

NS32CG160-15/NS32CG160-20/NSC32CG160-25 32-Bit Integrated System Processor

General Description

The NS32CG160 is a highly-integrated member of the Series 32000/EP™ family of National's Embedded System Processors™ specifically optimized for office imaging peripherals.

The NS32CG160 provides a 16 Mbyte linear external address space and a 16-bit external data bus. The CPU core incorporates a 32-bit ALU and instruction pipeline, an 8-byte prefetch queue and a 16 x 16-bit hardware multiplier.

Integrated on the same chip with the CPU is also a 2-channel DMA controller, a 16-function BitBLT Processing Unit (BPU), a 15-level Interrupt Control Unit (ICU) and 3 programmable 16-bit timers.

The on-chip hardware multiplier considerably enhances the multiply instruction performance over its predecessors making it better for graphics data calculations, such as outline fonts.


The NS32CG160 capabilities can be expanded by using an external floating point unit (FPU) which directly interfaces to the NS32CG160 using the slave protocol. The CPU-FPU cluster features high speed execution of the floating-point instructions.

The power and programmability of the 32-bit CPU Core combined with the I/O peripherals generally found in micro-controllers make the NS32CG160 an ideal single chip solution for embedded applications.

Features

- Software compatible with the NS32CG16 and other embedded system processors
- 32-bit architecture and implementation
- Special support for graphics applications
 - On-chip BITBLT processing unit supports very fast BITBLT operations
 - 18 graphics instructions
 - Binary compression/expansion capability for font storage using RLL encoding
 - Pattern magnification
 - Interface to external BITBLT processing units for multiple bitplane/color applications
- On-chip clock generator
- Three on-chip programmable 16-bit timers
- Two on-chip DMA channels with transfer rates up to 8 Mbyte/sec
- On-chip Interrupt Control Unit provides 15 levels of prioritized interrupts
- Floating-point support via the NS32081 or NS32181
- On-chip hardware multiplier
- Optimal interface to large memory arrays via the NS32CG821 and the DP84xx family of DRAM controllers
- Fast interrupt service and task switching for real-time applications
- Power save mode
- High-speed CMOS technology
- 84-pin PLCC package

Block Diagram



TL/EE/10752-1

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

Table of Contents

1.0 PRODUCT INTRODUCTION

1.1 NS32CG160 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 BITBLT Processing Unit (BPU) Registers

2.1.6 DMA Controller Registers

2.1.7 Interrupt Control Unit (ICU) Registers

2.1.8 Timer Registers

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

2.5 Graphics 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

3.1.2.4 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

3.0 FUNCTIONAL DESCRIPTION (Continued)

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 On-Chip Peripherals

3.4.1 BITBLT Processing Unit (BPU)

3.4.1.1 BPU Operation

3.4.2 DMA Controller

3.4.2.1 DMA Data Transfers

3.4.2.2 Single Transfer Operation

3.4.2.3 Double Buffer Operation

3.4.2.4 Auto Initialize Operation

3.4.2.5 Bus Arbitration

3.4.3 Interrupt Control Unit (ICU)

3.4.3.1 Interrupt-Acknowledge Processing

3.4.3.2 End-Of-Interrupt Processing

3.4.4 Timers

3.4.4.1 Processor Independent Mode (Mode 1)

3.4.4.2 External Event Counter Mode (Mode 2)

3.4.4.3 Input Capture Mode (Mode 3)

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 Special Bus Cycles

3.5.5.6 BPU Bus Cycles

3.5.5.7 DMA Controller Bus Cycles

3.5.5.8 Slave Processor Bus Cycles

3.5.5.9 Data Access Sequences

3.5.5.10 Bus Access Control

3.5.5.11 External Interrupt Requests

3.5.5.12 Instruction Status

Table of Contents (Continued)

