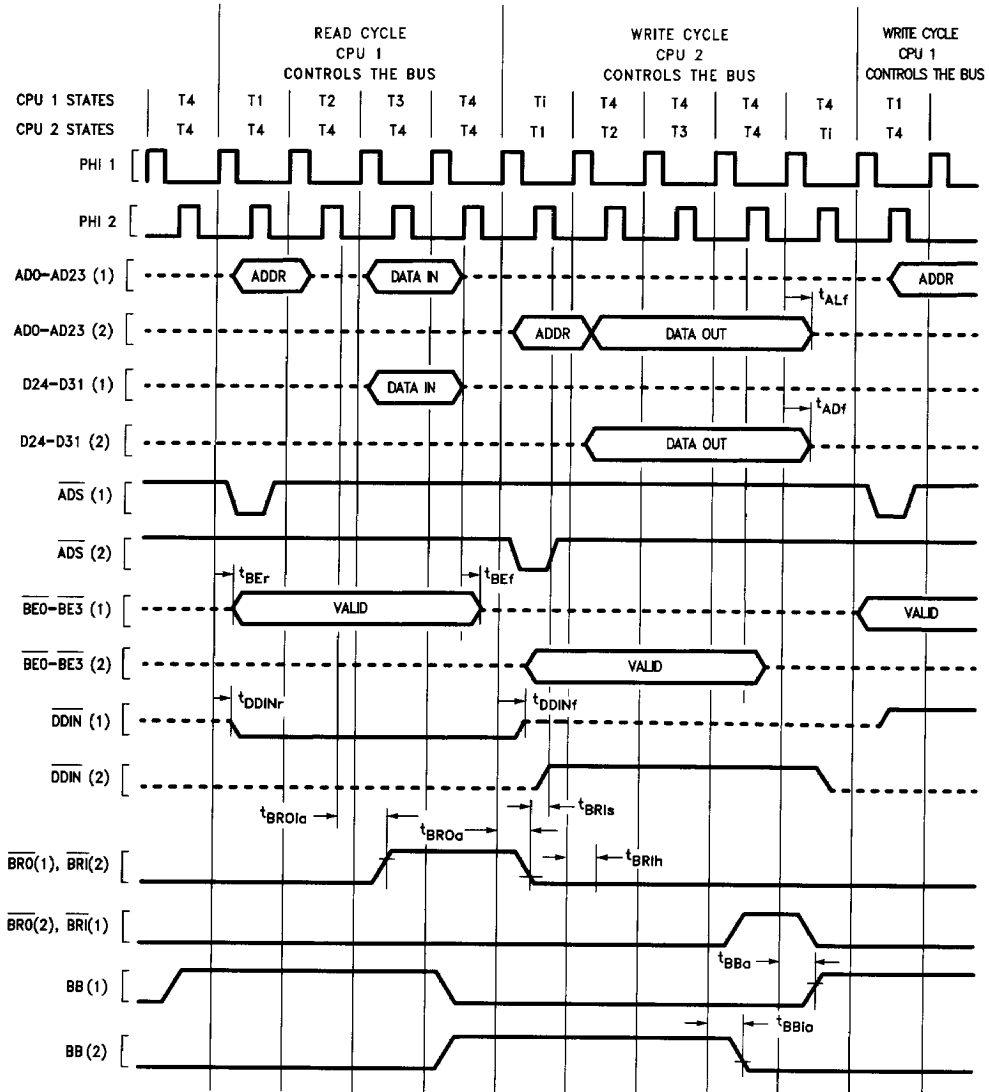


FIGURE 4-30. Dual Processor Bus Arbitration Timing

TL/EE/8583-72

Appendix A: Instruction Formats


NOTATIONS

- i= Integer Type Field
 - B = 00 (Byte)
 - W = 01 (Word)
 - D = 11 (Double Word)
 - f= Floating Point Type Field
 - F = 1 (Std. Floating: 32 bits)
 - L = 0 (Long Floating: 64 bits)
 - c= Custom Type Field
 - D = 1 (Double Word)
 - Q = 0 (Quad Word)
 - op= Operation Code

Valid encodings shown with each format.
 - gen, gen 1, gen 2= General Addressing Mode Field


