

NS32FX16-15/NS32FX16-20/NS32FX16-25 Imaging/Signal Processor

General Description

The NS32FX16 is a high-performance 32-bit member of the Series 32000/EP™ family of National's Embedded System Processors™ specifically optimized for CCITT Group 2 and Group 3 Facsimile Applications, Data Modems, Voice Mail Systems, Laser Printers, or any combination of the above.

It can perform all the computations and control functions required for a stand-alone Fax system, a PC add-in Fax/Data Modem card or a Laser/Fax system.

It also meets the performance requirements to implement 9600 and 7200 bps modems complying with CCITT V.29 and V.27 standards.

The NS32FX16 provides a 16 Mbyte Linear external address space and a 16-bit external data bus.

The CPU core, which is the same as that of the NS32CG16, incorporates a 32-bit ALU and instruction pipeline, and an 8-byte prefetch queue.

Also integrated on-chip with the CPU are a DSP Module and a 384-byte RAM Array. The DSP Module executes vector operations on complex variables and is specially designed to enhance performance in modem applications. The vector operations can also be used to efficiently implement FIR Filters and other DSP primitives. The on-chip RAM Array is used to store the coefficients of the various filters and can be accessed by both the CPU and the DSP Module.


The NS32FX16 capabilities can be expanded by using an external floating point unit (FPU) which directly interfaces to the NS32FX16 using the slave protocol. The CPU-FPU cluster features high speed execution of the floating-point instructions.

The NS32FX16 highly-efficient architecture combined with the NS32CG16 graphics instructions and the high-performance vector operation capability, makes the device the ideal choice for Postscript™ and Fax applications.

Features

- Software compatible with the Series 32000/EP processors
- Designed around the CPU core of the NS32CG16
- 32-bit architecture and implementation
- On-chip DSP Module for high-speed DSP operations
- Special support for graphics applications
 - 18 graphics instructions
 - Binary compression/expansion capability for font storage using RLL encoding
 - Pattern magnification
 - Interface to an external BITBLT processing units for fast color BITBLT operations
- 384-byte on-chip RAM array
- On-chip clock generator
- Floating-point support via the NS32081 or NS32181
- Optimal interface to large memory arrays via the NS32CG821 and the DP84xx family of DRAM controllers
- Power save mode
- High-speed CMOS technology
- 68-pin PLCC package

Block Diagram



TL/EE/10818-67

Series 32000® is a registered trademark of National Semiconductor Corporation.
 EP™ and Embedded System Processors™ are trademarks of National Semiconductor Corporation.

Table of Contents

1.0 PRODUCT INTRODUCTION

1.1 NS32FX16 Special Features

2.0 ARCHITECTURAL DESCRIPTION

2.1 Register Set

2.1.1 General Purpose Registers

2.1.2 Address Registers

2.1.3 Processor Status Register

2.1.4 Configuration Register

2.1.5 DSP Module Registers

2.1.6 RAM Array

2.2 Memory Organization

2.2.1 Address Mapping

2.3 Modular Software Support

2.4 Instruction Set

2.4.1 General Instruction Format

2.4.2 Addressing Modes

2.4.3 Instruction Set Summary

2.4 Graphic Support

2.5.1 Frame Buffer Addressing

2.5.2 BITBLT Fundamentals

2.5.2.1 Frame Buffer Architecture

2.5.2.2 BIT Alignment

2.5.2.3 Block Boundaries and Destination Masks

2.5.2.4 BITBLT Directions

2.5.2.5 BITBLT Variations

2.5.3 Graphics Support Instructions

2.5.3.1 BITBLT (BIT-aligned BLock Transfer)

2.5.3.2 Pattern Fill

2.5.3.3 Data Compression, Expansion and Magnify

2.5.3.3.1 Magnifying Compressed Data

3.0 FUNCTIONAL DESCRIPTION

3.1 Instruction Execution

3.1.1 Operating States

3.1.2 Instruction Endings

3.1.2.1 Completed Instructions

3.1.2.2 Suspended Instructions

3.1.2.3 Terminated Instructions

3.1.2.3 Partially Completed Instructions

3.1.3 Slave Processor Instructions

3.1.3.1 Slave Processor Protocol

3.1.3.2 Floating-Point Instructions

3.2 Exception Processing

3.2.1 Exception Acknowledge Sequence

3.2.2 Returning from an Exception Service Procedure

3.2.3 Maskable Interrupts

3.2.3.1 Non-Vectored Mode

3.2.3.2 Vectored Mode: Non-Cascaded Case

3.2.3.3 Vectored Mode: Cascaded Case

3.2.4 Non-Maskable Interrupt

3.2.5 Traps

3.2.6 Priority among Exceptions

3.2.7 Exception Acknowledge Sequences: Detailed Flow

3.2.7.1 Maskable/Non-Maskable Interrupt Sequence

3.2.7.2 SLAVE/ILL/SVC/DVZ/FLG/BPT/UND Trap Sequence

3.2.7.3 Trace Trap Sequence

3.3 Debugging Support

3.3.1 Instruction Tracing

3.4 DSP Module (DSPM)

3.4.1 DSPM Operation

3.4.2 Complex Number Representation

3.4.3 DSPM Instructions

3.4.4 Circular Buffers

3.5 System Interface

3.5.1 Power and Grounding

3.5.2 Clocking

3.5.3 Power Save Mode

3.5.4 Resetting

3.5.5 Bus Cycles

3.5.5.1 Bus Status

3.5.5.2 Basic Read and Write Cycles

3.5.5.3 Cycle Extension

3.5.5.4 Instruction Fetch Cycles

3.5.5.5 Interrupt Control Cycles

3.5.5.6 Special Bus Cycles

3.5.5.7 Slave Processor Bus Cycles

3.5.5.8 Data Access Sequences

3.5.5.9 Bus Access Control

3.5.5.10 Instruction Status

Table of Contents (Continued)

4.0 DEVICE SPECIFICATIONS

- 4.1 NS32FX16 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 Signal Requirements
 - 4.4.3 Timing Diagrams

Appendix A: INSTRUCTION FORMATS

Appendix B: INSTRUCTION EXECUTION TIMES

- B.1 Basic and Floating-Point Instructions
 - B.1.1 Equations
 - B.1.2 Notes on Table Use
 - B.1.3 Calculation of the Execution Time TEX for Basic Instructions
 - B.1.4 Calculation of the Execution Time TEX for Floating-Point Instructions
- B.2 Special Graphics Instructions
 - B.2.1 Execution Time Calculation for Special Graphics Instructions
- B.3 DSPM Instructions

List of Illustrations

CPU Block Diagram	1-1
NS32FX16 Internal Registers	2-1
Processor Status Register (PSR)	2-2
Configuration Register (CFG)	2-3
DSP Module Registers Address Map	2-4
CPTR Register Format	2-5
CTL Register Format	2-6
ST Register Format	2-7
On-Chip RAM Array Address Map	2-8
NS32FX16 Address Mapping	2-9
NS32FX16 Run-Time Environment	2-10
General Instruction Format	2-11
Index Byte Format	2-12
Displacement Encodings	2-13
Correspondence between Linear and Cartesian Addressing	2-14
32-Pixel by 32-Scan Line Frame Buffer	2-15
Overlapping BITBLT Blocks	2-16

List of Illustrations (Continued)

B B Instructions Format	2-17
BITWT Instruction Format	2-18
EXTBLT Instruction Format	2-19
MOVMPi Instruction Format	2-20
TBITS Instruction Format	2-21
SBITS Instruction Format	2-22
SBITPS Instruction Format	2-23
Bus Activity for a Simple BITBLT Operation	2-24
Operating States	3-1
Slave Processor Protocol	3-2
Slave Processor Status Word	3-3
Interrupt Dispatch Table	3-4
Exception Acknowledge Sequence	3-5
Return from Trap (RETTn) Instruction Flow	3-6
Return from Interrupt (RETI) Instruction Flow	3-7
Interrupt Control Unit Connections (16 Levels)	3-8
Cascaded Interrupt Control Unit Connections	3-9
Exception Processing Flowchart	3-10
Service Sequence	3-11
DSP Module Block Diagram	3-12
Memory Organization of a Complex Vector	3-13
Power and Ground Connections	3-14
Crystal Interconnections—30 MHz	3-15
Crystal Interconnections—40 MHz, 50 MHz	3-16
Recommended Reset Connections	3-17
Power-On Reset Requirements	3-18
General Reset Timings	3-19
Bus Connections	3-20
Read Cycle Timing	3-21
Write Cycle Timing	3-22
Cycle Extension of a Read Cycle	3-23
Special Bus Cycle Timing	3-24
Slave Processor Read Cycle	3-25
Slave Processor Write Cycle	3-26
NS32FX16 and FPU Interconnections	3-27
Memory Interface	3-28
HOLD Timing, Bus Initially Idle	3-29
HOLD Timing, Bus Initially Not Idle	3-30

List of Illustrations (Continued)

Connection Diagram	4-1
Output Signals Specification Standard	4-2
Input Signals Specification Standard	4-3
Read Cycle	4-4
Write Cycle	4-5
Special Bus Cycle	4-6
$\overline{\text{HOLD}}$ Acknowledge Timing (Bus Initially Not Idle)	4-7
$\overline{\text{HOLD}}$ Timing (Bus Initially Idle)	4-8
External DMA Controller Bus Cycle	4-9
Slave Processor Write Timing	4-10
Slave Processor Read Timing	4-11
$\overline{\text{SPC}}$ Timing	4-12
$\overline{\text{PFS}}$ Signal Timing	4-13
$\overline{\text{ILO}}$ Signal Timing	4-14
Clock Waveforms	4-15
$\overline{\text{INT}}$ Signal Timing	4-16
$\overline{\text{NITI}}$ Signal Timing	4-17
Power-On Reset	4-18
Non-Power-On Reset	4-19

List of Tables

NS32FX16 Addressing Modes	2-1
NS32FX16 Instruction Set Summary	2-2
'op' and 'i' Field Encodings	2-3
Floating-Point Instruction Protocols	3-1
Summary of Exception Processing	3-2
DSPM Instructions Summary	3-3
Circular Buffer Sizes	3-4
External Oscillator Specifications	3-5
Interrupt Sequences	3-6
Bus Cycle Categories	3-7
Data Access Sequences	3-8
Basic Instructions	B-1
Floating-Point Instructions: CPU Portion	B-2
Average Instruction Execution Times with No Wait-States	B-3
Average Instruction Execution Times with Wait-States	B-4
DSPM Instruction Execution Times	B-5

1.0 Product Introduction

The NS32FX16 is a high speed CMOS microprocessor in the Series 32000/EP family.

It includes two main execution units: the NS32CG16 compatible CPU core and the DSP Module. The CPU core is designed for general purpose computations and system control functions. The DSP Module is tuned to perform the DSP primitives needed in Voice Band Modems. The NS32FX16 also incorporates a 384-byte RAM Array as a shared resource for both the CPU core and the DSP Module.

The NS32FX16 is software-compatible with all other CPUs in the family.

The device incorporates all of the Series 32000 advanced architectural features, with the exception of the virtual memory capability.

Brief descriptions of the NS32FX16 features that are shared with other members of the family are provided 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-address machines. This means that each operand can be referenced by any one of the addressing modes provided.

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.

Large, Uniform Addressing. The NS32FX16 has 24-bit address pointers that can address up to 16 megabytes without any segmentation; this addressing scheme provides flexible memory management without add-on expense.

Modular Software Support. Any software package for the Series 32000 architecture 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 exten-

sions to the CPU. This concept of slave processors is unique to the Series 32000 architecture. 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

1.1 NS32FX16 SPECIAL FEATURES

In addition to the above Series 32000 features, the NS32FX16 provides features that make the device extremely attractive for a wide range of applications where graphics support, low chip count, and low power consumption are required.

The most relevant of these features are the enhanced Digital Signal Processing performance which makes the chip very attractive for facsimile applications, and the graphics support capabilities, that can be used in applications such as printers, CRT terminals, and other varieties of display systems, where text and graphics are to be handled.

Graphics support is provided by eighteen instructions that allow operations such as BITBLT, data compression/expansion, fills, and line drawing, to be performed very efficiently. In addition, the device can be easily interfaced to an external BITBLT Processing Unit (BPU) for high BITBLT performance.

The NS32FX16 allows systems to be built with a relatively small amount of random logic. The bus is highly optimized to allow simple interfacing to a large variety of DRAMs and peripheral devices. All the relevant bus access signals and clock signals are generated on-chip. The cycle extension logic is also incorporated on-chip.

The device is fabricated in a low-power, high speed CMOS technology. It also includes a power-save feature that allows the clock to be slowed down under software control, thus minimizing the power consumption. This feature can be used in those applications where power saving during periods of low performance demand is highly desirable.

The power save feature, the DSP Module and the Bus Characteristics are described in the "Functional Description" section. A general overview of BITBLT operations and a description of the graphics support instructions is provided in Section 2.5. Details on all the NS32FX16 graphics instructions can be found in the NS32CG16 Printer/Display Processor Programmer's Reference Supplement.

1.0 Product Introduction (Continued)

Below is a summary of the instructions that are directly applicable to graphics along with their intended use.

Instruction	Application
BBAND BBOR BBFOR BBXOR BBSTOD BITWT EXTBLT	The BITBLT group of instructions provide a method of quickly imaging characters, creating patterns, windowing and other block oriented effects.
MOVMP	Move Multiple Pattern is a very fast instruction for clearing memory and drawing patterns and lines.
TBITS	Test Bit String will measure the length of 1's or 0's in an image, supporting many data compression methods (RLL), TBITS may also be used to test for boundaries of images.
SBITS	Set Bit String is a very fast instruction for filling objects, outline characters and drawing horizontal lines. The TBITS and SBITS instructions support Group 3 and Group 4 CCITT standards for compression and decompression algorithms.
SBITPS	Set Bit Perpendicular String is a very fast instruction for drawing vertical, horizontal and 45° lines. In printing applications SBITS and SBITPS may be used to express portrait and landscape respectively from the same compressed font data. The size of the character may be scaled as it is drawn.
SBIT CBIT TBIT IBIT	The Bit group of instructions enable single pixels anywhere in memory to be set, cleared, tested or inverted.
INDEX	The INDEX instruction combines a multiply-add sequence into a single instruction. This provides a fast translation of an X-Y address to a pixel relative address.

2.0 Architectural Description

2.1 REGISTER SET

The NS32FX16 has 23 internal registers and a 384-byte RAM array. 17 of these registers belong to the CPU portion of the device and are addressed either implicitly by specific instructions or through the register addressing mode. The other 6 control the operation of the DSP Module, and are memory mapped. *Figure 2-1* shows the NS32FX16 internal registers.




FIGURE 2-1. NS32FX16 Internal Registers

2.0 Architectural Description (Continued)

2.1.1 General Purpose Registers

There are eight registers (R0–R7) used for satisfying the 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 32 bits in length. If a general purpose register is specified for an operand that is 8 or 16 bits long, only the low part of the register is used; the high part is not referenced or modified.

2.1.2 Address Registers

The seven address registers are used by the processor to implement specific address functions. Except for the MOD register that is 16 bits wide, all the others are 32 bits. A description of the address registers follows.

PC—Program Counter. The PC 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.

SP0, SP1—Stack Pointers. 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.

When a reference is made to the selected Stack Pointer (see PSR S-bit), the terms “SP Register” or “SP” are used. SP refers 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, SP refers to SP0. If the S bit in the PSR is 1 then SP refers to SP1.

Stacks in the Series 32000 architecture 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—Frame Pointer. The FP 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.

SB—Static Base. The SB 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.

INTBASE—Interrupt Base. The INTBASE register holds the address of the dispatch table for interrupts and traps (Section 3.2.1).

MOD—Module. The MOD register holds the address of the module descriptor of the currently executing software module. The MOD register is 16 bits long, therefore the module table must be contained within the first 64 kbytes of memory.

2.1.3 Processor Status Register

The Processor Status Register (PSR) holds status information for the 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.




FIGURE 2-2. Processor Status Register (PSR)

- 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 set to 1, a TRC trap is executed after every instruction (Section 3.3.1).
- 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.
- K** Reserved for use by the CPU.
- J** Reserved for use by the CPU.
- 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 processor is said to be in Supervisor Mode; when U=1 the processor 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 (Section 3.3.1). It may have a setting of 0 (no trace pending) or 1 (trace pending).

2.0 Architectural Description (Continued)

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

B Reserved for use by the CPU. This bit is set to 1 during the execution of the EXTBLT instruction and causes the \overline{BPU} signal to become active. Upon reset, B is set to zero and the \overline{BPU} signal is set high.

Note 1: When an interrupt is acknowledged, the B, I, P, S and U bits are set to zero and the \overline{BPU} signal is set high. A return from interrupt will restore the original values from the copy of the PSR register saved in the interrupt stack.

Note 2: If BITBLT (BB) or EXTBLT instructions are executed in an interrupt routine, the PSR bits J and K must be cleared first.

2.1.4 Configuration Register

The Configuration Register (CFG) is 8 bits wide, of which four bits are implemented. The implemented bits are used to declare the presence of certain external devices and to select the clock scaling factor. CFG is programmed by the SETCFG instruction. The format of CFG is shown in Figure 2-3. The various control bits are described below.




FIGURE 2-3. Configuration Register (CFG)

I Interrupt vectoring. This bit controls whether maskable interrupts are handled in nonvectored ($I=0$) or vectored ($I=1$) mode. Refer to Section 3.2.3 for more information.

F Floating-point instruction set. This bit indicates whether a floating-point unit (FPU) is present to execute floating-point instructions. If this bit is 0 when the CPU executes a floating-point instruction, a Trap (UND) occurs. If this bit is 1, then the CPU transfers the instruction and any necessary operands to the FPU using the slave-processor protocol described in Section 3.1.3.1.

M Clock scaling. This bit is used in conjunction with the C bit to select the clock scaling factor.

C Clock scaling. Same as the M bit above. Refer to Section 3.5.3 on "Power Save Mode" for details.

2.1.5 DSP Module Registers

The DSP Module (DSPM) contains 6 memory-mapped registers. All the registers, except ST, are readable and writable. ST is read-only.

Accesses to these registers must be aligned; word and double-word accesses must occur on word and double-word address boundaries respectively. Section 2.2. Failing to do so will cause unpredictable results. Figure 2-4 shows the address map of the DSP Module Registers.

Register Names	Register Addresses
A	FFFFD400
Y	FFFFD404
0 DPTR	FFFFD408
CPTR	FFFFD40C
CTL	FFFFD40E
ST	FFFFD410

FIGURE 2-4. DSP Module Registers Address Map

A—Accumulator. This 32-bit register holds one complex number, and is mapped into two consecutive 16-bit words, called A0 and A1.

Internally, A0 and A1 are 32-bit registers, however, only bits 15–30 (16 bits) are accessible. The rest of the bits are used for bigger dynamic range during intermediate calculations. Section 3.4.

Y—Multiplier Input Register. This 32-bit register holds one complex operand. The Y register is mapped into two consecutive words called Y0 and Y1.

DPTR—Data Pointer. 24-bit pointer pointing at the data vector in main memory. In order to implement circular buffers, only the least significant bits of DPTR are incremented. When the end of a buffer is reached, the least significant bits of DPTR are reloaded with zeros. The number of bits that are set to zero (which defines the size of the circular buffer) is controlled by the CTL register. The least-significant 16-bit word of DPTR is called DPTR0, and the most-significant byte is called DPTR1. Bits 24–31 are reserved and should be set to 0.

CPTR—Coefficient Memory Vector Pointer. This 16-bit register holds the address and length of the coefficient vector, stored in the on-chip RAM array. The format of CPTR is shown in Figure 2-5.




FIGURE 2-5. CPTR Register Format

STRT Coefficient's Vector Index. At the beginning of the DSP operation, STRT points to the first coefficient in the coefficient vector. Section 2.1.6.

LNGT Coefficient's Vector Length. This field specifies the number of coefficients in the coefficient vector.

STRT and LNGT must be in the ranges 0 to 95 and 1 to 96 respectively. The relation $LNGT + STRT < 97$ must also be satisfied. Specifying values violating any of these conditions will cause unpredictable results.

CTL Control Register. This register controls the modes of operation of the DSP Module. The format of CTL is shown in Figure 2-6. The reserved bits should be set to 0. See Section 3.4 for more details.




FIGURE 2-6. CTL Register Format

OPM Opcode Modifier. This field is used to further qualify the operation specified by OPC. See Section 3.4.3 for details.

DS Data Buffer Size. Specifies the number of complex variables (4 bytes each) in the data buffer in main memory.

DS = 00 → 8 Complex Variables

DS = 01 → 16 Complex Variables

DS = 10 → 32 Complex Variables

DS = 11 → 64 Complex Variables

2.0 Architectural Description (Continued)

OPC Operation Code. Specifies the vector operation to be performed.

OPC = 00 → VCMAD Vector Complex Multiply Add

OPC = 01 → VCMUL Vector Complex Multiply

OPC = 10 → VCMAC Vector Complex Multiply Accumulate

OPC = 11 → VCMAG Vector Complex Magnitude

ST—Status Register. ST is a read-only register that holds the status of the last vector operation. The format of the ST register is shown in *Figure 2-7*.




FIGURE 2-7. ST Register Format

OP0 Overflow occurred on calculation of A0.

OP1 Overflow occurred on calculation of A1.

OVF Overflow indication.

The ST register is cleared to 0 in the following cases.

- the user writes directly to either A0 or A1,
- the user writes to the CTL register,
- upon reset.

2.1.6 RAM Array

The on-chip RAM array provides 384 bytes of storage that is used to store up to 96 32-bit complex numbers. These numbers represent the coefficients C[0]–C[95] used by the DSP Module.

During a vector operation, the DSP Module accesses these coefficients sequentially starting with the coefficient indexed by the STRT field in the CPTR register.

The RAM array is not limited to coefficient storage only. It can be used as a fast, zero wait-state on-chip memory for instructions and data storage.

RAM array accesses must be word or double-word aligned. Failing to do so may cause unpredictable results. *Figure 2-8* shows the RAM array address map.




FIGURE 2-8. On-Chip RAM Array Address Map

2.2 MEMORY ORGANIZATION


The main memory of the NS32FX16 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. Un-

less 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, 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, 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 higher. 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 words. Therefore, words and double-words that are aligned to start at even addresses (multiples of two) are accessed more quickly than words and double-words that are not so aligned.

2.2.1 Address Mapping

The NS32FX16 supports the use of memory-mapped peripheral devices and coprocessors. Such memory-mapped devices can be located at arbitrary locations within the 16-Mbyte address range available externally.

The address range from 01000000 (hex) to FF800000 (hex) is not available in the present implementation of the NS32FX16, and should not be used. The top 8-Mbyte block is reserved by National Semiconductor Corporation, and only a few locations within this block are presently used to access the on-chip RAM array and DSP Module control registers. *Figure 2-9* shows the NS32FX16 address mapping.

2.0 Architectural Description (Continued)

Address (Hex)		
00000000		Memory and I/O
00FFFE00		Interrupt Control
01000000		Address Not Available Off-Chip (Do Not Use)
FF800000		Reserved
FFFFD000		On-Chip RAM Array
FFFFD200		Reserved
FFFFD4XX		DSP Module
		Reserved
FFFFFFFF		

FIGURE 2-9. NS32FX16 Address Mapping

2.3 MODULAR SOFTWARE SUPPORT

The NS32FX16 provides special support for software modules and modular programs.

Each module in a NS32FX16 software environment consists of three components:

1. Program Code Segment.

This segment contains the module's code and constant data.

2. Static Data Segment.

Used to store variables and data that may be accessed by all procedures within the module.

3. Link Table.

This component contains two types of entries: Absolute Addresses and Procedure Descriptors.

An Absolute Address is used in the external addressing mode, in conjunction with a displacement and the current MOD Register contents to compute the effective address of an external variable belonging to another module.

The Procedure Descriptor is used in the call external procedure (CXP) instruction to compute the address of an external procedure.

Normally, the linker program specifies the locations of the three components. The Static Data and Link Table typically reside in RAM; the code component can be either in RAM or in ROM. The three components can be mapped into non-contiguous locations in memory, and each can be independently relocated. Since the Link Table contains the absolute addresses of external variables, the linker need not assign absolute memory addresses for these in the module itself; they may be assigned at load time.

To handle the transfer of control from one module to another, the NS32FX16 uses a module table in memory and two registers in the CPU.

The Module Table is located within the first 64 bytes of memory. This table contains a Module Descriptor (also called a Module Table Entry) for each module in the address space of the program. A Module Descriptor has four 32-bit entries corresponding to each component of a module:

- The Static Base entry contains the address of the beginning of the module's static data segment.
- The Link Table Base points to the beginning of the module's Link Table.
- The Program Base is the address of the beginning of the code and constant data for the module.
- A fourth entry is currently unused but reserved.

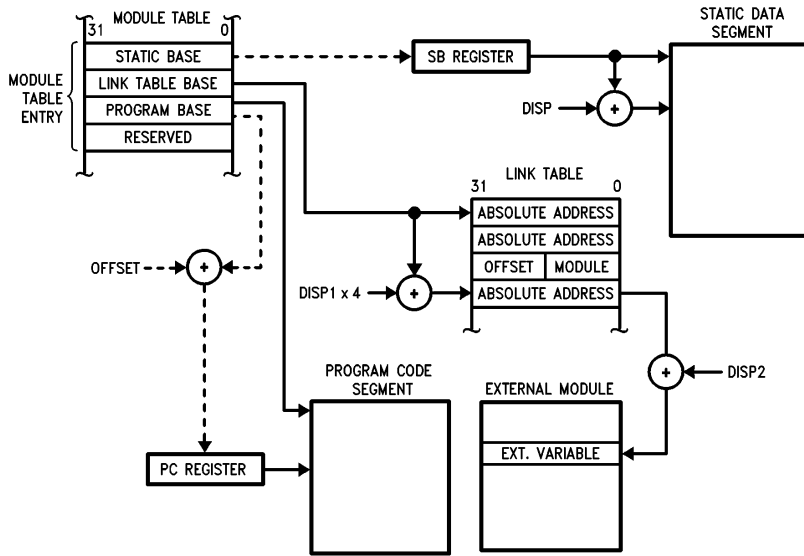
The MOD Register in the CPU contains the address of the Module Descriptor for the currently executing module.