4.0 DEVICE SPECIFICATIONS

- 4.1 NS32CG160 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

4.0 DEVICE SPECIFICATIONS (Continued)

- 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

List of Illustrations

CPU Block Diagram	1
NS32CG160 Internal Registers	2-1
Processor Status Register (PSR)	2-2
Configuration Register (CFG)	2-3
BPU Registers Address Map	2-4
Barrel Shifter Control Register	2-5
Right Mask Register (BRM)	2-6
Left Mask Register (BLM)	2-7
Function Select Register (BFS)	2-8
Control Register (BCNTL)	2-9
BPU Count Register (BC)	2-10
DMA Controller Registers Address Map	2-11
DMA Status Register (STAT)	2-12
DMA Interrupt Mask Register (IMSK)	2-13
DMA Mode Control Register (MODE)	2-14
Channel Control Register (CNTL)	2-15
ICU Registers Address Map	2-16
Interrupt Vector Register (IVCT)	2-17
Timer Registers Address Map	2-18
Timer Control Register (TCNTL)	2-19
NS32CG160 Address Mapping	2-20
NS32CG160 Run-Time Environment	2-21
General Instruction Format	2-22
Index Byte Format	2-23
Displacement Encodings	2-24
Correspondence between Linear and Cartesian Addressing	2-25
32-Pixel by 32-Scan Line Frame Buffer	2-26
Overlapping BITBLT Blocks	2-27
B B Instructions Formats	2-28
BITWT Instruction Format	2-29
EXTBLT Instruction Format	2-30
MOVMPi Instruction Format	2-31
TBITS Instruction Format	2-32
SBITS Instruction Format	2-33
SBITPS Instruction Format	2-34
Bus Activity for a Simple BITBLT Operation	2-35
Operating States	3-1
Slave Processor Protocol	3-2
Slave Processor Status Word	3-3
Interrupt Dispatch Table	3-4

List of Illustrations (Continued)

Exception Acknowledge Sequence: Direct-Exception Mode Disabled	3-5
Exception Acknowledge Sequence: Direct-Exception Mode Enabled	3-6
Return from Trap (RETT n) Instruction Flow: Direct-Exception Mode Disabled	3-7
Return from Interrupt (RETI) Instruction Flow: Direct-Exception Mode Disabled	3-8
Exception Processing Flowchart	3-9
Service Sequence	3-10
BPU Block Diagram	3-11
Timer Block Diagram for the Processor Independent and External Event Counter Modes	3-12
Timer Block Diagram for the Input Capture Mode	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
BPU Read Cycle	3-25
DMA Bus Cycle Timing	3-26
Slave Processor Read Cycle	3-27
Slave Processor Write Cycle	3-28
NS32CG160 and FPU Interconnections	3-29
Memory Interface	3-30
HOLD Timing, Bus Initially Idle	3-31
HOLD Timing, Bus Initially Not Idle	3-32
Interrupt Request Encoding Logic	3-33
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
On-Chip DMA Bus Cycle	4-7
HOLD Acknowledge Timing (Bus Initially Not Idle)	4-8
HOLD Timing (Bus Initially Idle)	4-9
External DMA Controller Bus Cycle	4-10
Slave Processor Write Timing	4-11
Slave Processor Read Timing	4-12
SPC Timing	4-13
PFS Signal Timing	4-14
ILO Signal Timing	4-15
Clock Waveforms	4-16
External Interrupt Requests	4-17
TXAn and TXBn Input Timing	4-18
TXAn Output Timing	4-19
TXAn Float/Non-Float Timing	4-20
Power-On Reset	4-21
Non-Power-On Reset	4-22

List of Tables

Right Mask Truth Table	2-1
Left Mask Truth Table	2-2
BITBLT Logical Operations	2-3
NS32CG160 Addressing Modes	2-4
NS32CG160 Instruction Set Summary	2-5
“OP” and “i” Field Encodings	2-6
Floating-Point Instruction Protocols	3-1
Summary of Exception Processing	3-2
Barrel Shifter Truth Table	3-3
Maximum Number of Back-to-Back DMA Transfers	3-4
Interrupt Sources and Priority Levels	3-5
Timer Modes	3-6
External Oscillator Specifications	3-7
Interrupt Sequences	3-8
Bus Cycle Categories	3-9
Data Access Sequences	3-10
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

1.0 Product Introduction

The NS32CG160 is a high speed CMOS microprocessor in the Series 32000/EP family. It is software compatible with all the 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 NS32CG160 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 NS32CG160 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 extensions 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 NS32CG160 SPECIAL FEATURES

In addition to the above Series 32000 features, the NS32CG160 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 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 multiple bitplane applications.