See Sec. 2.2 for encodings.
 - reg= General Purpose Register Number
 - cond= Condition Code Field
 - 0000 = EQual: Z = 1
 - 0001 = Not Equal: Z = 0
 - 0010 = Carry Set: C = 1
 - 0011 = Carry Clear: C = 0
 - 0100 = Higher: L = 1
 - 0101 = Lower or Same: L = 0
 - 0110 = Greater Than: N = 1
 - 0111 = Less or Equal: N = 0
 - 1000 = Flag Set: F = 1
 - 1001 = Flag Clear: F = 0
 - 1010 = LOver: L = 0 and Z = 0
 - 1011 = Higher or Same: L = 1 or Z = 1
 - 1100 = Less Than: N = 0 and Z = 0
 - 1101 = Greater or Equal: N = 1 or Z = 1
 - 1110 = (Unconditionally True)
 - 1111 = (Unconditionally False)
 - short= Short Immediate value. May contain
 - quick: Signed 4-bit value, in MOVQ, ADDQ, CMPQ, ACB.
 - cond: Condition Code (above), in Scond.
 - areg: CPU Dedicated Register, in LPR, SPR.
 - 0000 = US
 - 0001 - 0111 = (Reserved)
 - 1000 = FP
 - 1001 = SP
 - 1010 = SB
 - 1011 = (Reserved)
 - 1100 = (Reserved)
 - 1101 = PSR
 - 1110 = INTBASE
 - 1111 = MOD
- Options: in String Instructions
- | | | |
|-----|---|---|
| U/W | B | T |
|-----|---|---|
- T = Translated
 - B = Backward
 - U/W = 00: None
 - 01: While Match
 - 11: Until Match

Configuration bits, in SETCFG:




mreg: MMU Register number, in LMR, SMR.

- 0000 = BPR0
- 0001 = BPR1
- 0010 = (Reserved)
- 0011 = (Reserved)
- 0100 = PF0
- 0101 = PF1
- 0110 = (Reserved)
- 0111 = (Reserved)
- 1000 = SC
- 1001 = (Reserved)
- 1010 = MSR
- 1011 = BCNT
- 1100 = PTB0
- 1101 = PTB1
- 1110 = (Reserved)
- 1111 = EIA




Format 0

Bcond (BR)



Format 1


BSR	-0000	ENTER	-1000
RET	-0001	EXIT	-1001
CXP	-0010	NOP	-1010
RXP	-0011	WAIT	-1011
RETT	-0100	DIA	-1100
RETI	-0101	FLAG	-1101
SAVE	-0110	SVC	-1110
RESTORE	-0111	BPT	-1111



Format 2

ADDQ	-000	ACB	-100
CMPQ	-001	MOVQ	-101
SPR	-010	LPR	-110
Scond	-011		


Appendix A: Instruction Formats (Continued)



Format 3


CXPD	-0000	ADJSP	-1010
BICPSR	-0010	JSR	-1100
JUMP	-0100	CASE	-1110
BISPSR	-0110		

Trap (UND) on XXX1, 1000



Format 4


ADD	-0000	SUB	-1000
CMP	-0001	ADDR	-1001
BIC	-0010	AND	-1010
ADDC	-0100	SUBC	-1100
MOV	-0101	TBIT	-1101
OR	-0110	XOR	-1110



Format 5


MOVS	-0000	SETCFG	-0010
CMPS	-0001	SKPS	-0011

Trap (UND) on 1XXX, 01XX




Format 6

ROT	-0000	NEG	-1000
ASH	-0001	NOT	-1001
CBIT	-0010	Trap (UND)	-1010
CBITI	-0011	SUBP	-1011
Trap (UND)	-0100	ABS	-1100
LSH	-0101	COM	-1101
SBIT	-0110	IBIT	-1110
SBITI	-0111	ADDP	-1111



Format 7


MOVM	-0000	MUL	-1000
CMPM	-0001	MEI	-1001
INSS	-0010	Trap (UND)	-1010
EXTS	-0011	DEI	-1011
MOVXBW	-0100	QUO	-1100
MOVZBW	-0101	REM	-1101
MOVZID	-0110	MOD	-1110
MOVXID	-0111	DIV	-1111



TL/EE/8583-73


Format 8

EXT	-0 00	INDEX	-1 00
CVTP	-0 01	FFS	-1 01
INS	-0 10		
CHECK	-0 11		
MOVSU	-110, reg = 001		
MOVUS	-110, reg = 011		



Format 9

MOVif	-000	ROUND	-100
LFSR	-001	TRUNC	-101
MOVLF	-010	SFSR	-110
MOVFL	-011	FLOOR	-111




TL/EE/8583-74

Format 10

Trap (UND) Always

Appendix A: Instruction Formats (Continued)



Format 11


ADDf	-0000	DIVf	-1000
MOVf	-0001	Trap (Slave)	-1001
CMPf	-0010	Trap (UND)	-1010
Trap (Slave)	-0011	Trap (UND)	-1011
SUBf	-0100	MULf	-1100
NEGf	-0101	ABSf	-1101
Trap (UND)	-0110	Trap (UND)	-1110
Trap (UND)	-0111	Trap (UND)	-1111



TL/EE/8583-75

Format 12


Trap (UND) Always



TL/EE/8583-76


Format 13

Trap (UND) Always



Format 14

RDVAL	-0000	LMR	-1010
WRVAL	-0001	SMR	-1011
Trap (UND) on 01XX, 1XXX			



Operation Word


ID Byte

Format 15

(Custom Slave)