The Static Base Register (SB) contains a copy of the Static Base entry in the Module Descriptor of the currently executing module, i.e., it points to the beginning of the current module's static data area.

This register is implemented in the CPU for efficiency purposes. By having a copy of the static base entry or chip, the CPU can avoid reading it from memory each time a data item in the static data segment is accessed.

In an NS32FX16 software environment modules need not be linked together prior to loading. As modules are loaded, a linking loader simply updates the Module Table and fills the Link Table entries with the appropriate values. No modification of a module's code is required. Thus, modules may be stored in read-only memory and may be added to a system independently of each other, without regard to their individual addressing. *Figure 2-10* shows a typical NS32FX16 run-time environment.

2.0 Architectural Description (Continued)



TL/EE/10818-1

Note: Dashed lines indicate information copied to register during transfer of control between modules.

FIGURE 2-10. NS32FX16 Run-Time Environment

2.4 INSTRUCTION SET

2.4.1 General Instruction Format


Figure 2-11 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.

Following Index Bytes come any displacements (addressing constants) or immediate values associated with the selected addressing modes. Each Disp/Imm field may contain


one of two displacements, or one immediate value. The size of a Displacement field is encoded within the top bits of that field, as shown in Figure 2-13, 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 (Section 2.2).

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 (Section 2.4.3).



TL/EE/10818-3

FIGURE 2-12. Index Byte Format



TL/EE/10818-2

FIGURE 2-11. General Instruction Format

2.0 Architectural Description (Continued)

2.4.2 Addressing Modes

The NS32FX16 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."

Addressing modes in the NS32FX16 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.

NS32FX16 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.

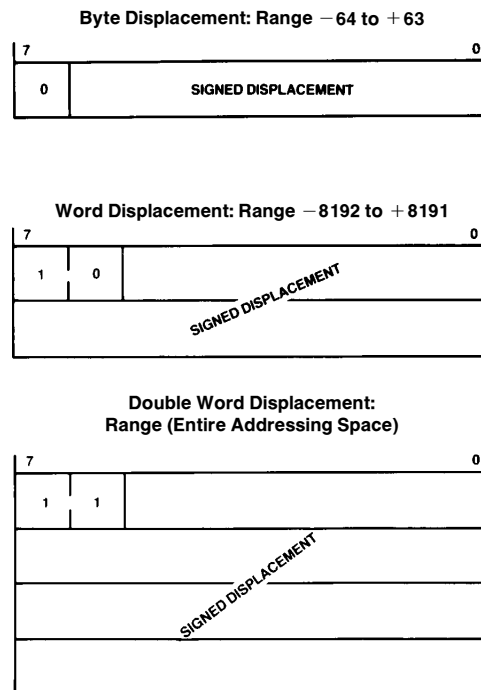
External: A pointer value is read from a specified entry of the current Link Table. To this pointer value is added a displacement, yielding the Effective Address of the operand.

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 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 Series 32000 Instruction Set Reference Manual.

In addition to the general modes, Register-Indirect with auto-increment/decrement and warps or pitch are available on several of the graphics instructions.



TL/EE/10818-4

FIGURE 2-13. Displacement Encodings

2.0 Architectural Description (Continued)

TABLE 2-1. NS32FX16 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	R6 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 Disp 1 + 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)

2.4.3 Instruction Set Summary

Table 2-2 presents a brief description of the NS32FX16 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 Series 32000 Instruction Set Reference Manual and the NS32CG16 Printer/Display Processor Programmer's Reference.

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 Processor Register: SP, SB, FP, INTBASE, MOD, PSR, US (bottom 8 PSR bits).

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

TABLE 2-2. NS32FX16 Instruction Set Summary

MOVES

Format	Operation	Operands	Description
4	MOV _i	gen,gen	Move a value.
2	MOVQ _i	short,gen	Extend and move a signed 4-bit constant.
7	MOV _M	gen,gen,disp	Move multiple: disp bytes (1 to 16).
7	MOVZ _{BW}	gen,gen	Move with zero extension.
7	MOVZ _{ID}	gen,gen	Move with zero extension.
7	MOVX _{BW}	gen,gen	Move with sign extension.
7	MOVX _{ID}	gen,gen	Move with sign extension.
4	ADDR	gen,gen	Move effective address.

INTEGER ARITHMETIC

Format	Operation	Operands	Description
4	ADD _i	gen,gen	Add.
2	ADDQ _i	short,gen	Add signed 4-bit constant.
4	ADD _C	gen,gen	Add with carry.
4	SUB _i	gen,gen	Subtract.
4	SUB _C	gen,gen	Subtract with carry (borrow).
6	NEG _i	gen,gen	Negate (2's complement).
6	ABS _i	gen,gen	Take absolute value.
7	MUL _i	gen,gen	Multiply.
7	QUO _i	gen,gen	Divide, rounding toward zero.
7	REM _i	gen,gen	Remainder from QUO.
7	DIV _i	gen,gen	Divide, rounding down.
7	MOD _i	gen,gen	Remainder from DIV (Modulus).
7	ME _i	gen,gen	Multiply to extended integer.
7	DE _i	gen,gen	Divide extended integer.

PACKED DECIMAL (BCD) ARITHMETIC

Format	Operation	Operands	Description
6	ADD _P	gen,gen	Add packed.
6	SUB _P	gen,gen	Subtract packed.

2.0 Architectural Description (Continued)

TABLE 2-2. NS32FX16 Instruction Set Summary (Continued)

INTEGER COMPARISON

Format	Operation	Operands	Description
4	CMPI	gen,gen	Compare.
2	CMPQi	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.

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.

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.

2.0 Architectural Description (Continued)

TABLE 2-2. NS32FX16 Instruction Set Summary (Continued)

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	Description
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.

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	Multiway 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)

2.0 Architectural Description (Continued)

TABLE 2-2. NS32FX16 Instruction Set Summary (Continued)

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.
12	POLYf	gen,gen	Polynomial Step.
12	DOTf	gen,gen	Dot Product.
12	SCALBf	gen,gen	Binary Scale.
12	LOGBf	gen,gen	Binary Log.

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.

GRAPHICS

Format	Operation	Operands	Description
5	BBOR	options*	Bit-aligned block transfer 'OR'.
5	BBAND	options	Bit-aligned block transfer 'AND'.
5	BBFOR		Bit-aligned block transfer fast 'OR'.
5	BBXOR	options	Bit-aligned block transfer 'XOR'.
5	BBSTOD	options	Bit-aligned block source to destination.
5	BITWT		Bit-aligned word transfer.
5	EXTBLT	options	External bit-aligned block transfer.
5	MOVMPi		Move multiple pattern.
5	TBITS	options	Test bit string.
5	SBITS		Set bit string.
5	SBITPS		Set bit perpendicular string.

BITS

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

*Note: Options are controlled by fields of the instruction, PSR status bits, or dedicated register values.

2.0 Architectural Description (Continued)

2.5 GRAPHICS SUPPORT

The following sections provide a brief description of the NS32FX16 graphics support capabilities. Basic discussions on frame buffer addressing and BITBLT operations are also provided. More detailed information on the NS32FX16 graphics support instructions can be found in the NS32CG16 Printer/Display Processor Programmer's Reference.

2.5.1 Frame Buffer Addressing

There are two basic addressing schemes for referencing pixels within the frame buffer: Linear and Cartesian (or x-y). Linear addressing associates a single number to each pixel representing the physical address of the corresponding bit in memory. Cartesian addressing associates two numbers to each pixel representing the x and y coordinates of the pixel relative to a point in the Cartesian space taken as the origin. The Cartesian space is generally defined as having the origin in the upper left. A movement to the right increases the x coordinate; a movement downward increases the y coordinate.

The correspondence between the location of a pixel in the Cartesian space and the physical (BIT) address in memory is shown in *Figure 2-14*. The origin of the Cartesian space ($x=0, y=0$) corresponds to the bit address 'ORG'. Incrementing the x coordinate increments the bit address by one. Incrementing the y coordinate increments the bit address by an amount representing the warp (or pitch) of the Cartesian space. Thus, the linear address of a pixel at location (x, y) in the Cartesian space can be found by the following expression.

$$\text{ADDR} = \text{ORG} + y * \text{WARP} + x$$

Warp is the distance (in bits) in the physical memory space between two vertically adjacent bits in the Cartesian space.

Example 1 below shows two NS32FX16 instruction sequences to set a single pixel given the x and y coordinates. Example 2 shows how to create a fat pixel by setting four adjacent bits in the Cartesian space.

Example 1: Set pixel at location (x, y)

Setup: R0 x coordinate
R1 y coordinate

Instruction Sequence 1:

```
MULD  WARP, R1      ; Y*WARP
ADDD  R0, R1        ; + X = BIT OFFSET
SBITD R1, ORG       ; SET PIXEL
```

Instruction Sequence 2:

```
INDEXD R1, (WARP-1), R0 ; Y*WARP + X
SBITD  R1, ORG          ; SET PIXEL
```

Example 2: Create fat pixel by setting bits at locations (x, y), (x+1, y), (x, y+1) and (x+1, y+1).

Setup: R0 x coordinate
R1 y coordinate


Instruction Sequence:

```
INDEXD R1, (WARP-1), R0 ; BIT ADDRESS
SBITD  41, ORG          ; SET FIRST PIXEL

ADDQD  1, R1            ; (X+1, Y)
SBITD  R1, ORG          ; SECOND PIXEL

ADDD   (WARP-1), R1     ; (X, Y+1)
SBITD  R1, ORG          ; THIRD PIXEL

ADDQD  1, R1            ; (X+1, Y+1)
SBITD  R1, ORG          ; LAST PIXEL
```



TL/EE/10818-5

FIGURE 2-14. Correspondence between Linear and Cartesian Addressing

2.5.2 BITBLT Fundamentals

BITBLT, BIT-aligned BLock Transfer, is a general operator that provides a mechanism to move an arbitrary size rectangle of an image from one part of the frame buffer to another. During the data transfer process a bitwise logical operation can be performed between the source and the destination data. BITBLT is also called RasterOp: operations on rasters. It defines two rectangular areas, source and destination, and performs a logical operation (e.g., AND, OR, XOR) between these two areas and stores the result back to the destination. It can be expressed in simple notation as:

Source op Destination → Destination
op: AND, OR, XOR, etc.

2.0 Architectural Description (Continued)

2.5.2.1 Frame Buffer Architecture

There are two basic types of frame buffer architectures: plane-oriented or pixel-oriented. BITBLT takes advantage of the plane-oriented frame buffer architecture's attribute of multiple, adjacent pixels-per-word, facilitating the movement of large blocks of data. The source and destination starting addresses are expressed as pixel addresses. The width and height of the block to be moved are expressed in terms of pixels and scan lines. The source block may start and end at any bit position of any word, and the same applies for the destination block.

2.5.2.2 Bit Alignment

Before a logical operation can be performed between the source and the destination data, the source data must first be bit aligned to the destination data. In *Figure 2-15*, the source data needs to be shifted three bits to the right in order to align the first pixel (i.e., the pixel at the top left corner) in the source data block to the first pixel in the destination data block.

2.5.2.3 Block Boundaries and Destination Masks

Each BITBLT destination scan line may start and end at any bit position in any data word. The neighboring bits (bits sharing the same word address with any words in the destination data block, but not a part of the BITBLT rectangle) of the BITBLT destination scan line must remain unchanged after the BITBLT operation.

Due to the plane-oriented frame buffer architecture, all memory operations must be word-aligned. In order to preserve the neighboring bits surrounding the BITBLT destination block, both a left mask and a right mask are needed for all the leftmost and all the rightmost data words of the destination block. The left mask and the right mask both remain the same during a BITBLT operation.

The following example illustrates the bit alignment requirements. In this example, the memory data path is 16 bits wide. *Figure 2-15* shows a 32 pixel by 32 scan line frame buffer which is organized as a long bit stream which wraps around every two words (32 bits). The origin (top left corner) of the frame buffer starts from the lowest word in memory (word address 00 (hex)).

Each word in the memory contains 16 bits, D0–D15. The least significant bit of a memory word, D0, is defined as the first displayed pixel in a word. In this example, BITBLT addresses are expressed as pixel addresses relative to the origin of the frame buffer. The source block starting address is 021 (hex) (the second pixel in the third word). The destination block starting address is 204 (hex) (the fifth pixel in the 33rd word). The block width is 13 (hex), and the height is 06 (hex) (corresponding to 6 scan lines). The shift value is 3.




FIGURE 2-15. 32-Pixel by 32-Scan Line Frame Buffer

TL/EE/10818-6

2.0 Architectural Description (Continued)




FIGURE 2-16. Overlapping BITBLT Blocks

The left mask and the right mask are 0000,1111,1111,1111 and 1111,1111,0000,0000 respectively.

Note 1: Zeros in either the left mask or the right mask indicate the destination bits which will not be modified.

Note 2: The BB(function) and EXTBLT instructions use different set up parameters, and techniques.

2.5.2.2 BITBLT Directions

A BITBLT operation moves a rectangular block of data in a frame buffer. The operation itself can be considered as a subroutine with two nested loops. The loops are preceded by setup operations. In the outer loop the source and destination starting addresses are calculated, and the test for completion is performed. In the inner loop the actual data movement for a single scan line takes place. The length of the inner loop is the number of (aligned) words spanned by each scan line. The length of the outer loop is equal to the height (number of scan lines) of the block to be moved. A skeleton of the subroutine representing the BITBLT operation follows.

BITBLT: calculate BITBLT setup parameters;
(once per BITBLT operation).
such as
width, height
bit misalignment (shift number)
left, right masks
horizontal, vertical directions
etc

-
-

OUTERLOOP: calculate source, dest addresses;
(once per scanline).

INNERLOOP: move data, (logical operation) and increment addresses;
(once per word).

UNTIL done horizontally

UNTIL done vertically

RETURN (from BITBLT).

Note: In the NS32FX16 only the setup operations must be done by the programmer. The inner and outer loops are automatically executed by the BITBLT instructions.

Each loop can be executed in one of two directions: the inner loop from left to right or right to left, the outer loop from top to bottom (down) or bottom to top (up).

The ability to move data starting from any corner of the BITBLT rectangle is necessary to avoid destroying the BITBLT source data as a result of destination writes when the source and destination are overlapped (i.e., when they share pixels). This situation is routinely encountered while panning or scrolling.

A determination of the correct execution directions of the BITBLT must be performed whenever the source and destination rectangles overlap. Any overlap will result in the destruction of source data (from a destination write) if the correct vertical direction is not used. Horizontal BITBLT direction is of concern only in certain cases of overlap, as will be explained below.

Figures 2-16(a) and (b) illustrate two cases of overlap. Here, the BITBLT rectangles are three pixels wide by five scan lines high; they overlap by a single pixel in (a) and a single column of pixels in (b). For purposes of illustration, the BITBLT is assumed to be carried out pixel-by-pixel. This convention does not affect the conclusions.

In *Figure 2-16(a)*, if the BITBLT is performed in the UP direction (bottom-to-top) one of the transfers of the bottom scan line of the source will write to the circled pixel of the destination. Due to the overlap, this pixel is also part of the uppermost scan line of the source rectangle. Thus, data needed later is destroyed. Therefore, this BITBLT must be performed in the DOWN direction. Another example of this oc-

2.0 Architectural Description (Continued)

curs any time the screen is moved in a purely vertical direction, as in scrolling text. It should be noted that, in both of these cases, the choice of horizontal BITBLT direction may be made arbitrarily.

Figure 2-16(b) demonstrates a case in which the horizontal BITBLT direction may not be chosen arbitrarily. This is an instance of purely horizontal movement of data (panning). Because the movement from source to destination involves data within the same scan line, the incorrect direction of movement will overwrite data which will be needed later. In this example, the correct direction is from right to left.

2.5.2.5 BITBLT Variations

The "classical" definition of BITBLT, as described in "Smalltalk-80 The Language and its Implementation", by Adele Goldberg and David Robson, provides for three operands: source, destination and mask/texture. This third operand is commonly used in monochrome systems to incorporate a stipple pattern into an area. These stipple patterns provide the appearance of multiple shades of gray in single-bit-per-pixel systems, in a manner similar to the "halftone" process used in printing.

Texture op1 Source op2 Destination → Destination

While the NS32FX16 and the external BPU (if used) are essentially two-operand devices, three-operand BITBLT operations can be implemented quite flexibly and efficiently by performing the two operations serially.

2.5.3 GRAPHICS SUPPORT INSTRUCTIONS

The NS32FX16 provides eleven instructions for supporting graphics oriented applications. These instructions are divided into three groups according to the operations they perform. General descriptions for each of them and the related formats are provided in the following sections.

2.5.3.1 BITBLT (BIT-aligned BLock Transfer)

This group includes seven instructions. They are used to move characters and objects into the frame buffer which will be printed or displayed. One of the instructions works in conjunction with an external BITBLT Processing Unit (BPU) to maximize performance. The other six are executed by the NS32FX16.

BIT-aligned BLock Transfer

Syntax: BB(function) Options

Setup: R0 base address, source data
 R1 base address, destination data
 R2 shift value
 R3 height (in lines)
 R4 first mask
 R5 second mask
 R6 source warp (adjusted)
 R7 destination warp (adjusted)
 0(SP) width (in words)

Function: AND, OR, XOR, FOR, STOD

Options: IA Increasing Address (default option).
 When IA is selected, scan lines are transferred in the increasing BIT/BYTE order.
 DA Decreasing Address.
 S True Source (default option).
 -S Inverted Source.

These five instructions perform standard BITBLT operations between source and destination blocks. The operations available include the following:

BBAND: src AND dst
 -src AND dst
BBOR: src OR dst
 -src OR dst
BBXOR: src XOR dst
 -src XOR dst
BBFOR: src OR dst
BBSTOD: src TO dst
 -src TO dst

'src' and '-src' stand for 'True Source' and 'Inverted Source' respectively; 'dst' stands for 'Destination'.

Note 1: For speed reasons, the BB instructions require the masks to be specified with respect to the source block. In Figure 2-15 masking was defined relative to the destination block.

Note 2: The options -S and DA are not available for the BBFOR instruction.


Note 3: BBFOR performs the same operation as BBOR with IA and S options.

Note 4: IA and DA are mutually exclusive and so are S and -S.

Note 5: The width is defined as the number of words of source data to read.

Note 6: An odd number of bytes can be specified for the source warp. However, word alignment of source scan lines will result in faster execution.

The horizontal and vertical directions of the BITBLT operations performed by the above instructions, with the exception of BBFOR, are both programmable. The horizontal direction is controlled by the IA and DA options. The vertical direction is controlled by the sign of the source and destination warps. Figure 2-17 and Table 2-3 show the format of the BB instructions and the encodings for the 'op' and 'i' fields.



- D is set when the DA option is selected
- S is set when the -S option is selected
- X is set for BBAND, and it is clear for all other BB instructions

FIGURE 2-17. BB Instructions Format

TABLE 2-3. 'op' and 'i' Field Encodings

Instruction	Options	'op' Field	'i' Field
BBAND	Yes	1010	11
BBOR	Yes	0110	01
BBXOR	Yes	1110	01
BBFOR	No	1100	01
BBSTOD	Yes	0100	01

BIT-aligned Word Transfer

Syntax: BITWT

Setup: R0 Base address, source word
 R1 Base address, destination double word
 R2 Shift value

The BITWT instruction performs a fast logical OR operation between a source word and a destination double word, stores the result into the destination double word and increments registers R0 and R1 by two. Before performing the OR operation, the source word is shifted left (i.e., in the direction of increasing bit numbers) by the value in register R2.

2.0 Architectural Description (Continued)

This instruction can be used within the inner loop of a block OR operation. Its use assumes that the source data is 'clean' and does not need masking. The BITWT format is shown in *Figure 2-18*.




FIGURE 2-18. BITWT Instruction Format

External BITBLT

Syntax: EXTBLT

Setup: R0 base addresses, source data
 R1 base address, destination data
 R2 width (in bytes)
 R3 height (in lines)
 R4 horizontal increment/decrement
 R5 temporary register (current width)
 R6 source warp (adjusted)
 R7 destination warp (adjusted)

Note 1: R0 and R1 are updated after execution to point to the last source and destination addresses plus related warps. R2, R3 and R5 will be modified. R4, R6, and R7 are returned unchanged.

Note 2: Source and destination pointers should point to word-aligned operands to maximize speed and minimize external interface logic.

This instruction performs an entire BITBLT operation in conjunction with an external BITBLT Processing Unit (BPU). The external BPU Control Register should be loaded by the software before the instruction is executed (refer to the DP8510 or DP8511 data sheets for more information on the BPU). The NS32FX16 generates a series of source read, destination read and destination write bus cycles until the entire data block has been transferred. The BITBLT operation can be performed in either horizontal direction. As controlled by the sign of the contents of register R4.

Depending on the relative alignment of the source and destination blocks, an extra source read may be required at the beginning of each scan line, to load the pipeline register in the external BPU. The L bit in the PSR register determines whether the extra source read is performed. If L is 1, no extra read is performed. The instructions CMPQB 2,1 or CMPQB 1,2 could be executed to provide the right setting for the L bit just before executing EXTBLT. *Figure 2-19* shows the EXTBLT format. The bus activity for a simple BITBLT operation is shown in *Figure 2-24*.



FIGURE 2-19. EXTBLT Instruction Format

2.5.3.2 Pattern Fill

Only one instruction is in this group. It is usually used for clearing RAM and drawing patterns and lines.

Move Multiple Pattern

Syntax: MOVMPi

Setup: R0 base address of the destination
 R1 pointer increment (in bytes)
 R2 number of pattern moves
 R3 source pattern

Note: R1 and R3 are not modified by the instruction. R2 will always be returned as zero. R0 is modified to reflect the last address into which a pattern was written.

This instruction stores the pattern in register R3 into the destination area whose address is in register R0. The pattern count is specified in register R2. After each store operation the destination address is changed by the contents of register R1. This allows the pattern to be stored in rows, in columns, and in any direction, depending on the value and sign of R1. The MOVMPi instruction format is shown in *Figure 2-20*.




FIGURE 2-20. MOVMPi Instruction Format

2.5.3.3 Data Compression, Expansion and Magnify

The three instructions in this group can be used to compress data and restore data from compression. A compressed character set may require from 30% to 50% less memory space for its storage.

The compression ratio possible can be 50:1 or higher depending on the data and algorithm used. TBITS can also be used to find boundaries of an object. As a character is needed, the data is expanded and stored in a RAM buffer. The expand instructions (SBITS, SBITPS) can also function as line drawing instructions.

Test Bit String

Syntax: TBITS option


Setup: R0 base address, source (byte address)
 R1 starting source bit offset
 R2 destination run length limited code
 R3 maximum value run length limit
 R4 maximum source bit offset

Option: 1 count set bits until a clear bit is found
 0 count clear bits until a set bit is found

Note: R0, R3 and R4 are not modified by the instruction execution. R1 reflects the new bit offset. R2 holds the result.

This instruction starts at the base address, adds a bit offset, and tests the bit for clear if "option" = 0 (and for set if "option" = 1). If clear (or set), the instruction increments to the next higher bit and tests for clear (or set). This testing for clear proceeds through memory until a set bit is found or until the maximum source bit offset or maximum run length value is reached. The total number of clear bits is stored in the destination as a run length value.

When TBITS finds a set bit and terminates, the bit offset is adjusted to reflect the current bit address. Offset is then ready for the next TBITS instruction with "option" = 0. After the instruction is executed, the F flag is set to the value of the bit previous to the bit currently being pointed to (i.e., the value of the bit on which the instruction completed execution). In the case of a starting bit offset exceeding the maximum bit offset ($R1 \geq R4$), the F flag is set if the option was 1 and clear if the option was 0. The L flag is set when the desired bit is found, or if the run length equalled the maximum run length value and the bit was not found. It is cleared otherwise. *Figure 2-21* shows the TBITS instruction format.



• S is set for 'TBITS 1' and clear for 'TBITS 0'.

FIGURE 2-21. TBITS Instruction Format

2.0 Architectural Description (Continued)

Set Bit String

Syntax: SBITS

Setup: R0 base address of the destination
 R1 starting bit offset (signed)
 R2 number of bits to set (unsigned)
 R3 address of string look-up table

Note: When the instruction terminates, the registers are returned unchanged.

SBITS sets a number of contiguous bits in memory to 1, and is typically used for data expansion operations. The instruction draws the number of ones specified by the value in R2, starting at the bit address provided by registers R0 and R1. In order to maximize speed and allow drawing of patterned lines, an external 1k byte lookup table is used. The lookup table is specified in the NS32CG16 Printer/Display Processor Programmer's Reference Supplement.

When SBITS begins executing, it compares the value in R2 with 25. If the value in R2 is less than or equal to 25, the F flag is cleared and the appropriate number of bits are set in memory. If R2 is greater than 25, the F flag is set and no other action is performed. This allows the software to use a faster algorithm to set longer strings of bits. Figure 2-22 shows the SBITS instruction format.




FIGURE 2-22. SBITS Instruction Format

Set BIT Perpendicular String

Syntax: SBITPS

Setup: R0 base address, destination (byte address)
 R1 starting bit offset
 R2 number of bits to set
 R3 destination warp (signed value, in bits)

Note: When the instruction terminates, the R0 and R3 registers are returned unchanged. R1 becomes the final bit offset. R2 is zero.

The SBITPS can be used to set a string of bits in any direction. This allows a font to be expanded with a 90 or 270 degree rotation, as may be required in a printer application. SBITPS sets a string of bits starting at the bit address specified in registers R0 and R1. The number of bits in the string is specified in R2. After the first bit is set, the destination warp is added to the bit address and the next bit is set. The process is repeated until all the bits have been set. A negative raster warp offset value leads to a 90 degree rotation. A positive raster warp value leads to a 270 degree rotation. If the R3 value is = (space warp + 1 or -1), then the result is a 45 degree line. If the R3 value is +1 or -1, a horizontal line results.