The NS32CG160 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 bus characteristics and the power save feature 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 2.4. Details on all the NS32CG160 graphics instructions can be found in the NS32CG16 Printer/Display Processor Programmer's Reference Supplement.

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	The BitBLT group of instructions provide a method of quickly imaging characters, creating patterns, windowing and other block oriented effects.
BITWT EXTBLT 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 enables 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 NS32CG160 has 52 internal registers. 17 of them belong to the CPU portion of the device and are addressed either implicitly by specific instructions or through the register addressing mode. The other 35 control the operation of the on-chip peripherals, and are memory mapped. *Figure 2-1* shows the NS32CG160 internal registers.

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.

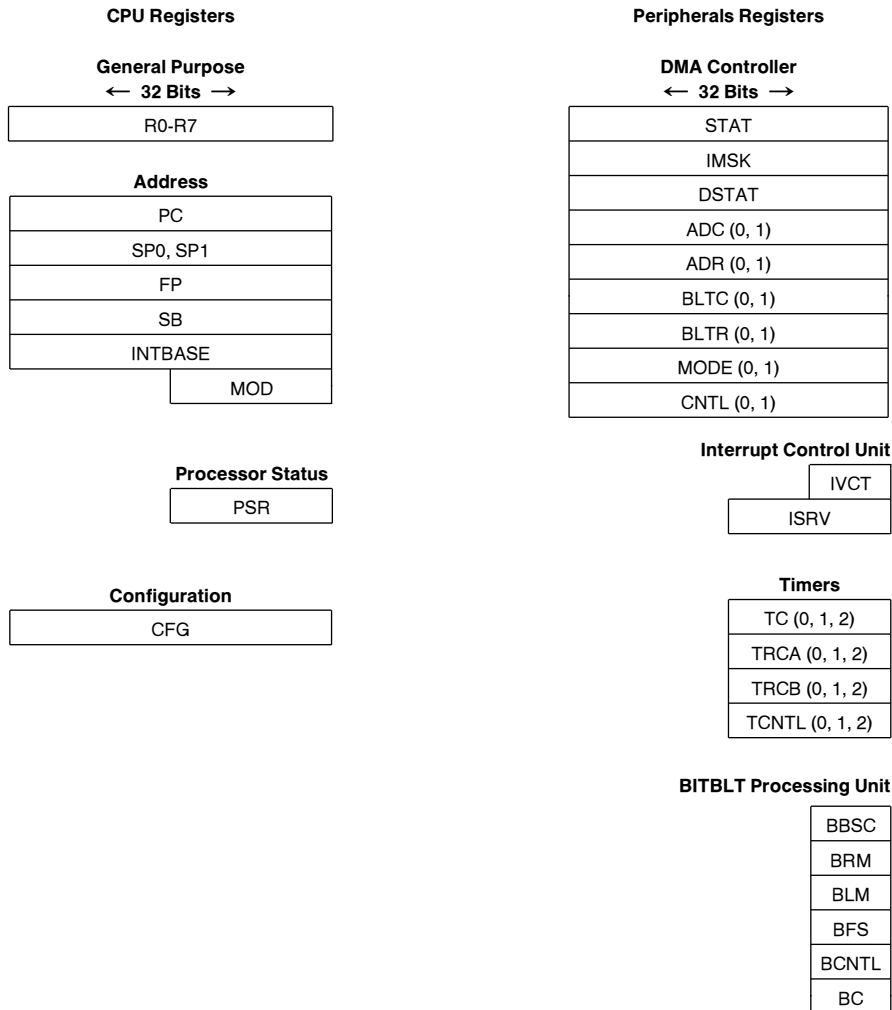


FIGURE 2-1. NS32CG160 Internal Registers

2.0 Architectural Description

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).
- 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.

2.0 Architectural Description (Continued)

2.1.4 Configuration Register

The Configuration Register (CFG) is 32 bits wide, of which 5 bits are implemented. The implemented bits enable various operating modes for the CPU, including vectoring of interrupts, execution of floating-point instructions, processing of exceptions and selection of clock scaling factor. The CFG is programmed by the SETCFG instruction. The format of CFG is shown in *Figure 2-3*. The various control bits are described below.

31	8	7					0
Reserved	DE	Res	C	M	F	I	

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.
- DE** Direct-Exception mode enable. This bit enables the Direct-Exception mode for processing exceptions. When this mode is selected, the CPU response time to interrupts and other exceptions is significantly improved. Refer to Section 3.2 for more information.