Operation Word Format

nnn




000

Format 15.0

CATST0	-0000	LCR	-1010
CATST1	-0001	SCR	-1011


Trap (UND) on all others



001

Format 15.1

CCV3	-000	CCV2	-100
LCSR	-001	CCV1	-101
CCV5	-010	SCSR	-110
CCV4	-011	CCV0	-111





101

Format 15.5

CCAL0	-0000	CCAL3	-1000
CMOV0	-0001	CMOV3	-1001
CCMP	-0010	Trap (UND)	-1010
CCMP1	-0011	Trap (UND)	-1011
CCAL1	-0100	CCAL2	-1100
CMOV2	-0101	CMOV1	-1101
Trap (UND)	-0110	Trap (UND)	-1110
Trap (UND)	-0111	Trap (UND)	-1111


If nnn = 010, 011, 100, 110, 111
then Trap (UND) Always

Appendix A: Instruction Formats (Continued)




Format 16

Trap (UND) Always



Format 17

Trap (UND) Always






Format 18

Trap (UND) Always

Format 19

Trap (UND) Always

Implied Immediate Encodings:



Appendix B. Interfacing Suggestions

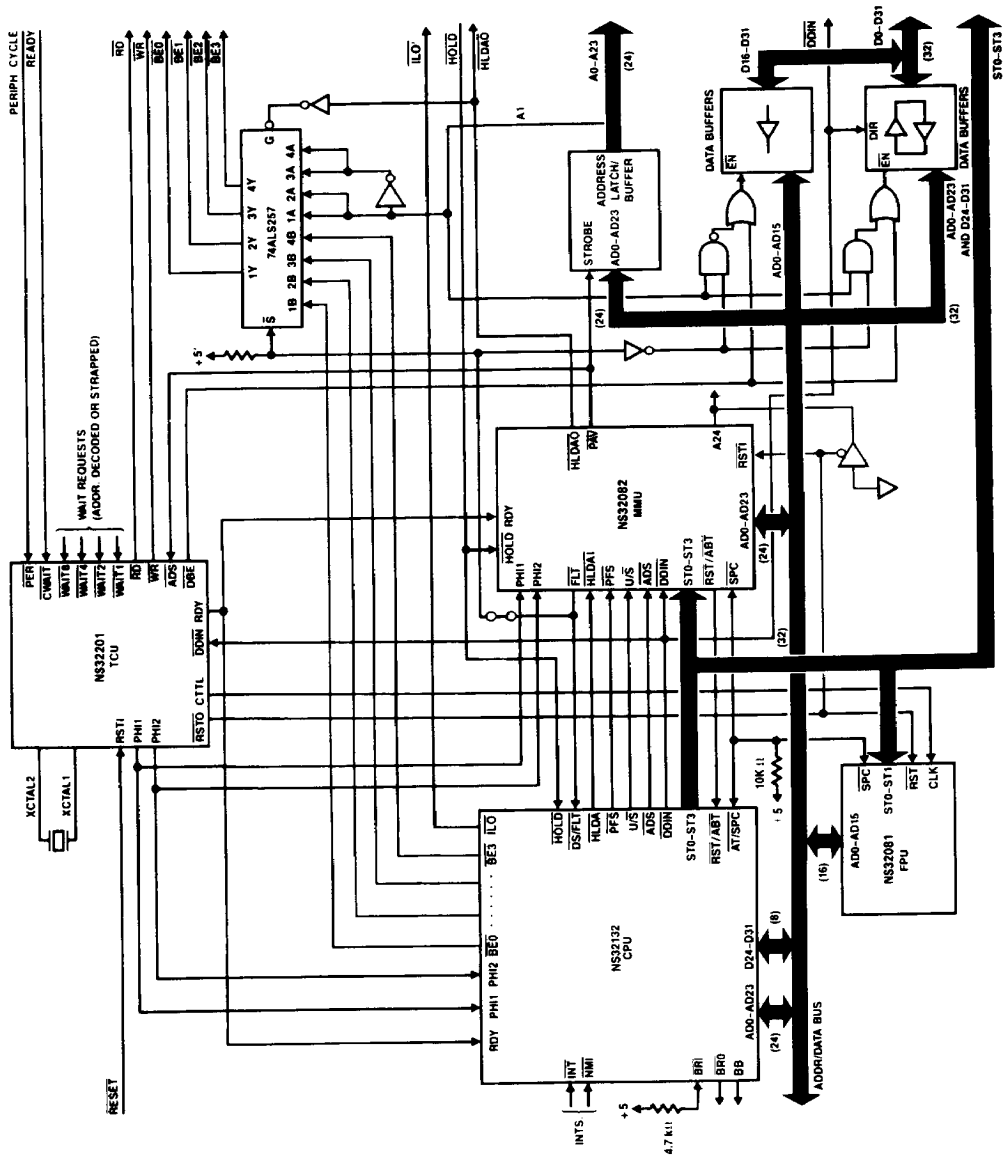


FIGURE B-1. Single Processor System Connection Diagram

TL/EE/8683-81