SBITS and SBITPS allow expansion on any 90 degree angle, giving portrait, landscape and mirror images from one font. Figure 2-23 shows the SBITPS instruction format.




FIGURE 2-23. SBITPS Instruction Format

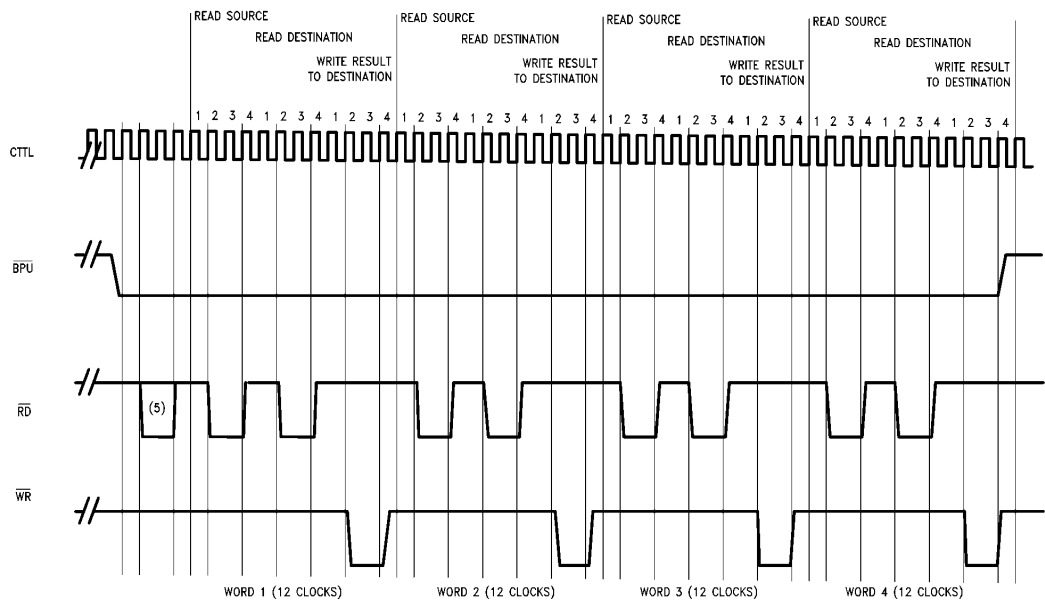


FIGURE 2-24. Bus Activity for a Simple BITBLT Operation

- Note 1:** This example is for a block 4 words wide and 1 line high.
- Note 2:** The sequence is common with all logical operations of the DP8510/DP8511 BPU.
- Note 3:** Mask values, shift values and number of bit planes do not affect the performance.
- Note 4:** Zero wait states are assumed throughout the BITBLT operation.
- Note 5:** The extra read is performed when the BPU pipeline register needs to be preloaded.

2.0 Architectural Description (Continued)

2.5.3.3.1 Magnifying Compressed Data

Restoring data is just one application of the SBITS and SBITPS instructions. Multiplying the "length" operand used by the SBITS and SBITPS instructions causes the resulting pattern to be wider, or a multiple of "length".

As the pattern of data is expanded, it can be magnified by 2x, 3x, 4x, . . . , 10x and so on. This creates several sizes of the same style of character, or changes the size of a logo. A magnify in both dimensions X and Y can be accomplished by drawing a single line, then using the MOVS (Move String) or the BB instructions to duplicate the line, maintaining an equal aspect ratio.

More information on this subject is provided in the NS32CG16 Printer/Display Processor Programmer's Reference Supplement.

3.0 Functional Description

This chapter provides details on the functional characteristics of the NS32FX16 microprocessor.

The chapter is divided into five main sections:

Instruction Execution, Exception Processing, Debugging, DSP Module and System Interface.

3.1 INSTRUCTION EXECUTION

To execute an instruction, the NS32FX16 performs the following operations:

- Fetch the Instruction
- Read Source Operands, if Any (1)
- Calculate Results
- Write Result Operands, if Any
- Modify Flags, if Necessary
- Update the Program Counter

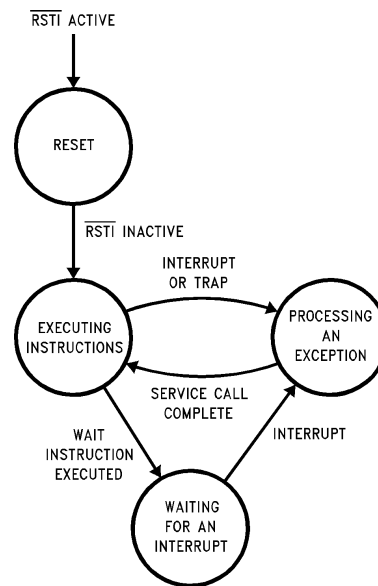
Under most circumstances, the CPU can be conceived to execute instructions by completing the operations above in strict sequence for one instruction and then beginning the sequence of operations for the next instruction. However, due to the internal instruction pipelining, as well as the occurrence of exceptions, the sequence of operations performed during the execution of an instruction may be altered. Furthermore, exceptions also break the sequentiality of the instructions executed by the CPU.

Note 1: In this and following sections, memory locations read by the CPU to calculate effective addresses for Memory-Relative and External addressing modes are considered like source operands, even if the effective address is being calculated for an operand with access class of write.

3.1.1 Operating States

The CPU has four operating states regarding the execution of instructions and the processing of exceptions: Reset, Executing Instructions, Processing An Exception and Waiting-For-An-Interrupt. The various states and transitions between them are shown in *Figure 3-1*.

Whenever the \overline{RSTI} signal is asserted, the CPU enters the reset state. The CPU remains in the reset state until the \overline{RSTI} signal is driven inactive, at which time it enters the Executing-Instructions state. In the Reset state the contents of certain registers are initialized. Refer to Section 3.5.4 for details.



TL/EE/10818-10

FIGURE 3-1. Operating States

In the Executing-Instructions state, the CPU executes instructions. It will exit this state when an exception is recognized or a WAIT instruction is encountered. At which time it enters the Processing-An-Exception state or the Waiting-For-An-Interrupt state respectively.

While in the Processing-An-Exception state, the CPU saves the PC, PSR and MOD register contents on the stack and reads the new PC and module linkage information to begin execution of the exception service procedure.

Following the completion of all data references required to process an exception, the CPU enters the Executing-Instructions state.

In the Waiting-For-An-Interrupt state, the CPU is idle. A special status identifying this state is presented on the system interface (Section 3.5). When an interrupt is detected, the CPU enters the Processing-An-Exception State.

3.1.2 Instruction Endings

The NS32FX16 checks for exceptions at various points while executing instructions. Certain exceptions, like interrupts, are in most cases recognized between instructions. Other exceptions, like Divide-By-Zero Trap, are recognized during execution of an instruction. When an exception is recognized during execution of an instruction, the instruction ends in one of four possible ways: completed, suspended, terminated, or partially completed. Each type of exception causes a particular ending, as specified in Section 3.2.

3.0 Functional Description (Continued)

3.1.2.1 Completed Instructions

When an exception is recognized after an instruction is completed, the CPU has performed all of the operations for that instruction and for all other instructions executed since the last exception occurred. Result operands have been written, flags have been modified, and the PC saved on the Interrupt Stack contains the address of the next instruction to execute. The exception service procedure can, at its conclusion, execute the RETT instruction (or the RETI instruction for maskable interrupts), and the CPU will begin executing the instruction following the completed instruction.

3.1.2.2 Suspended Instructions

An instruction is suspended when one of several trap conditions is detected during execution of the instruction. A suspended instruction has not been completed, but all other instructions executed since the last exception occurred have been completed. Result operands and flags due to be affected by the instruction may have been modified, but only modifications that allow the instruction to be executed again and completed can occur. For certain exceptions (Trap (UND)) the CPU clears the P-flag in the PSR before saving the copy that is pushed on the Interrupt Stack. The PC saved on the Interrupt Stack contains the address of the suspended instruction.

To complete a suspended instruction, the exception service procedure takes either of two actions:

1. The service procedure can simulate the suspended instruction's execution. After calculating and writing the instruction's results, the flags in the PSR copy saved on the Interrupt Stack should be modified, and the PC saved on the Interrupt Stack should be updated to point to the next instruction to execute. The service procedure can then execute the RETT instruction, and the CPU begins executing the instruction following the suspended instruction. This is the action taken when floating-point instructions are simulated by software in systems without a hardware floating-point unit.
2. The suspended instruction can be executed again after the service procedure has eliminated the trap condition that caused the instruction to be suspended. The service procedure should execute the RETT instruction at its conclusion; then the CPU begins executing the suspended instruction again. This is the action taken by a debugger when it encounters a BPT instruction that was temporarily placed in another instruction's location in order to set a breakpoint.

Note 1: It may be necessary for the exception service procedure to alter the P-flag in the PSR copy saved on the Interrupt Stack: If the exception service procedure simulates the suspended instruction and the P-flag was cleared by the CPU before saving the PSR copy, then the saved P-flag must be copied to the saved T-flag (like the floating-point instruction simulation described above). Or if the exception service procedure executes the suspended instruction again and the P-flag was not cleared by the CPU before saving the PSR copy, then the saved P-flag must be cleared (like the breakpoint trap described above). Otherwise, no alteration to the saved P-flag is necessary.

3.1.2.3 Terminated Instructions

An instruction being executed is terminated when reset occurs. Any result operands and flags due to be affected by the instruction are undefined, as is the contents of the PC.

3.1.2.4 Partially Completed Instructions

When an interrupt condition is recognized during execution of a string instruction, the instruction is said to be partially completed. A partially completed instruction has not completed, but all other instructions executed since the last exception occurred have been completed. Result operands and flags due to be affected by the instruction may have been modified, but the values stored in the string pointers and other general-purpose registers used during the instruction's execution allow the instruction to be executed again and completed.

The CPU clears the P-flag in the PSR before saving the copy that is pushed on the Interrupt Stack. The PC saved on the Interrupt Stack contains the address of the partially completed instruction. The exception service procedure can, at its conclusion, simply execute the RETT instruction (or the RETI instruction for maskable interrupts), and the CPU will resume executing the partially completed instruction.

3.1.3 Slave Processor Instructions

The NS32FX16 supports only one group of instructions, the floating-point instruction set, as being executable by a slave processor. The floating-point instruction set is validated by the F-bit in the CFG register.

If a floating-point instruction is encountered and the F-bit in the CFG register is not set, a Trap (UND) will result, without any slave processor communication attempted by the CPU. This allows software emulation in case an external floating-point unit (FPU) is not used.

3.1.3.1 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 3-2*. While applying Status Code 1111 (Broadcast ID, Section 3.5.5.1), the CPU transfers the ID Byte on the least-significant half 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.

3.0 Functional Description (Continued)

The CPU next sends the Operation Word while applying Status Code 1101 (Transfer Slave Operand, Section 3.5.5.1). 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 operands 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 for memory accesses, these extensions are not sent to the Slave Processor. The Status Code applied is 1101 (Transfer Slave Processor Operand, Section 3.5.5.1).

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.

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).

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). This word has the format shown in Figure 3-3. 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.

Status Combinations:

Send ID (ID): Code 1111

Xfer Operand (OP): Code 1101

Read Status (ST): Code 1110

Step	Status	Action
1	ID	CPU Sends ID Byte
2	OP	CPU Sends Operation Word
3	OP	CPU Sends Required Operands
4	—	Slave Starts Execution. CPU Pre-Fetches.
5	—	Slave Pulses \overline{SPC} Low
6	ST	CPU Reads Status Word. (Trap? Alter Flags?)
7	OP	CPU Reads Results (If Any).

FIGURE 3-2. Slave Processor Protocol

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).

3.1.3.2 Floating-Point Instructions
Table 3-1 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.

TABLE 3-1. 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
POLYf	read.f	read.f	f	f	f to F0	none
DOTf	read.f	read.f	f	f	f to F0	none
SCALBf	read.f	rmw.f	f	f	f to Op. 2	none
LOGBf	read.f	write.f	f	N/A	f to Op. 2	none

Notes:

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)

The Operand class columns give the Access Class for each general operand, defining how the addressing modes are interpreted (see Series 32000 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 3-3).



TL/EE/10818-11

FIGURE 3-3. Slave Processor Status Word

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.

3.2 EXCEPTION PROCESSING

Exceptions are special events that alter the sequence of instruction execution. The CPU recognizes two basic types of exceptions: interrupts and traps.

An interrupt occurs in response to an event signalled by activating the NMI or INT input signals. Interrupts are typically requested by peripheral devices that require the CPU's attention.

Traps occur 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., supervisor call instruction).

When an exception is recognized, the CPU saves the PC, PSR and the MOD register contents on the interrupt stack and then it transfers control to an exception service procedure.

Details on the operations performed in the various cases by the CPU to enter and exit the exception service procedure are given in the following sections.

It is to be noted that the reset operation is not treated here as an exception. Even though, like any exception, it alters the instruction execution sequence.

The reason being that the CPU handles reset in a significantly different way than it does for exceptions.

Refer to Section 3.5.4 for details on the reset operation.

3.2.1 Exception Acknowledge Sequence

When an exception is recognized, the CPU goes through three major steps:

- 1) Adjustment of Registers.

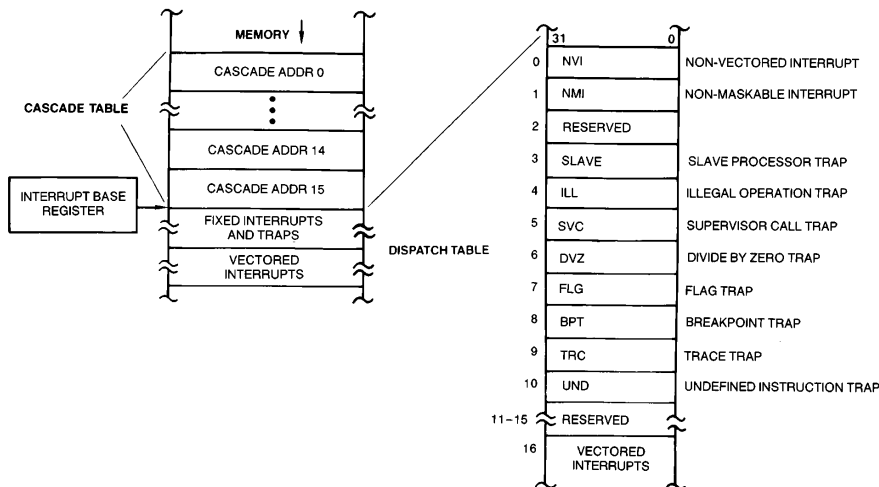
Depending on the source of the exception, the CPU may restore and/or adjust the contents of the Program Counter (PC), the Processor Status Register (PSR) and the currently-selected Stack Pointer (SP). A copy of the PSR is made, and the PSR is then set to reflect Supervisor Mode and selection of the Interrupt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

- 2) Vector Acquisition.

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

- 3) 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 3-4. 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.



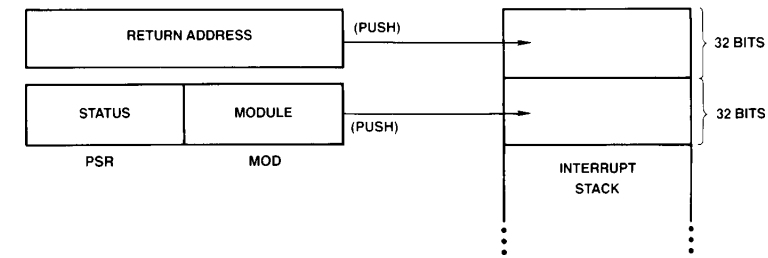
TL/EE/10818-12

FIGURE 3-4. Interrupt Dispatch and Cascade Tables

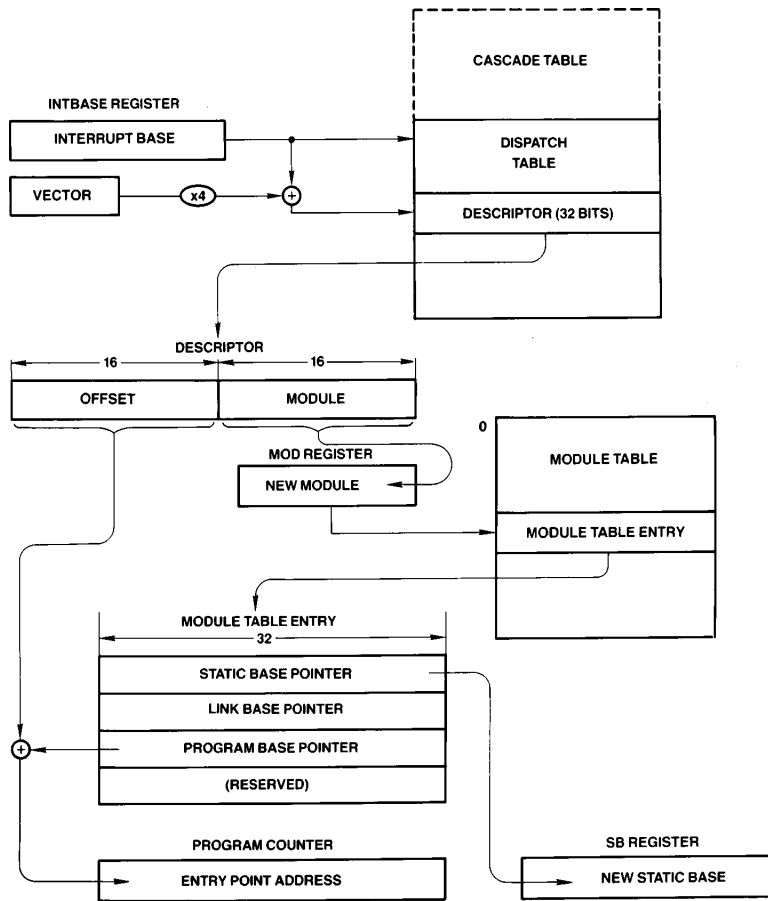
3.0 Functional Description (Continued)

This process is illustrated in *Figure 3-15*, from the viewpoint of the programmer.

Details on the sequences of events in processing interrupts and traps are given in the following sections.



TL/EE/10818-13



TL/EE/10818-14

FIGURE 3-5. Exception Acknowledge Sequence

3.0 Functional Description (Continued)

3.2.2 Returning from an Exception Service Procedure

To return control to an interrupted program, one of two instructions can be used: RETT (Return from Trap) and RETI (Return from Interrupt).

RETT is used to return from any trap or a non-maskable interrupt service procedure. Since some traps are often used deliberately as a call mechanism for supervisor mode procedures, RETT can also adjust the Stack Pointer (SP) to discard a specified number of bytes from the original stack as surplus parameter space.

RETI is used to return from a maskable interrupt service procedure. A difference of RETT, RETI also informs any external interrupt control units that interrupt service has completed. Since interrupts are generally asynchronous external events, RETI does not discard parameters from the stack.

Both of the above instructions always restore the PSR, MOD, PC and SB registers to their previous contents.

3.2.3 Maskable Interrupts

The $\overline{\text{INT}}$ 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 $\overline{\text{INT}}$ or $\overline{\text{NMI}}$ request, and is restored to its original setting upon return from the interrupt service routine via the RETT or RETI instruction.

The $\overline{\text{INT}}$ pin may be configured via the SETCFG instruction as either Non-Vectored (CFG Register bit I=0) or Vectored (bit I=1).

3.2.3.1 Non-Vectored Mode

In the Non-Vectored mode, an interrupt request on the $\overline{\text{INT}}$ 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.

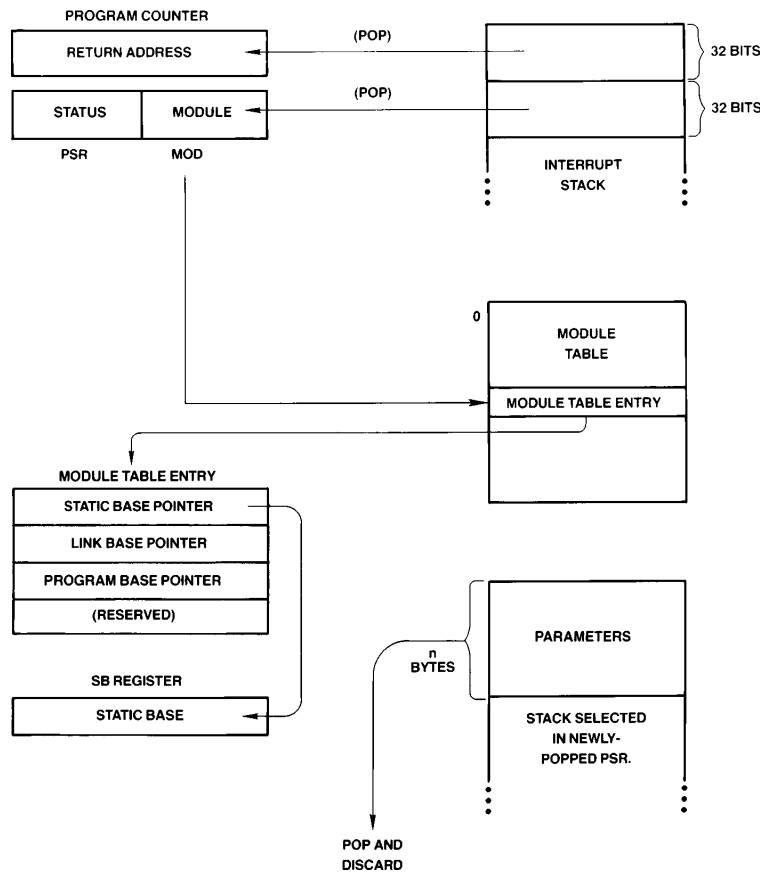
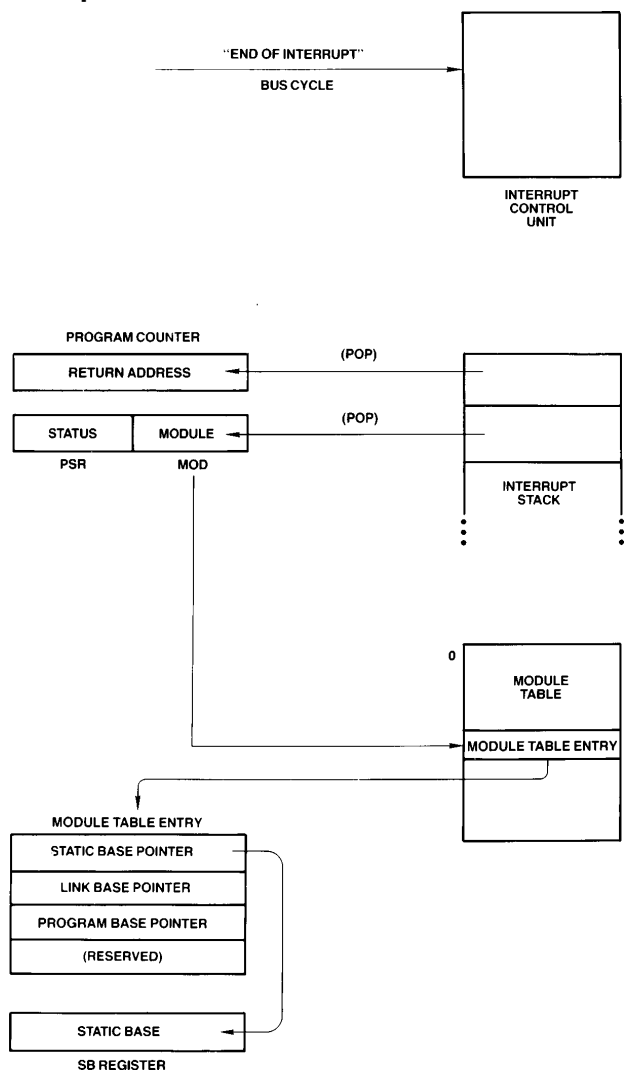


FIGURE 3-6. Return from Trap (RETT n) Instruction Flow

TL/EE/10818-15

3.0 Functional Description (Continued)



TL/EE/10818-16

FIGURE 3-7. Return from Interrupt (RETI) Instruction Flow

3.2.3.2 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 \overline{INT} pin, the CPU performs an "Interrupt Acknowledge, Master" bus cycle 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.

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).

Note: During a return from interrupt, the CPU looks at Bit 7 of the vector number from the master ICU. If Bit 7 is 0, bits 0 through 6 are ignored.

3.2.3.3 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

3.0 Functional Description (Continued)

Unit (ICU) to transparently support cascading. *Figure 3-9* 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 \overline{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.

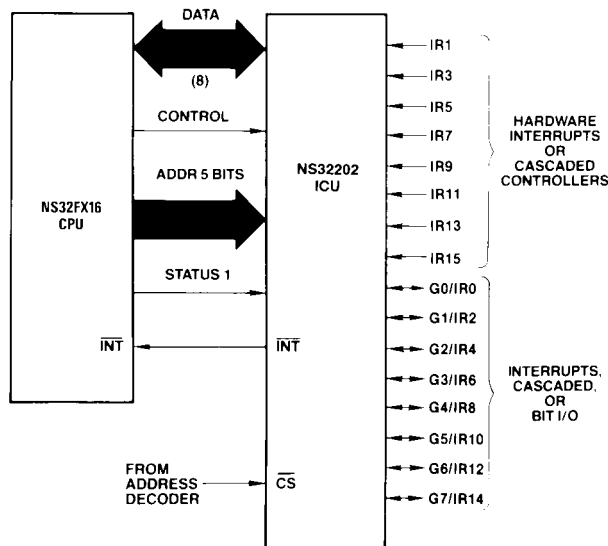
Figure 3-4 illustrates the position of the Cascade Table. To find the Cascade Table entry for a Cascaded ICU, take its Master ICU line number (0 to 15) and subtract 16 from it, giving an index in the range -16 to -1 . Multiply this value by 4, and add the resulting negative number to the contents of 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 neg-

ative 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, 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, whereupon the Master ICU again provides the negative Cascaded 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, 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 \overline{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/10818-17

FIGURE 3-8. Interrupt Control Unit Connections (16 Levels)

3.0 Functional Description (Continued)




FIGURE 3-9. Cascaded Interrupt Control Unit Connections

TL/EE/10818-18

3.0 Functional Description (Continued)

3.2.4 Non-Maskable Interrupt

The Non-Maskable Interrupt is triggered whenever a falling edge is detected on the $\overline{\text{NMI}}$ pin. The CPU performs an "Interrupt Acknowledge" bus cycle from Address FFFF00_{16} when processing of this interrupt actually begins. 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.