2.1.5 BITBLT Processing Unit (BPU) Registers

The BITBLT Processing Unit contains 6 8-bit registers that must be programmed before executing the EXTBLT instruction.

These registers are memory mapped and are accessible for both read and write. An operand length of 8 bits must be specified when accessing them. The address map of the BPU registers is shown in *Figure 2-4*.

Register Names	Register Addresses
BBSC	FFFFE000
BRM	FFFFE004
BLM	FFFFE008
BFS	FFFFE00C
BCNTL	FFFFE010
BC	FFFFE014

FIGURE 2-4. BPU Registers Address Map

BBSC—Barrel Shifter Control. This register controls the operation of the BPU barrel shifter. The format of BBSC is shown in *Figure 2-5*.

7	5	4	3	0
Reserved	BIS	SN		

FIGURE 2-5. Barrel Shifter Control Register (BBSC)

SN Shift Number. This field determines the number of bit positions shifted by the barrel shifter.

BIS Barrel Shifter Input Select. This bit controls the routing of the two 16-bit input words into the least significant and most significant halves of the 32-bit shift register inside the barrel shifter. Refer to Section 3.4.1 for more details.

BRM—Right Mask. This register controls the right mask. The format of BRM is shown in *Figure 2-6*. The value of the RM field defines the right mask bit pattern as shown in Table 2-1.

7	4	3	0
Reserved	RM		

FIGURE 2-6. Right Mask Register (BRM)

2.0 Architectural Description (Continued)

TABLE 2-1. Right Mask Truth Table

RM Field	Right Mask Patterns
0	1000000000000000
1	1100000000000000
2	1110000000000000
3	1111000000000000
4	1111100000000000
5	1111110000000000
6	1111111000000000
7	1111111100000000
8	1111111110000000
9	1111111111000000
10	1111111111100000
11	1111111111110000
12	1111111111111000
13	1111111111111100
14	1111111111111110
15	1111111111111111

1 = Enable Logical Operation
0 = Disable Logical Operation

BLM—Left Mask. This register controls the left mask. The format of BLM is shown in *Figure 2-7*. The value of the LM field defines the left mask bit pattern as shown in Table 2-2.

7	4	3	0
Reserved		LM	

FIGURE 2-7. Left Mask Register (BLM)

TABLE 2-2. Left Mask Truth Table

LM Field	Left Mask Patterns
0	1111111111111111
1	0111111111111111
2	0011111111111111
3	0001111111111111
4	0000111111111111
5	0000011111111111
6	0000001111111111
7	0000000111111111
8	0000000011111111
9	0000000001111111
10	0000000000111111
11	0000000000011111
12	0000000000001111
13	0000000000000111
14	0000000000000011
15	0000000000000001

1 = Enable Logical Operation
0 = Disable Logical Operation

Note: The leftmost bit in the mask patterns is the least significant one.

BFS—Function Select. The BSF register is used to select the BITBLT logical operation performed by the on-chip BPU between the source and destination data blocks. The format of BFS is shown in *Figure 2-8*. The value of the FS field selects 1 out of 16 logical operations as shown in Table 2-3.

In multi-plane systems that use one BPU per plane, each BPU can have its logic function programmed independently. This feature allows the NS32CG160 to BITBLT data to all bitplanes concurrently while maintaining the flexibility of having unique logical operations being performed on each plane.

7	4	3	0
Reserved		FS	

FIGURE 2-8. Function Select Register (BFS)

TABLE 2-3. BITBLT Logical Operations

FS Field				Operation Performed
0	0	0	0	0
0	0	0	1	s and d
0	0	1	0	s and -d
0	0	1	1	s
0	1	0	0	-s and d
0	1	0	1	d
0	1	1	0	s xor d
0	1	1	1	s or d
1	0	0	0	-s and -d
1	0	0	1	s xnor d
1	0	1	0	-d
1	0	1	1	s or -d
1	1	0	0	-s
1	1	0	1	-s or d
1	1	1	0	-s or -d
1	1	1	1	1

d: Destination Data
s: Source Data

BCNTL—Control Register. Controls the operation of the on-chip BPU. The format of BCNTL is shown in *Figure 2-9*. Upon reset, all the implemented bits in BCNTL are set to 0.