3.2.5 Traps

Traps are processing exceptions that are generated as direct results of the execution of an instruction.

The return address saved on the stack by any trap except Trap (TRC) is the address of the first byte of the instruction during which the trap occurred.

When a trap is recognized, maskable interrupts are not disabled.

There are 8 trap conditions recognized by the NS32FX16 as described below.

Trap (SLAVE): An exceptional condition was detected by the Floating-Point Unit during the execution of a Slave Instruction. This trap is requested via the Status Word returned as part of the Slave Processor Protocol (Section 3.1.3.1).

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 PSR F-bit.

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

Trap (TRC): The instruction just completed is being traced. Refer to Section 3.3.1 for details.

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

3.2.6 Priority among Exceptions

The CPU checks for specific exceptions at various points while executing an instruction. It is possible that several exceptions occur simultaneously. In that event, the CPU responds to the exception with highest priority.

Figure 3-10 shows an exception processing flowchart.

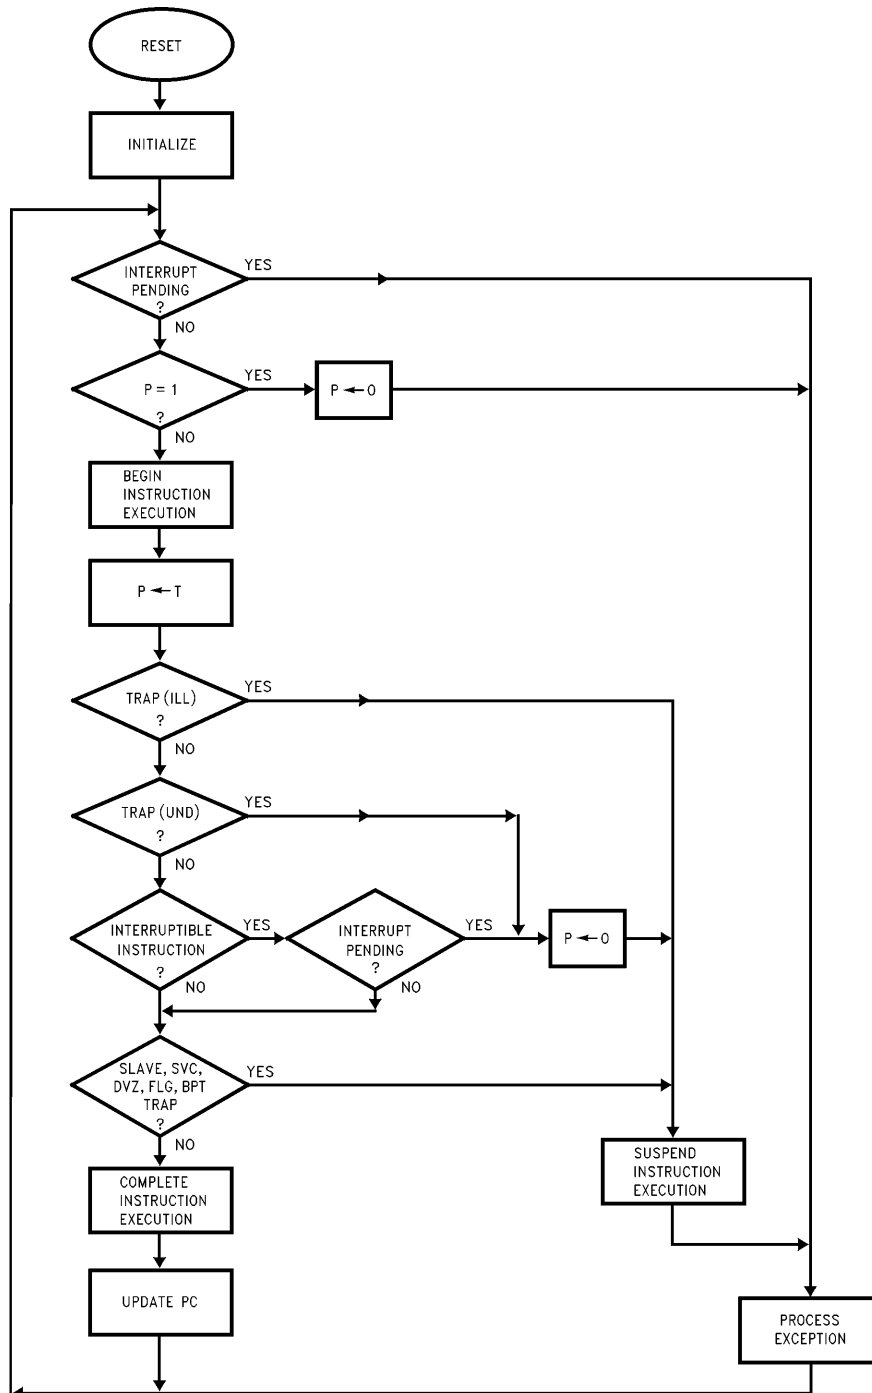
Before executing an instruction, the CPU checks for pending interrupts, or Trap (TRC). The CPU responds to any pending interrupt requests; nonmaskable interrupts are recognized with higher priority than maskable interrupts. If no interrupts are pending, then the CPU checks the P-flag in the PSR to determine whether a Trap (TRC) is pending. If the P-flag is 1, a Trap (TRC) is processed. If no interrupt or Trap (TRC) is pending, the CPU begins executing the instruction.

While executing an instruction, the CPU may recognize up to two exceptions:

1. Interrupt, if the instruction is interruptible.
2. One of 7 mutually exclusive traps: SLAVE, ILL, SVC, DVZ, FLG, BPT, UND

If no exception is detected while the instruction is executing, then the instruction is completed and the PC is updated to point to the next instruction.

3.0 Functional Description (Continued)



TL/EE/10818-19

FIGURE 3-10. Exception Processing Flowchart

3.0 Functional Description (Continued)

3.2.7 Exception Acknowledge Sequences:

Detailed Flow

For purposes of the following detailed discussion of exception acknowledge sequences, a single sequence called “service” is defined in *Figure 3-11*.

Upon detecting any interrupt request or trap condition, the CPU first performs a sequence dependent upon the type of exception. This sequence will include saving a copy of the Processor Status Register and establishing a vector and a return address. The CPU then performs the service sequence.

3.2.7.1 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, or Graphics instructions which have interior loops (BBOR, BBXOR, BBAND, BBFOR, EXTBLT, MOVMP, SBITPS, TBITS), at the next interruptible point during its execution. The graphics instructions are interruptible.

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 FFFF00_{16} , applying Status Code 0100 (Interrupt Acknowledge, Master: Section 3.4.1). 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 FFFE00_{16} , applying Status Code 0100 (Interrupt Acknowledge, Master: Section 3.4.1). Discard the byte read.
 - b. Set “Vector” to 0.
 - c. Go to Step 8.
5. Here the interrupt is Vectored. Read “Byte” from address FFFE00_{16} , applying Status Code 0100 (Interrupt Acknowledge, Master: Section 3.4.1).
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: Section 3.4.1).
8. Perform Service (Vector, Return Address), *Figure 3-11*.

3.2.7.2 SLAVE/ILL/SVC/DVZ/FLG/BPT/UND

Trap Sequence

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. If Trap (UND)

- a. Clear the Processor Status Register P Bit.
4. Copy the Processor Status Register (PSR) into a temporary register, then clear PSR bits T, U, S, and P.
5. Set “Return Address” to the address of the first byte of the trapped instruction.
6. Perform Service (Vector, Return Address), *Figure 3-11*.

3.2.7.3. 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. Set “Vector” to 9.
4. Set “Return Address” to the address of the next instruction.
5. Perform Service (Vector, Return Address), *Figure 3-11*.

Service (Vector, Return Address):

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

FIGURE 3-11. Service Sequence
Invoked during All Interrupt/Trap Sequences

3.0 Functional Description (Continued)

TABLE 3-2. Summary of Exception Processing

Exception	Instruction Ending	Cleared before Saving PSR	Cleared after Saving PSR
Interrupt	Before Instruction	None /P*	TUSPI
UND	Suspended	P	TUS
SLAVE, SVC, DVZ, FLG, BPT, ILL	Suspended	None	TUSP
TRC	Before Instruction	P	TUS

3.3 DEBUGGING SUPPORT

The NS32FX16 provides features to assist in program debugging.

Besides the Breakpoint (BPT) instruction that can be used to generate soft breaks, the CPU also provides the instruction tracing capability.

3.3.1 Instruction Tracing

Instruction tracing is a very useful feature that can be used during debugging to single-step through selected portions of a program. Tracing is enabled by setting the T-bit in the PSR Register. When enabled, the CPU generates a Trace Trap (TRC) after the execution of each instruction.

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 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.

The beginning of the execution of a TRAP(UND) is not considered to be a beginning of an instruction, and hence the T-bit is not copied into the P-bit.

Due to the fact that some instructions can clear the T- and P-bits in the PSR, in some cases a Trace Trap may not occur at the end of the instruction. This happens when one of the privileged instructions BICPSRW or LPRW PSR is executed.

In other cases, it is still possible to guarantee that a Trace Trap occurs at the end of the instruction, provided that special care is taken before returning from the Trace Trap Service Procedure. In case a BICPSRB instruction has been executed, the service procedure should make sure that the T-bit in the PSR copy saved on the Interrupt Stack is set before executing the RETT instruction to return to the program being traced. If the RETT or RETI instructions have to be traced, the Trace Trap Service Procedure should set the P- and T-bits in the PSR copy on the Interrupt Stack that is going to be restored in the execution of such instructions.

While debugging the NS32FX16 instructions which have interior loops (BBOR, BBXOR, BBAND, BBFOR, EXTBLT, MOVMP, SBITPS, TBITS), special care must be taken with the single-step trap. If an interrupt occurs during a single-step of one of the graphics instructions, the interrupt will be serviced. Upon return from the interrupt service routine, the new NS32FX16 instruction will not be re-entered, due to a single-step trap. Both the NMI and INT interrupts will cause this behavior. Another single-step operation (S command in DBG16/MONCG) will resume from where the instruction was interrupted. There are no side effects from this early termination, and the instruction will complete normally.

For all other Series 32000 instructions, a single-step operation will complete the entire instruction before trapping back to the debugger. On the instructions mentioned above, several single-step commands may be required to complete the instruction, ONLY when interrupts are occurring.

There are some methods to give the appearance of single-stepping for these NS32FX16 instructions.

1. MON16/MONCG monitors the return from single-step trap vector, PC value. If the PC has not changed since the last single-step command was issued, the single-step operation is repeated. It is also advisable to ensure that one of the NS32FX16 instructions is being single-stepped, by inspecting the first byte of the address pointed to by the PC register. If it is 0x0E, then the instruction is an NS32FX16-specific instruction.
2. A breakpoint following the instruction would also trap after the instruction had completed.

Note: If instruction tracing is enabled while the WAIT instruction is executed, the Trap (TRC) occurs after the next interrupt, when the interrupt service procedure has returned.

3.4 DSP MODULE (DSPM)

The DSP Module is specifically designed to execute high-speed vector operations on complex numbers. This is especially needed to meet the performance requirements in modern applications as well as to efficiently implement digital filters and other DSP primitives.

Detailed descriptions of the operational characteristics are provided in the following sections.

3.4.1 DSPM Operation

A block diagram of the DSP Module including the RAM array is shown in *Figure 3-12*. In order to maximize performance, an internal two-stage pipeline is provided. This allows to overlap operand fetches and multiply-accumulate operations for different vector elements.

Two data elements at a time can be fetched: one from main memory and the other from the on-chip RAM Array.

While fetching operands for one vector element, the DSPM performs the multiplication and additions on the previous vector element. Each complex multiply and accumulate operation requires two operand fetches, four multiplications and four additions. The DSPM pipeline allows a maximal throughput of a complex multiply accumulate operation in 8 clock cycles.

The DSPM uses the full bandwidth of the external bus during VCMAD, VCMUL or VCMAC operations. See Section 3.4.3.

While executing the VCMAG instruction, the external bus is free as no external operands are required. In this case the CPU can execute instructions in parallel with the DSPM.

3.0 Functional Description (Continued)

During a VCMAD, VCMUL or VCMAC instruction, external HOLD requests will be granted at the end of each memory access. Interrupt requests can only be acknowledged at the end of the DSPM instruction.

CPU accesses to any of the DSPM registers while a vector instruction is in progress are also delayed until the end of the instruction.




FIGURE 3-12. DSP Module Block Diagram

TL/EE/10818-21

3.0 Functional Description (Continued)

3.4.2 Complex Number Representation

The complex numbers processed by the DSP Module are given by pairs of 16-bit Fixed-Point values representing the real and imaginary parts of the number. The range of these values is defined by the interval $[x: -1 \leq x < 1]$.

The values are represented in 2's complement notation and the decimal point is between bits 14 and 15.

The intermediate results of vector operations are stored in temporary 32-bit registers in order to maintain full accuracy. Whenever data is transferred from the accumulator, the DSP module returns both the real and imaginary parts in 16-bit representation by rounding (to nearest) the least significant bits of the temporary registers.

An overflow is detected whenever a temporary result is outside the range given above. When an overflow is detected, the ST register OVF bit and either OP0 or OP1 is set to 1.

Complex numbers are stored in memory as two consecutive 16-bit words, with the word at the lower address representing the real part and the word at the higher address representing the imaginary part of the number.

Complex vectors consists of arrays of complex numbers stored in consecutive memory locations. Complex vectors must be aligned to double-word boundaries. *Figure 3-13* shows the memory organization of a complex vector.

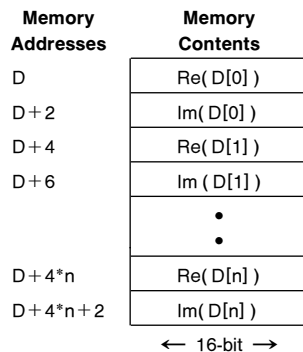


FIGURE 3-13. Memory Organization of a Complex Vector

3.4.3 DSPM Instructions

The DSP Module can execute the following vector instructions, in addition to the basic CPU load and store operations on its internal registers.

- VCMAD** Vector Complex Multiply Add
- VCMUL** Vector Complex Multiply
- VCMAC** Vector Complex Multiply Accumulate
- VCMAG** Vector Complex Magnitude

The following terms are used in the description of the operations:

- C[i]** Coefficient memory element. Entry [i] can be selected by the address generator or directly accessed by the CPU.

- D[i]** Data from external memory fetched using the address generator.
- Y** Complex Multiplier input register.
- D[i]•** The conjugate of D[i].
- A** Complex Accumulator.

Each DSP Module instruction is controlled by the CTL register OPC and OPM fields. OPC is the basic opcode, while OPM is an opcode modifier whose function is to further qualify the operation specified by OPC.

Table 3-3 provides a summary of the various vector instructions executed by the DSP module.

A DSPM instruction starts whenever the software writes into the CTL register.

Note that all the operands are complex numbers.

Thus,

$$A = \sum (C[i] \times D[i]) \text{ breaks down to:}$$

$$\text{Re}(A) = \sum (\text{Re}(C[i]) \times \text{Re}(D[i]) - \text{Im}(C[i]) \times \text{Im}(D[i]))$$

$$\text{Im}(A) = \sum (\text{Re}(C[i]) \times \text{Im}(D[i]) + \text{Im}(C[i]) \times \text{Re}(D[i]))$$

Note: The accumulator A, the multiplier input register Y, the external data pointer DPTR and the coefficient pointer CPTR registers are used as temporary registers during vector instructions. The values previously stored in these registers are destroyed.

TABLE 3-3. DSPM Instructions Summary

Instruction	OPC	OPM	Operation
VCMAD	00	0 0	$C[i] \leq C[i] + Y \times D[i]$
	00	0 1	$C[i] \leq C[i] + Y \times D[i] \bullet$
	00	1 0	$C[i] \leq Y \times D[i]$
	00	1 1	$C[i] \leq Y \times D[i] \bullet$
VCMUL	01	0 0	$C[i] \leq C[i] \times (1 + D[i])$
	01	0 1	$C[i] \leq C[i] \times (1 + D[i] \bullet)$
	01	1 0	$C[i] \leq C[i] \times D[i]$
	01	1 1	$C[i] \leq C[i] \times D[i] \bullet$
VCMAC	10	0 0	$A \leq A + \sum (C[i] \times D[i])$
	10	0 1	$A \leq A + \sum (C[i] \times D[i] \bullet)$
	10	1 0	$A \leq \sum (C[i] \times D[i])$
	10	1 1	$A \leq \sum (C[i] \times D[i] \bullet)$
VCMAG	11	0 0	$A \leq A + \sum (C[i] \times C[i])$
	11	0 1	$A \leq A + \sum (C[i] \times C[i] \bullet)$
	11	1 0	$A \leq \sum (C[i] \times C[i])$
	11	1 1	$A \leq \sum (C[i] \times C[i] \bullet)$

3.4.4 Circular Buffers

The DSP Module accesses arrays of data in external memory using the DPTR register as an address pointer. The DS field in the CTL register controls the size of the array. The DSPM handles the data in the external array in a circular fashion. Only the appropriate number of least significant address bits in the DPTR register are incremented after each memory access. The upper bits remain unchanged. Table 3-4 shows which bits are incremented for various buffer sizes.

3.0 Functional Description (Continued)

TABLE 3-4. Circular Buffer Sizes

DS Field	External Buffer Size (DW)	Constant Address bits	Incremented Address bits
00	8	A0, A5–A23	A1–A4
01	16	A0, A6–A23	A1–A5
10	32	A0, A7–A23	A1–A6
11	64	A0, A8–A23	A1–A7

3.5 SYSTEM INTERFACE

This section provides general information on the NS32FX16 interface to the external world. Descriptions of the CPU requirements as well as the various bus characteristics are provided here. Details on other device characteristics including timing are given in Chapter 4.

3.5.1 Power and Grounding

The NS32FX16 requires a single 5V power supply, applied on the V_{CC} pins. These pins should be connected together by a power (V_{CC}) plane on the printed circuit board.

The grounding connections are made on the GND pins. These pins should be connected together by a ground (GND) plane on the printed circuit board.

Both power and ground connections are shown in *Figure 3-14*.

For optimal noise immunity, the power and ground pins should be connected to V_{CC} and ground planes respectively. If V_{CC} and ground planes are not used, single conductors should be run directly from each V_{CC} pin to a power point, and from each GND pin to a ground point. Daisy-chained connections should be avoided.

Decoupling capacitors should also be used to keep the noise level to a minimum. Standard 0.1 μF ceramic capacitors can be used for this purpose. They should attach to V_{CC} , GND pins as close as possible to the NS32FX16.


During prototype using wire-wrap or similar methods, the capacitors should be soldered directly to the power pins of the NS32FX16 socket, or as close as possible, with very short leads.

Design Notes

When constructing a board using high frequency clocks with multiple lines switching, special care should be taken to avoid resonances on signal lines. A separate power and ground layer is recommended. This is true when designing boards for the NS32FX16. Switching times of under 5 ns on some lines are possible. Resonant frequencies should be maintained well above the 200 MHz frequency range on signal paths by keeping traces short and inductance low. Loading capacitance at the end of a transmission line contributes to the resonant frequency and should be minimized if possible. Capacitors should be located as close as possible across each power and ground pair near the NS32FX16. Power and ground connections are shown in *Figure 3-14*.

3.5.2 Clocking

The NS32FX16 provides an internal oscillator that interacts with an external clock source through two signals; OSCIN and OSCOUT.




TL/EE/10818-22

FIGURE 3-14. Power and Ground Connections


Either an external single-phase clock signal or a crystal can be used as the clock source. If a single-phase clock source is used, only the connection on OSCIN is required; OSCOUT should be left unconnected or loaded with no more than 5 pF of stray capacitance. The voltage level requirements specified in Section 4.3 must also be met for proper operation.

When operation with a crystal is desired, special care should be taken to minimize stray capacitances and inductances. The crystal, as well as the external components, should be placed in close proximity to the OSCIN and OSCOUT pins to keep the printed circuit trace lengths to an absolute minimum. *Figure 3-15* and *3-16* show the external crystal interconnections. Table 3-5 provides the crystal characteristics and the values of the R, C, and L components, including stray capacitance, required for various frequencies.



TL/EE/10818-23

FIGURE 3-15. Crystal Interconnections—30 MHz



TL/EE/10818-24

FIGURE 3-16. Crystal Interconnections, 40 MHz, 50 MHz

3.0 Functional Description (Continued)




FIGURE 3-17. Recommended Reset Connections

TL/EE/10818-25

TABLE 3-5. External Oscillator Specifications Crystal Characteristics

Type	AT-Cut
Tolerance	0.005% at +25°C
Stability	0.01% from 0°C to +70°C
Resonance	
30 MHz:	Fundamental (Parallel)
40 MHz or 50 MHz:	Third Overtone (Parallel)
Maximum Series Resistance	50Ω
Maximum Shunt Capacitance	7 pF

R, C and L Values

Frequency (MHz)	R1 (kΩ)	R2 (Ω)	C1 (pF)	C2 (pF)	C3 (pF)	L (μH)
30	180	51	20	20		
30	180	51	20	20	800-1300	3.3
40	150	51	20	20	800-1300	1.8
50	150	51	20	20	800-1300	1.1

3.5.3 Power Save Mode

The NS32FX16 provides a power save feature that can be used to significantly reduce the power consumption at times when the computational demand decreases. The device uses the clock signal at the OSCIN pin to derive the internal clock as well as the external signals CTTL and FCLK. The frequency of these clock signals is affected by the clock scaling factor. Scaling factors of 1, 2, 4, or 8 can be selected by properly setting the C- and M-bits in the CFG register. The power save mode should not be used to reduce the clock frequency below the minimum frequency required by the CPU.

Upon reset, both C and M are set to zero, thus maximum clock rate is selected.

Due to the fact that the C- and M-bits are programmed by the SETCFG instruction, the power save feature can only be controlled by programs running in supervisor mode.

The following table shows the C- and M-bit settings for the various scaling factors, and the resulting supply current for a crystal frequency of 50 MHz.

Clock Scaling Factor vs Supply Current


C	M	Scaling Factor	CPU Clock Frequency	Typical I _{CC} at +5V
0	0	1	25 MHz	170 mA
0	1	2	12.5 MHz	100 mA
1	0	4	6.25 MHz	65 mA
1	1	8	3.13 MHz	45 mA

3.5.4 Resetting

The \overline{RSTI} input pin is used to reset the NS32FX16. The CPU samples \overline{RSTI} on the falling edge of CTTL.


Whenever a low level is detected, the CPU responds immediately. Any instruction being executed is terminated; any results that have not yet been written to memory are discarded; and any pending interrupts and traps are eliminated. The internal latch for the edge-sensitive \overline{NMI} signal is cleared. The DSP module ST register is set to 0.

On application of power, \overline{RSTI} must be held low for at least 50 μs 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 active for not less than 64 CTTL cycles. See Figures 3-18 and 3-19.



TL/EE/10818-26

FIGURE 3-18. Power-On Reset Requirements



TL/EE/10818-27

FIGURE 3-19. General Reset Timing

While in the Reset state, the CPU drives the signals \overline{ADS} , \overline{IAS} , \overline{RD} , \overline{WR} , \overline{DBE} , \overline{TSO} , \overline{BPU} , and \overline{DDIN} inactive. $\overline{AD0-AD15}$, $\overline{A16-A23}$ and \overline{SPC} are floated, \overline{ALE} is HIGH and the state of all other output signals is undefined.

The internal CPU clock and CTTL run at half the frequency of the signal on the OSCIN pin.

The \overline{HOLD} signal must be kept inactive. After the \overline{RSTI} signal is driven high, the CPU will stay in the reset condition for approximately 8 clock cycles and then it will begin execution at address 0.

The PSR is reset to 0. The CFG C- and M-bits are reset to 0. FCLK runs at the same frequency as OSCIN. \overline{NMI} is enabled to allow Non-Maskable Interrupts. The following conditions are present after reset due to the PSR being reset to 0:

3.0 Functional Description (Continued)

Tracing is disabled.

Supervisor mode is enabled.

Supervisor stack space is used when the TOS addressing mode is indicated.

No trace traps are pending.

Only $\overline{\text{NMI}}$ is enabled. Maskable interrupts are disabled.

$\overline{\text{BPU}}$ is inactive high.

The Clock Scaling Factor is set to 1, refer to Section 3.5.3.

Note that vector/non-vectorized interrupts have not been selected. While interrupts are disabled, a SETCFG [I] instruction must be executed to enable vectored interrupts. If non-vectorized interrupts are required, a SETCFG without the [I] must be executed.

The presence/absence of the NS32081, NS32181, or NS32381 has also not been declared. If there is a Floating-Point Unit, a SETCFG [F] instruction must be executed. If there is no floating-point unit, a SETCFG without the [F] must be executed.

In general, a SETCFG instruction must be executed in the reset routine, in order to properly configure the CPU. The options should be combined, and executed in a single instruction. For example, to declare vectored interrupts, a Floating-Point unit installed, and full CPU clock rate, execute a SETCFG [F, I] instruction. To declare non-vectorized interrupts, no FPU, and full CPU clock rate, execute a SETCFG [] instruction.

3.5.5 Bus Cycles

The NS32FX16 will perform bus cycles for one of the following reasons:

1. To fetch instructions from memory.
2. To write or read data to or from memory or external peripheral devices.
3. To acknowledge an interrupt, or to acknowledge completion of an interrupt service routine.
4. To notify external logic of any accesses to the on-chip peripheral device registers or internal RAM.
5. To transfer information to or from a Slave Processor.

3.5.5.1 Bus Status

The NS32FX16 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, they why it is idle.

The Bus Status pins are interpreted as a 4-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 — DSP Module Data Transfer.
- 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 an interrupt request. See Section 3.2.3.

- 0101 — Interrupt Acknowledge, Cascaded.

The CPU is reading an interrupt vector to acknowledge a maskable interrupt request from a Cascaded Interrupt Control Unit.

- 0110 — End of Interrupt, Master.

The CPU is performing a Read cycle to indicate that it is executing a Return from Interrupt (RETI) instruction at the completion of an interrupt's service procedure.

- 0111 — End of Interrupt, Cascaded.

The CPU is performing a read cycle from a Cascaded Interrupt Control Unit to indicate that it is executing a Return from Interrupt (RETI) instruction at the completion of an interrupt's service procedure.

- 1000 — Sequential Instruction Fetch.

The CPU is reading the next sequential word from the instruction stream into the Instruction 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, 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. The write cycle of RMW will have a "write" status.

- 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.

- 1110 — Read Slave Processor Status.

The CPU is reading a Status Word from a Slave Processor after the Slave Processor has signalled completion of an instruction.

- 1111 — Broadcast Slave ID.

The CPU is initiating the execution of a Slave Processor instruction by transferring the first byte of the instruction, which represents the slave processor identification.

3.5.5.2 Basic Read and Write Cycles

The sequence of events occurring during a CPU access to either memory or peripheral device is shown in *Figure 3-21* for a read cycle, and *Figure 3-22* for a write cycle.

3.0 Functional Description (Continued)

The cases shown assume that the selected memory or peripheral device is capable of communicating with the CPU at full speed. If not, then cycle extension may be requested through $\overline{\text{CWAIT}}$ and/or $\overline{\text{WAIT1-2}}$.

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

During T1, the CPU applies an address on pins AD0-AD15 and A16-A23 and provides a low-going pulse on the $\overline{\text{ADS}}$ 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 0-15 from the AD0-AD15 pins. It also deasserts the ALE signal, which eliminates the need to invert $\overline{\text{ADS}}$ to generate the strobe for the address latches. See Figure 3-20. During this time also the status signals $\overline{\text{DDIN}}$, indicating the direction of the transfer, and $\overline{\text{HBE}}$, indicating whether the high byte (AD8-AD15) is to be referenced, become valid.

During T2 the CPU switches the Data Bus, AD0-AD15, to either accept or present data. Note that the signals A16-A23 remain valid, and need not be latched.




FIGURE 3-20. Bus Connections

TL/EE/10818-28

3.0 Functional Description (Continued)

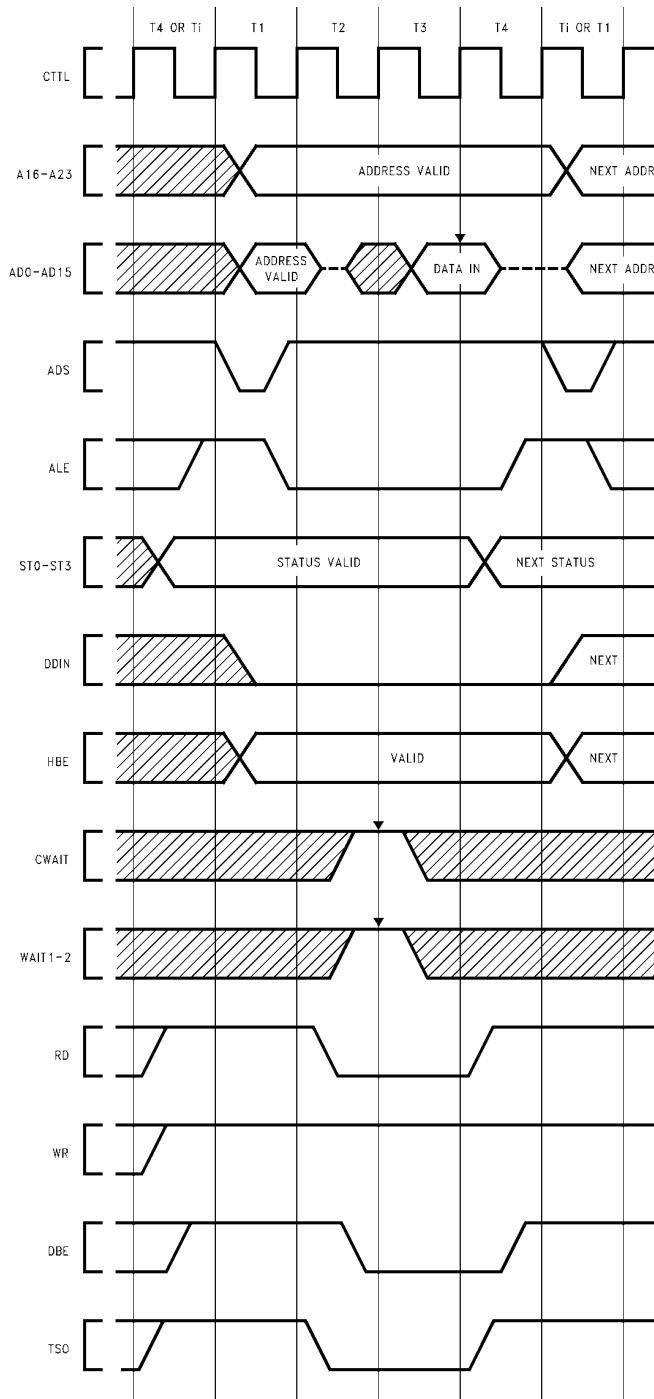
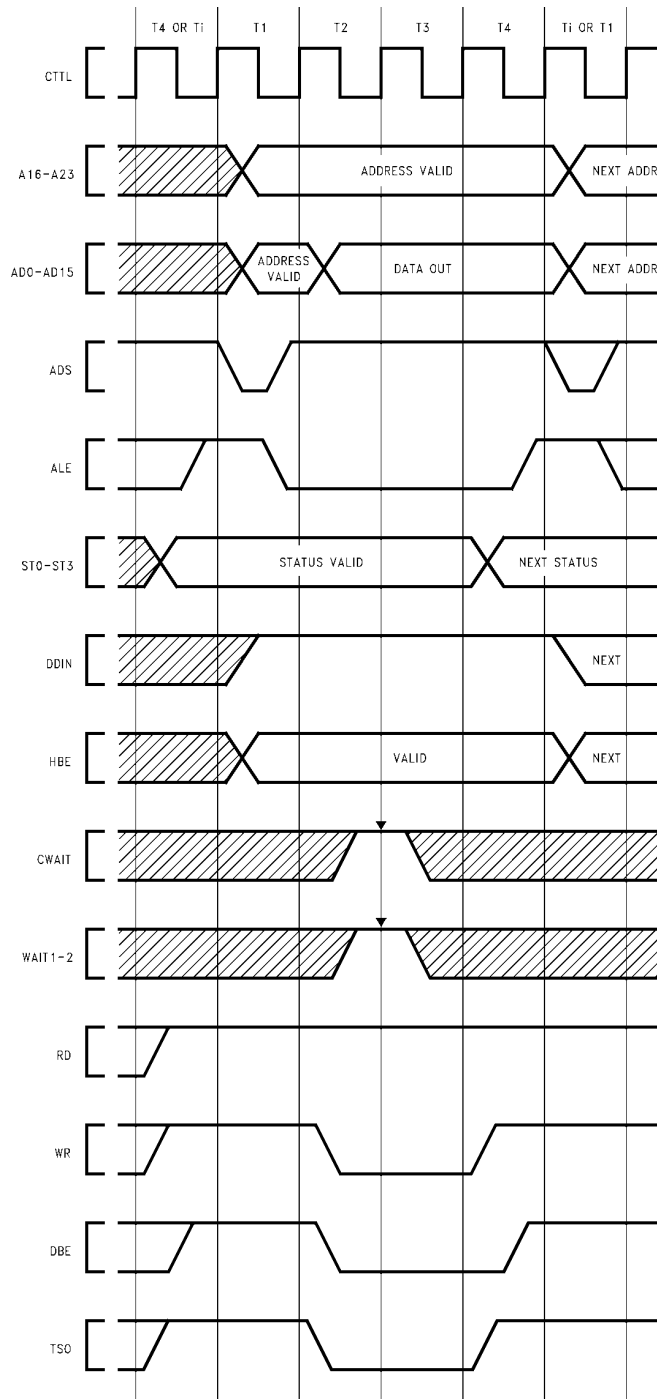


FIGURE 3-21. Read Cycle Timing

TL/EE/10818-29

3.0 Functional Description (Continued)



TL/EE/10818-30

FIGURE 3-22. Write Cycle Timing

3.0 Functional Description (Continued)

At this time the signals $\overline{\text{TSO}}$ (Timing State Output), $\overline{\text{DBE}}$ (Data Buffer Enable) and either $\overline{\text{RD}}$ (Read Strobe) or $\overline{\text{WR}}$ (Write Strobe) will also be activated.

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 rising edge of CTTL, the $\overline{\text{CWAIT}}$ and $\overline{\text{WAIT1-2}}$ signals are sampled to determine whether the bus cycle will be extended. See Section 3.5.5.3.

If the CPU is performing a read cycle, the data bus (AD0–AD15) is sampled at the beginning of T4 on the rising edge of CTTL. Data must, however, be held a little longer to meet the data hold time requirements. The $\overline{\text{RD}}$ signal is guaranteed not to go inactive before this time, so its rising edge can be safely used to disable the device providing the input data.

The T4 state finishes the bus cycle. At the beginning of T4, the $\overline{\text{RD}}$ or $\overline{\text{WR}}$, and $\overline{\text{TSO}}$ signals go inactive, and on the falling edge of CTTL, $\overline{\text{DBE}}$ 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).

3.5.5.3 Cycle Extension

To allow sufficient access time for any speed of memory or peripheral device, the NS32FX16 provides for extension of a bus cycle. Any type of bus cycle except a Slave Processor cycle and a special bus cycle can be extended.

In *Figures 3-21* and *3-22*, note that during T3 all bus control signals from the CPU are flat. Therefore, a bus cycle can be cleanly extended by causing the T3 state to be repeated. This is the purpose of the $\overline{\text{WAIT1-2}}$ and $\overline{\text{CWAIT}}$ input signals.

At the end of state T2, on the rising edge of CTTL, $\overline{\text{WAIT1-2}}$ and $\overline{\text{CWAIT}}$ are sampled.

If any of these signals are active, the bus cycle will be extended by at least one clock cycle. Thus, one or more addi-

tional T3 state (also called wait state) will be inserted after the next T-State. Any combination of the above signals can be activated at one time. However, the $\overline{\text{WAIT1-2}}$ inputs are only sampled by the CPU at the end of state T2. They are ignored at all other times.

The $\overline{\text{WAIT1-2}}$ inputs are binary weighted, and can be used to insert up to 3 wait states, according to the following table.

$\overline{\text{WAIT2}}$	$\overline{\text{WAIT1}}$	Number of Wait States
HIGH	HIGH	0
HIGH	LOW	1
LOW	HIGH	2
LOW	LOW	3

$\overline{\text{CWAIT}}$ causes wait states to be inserted continuously as long as it is sampled active. It is normally used when the number of wait states to be inserted in the CPU bus cycle is not known in advance.

The following sequence shows the CPU response to the $\overline{\text{WAIT1-2}}$ and $\overline{\text{CWAIT}}$ inputs.

1. Start bus cycle.
2. Sample $\overline{\text{WAIT1-2}}$ and $\overline{\text{CWAIT}}$ at the end of state T2.
3. If the $\overline{\text{WAIT1-2}}$ inputs are both inactive, then go to step 6.
4. Insert the number of wait states selected by $\overline{\text{WAIT1-2}}$.
5. Sample $\overline{\text{CWAIT}}$ again.
6. If $\overline{\text{CWAIT}}$ is not active, then go to step 8.
7. Insert one wait state and then go to step 5.
8. Complete bus cycle.

Figure 3-23 shows a bus cycle extended by three wait states, two of which are due to $\overline{\text{WAIT2}}$, and one is due to $\overline{\text{CWAIT}}$.

3.0 Functional Description (Continued)

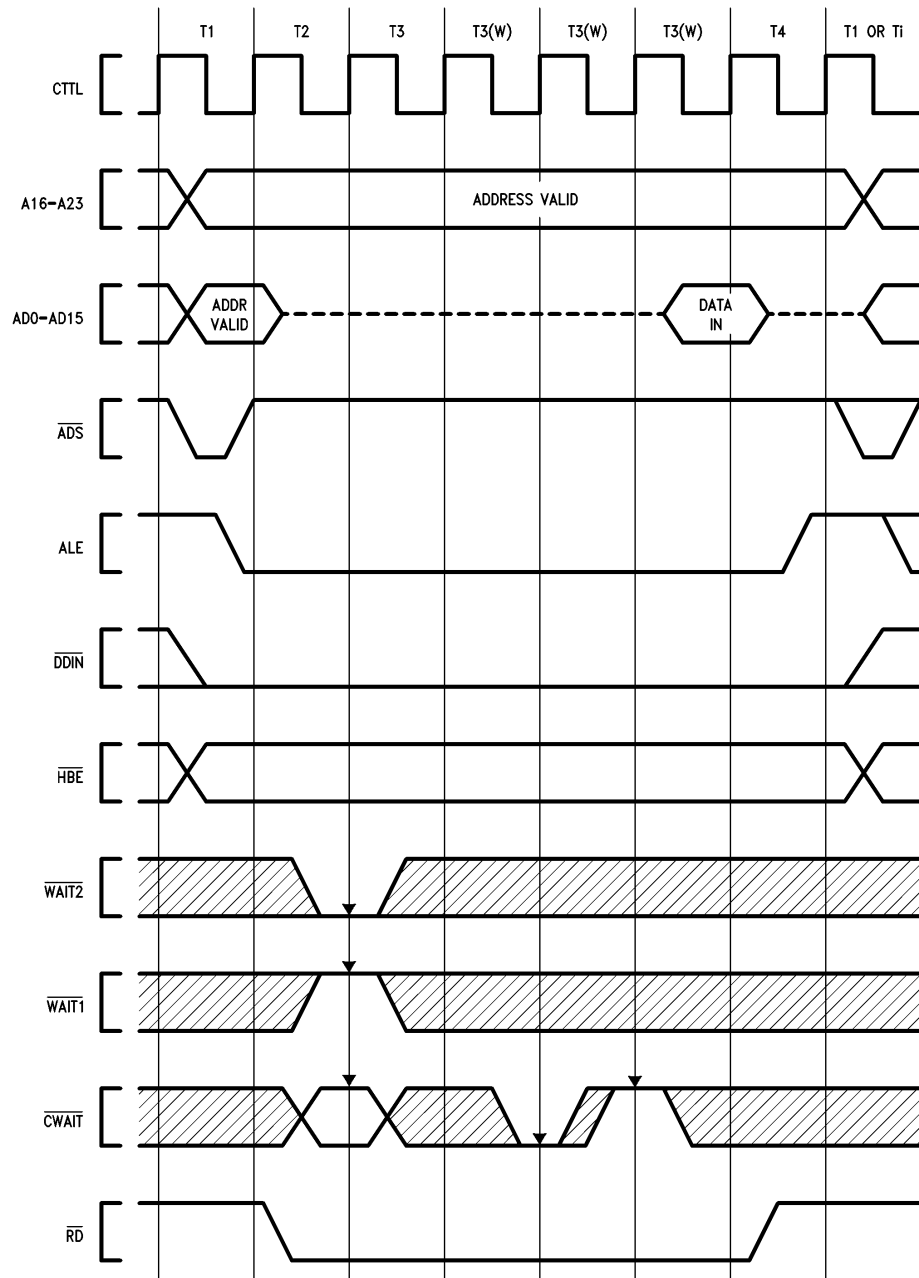


FIGURE 3-23. Cycle Extension of a Read Cycle

TL/EE/10818-31

3.0 Functional Description (Continued)

3.5.5.4 Instruction Fetch Cycles

Instructions for the NS32FX16 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 (Section 3.5.5.1).

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 Even Word Read cycles (Table 3-7).

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 is either an Even Word Read or an Odd Byte Read, depending on whether the destination address is even or odd.

3.5.5.5 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.

Table 3-6 shows the Interrupt Control sequences associated with each interrupt and with the return from its service routine. For full details of the NS32FX16 interrupt structure, see Section 3.2.

3.0 Functional Description (Continued)

TABLE 3-6. Interrupt Sequences

Cycle	Status	Address	\overline{DDIN}	\overline{HBE}	A0	High Bus	Low Bus
A. Non-Maskable Interrupt Control Sequence							
Interrupt Acknowledge							
1	0100	FFFF00 ₁₆	0	1	0	Don't Care	Don't Care
Interrupt Return							
None: Performed through Return from Trap (RETT) instruction.							
B. Non-Vectored Interrupt Control Sequence							
Interrupt Acknowledge							
1	0100	FFFE00 ₁₆	0	1	0	Don't Care	Don't Care
Interrupt Return							
None: Performed through Return from Trap (RETT) instruction.							
C. Vectored Interrupt Sequence: Non-Cascaded							
Interrupt Acknowledge							
1	0100	FFFE00 ₁₆	0	1	0	Don't Care	Vector: Range: 0–127
Interrupt Return							
1	0110	FFFE00 ₁₆	0	1	0	Don't Care	Vector: Same as in Previous Int. Ack. Cycle
D. Vectored Interrupt Sequence: Cascaded							
Interrupt Acknowledge							
1	0100	FFFE00 ₁₆	0	1	0	Don't Care	Cascade Index: range – 16 to – 1
(The CPU here uses the Cascade Index to find the Cascade Address.)							
2	0101	Cascade Address	0	1 or 0*	0 or 1*	Vector, range 0–255; on appropriate half or Data Bus for even/odd address	
Interrupt Return							
1	0110	FFFE00 ₁₆	0	1	0	Don't Care	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	1 or 0*	0 or 1*	Don't Care	Don't Care

* If the Cascaded ICU Address is Even (A0 is low), then the CPU applies \overline{HBE} high and reads the vector number from bits 0–7 of the Data Bus.

If the address is Odd (A0 is high), then the CPU applies \overline{HBE} low and reads the vector number from bits 8–15 of the Data Bus. The vector number may be in the range 0–225.

3.0 Functional Description (Continued)

3.5.5.6. Special Bus Cycles

Special bus cycles are performed during CPU accesses to the DSP Module (DSPM) registers or internal RAM. These cycles may be used by external logic to track CPU activities involving on-chip bus transactions.

A special bus cycle starts with the assertion of the special output signal \overline{IAS} . The ALE signal stays high during the entire cycle, and the signals \overline{ADS} , \overline{TSO} , \overline{DBE} , \overline{RD} and \overline{WR} are not activated. \overline{CWAIT} and $\overline{WAIT1-2}$ are ignored.

A CPU access to a DSP Module register or internal RAM occurring while a vector operation is being executed, is delayed until the end of the vector operation. This delay cannot be observed externally.

The CPU drives the data bus with the same data that is being written into the on-chip register or RAM during a special write cycle, and ignores the data placed on the data bus during a special read cycle. The 24 least significant address bits of the DSPM register being accessed are output on the AD0-AD15 and A16-A23 signals. *Figure 3-24.* shows the timing for special read and write cycles.

3.5.5.7 Slave Processor Bus Cycles

A Slave Processor bus cycle always takes exactly two clock cycles, labeled T1 and T4 (see *Figures 3-25* and *3-26*). During a Read cycle \overline{SPC} is active from the beginning of T1 to the beginning of T4, and the data is sampled at the end of T1. The Cycle Status pins lead the cycle by one clock period, and are sampled on the leading edge of \overline{SPC} . During a




FIGURE 3-24. Special Bus Cycle Timing

TL/EE/10818-32

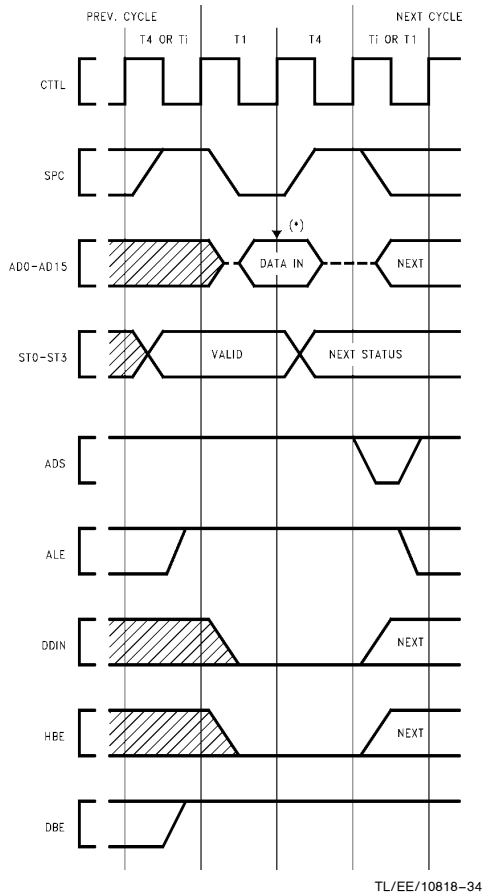
3.0 Functional Description (Continued)

Write cycle, the CPU applies data and activates \overline{SPC} at T1, removing \overline{SPC} at T4. The Slave Processor latches the status on the leading edge of \overline{SPC} and latches data on the trailing edge.

The CPU does not pulse the Address Strobe (\overline{ADS}), and no bus signals are generated. 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.


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 the entire bus. 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.

Figure 3-27 shows the NS32FX16 and FPU connection diagram.



Note: CPU samples Data Bus here.

FIGURE 3-25. Slave Processor Read Cycle



TL/EE/10818-35

*Note: Slave Processor samples Data Bus here.

FIGURE 3-26. Slave Processor Write Cycle


3.5.5.8 Data Access Sequences

The 24-bit address provided by the NS32FX16 is a byte address; that is, it uniquely identifies one of up to 16,777,216 8-bit memory locations. An important feature of the NS32FX16 is that the presence of a 16-bit data bus imposes no restrictions on data alignment; any data item, regardless of size, may be placed starting at any memory address. The NS32FX16 provides a special control signal, High Byte Enable (\overline{HBE}), which facilitates individual byte addressing on a 16-bit bus.


Memory is organized as two 8-bit banks, each bank receiving the word address (A1-A23) in parallel. One bank, connected to Data Bus pins AD0-AD7, is enabled to respond to even byte addresses; i.e., when the least significant address bit (A0) is low. The other bank, connected to Data Bus pins AD8-AD15, is enabled when \overline{HBE} is low. See Figure 3-28.

Any bus cycle falls into one of three categories: Even Byte Access, Odd Byte Access, and Even Word Access. All accesses to any data type are made up of sequences of these cycles. Table 3-7 gives the state of A0 and \overline{HBE} for each category.

3.0 Functional Description (Continued)



TL/EE/10818-33
FIGURE 3-27. NS32FX16 and FPU Interconnections



TL/EE/10818-36
FIGURE 3-28. Memory Interface

TABLE 3-7. Bus Cycle Categories

Category	HBE	A0
Even Byte	1	0
Odd Byte	0	1
Even Word	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 (i.e., whether it starts on an even byte address or an odd byte address). Table 3-8 lists the bus cycles performed for each situation. For the timing of A0 and HBE, see Section 3.5.5.2.

3.0 Functional Description (Continued)

TABLE 3-8. Data Access Sequences

Cycle	Type	Address	HBE	A0	High Bus	Low Bus						
A. Odd Word Access Sequence												
					Byte 1	Byte 0						
1	Odd Byte	A	0	1	Byte 0	Don't Care						
2	Even Byte	A + 1	1	0	Don't Care	Byte 1						
B. Even Double-Word Access Sequence												
					Byte 3	Byte 2	Byte 1	Byte 0				
1	Even Word	A	0	0	Byte 1	Byte 0						
1	Even Word	A + 2	0	0	Byte 3	Byte 2						
C. Odd Double-Word Access Sequence												
					Byte 3	Byte 2	Byte 1	Byte 0				
1	Odd Byte	A	0	1	Byte 0	Don't Care						
2	Even Word	A + 1	0	0	Byte 2	Byte 1						
3	Even Byte	A + 3	1	0	Don't Care	Byte 3						
D. Even Quad-Word Access Sequence												
					Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0
1	Even Word	A	0	0	Byte 1	Byte 0						
2	Even Word	A + 2	0	0	Byte 3	Byte 2						
Other Bus Cycles (Instruction Prefetch or Slave) can occur here.												
3	Even Word	A + 4	0	0	Byte 5	Byte 4						
4	Even Word	A + 6	0	0	Byte 7	Byte 6						
E. Odd Quad-Word Access Sequence												
					Byte 7	Byte 6	Byte 5	Byte 4	Byte 3	Byte 2	Byte 1	Byte 0
1	Odd Byte	A	0	1	Byte 0	Don't Care						
2	Even Word	A + 1	0	0	Byte 2	Byte 1						
3	Even Byte	A + 3	1	0	Don't Care	Byte 3						
Other Bus Cycles (Instruction Prefetch or Slave) can occur here.												
4	Odd Byte	A + 4	0	1	Byte 4	Don't Care						
5	Even Word	A + 5	0	0	Byte 6	Byte 5						
6	Even Byte	A + 7	1	0	Don't Care	Byte 7						

3.0 Functional Description (Continued)

3.5.5.9 Bus Access Control

The NS32FX16 CPU has the capability of relinquishing its control of the bus upon request from a DMA controller or another CPU. This capability is implemented by means of the $\overline{\text{HOLD}}$ (Hold Request) and $\overline{\text{HLD\bar{A}}}$ (Hold Acknowledge) pins. By asserting $\overline{\text{HOLD}}$ low, an external device requests access to the bus. On receipt of $\overline{\text{HLD\bar{A}}}$ from the CPU, the device may perform bus cycles, as the CPU at this point has set AD0--AD15 , A16--A23 and $\overline{\text{HBE}}$ to the TRI-STATE® condition and has switched $\overline{\text{ADS}}$ and $\overline{\text{DDIN}}$ to the input mode. ALE is asserted in T4 , and stays high during the time the bus is granted. The CPU now monitors $\overline{\text{ADS}}$ and $\overline{\text{DDIN}}$ from the external device to generate the relevant strobe signals (i.e., $\overline{\text{TSO}}$, $\overline{\text{DBE}}$, $\overline{\text{RD}}$ or $\overline{\text{WR}}$). To return control of the bus to the CPU, the device sets $\overline{\text{HOLD}}$ inactive, and the CPU acknowledges it by setting $\overline{\text{HLD\bar{A}}}$ inactive.

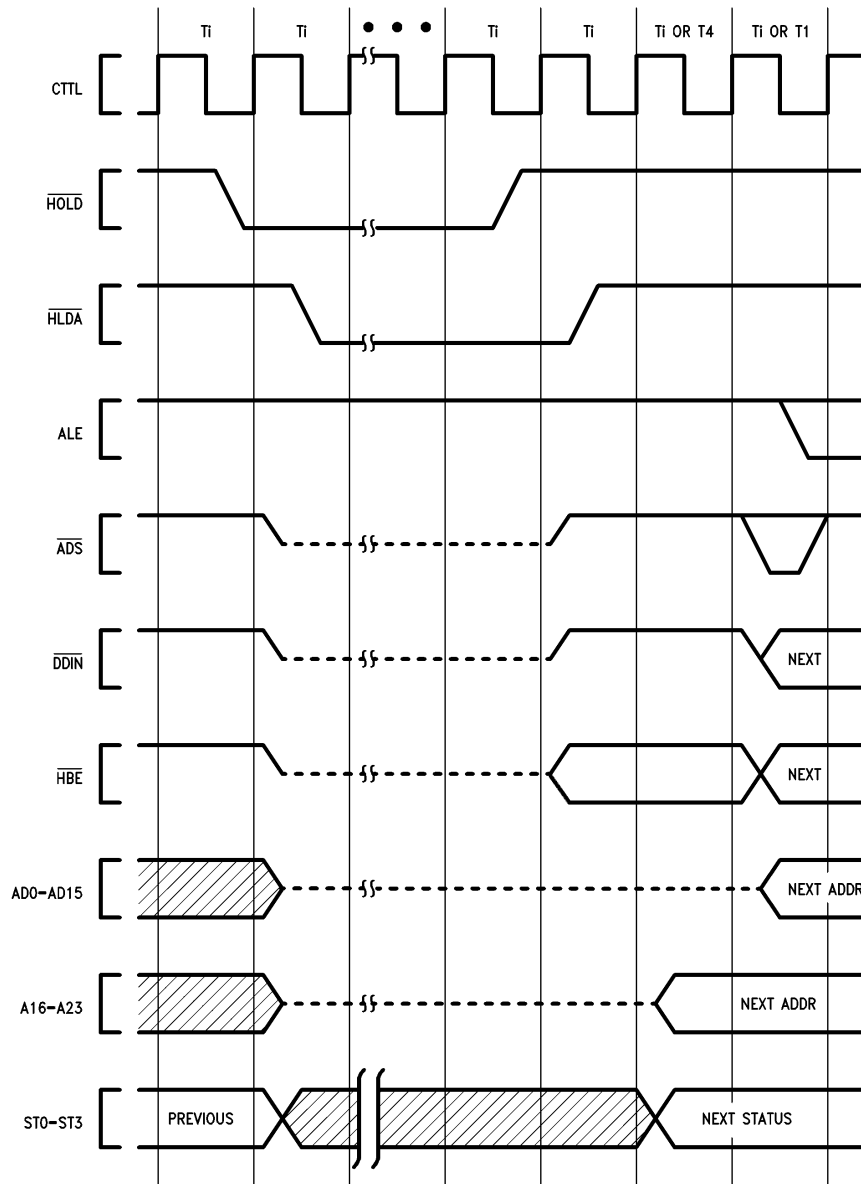
How quickly the CPU releases the bus depends on whether it is idle on the bus at the time the $\overline{\text{HOLD}}$ request is made,

as the CPU must always complete the current bus cycle. *Figure 3-29* shows the timing sequence when the CPU is idle. In this case, the CPU grants the bus during the immediately following clock cycle. *Figure 3-30* shows the sequence when the CPU is using the bus at the time the $\overline{\text{HOLD}}$ 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.

Note 1: During DMA cycles the $\overline{\text{WAIT1--2}}$ signals should be kept inactive, unless they are also monitored by the DMA controller. If wait states are required, $\overline{\text{CWAIT}}$ should be used.

Note 2: The logic value of the status pins, ST0--3 , is undefined during DMA activity.


3.0 Functional Description (Continued)



TL/EE/10818-37

FIGURE 3-29. $\overline{\text{HOLD}}$ Timing, Bus Initially Idle

3.0 Functional Description (Continued)



TL/EE/10818-38

FIGURE 3-30. $\overline{\text{HOLD}}$ Timing, Bus Initially Not Idle

3.0 Functional Description (Continued)

3.5.5.10 Instruction Status

In addition to the four bits of Bus Cycle status (ST0–3), the NS32FX16 CPU also presents Instruction Status information on three separate pins. These pins differ from ST0–3 in that they are synchronous to the CPU's internal instruction execution section rather than to its bus interface section.

\overline{PFS} (Program Flow Status) is pulsed low as each instruction begins execution. It is intended for debugging purposes.

U/\overline{S} originates from the U-bit of the Processor Status Register, and indicates whether the CPU is currently running in User or Supervisor mode. Although it is not synchronous to bus cycles, there are guarantees on its validity during any given bus cycle. See the Timing Specifications in Section 4.

\overline{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. \overline{ILO} is guaranteed to be active during the operand accesses performed by the interlocked instructions.

Note: The acknowledgement of \overline{HOLD} is on a cycle by cycle basis. Therefore, it is possible to have HLDA active when an interlock operation is in progress. In this case, \overline{ILO} remains low and the interlocked instruction continues only after \overline{HOLD} is de-asserted.

4.0 Device Specifications

4.1 NS32FX16 PIN DESCRIPTIONS

The following is a brief description of all NS32FX16 pins. The descriptions reference portions of the Function Description, Section 3.

Unless otherwise indicated, reserved pins should be left open.

Note: An asterisk next to the signal name indicates a TRI-STATE condition for that signal during \overline{HOLD} acknowledgement.

4.1.1 Supplies

V_{CC}	Power. +5V positive supply.
GND	Ground. Ground reference for both on-chip logic and output drivers.

4.1.2 Input Signals

RSTI	Reset Input. Schmitt triggered, asynchronous signal used to generate a CPU reset. See Section 3.5.4. Note: The reset signal is a true asynchronous input. Therefore, no external synchronizing circuit is needed.
-------------	---

\overline{HOLD}	Hold Request. When active, causes the CPU to release the bus for DMA or multiprocessing purposes. See Section 3.5.5.9.
-------------------------------------	--

Note: If the \overline{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 HLDA latency. This is to avoid speed degradations in cases of heavy \overline{HOLD} activity (i.e., DMA controller cycles interleaved with CPU cycles).

\overline{INT}	Interrupt. A low level on this pin requests a maskable interrupt. \overline{INT} must be kept asserted until the interrupt is acknowledged.
------------------------------------	---

\overline{NMI}	Non-Maskable Interrupt. A High-to-Low transition on this signal requests a non-maskable interrupt.
------------------------------------	--

Note: \overline{INT} and \overline{NMI} are true asynchronous inputs. Therefore, no external synchronizing circuit is needed.

\overline{CWAIT}	Continuous Wait. Causes the CPU to insert continuous wait states if sampled low at the end of T2 and each following T-State. See Section 3.5.5.3.
--------------------------------------	---

$\overline{WAIT1-2}$	Two-Bit Wait State Inputs These inputs, collectively called $\overline{WAIT1-2}$, allow from zero to three wait states to be specified. They are binary weighted. See Section 3.5.5.3.
--	---

Note: During a DMA cycle, $\overline{WAIT1-2}$ should be kept inactive unless they are also monitored by the DMA Controller. Wait states, in this case, should be generated through \overline{CWAIT} .

OSCIN	Crystal/External Clock Input. Input from a crystal or an external clock source. See Section 3.5.2.
--------------	--

4.1.3 Output Signals

A16–A23	*High-Order Address Bits. These are the most significant 8 bits of the memory address bus.
----------------	--

\overline{HBE}	*High Byte Enable. Status signal used to enable data transfers on the most significant byte of the data bus.
------------------------------------	--

ST0–3	Status. Bus cycle status code; ST0 is the least significant. Encodings are: 0000— Idle: CPU Inactive on Bus. 0001— Idle: WAIT Instruction. 0010— DSP Module Data Transfer. 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.
--------------	--

4.0 Device Specifications (Continued)

1100— Read for Effective Address.
1101— Transfer Slave Operand.
1110— Read Slave Status Word.
1111— Broadcast Slave ID.

$\overline{U/\overline{S}}$ User/Supervisor.
User or Supervisor Mode status. High indicates User Mode; low indicates Supervisor Mode.

\overline{ILO} Interlocked Operation.
When active, indicates that an interlocked operation is being executed.

\overline{HLDA} Hold Acknowledge.
Activated by the CPU in response to the \overline{HOLD} input to indicate that the CPU has released the bus.

\overline{PFS} Program Flow Status.
A pulse on this signal indicates the beginning of execution of an instruction.

\overline{BPU} BPU Cycle.
This signal is activated during a bus cycle to enable an external BITBLT processing unit. The EXTBLT instruction activates this signal.

Note: \overline{BPU} is low (Active) only during bus cycles involving pre-fetching instructions and execution of EXTBLT operands. It is recommended that \overline{BPU} , \overline{ADS} and status lines (ST0–ST3) be used to qualify BPU bus cycles. If a DMA circuit exists in the system, the \overline{HLDA} signal should be used to further qualify BPU cycles. \overline{BPU} may become active during T4 of a non-BPU bus cycle, and may become inactive during T4 of a BPU bus cycle. \overline{BPU} must be qualified by \overline{ADS} and status lines (ST0–ST3) to be used as an external gating signal.

\overline{RSTO} Reset Output.
This signal becomes active when \overline{RSTI} is low, initiating a system reset.

\overline{RD} Read Strobe.
Activated during CPU or DMA read cycles to enable reading of data from memory or peripherals. See Section 3.5.5.2.

\overline{WR} Write Strobe.
Activated during CPU or DMA write cycles to enable writing of data to memory or peripherals.

$\overline{T\overline{SO}}$ Timing State Output.
The falling edge of $\overline{T\overline{SO}}$ identifies the beginning of state T2 of a bus cycle. The rising edge identifies the beginning of state T4.

\overline{DBE} Data Buffers Enable.
Used to control external data buffers. It is active when the data buffers are to be enabled.

OSCOUT Crystal Output.
This line is used as the return path for the crystal (if used). When an external clock source is used, OSCOUT should be left unconnected or loaded with no more than 5 pF of stray capacitance.

\overline{IAS} Special Cycle Address Strobe.
Signals the beginning of a special bus cycle.

CTTL1–2 System Clock.
Output clock for bus timing. CTTL1 and CTTL2 must be externally connected together.

FCLK Fast Clock.
This clock is derived from the clock waveform on OSCIN. Its frequency is either the same as OSCIN or is lower, depending upon the scale factor programmed into the CFG register.

ALE Address Latch Enable.
Active high signal that can be used to control external address latches.

4.1.4 Input-Output Signals


AD0–15 *Address/Data Bus.
Multiplexed Address/Data Information. Bit 0 is the least significant bit of each.

\overline{SPC} Slave Processor Control.
Used by the CPU as the data strobe output for slave processor transfers; used by a slave processor to acknowledge completion of a slave instruction. See Section 3.5.5.7.

\overline{DDIN} *Data Direction.
Status signal indicating the direction of the data transfer during a bus cycle. During \overline{HOLD} acknowledge this signal becomes an input and determines the activation of \overline{RD} or \overline{WR} .

\overline{ADS} *Address Strobe
Controls address latches; signals the beginning of a bus cycle. During \overline{HOLD} acknowledge this signal becomes an input and the CPU monitors it to detect the beginning of a DMA cycle and generate the relevant strobe signals. When a DMA is used, \overline{ADS} should be pulled up to V_{CC} through a 10 k Ω resistor.

4.0 Device Specifications (Continued)



TL/EE/10818-39

Bottom View

**Order Number NS32FX16V-15, NS32FX16V-20 or NS32FX16V-25
NS Package Number V68A**

FIGURE 4-1. Connection Diagram

4.0 Device Specifications (Continued)

4.2 ABSOLUTE MAXIMUM RATINGS

If Military/Aerospace specified devices are required, please contact the National Semiconductor Sales Office/Distributors for availability and specifications.

Temperature under Bias 0°C to +70°C
Storage Temperature -65°C to +150°C

All Input or Output Voltages
with Respect to GND -0.5V to +7V

Note: Absolute maximum ratings indicate limits beyond which permanent damage may occur. Continuous operation at these limits is not intended; operation should be limited to those conditions specified under Electrical Characteristics.

4.3 ELECTRICAL CHARACTERISTICS $T_A = 0^\circ\text{C}$ to $+70^\circ\text{C}$, $V_{CC} = 5V \pm 10\%$, $GND = 0V$

Symbol	Parameter	Conditions	Min	Typ	Max	Units
V_{IH}	High Level Input Voltage		2.0		$V_{CC} + 0.5$	V
V_{IL}	Low Level Input Voltage		-0.5		0.8	V
V_{T+}	\overline{RSTI} Rising Threshold Voltage	$V_{CC} = 5.0V$	2.5		3.5	V
V_{XL}	OSCIN Input Low Voltage				0.5	V
V_{XH}	OSCIN Input High Voltage		4.5			V
V_{OH}	High Level Output Voltage	$I_{OH} = -400 \mu\text{A}$	2.4			V
V_{OL}	Low Level Output Voltage	$I_{OL} = 4 \text{ mA}$			0.45	V
I_{ILS}	\overline{SPC} Input Current (Low)	$V_{IN} = 0.4V$, \overline{SPC} in Input Mode			1.0	mA
I_I	Input Load Current	$0 \leq V_{IN} \leq V_{CC}$, All Inputs except \overline{SPC}	-20		20	μA
I_L	Leakage Current Output and I/O Pins in TRI-STATE or Input Mode	$0.4 \leq V_{OUT} \leq V_{CC}$	-20		20	μA
I_{CC}	Active Supply Current	$I_{OUT} = 0$, $T_A = 25^\circ\text{C}$ (Note 2)		170	240	mA


Note 1: Care should be taken by designers to provide a minimum inductance path between the GND pins and system ground in order to minimize noise.

Note 2: I_{CC} is affected by the clock scaling factor selected by the C- and M-bits in the CFG register, see Section 3.5.3.

4.4 SWITCHING CHARACTERISTICS

4.4.1 Definitions

All the timing specifications given in this section refer to 0.8V or 2.0V on the rising or falling edges of all the signals as illustrated in Figures 4-2 and 4-3 unless specifically stated otherwise. The capacitive load is assumed to be 100 pF on CTTL and 50 pF on all the other output signals.




TL/EE/10818-40

FIGURE 4-2. Output Signals Specification Standard

Abbreviations:

L.E.— Leading Edge
T.E.— Trailing Edge

R.E.— Rising Edge
F.E.— Falling Edge



TL/EE/10818-41

FIGURE 4-3. Input Signals Specification Standard

4.0 Device Specifications (Continued)

4.4.2 Timing Tables

4.4.2.1 Output Signals: Internal Propagation Delays, NS32FX16-15, NS32FX16-20, NS32FX16-25

- The output to input timings (e.g., address to data-in) are at least 2 ns better than the worst case values calculated from the output valid and input setup times relative to CTTL.

Symbol	Figure	Description	Reference/ Conditions	NS32FX16-15		NS32FX16-20		NS32FX16-25		Units
				Min	Max	Min	Max	Min	Max	
t _{CTp}	4-15	CTTL Clock Period	R.E., CTTL to Next R.E., CTTL	66	1000	50	1000	40	1000	ns
t _{CTh}	4-15	CTTL High Time	At 2.0V (Both Edges)	0.5 t _{CTp} - 6 ns		0.5 t _{CTp} - 5 ns		0.5 t _{CTp} - 5 ns		
t _{CTl}	4-15	CTTL Low Time	At 0.8V (Both Edges)	0.5 t _{CTp} - 6 ns		0.5 t _{CTp} - 5 ns		0.5 t _{CTp} - 4 ns		
t _{CTr}	4-15	CTTL Rise Time	0.8V to 2.0V on R.E., CTTL		6		5		4	ns
t _{CTf}	4-15	CTTL Fall Time	2.0V to 0.8V on F.E., CTTL		6		5		4	ns
t _{XCTd}	4-15	OSCIN to CTTL Delay	4.2V on R.E., OSCIN to R.E., CTTL		35		29		25	ns
t _{XFr}	4-15	OSCIN to FCLK R.E. Delay	4.2V on R.E., OSCIN to R.E., FCLK		25		20		15	ns
t _{FCr}	4-15	FCLK to CTTL R.E. Delay	R.E., FCLK to R.E., CTTL		10		10		10	ns
t _{FCf}	4-15	FCLK to CTTL F.E. Delay	R.E., FCLK to F.E., CTTL		10		10		10	ns
t _{ALv}	4-4	AD0-AD15 Valid (Note 5)	After R.E., CTTL T1		14		13		12	ns
t _{ALh}	4-4	AD0-AD15 Hold	After R.E., CTTL T2	0		0		0		ns
t _{AHv}	4-4	A16-A23 Valid (Note 5)	After R.E., CTTL T1		14		13		12	ns
t _{AHh}	4-4	A16-A23 Hold	After R.E., CTTL Next T1 or Ti	0		0		0		ns
t _{ALfr}	4-4	AD0-AD15 Floating (during Read)	After R.E., CTTL T2		14		13		12	ns
t _{ALf}	4-7	AD0-AD15 Floating	After R.E., CTTL Ti		14		13		12	ns
t _{AHf}	4-7	A16-A23 Floating	After R.E., CTTL Ti		14		13		12	ns
t _{Dv}	4-5	Data Valid (Write Cycle)	After R.E., CTTL T2 or T1		14		13		12	ns
t _{Dh}	4-5	Data Hold	After R.E., CTTL Next T1 or Ti	0		0		0		ns
t _{ADSa}	4-4	\overline{ADS} Signal Active	After R.E., CTTL T1		14		13		12	ns
t _{ADSi}	4-4	\overline{ADS} Signal Inactive (Note 4)	After R.E., CTTL T1	0.5 t _{CTp} - 6 ns	0.5 t _{CTp} + 16 ns	0.5 t _{CTp} - 6 ns	0.5 t _{CTp} + 15 ns	0.5 t _{CTp} - 6 ns	0.5 t _{CTp} + 14 ns	
t _{ADSw}	4-5	\overline{ADS} Pulse Width	At 0.8V (Both Edges)	20		15		10		ns
t _{ADsf}	4-7	\overline{ADS} Floating	After R.E., CTTL Ti		14		13		12	ns
t _{ALADSS}	4-4	AD0-AD15 Setup	Before \overline{ADS} T.E.	10		10		10		ns
t _{HBEv}	4-4	HBE Signal Valid	After R.E., CTTL T1		14		13		12	ns
t _{HBEh}	4-4	HBE Signal Hold	After R.E., CTTL Next T1 or Ti	0		0		0		ns
t _{HBEf}	4-7	HBE Signal Floating	After R.E., CTTL Ti		14		13		12	ns

4.0 Device Specifications (Continued)

4.4.2 Timing Tables (Continued)

4.4.2.1 Output Signals: Internal Propagation Delays, NS32FX16-15, NS32FX16-20, NS32FX16-25

Symbol	Figure	Description	Reference/ Conditions	NS32FX16-15		NS32FX16-20		NS32FX16-25		Units
				Min	Max	Min	Max	Min	Max	
t_{DDINv}	4-4	\overline{DDIN} Signal Valid	After R.E., CTTL T1		14		13		12	ns
t_{DDINh}	4-4	\overline{DDIN} Signal Hold	After R.E., CTTL Next T1 or Ti	0		0		0		ns
t_{DDINF}	4-7	\overline{DDIN} Floating	After R.E., CTTL Ti		14		13		12	ns
t_{SPCa}	4-10	\overline{SPC} Output Active	After R.E., CTTL T1		14		13		12	ns
t_{SPCia}	4-10	\overline{SPC} Output Inactive	After R.E., CTTL T4		14		13		12	ns
t_{HLDAa}	4-7	\overline{HLDA} Signal Active	After R.E., CTTL Ti		14		13		12	ns
t_{HLDAia}	4-8	\overline{HLDA} Signal Inactive	After R.E., CTTL Ti		14		13		12	ns
t_{STv}	4-4	Status ST0–ST3 Valid	After R.E., CTTL T4 (Before T1, see Note 1)		14		13		12	ns
t_{STh}	4-4	Status ST0–ST3 Hold	After R.E., CTTL T4	0		0		0		ns
t_{BPUv}	4-4	\overline{BPU} Signal Valid	After R.E., CTTL T4 or Ti		14		13		12	ns
t_{BPUh}	4-4	\overline{BPU} Signal Hold	After R.E., CTTL T4 or Ti	0		0		0		ns
t_{TSOa}	4-4	\overline{TSO} Signal Active	After R.E., CTTL T2		14		13		12	ns
t_{TSOia}	4-4	\overline{TSO} Signal Inactive	After R.E., CTTL T4		14		13		12	ns
t_{RDa}	4-4	\overline{RD} Signal Active	After R.E., CTTL T2		14		13		12	ns
t_{RDia}	4-4	\overline{RD} Signal Inactive	After R.E., CTTL T4		14		13		12	ns
t_{WRa}	4-5	\overline{WR} Signal Active	After R.E., CTTL T2		14		13		12	ns
t_{WRia}	4-5	\overline{WR} Signal Inactive	After R.E., CTTL T4		14		13		12	ns
$t_{DBEa(R)}$	4-4	\overline{DBE} Active (Read Cycle) (Note 4)	After R.E., CTTL T2	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 16 ns	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 15 ns	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 14 ns	
$t_{DBEa(W)}$	4-5	\overline{DBE} Active (Write Cycle)	After R.E., CTTL T2		14		13		12	ns
t_{DBEia}	4-5, 4-6	\overline{DBE} Inactive (Note 4)	After R.E., CTTL T4	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 16 ns	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 15 ns	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 14 ns	
t_{USv}	4-4	U/\overline{S} Signal Valid	After R.E., CTTL T4		14		13		12	ns
t_{USh}	4-4	U/\overline{S} Signal Hold	After R.E., CTTL T4	0		0		0		ns
t_{PFSa}	4-13	\overline{PFS} Signal Active (Note 4)	After R.E., CTTL	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 16 ns	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 15 ns	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 14 ns	
t_{PFSia}	4-13	\overline{PFS} Signal Inactive (Note 4)	After R.E., CTTL	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 16 ns	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 15 ns	$0.5 t_{CTP}$ – 3 ns	$0.5 t_{CTP}$ + 14 ns	
t_{ALEa}	4-5	ALE Signal Active (Note 4)	After R.E., CTTL T4	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 16 ns	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 15 ns	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 14 ns	
t_{ALEia}	4-5	ALE Signal Inactive (Note 4)	After R.E., CTTL T1	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 16 ns	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 15 ns	$0.5 t_{CTP}$ – 6 ns	$0.5 t_{CTP}$ + 14 ns	
T_{ALALEs}	4-5	AD0–AD15 Setup	Before ALE T.E.	10		10		10		ns

4.0 Device Specifications (Continued)

4.4.2 Timing Tables (Continued)

4.4.2.1 Output Signals: Internal Propagation Delays, NS32FX16-15, NS32FX16-20, NS32FX16-25

Symbol	Figure	Description	Reference/ Conditions	NS32FX16-15		NS32FX16-20		NS32FX16-25		Units
				Min	Max	Min	Max	Min	Max	
t_{IASa}	4-6	\overline{IAS} Signal Active	After R.E., CTTL T1		14		13		12	ns
t_{IASia}	4-6	\overline{IAS} Signal Inactive (Note 4)	After R.E., CTTL T1	$0.5 t_{CTP}$ - 6 ns	$0.5 t_{CTP}$ + 16 ns	$0.5 t_{CTP}$ - 6 ns	$0.5 t_{CTP}$ + 15 ns	$0.5 t_{CTP}$ - 6 ns	$0.5 t_{CTP}$ + 14 ns	
t_{ASw}	4-6	\overline{IAS} Pulse Width	At 0.8V (Both Edges)	20		15		10		ns
t_{AIASs}	4-6	AD0-AD15 Setup	Before \overline{IAS} T.E.	10		10		10		ns
t_{ILOa}	4-14	\overline{ILO} Signal Active	After R.E., CTTL		14		13		12	ns
t_{ILOia}	4-14	\overline{ILO} Signal Inactive	After R.E., CTTL		14		13		12	ns
t_{RSTOa}	4-19	\overline{RSTO} Signal Active	After R.E., CTTL		14		13		12	ns
t_{RSTOia}	4-19	\overline{RSTO} Signal Inactive	After R.E., CTTL		14		13		12	ns
t_{RTOI}	4-19	Reset to Idle (Note 3)	After F.E. of \overline{RSTO}		10		10		10	t_{CTP}

Note 1: Every memory cycle starts with T4, during which Cycle Status is applied. If the CPU was idling, the sequence will be "... Ti, T4, T1 ...". If the CPU was not idling, the sequence will be "... T4, T1 ...".

Note 2: The parameters related to the "floating/not floating" conditions are guaranteed by characterization. Due to tester conditions, these parameters are not 100% tested.

Note 3: Not tested, guaranteed by design.

Note 4: Minimum values not tested, guaranteed by design.

Note 5: When the load on AD0-15 is increased to 90 pF the value of t_{ALV} is increased by no more than 5 ns. When the load on A16-23 is increased to 90 pF the value of t_{AHV} is increased by no more than 5 ns.

4.4.2.2 Input Signal Requirements: NS32FX16-15, NS32FX16-20 and NS32FX16-25

Symbol	Figure	Description	Reference/ Conditions	NS32FX16-15		NS32FX16-20		NS32FX16-25		Units
				Min	Max	Min	Max	Min	Max	
t_{xp}	4-15	OSCIN Clock Period	R.E., OSCIN to Next R.E., OSCIN	33	500	25	500	20	500	ns
t_{xh}	4-15	OSCIN High Time (External Clock)	At 4.2V (Both Edges)	$0.5 t_{xp}$ - 5 ns		$0.5 t_{xp}$ - 4 ns		$0.5 t_{xp}$ - 3 ns		
t_{xl}	4-15	OSCIN Low Time	At 1.0V (Both Edges)	$0.5 t_{xp}$ - 5 ns		$0.5 t_{xp}$ - 4 ns		$0.5 t_{xp}$ - 3 ns		
t_{DIs}	4-4, 4-11	Data In Setup	Before R.E., CTTL T4	15		14		10		ns
t_{DIh}	4-4, 4-11	Data In Hold (Note 1)	After R.E., CTTL T4	2		2		2		ns
t_{CWs}	4-4, 4-5	\overline{CWAIT} Signal Setup	Before R.E., CTTL T3 or T3(w)	22		18		10		ns
t_{CWh}	4-4, 4-5	\overline{CWAIT} Signal Hold	After R.E., CTTL T3 or T3(w)	2		2		2		ns
t_{Ws}	4-4, 4-5	\overline{WAITn} Signals Setup	Before R.E., CTTL T3 or T3(w)	22		21		20		ns
t_{Wh}	4-4, 4-5	\overline{WAITn} Signals Hold	After R.E., CTTL T3 or T3(w)	2		2		2		ns
t_{HLDs}	4-7, 4-8	\overline{HOLD} Setup Time	Before R.E., CTTL T2 or Ti	16		15		14		ns
t_{HLDh}	4-7, 4-8	\overline{HOLD} Hold Time	After R.E., CTTL Ti	2		2		2		ns

4.0 Device Specifications (Continued)

4.4.2.2 Input Signal Requirements: NS32FX16-15, NS32FX16-20 and NS32FX16-25 (Continued)

Symbol	Figure	Description	Reference/ Conditions	NS32FX16-15		NS32FX16-20		NS32FX16-25		Units
				Min	Max	Min	Max	Min	Max	
t _{PWR}	4-18	Power Stable to $\overline{\text{RSTI}}$ R.E. (Note 2)	After V _{CC} Reaches 4.5V	50		40		30		μs
t _{RSTW}	4-19	$\overline{\text{RSTI}}$ Pulse Width	At 0.8V (Both Edges)	64		64		64		t _{CTP}
t _{INTH}	4-16	$\overline{\text{INT}}$ Signal Hold	After R.E., CTTL T2 of Interrupt Acknowledge Cycle	0		0		0		ns
t _{NMIS}	4-17	$\overline{\text{NMI}}$ Setup Time	Before F.E., CTTL	15		14		12		ns
t _{NMIH}	4-17	$\overline{\text{NMI}}$ Hold Time	After F.E., CTTL	2		2		2		ns
t _{SPCd}	4-12	SPC Pulse Delay from Slave (Note 2)	After F.E., CTTL T4	2		2		2		t _{CTP}
t _{SPCs}	4-12	$\overline{\text{SPC}}$ Input Setup	Before R.E., CTTL	22		21		20		ns
t _{SPCh}	4-12	SPC Hold Time	After R.E., CTTL	2		2		2		ns
t _{ADSS}	4-9	$\overline{\text{ADS}}$ Input Setup	Before F.E., CTTL	15	t _{CTP} -3	14	t _{CTP} -3	12	t _{CTP} -3	ns
t _{ADSh}	4-9	$\overline{\text{ADS}}$ Input Hold (Note 3)	After F.E., CTTL T1	2		2		2		ns
t _{DDINs}	4-9	$\overline{\text{DDIN}}$ Input Setup	Before F.E., CTTL	15		14		12		ns
t _{DDINih}	4-9	$\overline{\text{DDIN}}$ Input Hold	After R.E., CTTL T4	2		2		2		ns

Note 1: t_{Dih} is always less than or equal to t_{RDia}.

Note 2: Not tested, guaranteed by design.

Note 3: $\overline{\text{ADS}}$ must be deasserted before state T4 of the DMA controller cycle.

4.0 Device Specifications (Continued)

4.4.3 Timing Diagrams




FIGURE 4-4. Read Cycle

TL/EE/10818-42

4.0 Device Specifications (Continued)




FIGURE 4-5. Write Cycle

TL/EE/10818-43

4.0 Device Specifications (Continued)


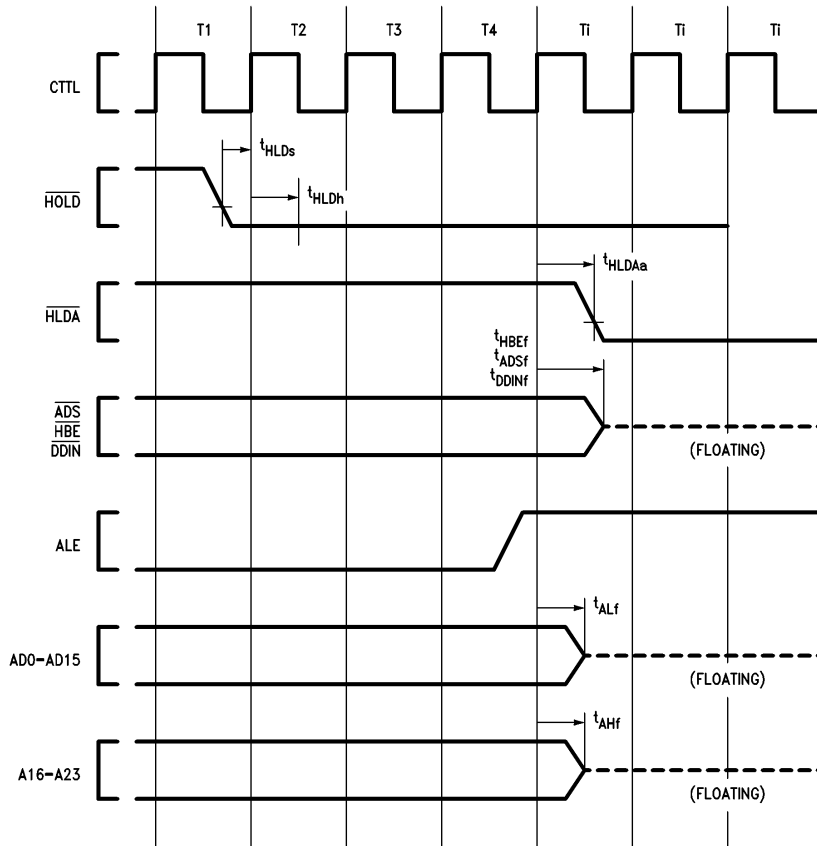


FIGURE 4-6. Special Bus Cycle

TL/EE/10818-44

4.0 Device Specifications (Continued)



TL/EE/10818-45

Note: When the bus is not idle, \overline{HOLD} must be asserted before the rising edge of CCTL of the timing state that precedes state T4 in order for the request to be acknowledged.

FIGURE 4-7. \overline{HOLD} Acknowledge Timing (Bus Initially Not Idle)

4.0 Device Specifications (Continued)


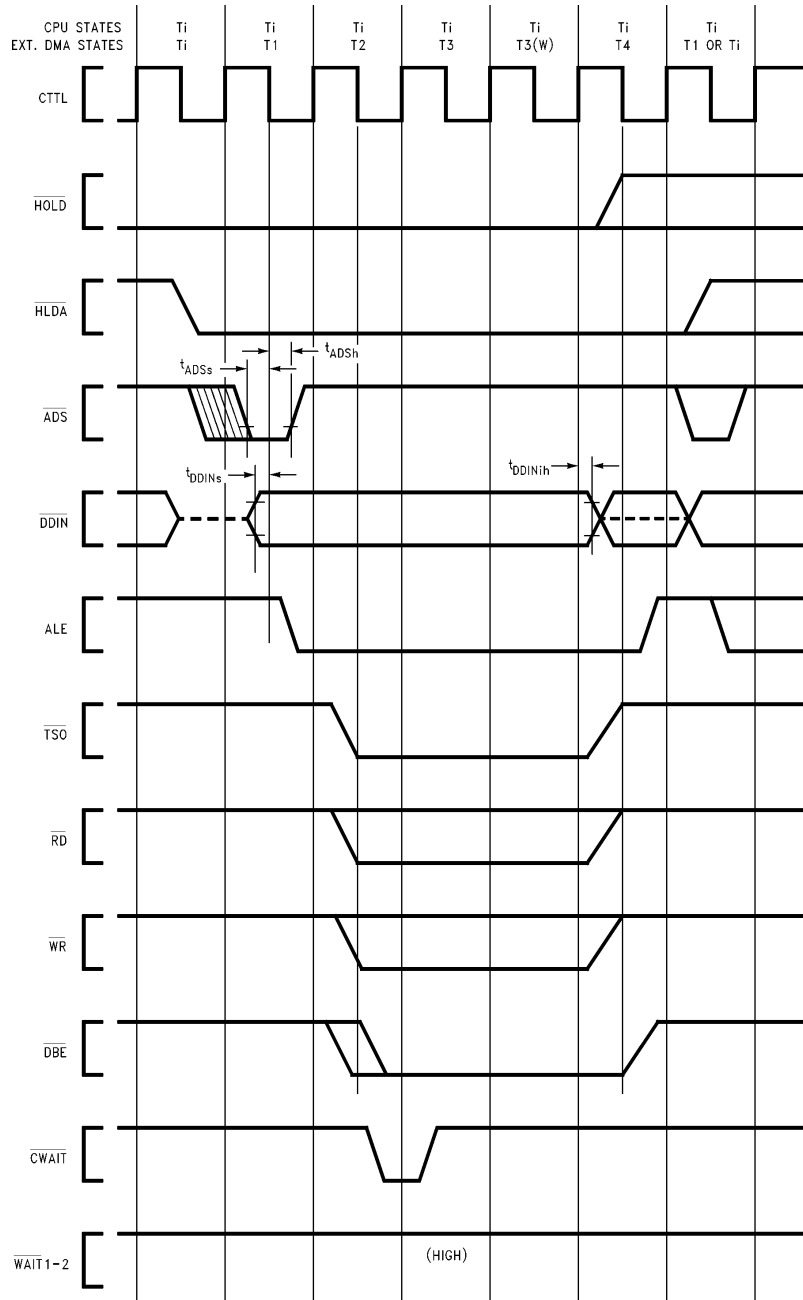


FIGURE 4-8. $\overline{\text{HOLD}}$ Timing (Bus Initially Idle)

TL/EE/10818-46

4.0 Device Specifications (Continued)



TL/EE/10818-47

Note 1: $\overline{\text{ADS}}$ must be deactivated before state T₄ of the external DMA controller cycle.

Note 2: During an external DMA cycle $\overline{\text{WAIT1-2}}$ must be kept inactive unless they are monitored by the DMA Controller. An external DMA cycle is similar to a CPU cycle. The NS32FX16 generates $\overline{\text{TSO}}$, $\overline{\text{RD}}$, $\overline{\text{WR}}$, ALE and DBE. The external DMA controller drives the address/data lines $\overline{\text{HBE}}$, $\overline{\text{ADS}}$ and DDIN.

Note 3: During an external DMA cycle, if the $\overline{\text{ADS}}$ signal is pulsed in order to initiate a bus cycle, the $\overline{\text{HOLD}}$ signal must remain asserted until state T₄ of the DMA cycle.

FIGURE 4-9. External DMA Controller Bus Cycle

4.0 Device Specifications (Continued)




FIGURE 4-10. Slave Processor Write Timing

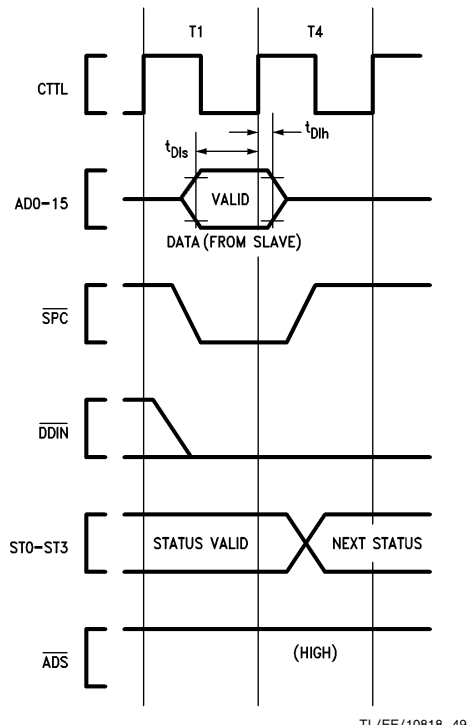




FIGURE 4-11. Slave Processor Read Timing



After transferring the last operand to the FPU, the CPU turns OFF the output driver and holds \overline{SPC} high with an internal 5 k Ω pullup.


FIGURE 4-12. \overline{SPC} Timing

4.0 Device Specifications (Continued)



TL/EE/10818-51


FIGURE 4-13. $\overline{\text{PFS}}$ Signal Timing



TL/EE/10818-52

Note: $\overline{\text{ILO}}$ may be asserted more than one clock cycle before the beginning of an interlocked access.

FIGURE 4-14. $\overline{\text{ILO}}$ Signal Timing



TL/EE/10818-53

FIGURE 4-15. Clock Waveforms

4.0 Device Specifications (Continued)




FIGURE 4-16. $\overline{\text{INT}}$ Signal Timing

TL/EE/10818-54

Note 1: Once $\overline{\text{INT}}$ is asserted, it must remain asserted until it is acknowledged.

Note 2: $\overline{\text{INTA}}$ is the Interrupt Acknowledge bus cycle (not a CPU signal). Refer to Section 3.4.1 and Table 3.4.




FIGURE 4-17. $\overline{\text{NMI}}$ Signal Timing

TL/EE/10818-55





FIGURE 4-18. Power-On Reset

TL/EE/10818-56

4.0 Device Specifications (Continued)



TL/EE/10818-57

Note 1: During Reset the $\overline{\text{HOLD}}$ signal must be kept high.

Note 2: After $\overline{\text{RSTi}}$ is deasserted the first bus cycle will be an instruction fetch at address zero.

FIGURE 4-19. Non-Power-On Reset

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)
 op = Operation Code
 Valid encodings shown with each format.
 gen, gen 1, gen 2 = General Addressing Mode Field
 See Section 2.4.2 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 = UPSR
 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 instruction:




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		



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


Appendix A: Instruction Formats (Continued)



Format 5


MOVS	-0000	BITWT	-1000
CMPS	-0001	TBITS	-1001
SETCFG	-0010	BBAND	-1010
SKPS	-0011	SBITPS	-1011
BBSTOD	-0100	BBFOR	-1100
EXTBLT	-0101	SBITS	-1101
BBOR	-0110	BBXOR	-1110
MOVMP	-0111		

No Operation on 1111




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/10818-58

Format 8


EXT	-0 00	INDEX	-1 00
CVTP	-0 01	FFS	-1 01
INS	-0 10		
CHECK	-0 11		

Trap (UND) on -1 10 and -1 11



Format 9


MOVf	-000	ROUND	-100
LFSSR	-001	TRUNC	-101
MOVLf	-010	SFSSR	-110
MOVFL	-011	FLOOR	-111



TL/EE/10818-59


Format 10

Trap (UND) Always



Format 11


ADDf	-0000	DIVf	-1000
MOVf	-0001	(Note 1)	-1001
CMPf	-0010	Trap (UND)	-1010
(Note 3)	-0011	Trap (UND)	-1011
SUBf	-0100	MULf	-1100
NEGf	-0101	ABSf	-1101
Trap (UND)	-0110	Trap (UND)	-1110
Trap (UND)	-0111	Trap (UND)	-1111



Format 12

(Note 2)	-0000	(Note 2)	-1000
(Note 1)	-0001	(Note 1)	-1001
POLYf	-0010	Trap (UND)	-1010
DOTf	-0011	Trap (UND)	-1011
SCALBf	-0100	(Note 2)	-1100
LOGBf	-0101	(Note 1)	-1101
Trap (UND)	-0110	Trap (UND)	-1110
Trap (UND)	-0111	Trap (UND)	-1111


*Instructions with Format 12 are available only when the NS32381 is used.



TL/EE/10818-60

Format 13


Trap (UND) Always




TL/EE/10818-61

Appendix A: Instruction Formats (Continued)


Trap (UND) **Format 14**
Always



Trap (UND) **Format 18**
Always




Trap (UND) **Format 15**
Always





Trap (UND) **Format 19**
Always



Implied Immediate Encodings:



Trap (UND) **Format 16**
Always



Trap (UND) **Format 17**
Always



Note 1: Opcode not defined; CPU treats like MOVf. First operand has access class of read; second operand has access class of write; f-field selects 32-bit or 64-bit data.

Note 2: Opcode not defined; CPU treats like ADDf. First operand has access class of read; second operand has access class of read-modify-write. f-field selects 32-bit or 64-bit data.

Note 3: Reserved opcode; execution of this opcode will generate an undefined result.

Appendix B: Instruction Execution Times

This section provides the necessary information to calculate the instruction execution times for the NS32FX16.

The following assumptions are made:

- The entire instruction, with all displacements and immediate operands, is assumed to be present in the instruction queue when needed.
- Interference from instruction prefetches, which is very dependent upon the preceding instruction(s), is ignored. This assumption will tend to affect the timing estimate in an optimistic direction.
- It is assumed that all memory operand transfers are completed before the next instruction begins execution. In the case of an operand of access class *rmw* in memory, this is pessimistic, as the Write transfer occurs in parallel with the execution of the next instruction.
- It is assumed that there is no overlap between the fetch of an operand and the following sequences of microcode. This is pessimistic, as the fetch of Operand 1 will generally occur in parallel with the effective address calculation of Operand 2, and the fetch of Operand 2 will occur in parallel with the execution phase of the instruction.
- Where possible, the values of operands are taken into consideration when they affect instruction timing, and a range of times is given. Where this is not done, the worst case is assumed.

B.1 BASIC AND FLOATING-POINT INSTRUCTIONS

Execution times for basic and floating-point instructions are given in Tables B-1 and B-2. The parameters needed for the various calculations are defined below.

TEA— The time required to calculate an operand's Effective Address. For a Register or Immediate operand, this includes the fetch of that operand.

TEA1— TEA value for the GEN or GEN1 operand.

TEA2— TEA value for the GEN2 operand.

TOPB— The time needed to read or write a memory byte.

TOPW— The time needed to read or write a memory word.

TOPD— The time needed to read or write a memory double-word.

TOPi— The time needed to read or write a memory operand, where the operand size is given by the operation length of the instruction. It is always equivalent to either TOPB, TOPW or TOPD.

TCY— Internal processing overhead, in clock cycles.

L— Internal processing whose duration depends on the operation length. The number of clock cycles is derived by multiplying this value by the number of bytes in the operation length.

NCYC— Number of bus cycles performed by the CPU to fetch or store an operand. NCYC depends on the operand size and alignment.

TPR— CPU processing (in clock cycles) performed in parallel with the FPU.

TFPU— Processing time required by the FPU to execute the instruction. This is the time from the last data sent to the FPU, until done is issued. TFPU can be found in the FPU data sheets.

f— This parameter is related to the floating-point operand size.

Tf— The time required to transfer 32 bits of floating point value to or from the FPU.

Ti— The time required to transfer an integer value to or from the FPU.

B.1.1 Equations

The following equations assume that:

- Memory accesses occur at full speed.
- Any wait states should be reflected in the calculations of TOPB, TOPW and TOPD.

Note: When multiple writes are performed during the execution of an instruction, wait states occurring during intermediate write transactions may be partially hidden by the internal execution. Therefore, a certain number of wait states can be inserted with no effect on the execution time. For example, in the case of the MOVSi instructions each wait state on write operations subtracts 1 clock cycle per write bus access, from the TCY of the instruction, since updating the pointers occurs in parallel with the write operation. This means that wait states can be added to write cycles without changing the execution time of the instruction, up to a maximum of 13 wait states on writes for MOVSB and MOVSW, and 4 wait states on writes for MOVSD.

TEA— TEA values for the various addressing modes are provided in the following table.

TEA TABLE

Addressing Mode	TEA Value	Notes
IMMEDIATE, ABSOLUTE	4	
EXTERNAL	$11 + 2 * TOPD$	
MEMORY RELATIVE	$7 + TOPD$	
REGISTER	2	
REGISTER RELATIVE, MEMORY SPACE	5	
TOP OF STACK	4	Access Class Write
	2	Access Class Read
	3	Access Class RMW
SCALED INDEXED	$TI1 + TI2$	

TI1 = TEA of the basemode except:

if basemode is REGISTER then $TI1 = 5$

if basemode is TOP OF STACK then $TI1 = 4$

TI2 depends on the scale factor:

if byte indexing $TI2 = 5$

if word indexing $TI2 = 7$

if double-word indexing $TI2 = 8$

if quad-word indexing $TI2 = 10$

TOPB— If operand is in a register or is immediate then $TOPB = 0$

else $TOPB = 3$

TOPW— If operand is in a register or is immediate then $TOPW = 0$

else $TOPW = 4 * NCYC - 1$

TOPD— If operand is in a register or is immediate then $TOPD = 0$

else $TOPD = 4 * NCYC - 1$

Appendix B: Instruction Execution Times

TOPI— If operand is in a register or is immediate then
 $TOPI = 0$
 else if $i = \text{byte}$ then $TOPI = TOPB$
 else if $i = \text{word}$ then $TOPI = TOPW$
 else ($i = \text{double-word}$) then $TOPI = TOPD$

L— If i (operation length) = byte then $L = 1$
 else if $i = \text{word}$ then $L = 2$
 else ($i = \text{double-word}$) $L = 4$

f— If standard floating (32 bits): $f = 1$
 If long floating (64 bits): $f = 2$

Tf— $Tf = 4$

Ti— If integer = byte or word, then $Ti = 2$
 If integer = double-word, then $Ti = 4$

B.1.2 Notes on Table Use

Values in the #TEA1 and #TEA2 columns indicate whether effective addresses need to be calculated.

A value of 1 indicates that address calculation time is required for the corresponding operand. A 0 indicates that the operand is either missing, or it is in a register and the instruction has an optimized form which eliminates the TEA calculation for it.

In the L column, multiply the entry by the operation length in bytes (1, 2 or 4).

In the TCY column, special notations sometimes appear:

$n1 \rightarrow n2$ means $n1$ minimum, $n2$ maximum

$n1\%n2$ means that the instruction flushes the instruction queue after $n1$ clock cycles and nonsequentially fetches the next instruction. The value $n2$ indicates the number of clock cycles for the internal execution of the instruction (including $n1$).

The effective number of cycles (TCY) must take into account the time (T_{fetch}) required to fetch the portion of the next instruction including the basic encoding and the index bytes. This time depends on the size and the alignment of this portion.

If only one memory cycle is required, then:

$$TCY = n1 + 6 + T_{\text{fetch}}$$

If more than one memory cycle is required, then:

$$TCY = n1 + 5 + T_{\text{fetch}}$$

In the notes column, notations held within angle brackets < > indicate alternatives in the operand addressing modes which affect the execution time. A table entry which is affected by the operand addressing may have multiple values, corresponding to the alternatives. These addressing notations are:

<I> Immediate

<R> CPU Register

<M> Memory

<F> FPU Register, either 32 or 64 Bits

<x> Any Addressing Mode

<ab> a and b represent the addressing modes of operand 1 and 2 respectively. Both a and b can be any addressing mode (e.g., <MR> means memory to CPU register).

Note: Unless otherwise specified the TCY value for immediate addressing is the same as for CPU register addressing.

B.1.3. Calculation of the Execution Time TEX for Basic Instructions

The execution time for a basic instruction is obtained by performing the following steps:

1. Find the desired instruction in Table B-1.
2. Calculate the values of TEA, TOPB, etc. using the numbers in the table and the equations given in the previous sections.
3. The result derived by adding together these values is the execution time TEX in clock cycles.

EXAMPLE

Calculate TEX for the instruction CMPW R0, TOS.

Operand 1 is in a register; Operand 2 is in memory. This means that we must use the table values corresponding to the <xM> case as given in the Notes column.

Only the #TEA1, #TEA2, #TOPI and TCY columns have values assigned for the CMPi instruction. Therefore, they are they only ones that need to be calculated to find TEX. The blank columns are irrelevant to this instruction.

Both #TEA1 and #TEA2 columns contain 1 for the <xM> case. This means that effective address times have to be calculated for both operands. (For the <MR> case, the Register operand would have required no TEA time, therefore only the Memory operand TEA would have been necessary.) From the equations:

$$TEA1 (\text{Register mode}) = 2.$$

$$TEA2 (\text{Top of Stack mode, access class read}) = 2.$$

The #TOPI column represents potential operand transfers to or from memory. For a Compare instruction, each operand is read once, for a total of two operand transfers.

$$TOPI (\text{Word, Register}) = 0,$$

$$TOPI (\text{Word, TOS}) = 3 (\text{assuming the operand aligned})$$

$$\text{Total TOPI} = 3$$

TCY is the time required for internal operation within the CPU. The TCY value for this case is 3.

$$TEX = TEA1 + TEA2 + TOPI + TCY = 2 + 2 + 3 + 3 = 10 \text{ machine cycles.}$$

If the CPU is running at 20 MHz then a machine cycle (clock cycle) is 50 ns. Therefore, this instruction would take 10×50 ns, or 0.5 μ s, to execute.

B.1.4 Calculation of the Execution Time TEX for Floating-Point Instructions

The execution time for a floating-point instruction is obtained by performing the following steps:

1. Find the desired instruction in Table B-2.
2. Calculate the values of TEA1, TEA2, TOPB, etc., using the numbers in the table, and the equations given in the previous sections.
3. Get the floating-point instruction execution time TFPU from the appropriate FPU data sheet.
4. Choose the higher value between TPR and $TFPU + 3$.
5. The result derived by adding together these values is the execution time TEX in clock cycles.

EXAMPLE 1

Calculate TEX for the instruction MOVLF F0,@h'3000.

Assumptions:

- The FPU being used is the NS32181.
- Write cycles are performed with no wait states.

Appendix B: Instruction Execution Times (Continued)

TEX Calculation:

Operand 1 is in a register, operand 2 is in memory. This means that we have to use the table values for the <FM> case.

The following parameter values are obtained from Table B-2 and the equations in the previous sections.

$$\text{TEA2 (Absolute Mode)} = 4$$

$$\text{TOPD (Memory Write)} = 7 \text{ (Operand aligned, no waits)}$$

$$T_f = 4$$

$$\text{TCY} = 32$$

$$\text{TPR} = \text{TEA2} + 6 = 4 + 6 = 10$$

From the FPU Execution Timing table in the NS32181 data sheet we get a TFPU for MOVLF of 19 clock cycles.

The higher value between TPR and TFPU + 3 is 22. The total execution time in clock cycles is:

$$\text{TEX} = \text{TEA2} + \text{TOPD} + T_f + \text{TCY} + 22 = 65$$

EXAMPLE 2

Calculate TEX for the instruction MULF 20(R0), 4(10(FP))

Assumptions:

- The FPU being used is the NS32181.
- 20(R0) is an aligned read with one wait state.
- 10(FP) is an aligned read with no wait states.
- 4(10 (FP)) is an unaligned rmw with two wait states.

TEX Calculation:

Operand 1 and operand 2 are both in memory. Therefore, the table values for the <MM> case must be used.

The parameter values obtained from Table B-2 and the equations in the previous sections are as follows:

$$\text{TEA1 (Register Relative Mode)} = 5$$

$$\text{TEA2 (Memory Relative Mode)} = 8 + \text{TOPD} = 15$$

(TOPD = 7 (Operand Aligned, No Wait))

$$\text{TOPD}_1 \text{ (Read from GEN1)} = 7 + 2 = 9 \text{ (Operand Aligned, One Wait)}$$

$$\text{TOPD}_2 \text{ (RMW from GEN2)} = 11 + 6 = 17 \text{ (Operand Unaligned, Two Waits)}$$

$$T_f = 4$$

$$\text{TCY} = 22 \rightarrow 28$$

$$\text{TPR} = 0$$

From the FPU Execution Timing Table in the NS32181 data sheet we get a TFPU for MULF of 33 clock cycles.

The higher value between TPR and TFPU + 3 is 36. The total execution time in clock cycles is:

$$\text{TEX} = \text{TEA1} + \text{TEA2} + \text{TOPD}_1 + \text{TOPD}_2 + 3 \cdot T_f + \text{TCY} + 36 = 5 + 15 + 9 + 17 + (22 \rightarrow 28) + 36 = 133 \rightarrow 140$$

TABLE B-1. Basic Instructions

Mnemonic	#TEA1	#TEA2	#TOPB	#TOPW	#TOPD	#TOPI	#L	TCY	Notes
ABSi	1	1	—	—	—	2	—	9	SCR < 0
	1	1	—	—	—	2	—	8	SCR > 0
ACBi	1	—	—	—	—	2	—	16	<M> no branch
	1	—	—	—	—	2	—	15%20	<M> branch
	—	—	—	—	—	—	—	18	<R> no branch
	—	—	—	—	—	—	—	17%22	<R> branch
ADDi	1	1	—	—	—	3	—	3	<xM>
	1	—	—	—	—	1	—	4	<MR>
	—	—	—	—	—	—	—	4	<RR>
ADDCi	1	1	—	—	—	3	—	3	<xM>
	1	—	—	—	—	1	—	4	<MR>
	—	—	—	—	—	—	—	4	<RR>
ADDPi	1	1	—	—	—	3	—	16	No Carry
	1	1	—	—	—	3	—	18	Carry
ADDQi	—	1	—	—	—	2	—	6	<M>
	—	—	—	—	—	—	—	4	<R>
ADDR	1	1	—	—	1	—	—	2	<xM>
	1	—	—	—	—	—	—	3	<xR>
ADJSPi	1	—	—	—	—	1	—	6	
ANDi	1	1	—	—	—	3	—	3	<xM>
	1	—	—	—	—	1	—	4	<MR>
	—	—	—	—	—	—	—	4	<RR>
ASHi	1	1	1	—	—	2	—	14 → 45	
Bcond	—	—	—	—	—	—	—	7	no branch
	—	—	—	—	—	—	—	6%10	branch
BICi	1	1	—	—	—	3	—	3	<xM>
	1	—	—	—	—	1	—	4	<MR>
	—	—	—	—	—	—	—	4	<RR>

Appendix B: Instruction Execution Times (Continued)

TABLE B-1. Basic Instructions (Continued)

Mnemonic	#TEA1	#TEA2	#TOPB	#TOPW	#TOPD	#TOPi	#L	TCY	Notes
BICPSRB	1	—	1	—	—	—	—	18%22	
BICPSRW	1	—	—	1	—	—	—	30%34	
BISPSRB	1	—	1	—	—	—	—	18%22	
BISPSRW	1	—	—	1	—	—	—	30%34	
BPT	—	—	—	2	4	—	—	40	
BR	—	—	—	—	—	—	—	4%10	
BSR	—	—	—	—	1	—	—	6%16	
CASEi	1	—	—	—	—	1	—	4%9	
CBITi	1	1	2	—	—	1	—	15	<xM>
	1	—	—	—	—	1	—	7	<xR>
CBITi	1	1	2	—	—	1	—	15	<xM>
	1	—	—	—	—	1	—	7	<xR>
CHECKi	1	1	—	—	—	3	—	7	high
	1	1	—	—	—	3	—	10	low
	1	1	—	—	—	3	—	11	ok
CMPI	1	1	—	—	—	2	—	3	<xM>
	1	—	—	—	—	1	—	3	<MR>
	—	—	—	—	—	—	—	3	<RR>
CMPMi	1	1	—	—	—	2 * n	—	9 * n + 24	n = # of elements in block
CMPQi	1	—	—	—	—	1	—	3	<M>
	—	—	—	—	—	—	—	3	<R>
CMPSi	—	—	—	—	—	2 * n	—	35 * n + 53	n = # of elements, not Translated
CMPST	—	—	n	—	—	2 * n	—	38 * n + 53	Translated
COMi	1	1	—	—	—	2	—	7	
CVTP	1	1	—	—	1	—	—	7	
CXP	—	—	—	3	4	—	—	16%21	
CXPD	1	—	—	3	3	—	—	13%18	
DEIi	1	1	—	—	—	5	16	38	<xM>
	1	—	—	—	—	1	16	31	<xR>
DIA	—	—	—	—	—	—	—	3%7	
DIVi	1	1	—	—	—	3	16	58 → 68	
ENTER	—	—	—	—	n + 1	—	—	4 * n + 18	n = # of general registers saved
EXIT	—	—	—	—	n + 1	—	—	5 * n + 17	n = # of general registers restored
EXTi	1	1	—	—	1	1	—	19 → 29	field in memory
	1	1	—	—	—	1	—	17 → 51	field in register
EXTSi	1	1	—	—	1	1	—	26 → 36	
FFSi	1	1	2	—	—	1	24	24 → 28	
FLAG	—	—	—	—	—	—	—	6	no trap
	—	—	—	4	3	—	—	44	trap
IBITi	1	1	2	—	—	1	—	17	<xM>
	1	—	—	—	—	—	—	9	<xR>

Appendix B: Instruction Execution Times (Continued)

TABLE B-1. Basic Instructions (Continued)

Mnemonic	# TEA1	# TEA2	# TOPB	# TOPW	# TOPD	# TOPI	# L	TCY	Notes
INDEXi	1	1	—	—	—	2	16	25	
INSi	1	1	—	—	2	1	—	29 → 39	field in memory
	1	—	—	—	—	1	—	28 → 96	field in register
INSSi	1	1	—	—	2	1	—	39 → 49	
JSR	1	—	—	—	1	1	—	5%15	
JUMP	1	—	—	—	—	—	—	2%6	
LPRi	1	—	—	—	—	1	—	19 → 33	
LSHi	1	1	1	—	—	2	—	14 → 45	
MEIi	1	1	—	—	—	4	16	23	
MODi	1	1	—	—	—	3	16	54 → 73	
MOVi	1	1	—	—	—	2	—	1	<xM>
	1	—	—	—	—	1	—	3	<MR>
	—	—	—	—	—	—	—	3	<RR>
MOVMi	1	1	—	—	—	2 * n	—	3 * n + 20	n = # of elements in block
MOVQi	1	—	—	—	—	1	—	2	<M>
	—	—	—	—	—	—	—	3	<R>
MOVSB, W	—	—	—	—	—	2 * n	—	14 * n + 59	n = # elements no options
	—	—	—	—	—	2 * n	—	24 * n + 54	B, W and/or U option in effect
MOVSD	—	—	—	—	—	2 * n	—	10 * n + 59	n = # of elements no options
	—	—	—	—	—	2 * n	—	24 * n + 54	B, W and/or U option in effect
MOVST	—	—	n	—	—	2 * n	—	27 * n + 54	Translated
MOVXBD	1	1	1	—	1	—	—	6	
MOVXBW	1	1	1	1	—	—	—	6	
MOVXWD	1	1	—	1	1	—	—	6	
MOVZBD	1	1	1	—	1	—	—	5	
MOVZBW	1	1	1	1	—	—	—	5	
MOVZWD	1	1	—	1	1	—	—	5	
MULi	1	1	—	—	—	3	16	15	
NEGi	1	1	—	—	—	2	—	5	
NOP	—	—	—	—	—	—	—	3	
NOTi	1	1	—	—	—	2	—	5	
ORi	1	1	—	—	—	3	—	3	<xM>
	1	—	—	—	—	1	—	4	<MR>
	—	—	—	—	—	—	—	4	<RR>
QUOi	1	1	—	—	—	3	16	49 → 55	

Appendix B: Instruction Execution Times (Continued)

TABLE B-1. Basic Instructions (Continued)

Mnemonic	#TEA1	#TEA2	#TOPB	#TOPW	#TOPD	#TOPi	#L	TCY	Notes
REMi	1	1	—	—	—	3	16	57 → 62	
RESTORE	—	—	—	—	n	—	—	5 * n + 12	n = # of general registers restored
RET	—	—	—	—	1	—	—	2%8	
RETI	—	—	1	2	2	—	—	60	Non-Cascaded
	—	—	2	2	3	—	—	60	Cascaded
RETT	—	—	—	2	2	—	—	45	
ROTi	1	1	1	—	—	2	—	14 → 45	
RXP	—	—	—	1	2	—	—	2%6	
Scondi	1	—	—	—	—	1	—	9	False
	1	—	—	—	—	1	—	10	True
SAVE	—	—	—	—	n	—	—	4 * n + 13	n = # of general registers saved
SBITi	1	1	2	—	—	1	—	15	<xM>
	1	—	—	—	—	1	—	7	<xR>
SBITii	1	1	2	—	—	1	—	15	<xM>
	1	—	—	—	—	1	—	7	<xR>
SETCFG	—	—	—	—	—	—	—	15	
SKPSi	—	—	—	—	—	n	—	27 * n + 51	n = # of elements, not Translated
SKPST	—	—	n	—	—	n	—	30 * n + 51	Translated
SPRi	1	—	—	—	—	1	—	21 → 27	
SUBi	1	1	—	—	—	3	—	3	<xM>
	1	—	—	—	—	1	—	4	<MR>
	—	—	—	—	—	—	—	4	<RR>
SUBCi	1	1	—	—	—	3	—	3	<xM>
	1	—	—	—	—	1	—	4	<MR>
	—	—	—	—	—	—	—	4	<RR>
SUBPi	1	1	—	—	—	3	—	16	no carry
	1	1	—	—	—	3	—	18	carry
SVC	—	—	—	2	4	—	—	40	
TBITi	1	1	1	—	—	1	—	14	<xM>
	1	—	—	—	—	1	—	4	<xR>
WAIT	—	—	—	—	—	—	—	6 → ?	? = until an interrupt/reset
XORi	1	1	—	—	—	3	—	3	<xM>
	1	—	—	—	—	1	—	4	<MR>
	—	—	—	—	—	—	—	4	<RR>

Appendix B: Instruction Execution Times (Continued)

TABLE B-2. Floating-Point Instructions: CPU Portion

Mnemonic	#TEA1	#TEA2	#TOPD	#TOPI	#Ti	#Tf	TCY	TPR	Notes
ADDf, SUBf, MULf, DIVf	— 1 — — 1	— — — 1 1	— f — 2f 2f 3f	— — — — —	— — — — —	— f f 2f 3f	17 (14 → 17) + 3f 24 + f (25 → 29) + 6f (27 → 30) + 3f (13 → 19) + 9f	8 0 0 0 0	<FF> <MF> <IF> <FM> <IM> <MM>
MOVf, ABSf, NEGf	— 1 — — 1	— — — — —	— f — f f 2f	— — — — —	— — — — —	— f f 2f 2f	17 (14 → 17) + 3f 24 + f 23 + 3f 33 + f (20 → 23) + 6f	6 0 0 6 + TEA2 TEA2 - 2 - f TEA2 - 3	<FF> <MF> <IF> <FM> <IM> <MM>
MOVFL	— 1 — — — 1	— — — — — —	— 1 — 2 2 3	— — — — — —	— — — — — —	— 1 1 2 3 3	17 17 → 20 25 35 43 35 → 38	8 0 0 6 + TEA2 TEA2 - 3 TEA2 - 3	<FF> <MF> <IF> <FM> <IM> <MM>
MOVLf	— 1 — — — 1	— — — — — —	— 2 — 1 1 3	— — — — — —	— — — — — —	— 2 2 1 3 3	16 20 → 23 26 32 42 35 → 38	8 0 0 TEA2 + 6 TEA2 - 4 TEA2 - 3	<FF> <MF> <IF> <FM> <IM> <MM>
TRUNCfi, FLOORfi, ROUNDfi	— 1 — — — 1	— — — — — —	— f — — — f	— — — 1 1 1	— — — 1 1 1	— f f — f f	20 (17 → 20) + 3f 25 + f 20 26 + f (16 → 19) + 4f	9 0 0 TEA2 + 6 TEA2 - 2 TEA2 - 2 - f	<FR> <MR> <IR> <FM> <IM> <MM>
MOVif	— 1 — — — 1	— — — 1 1 1	— — — f f f	— — — — — 1	— — — — — 1	— — — 1 1 1	25 - f 18 26 20 + 4f 22 + 5f (10 → 13) + 5f	0 0 0 0 0 0	<RF> <MF> <IF> <RM> <IM> <MM>
CMPf	— 1 — — — 1 — — 1	— — — 1 1 — — — —	— f — f f 2f — — f —	— — — — — — — — —	— — — — — — — — —	— f f 2f 2f f 2f 2f	23 (20 → 23) + 3f 31 + f (27 → 30) + 3f 29 (15 → 21) + 6f 37 + f (21 → 29) + 8f 35 + 2f	13 7 7 0 0 0 0 0 0	<FF> <MF> <IF> <FM> <IM> <MM> <FI> <MI> <II>
SFSR	— 1	— —	— 1	— —	— —	— 1	19 20	7 TEA1 + 4	<R> <M>
LFSR	— 1	— —	— 1	— —	— —	— 1	23 18 → 21	0 0	<R> <M>

Appendix B: Instruction Execution Times (Continued)

B.2 SPECIAL GRAPHICS INSTRUCTIONS

This section provides the execution times for the special graphics instructions. Table B-3 lists the average instruction execution times for different shift values and for a no-wait-state system design. The “No Option” of each instruction is used. The effect of wait states on the execution time is rather difficult to evaluate due to the pipelined nature of the read and write operations.

Instructions that have *shift* amounts, such as BBOR, BBXOR, BBAND, BBFOR and BITWT, make use of the parallel nature of the Series 32000®/EP processors by doing the actual *shift* during the reading of the double-word destination data. This means that if there are wait states on read operations, these instructions are able to *shift* further, without impacting the overall execution time. For example, the total execution time for a BBFOR operation, *shifting* 8 bits, with 2 wait states on read operations, is the same as for a BBFOR operation *shifting* by 12 bits. This is because a destination read takes 4 clock cycles longer than a no-wait-state double-word read does. Note that this effect is not valid for more than 4 wait states because at 4 wait states, all possible *shift* values (0–15) are “hidden” during the destination read.

Table B-4 shows the average execution times with wait states, assuming a shift value of eight unless stated otherwise. The parameters used in the execution time equations are defined below.

- Twaitrd The number of wait states applied for a Read operation.
- Twaitr The number of wait states applied for a Write operation.
- Twaitrds The number of wait states applied for a Read operation on source data. This also refers to the number of wait states applied for a table memory access (in the SBITS instruction, for example).

- Twaitrdd The number of wait states applied for a Read operation on destination data.
- Twaitrwd The number of wait states applied for a Write operation on destination data.
- Twaitrbd $Twaitrds + Twaitrdd * 2 + Twaitrwd * 2$, the value used for BITBLT timing.
- width The width of a BITBLT operation, in words.
- height The height of a BITBLT operation, in scan lines.
- shift The number of bits of shift applied.

B.2.1 Execution Time Calculation for Special Graphics Instructions

The execution time for a special graphics instruction is obtained by inserting the appropriate parameters to the equation for that instruction and evaluating it.

For example, to calculate the execution time of the BBOR instruction applied to a 10-word wide and 5-line high data block, assuming a shift count of 15 and a no-wait-state system, the following equation from Table B-3 is used.

$$42 + (107 + 44 * (\text{width} - 2)) * \text{height} + ((\text{shift} - 8) * \text{width} * \text{height})$$

Substituting the appropriate values to the shift, width and height parameters yields:

$$45 + (107 + 44 * (10 - 2)) * 50 + ((15 - 8) * 10 * 50)$$

or

$$42 + (107 + 352) * 50 + (7 * 500) = 26,492 \text{ clocks or } 1.77 \text{ ms @ } 15 \text{ MHz}$$

This represents the “worst case” time for this instruction, since a *shift* of greater than 15 bits can be handled by moving the source and destination pointers by 2 bytes and adjusting the *shift* amount.

The “best case” and “average case” times for most instructions are the same, due to reading the destination data during the *shifting* of the source data.

TABLE B-3. Average Instruction Execution Times with No Wait-States

Instruction	Number of Clock Cycles	Notes
BBOR	$42 + (107 + 44 * (\text{width} - 2)) * \text{height}$ $42 + (107 + 44 * (\text{width} - 2)) * \text{height}$ $+ ((\text{shift} - 8) * \text{width} * \text{height})$	Shift = 0 → 8 Shift > 8
BBXOR	$44 + (107 + 44 * (\text{width} - 2)) * \text{height}$ $44 + (107 + 44 * (\text{width} - 2)) * \text{height}$ $+ ((\text{shift} - 8) * \text{width} * \text{height})$	Shift = 0 → 8 Shift > 8
BBAND	$45 + (111 + 44 * (\text{width} - 2)) * \text{height}$ $45 + (111 + 44 * (\text{width} - 2)) * \text{height}$ $+ ((\text{shift} - 8) * \text{width} * \text{height})$	Shift = 0 → 8 Shift > 8
BBFOR	$48 + (61 + 25 * (\text{width} - 2)) * \text{height}$ $48 + (74 + 32 * (\text{width} - 2)) * \text{height}$ $48 + (74 + 32 * (\text{width} - 2)) * \text{height} +$ $((\text{shift} - 8) * \text{width} * \text{height})$	Shift = 0 Shift = 1 → 8 Shift > 8
BBSTOD	$66 + (170 + 60 * (\text{width} - 2)) * \text{height}$ $66 + (170 + 60 * (\text{width} - 2)) * \text{height}$ $+ ((\text{shift} - 8) * \text{width} * \text{height})$	Shift = 0 → 8 Shift > 8

Appendix B: Instruction Execution Times (Continued)

TABLE B-3. Average Instruction Execution Times with No Wait-States (Continued)

Instruction	Number of Clock Cycles	Notes
BITWT	16 28 $28 + (\text{shift} - 8)$	Shift = 0 Shift = 1 → 8 Shift > 8
EXTBLT	$35 + (19 + 12 * \text{width}) * \text{height}$ $35 + (13 + 12 * \text{width}) * \text{height}$ $35 + (17 + 13 * \text{width}) * \text{height}$ $35 + (11 + 13 * \text{width}) * \text{height}$	Shift = 0 → 8, Pre-Read Shift = 0 → 8, No Pre-Read Shift > 8, Pre-Read Shift > 8, No Pre-Read
MOVMPB,W	$16 + 7 * R2$	
MOVMPD,W	$16 + 8 * R2$	
SBITS	39 42	$R2 \leq 25$ $R2 > 25$
SBITP	$8 + (34 * R2)$	

TABLE B-4. Average Instruction Execution Times with Wait-States

Instruction	Number of Clock Cycles	Notes
BBOR	$42 + ((107 + 2 * \text{Twaitblt}) + (44 + \text{Twaitblt}) * (\text{width} - 2)) * \text{height}$	
BBXOR	$44 + ((107 + 2 * \text{Twaitblt}) + (44 + \text{Twaitblt}) * (\text{width} - 2)) * \text{height}$	
BBAND	$45 + ((111 + 2 * \text{Twaitblt}) + (44 + \text{Twaitblt}) * (\text{width} - 2)) * \text{height}$	
BBFOR	$48 + ((74 + 2 * \text{Twaitblt}) + (32 + \text{Twaitblt}) * (\text{width} - 2)) * \text{height}$	
BBSTOD	$66 + ((170 + 2 * \text{Twaitblt}) + (60 + \text{Twaitblt}) * (\text{width} - 2)) * \text{height}$	
BITWIT	$16 + \text{Twaitrds} + \text{Twaitrdd} + \text{Twait wrd}$ $28 + \text{Twaitblt}$	Shift = 0 Shift = 1 → 8
EXTBLT	$35 + (19 + (12 + (\text{Twaitrds} + \text{Twaitrdd} + \text{Twait wrd})) * \text{width}) * \text{height}$ $35 + (13 + (12 + (\text{Twaitrds} + \text{Twaitrdd} + \text{Twait wrd})) * \text{width}) * \text{height}$	Pre-Read No Pre-Read
MOVMPB,W	$16 + 7 * R2 + (\text{Twaitwr} - 1) * R2$ $16 + 7 * R2$	$\text{Twaitwr} > 1$ $\text{Twaitwr} \leq 1$
MOVMPD	$16 + 8 * R2 + \text{Twaitwr} * R2$	
SBITS	$39 + (2 * \text{Twaitrdd} + 2 * \text{Twait wrd} + 2 * \text{Twaitrds})$ $42 + (2 * \text{Twaitrdd} + 2 * \text{Twaitrds})$	$R2 \leq 25$ $R2 > 25$
SBITP	$8 + (34 * R2) + ((\text{Twaitrdd} + \text{Twait wrd}) * R2)$	

B.3 DSPM INSTRUCTIONS


A DSPM instruction execution starts with the CPU core writing to the CTL register. The execution time is counted from state T3 of this transaction until all the results are ready, either in the accumulator or in the coefficient RAM array. The execution times, in clock cycles, for the various DSPM instructions are listed in Table B-5.

It is assumed that External Hold Requests do not occur in the middle of a VCMAD, VCMUL or VCMAC instruction.


The parameters *n* and *w* represent the number of elements in the vector instruction and the number of wait states applied to each DSPM bus transaction respectively.

TABLE B-5. DSPM Instruction Execution Times

Instruction	Number of Clock Cycles
VCMAD	$9 + 8 * n + 2 * n * w$
VCMUL	$9 + 8 * n + 2 * n * w$
VCMAC	$6 + 8 * n + 2 * n * w$
VCMAG	$5 + 8 * n$



Physical Dimensions inches (millimeters)



68-Pin Plastic Leaded Chip Carrier (V)
Order Number NS32FX 16V-15, NS32FX 16V-20 or NS32FX 16V-25
NS Package Number V68A

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.



National Semiconductor Corporation
 1111 West Bardin Road
 Arlington, TX 76017
 Tel: 1(800) 272-9959
 Fax: 1(800) 737-7018

National Semiconductor Europe
 Fax: (+49) 0-180-530 85 86
 Email: cnjwge@tevm2.nsc.com
 Deutsch Tel: (+49) 0-180-530 85 85
 English Tel: (+49) 0-180-532 78 32
 Français Tel: (+49) 0-180-532 93 58
 Italiano Tel: (+49) 0-180-534 16 80

National Semiconductor Hong Kong Ltd.
 19th Floor, Straight Block,
 Ocean Centre, 5 Canton Rd.
 Tsimshatsui, Kowloon
 Hong Kong
 Tel: (852) 2737-1600
 Fax: (852) 2736-9960

National Semiconductor Japan Ltd.
 Tel: 81-043-299-2309
 Fax: 81-043-299-2408

National does not assume any responsibility for use of any circuitry described, no circuit patent licenses are implied and National reserves the right at any time without notice to change said circuitry and specifications.