7	6	5	4	3	2	1	0
BRME	BLME	ERME	ELME	Res	EBPU	SPRD	Res

FIGURE 2-9. Control Register (BCNTL)

2.0 Architectural Description (Continued)

- SPRD** Source Preread. Specifies whether one or two operands are preread at the beginning of each scan line.
 SPRD = 0 → One Operand is Preread
 SPRD = 1 → Two Operands are Preread
- EBPU** BPU Enable.
 EBPU = 0 → BPU Disabled
 EBPU = 1 → BPU Enabled
- ELME** End Left Mask Enable. Enables the left mask at the end of each scan line.
 ELME = 0 → End Left Mask Disabled
 ELME = 1 → End Left Mask Enabled
- ERME** End Right Mask Enable. Enables the right mask at the end of each scan line.
 ERME = 0 → End Right Mask Disabled
 ERME = 1 → End Right Mask Enabled
- BLME** Begin Left Mask Enable. Enables the left mask at the beginning of each scan line.
 BLME = 0 → Begin Left Mask Disabled
 BLME = 1 → Begin Left Mask Enabled
- BRME** Begin Right Mask Enable. Enables the right mask at the beginning of each scan line.
 BRME = 0 → Begin Right Mask Disabled
 BRME = 1 → Begin Right Mask Enabled

BC—BPU Count Register. Specifies the width of the destination data block. The format of BC is shown in *Figure 2-10*. The value of WCNT is determined by the relation $WCNT = 256 - N$. N is the number of 16-bit words in each scan line, and must be in the range 1 to 255.




FIGURE 2-10. BPU Count Register (BC)

2.1.6 DMA Controller Registers

The DMA Controller contains 15 32-bit memory mapped registers that are both readable and writable by software. Three of these registers are common to both channels and are used to report various conditions as well as selectively enable or disable interrupt generation corresponding to each condition. The other 12 registers are divided into two sets of six, with each set associated to one channel.

All of the registers appear as memory locations, and must be accessed by specifying an operand length of 32 bits. Accesses specifying a length other than 32 bits may cause unpredictable results. The registers DSTAT, ADC, BLTC and MODE must not be written into while the associated channel is enabled. Upon reset, the registers STAT, IMSK, BLTC, BLTR, MODE and CNTL are cleared. *Figure 2-11* shows the address map of the DMA registers.

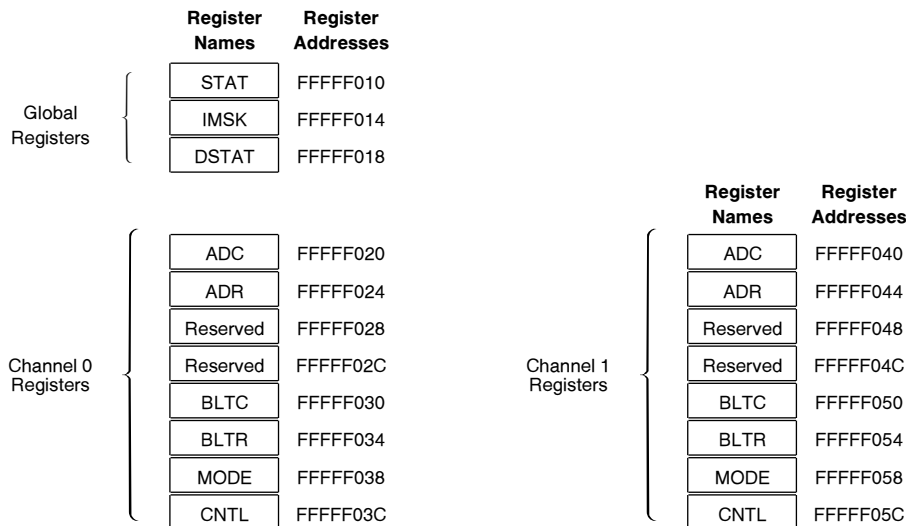


FIGURE 2-11. DMA Controller Registers Address Map

2.0 Architectural Description (Continued)

STAT—Status Register. This register contains status information for the two DMA channels. Its format is shown in *Figure 2-12*.




FIGURE 2-12. DMA Status Register (STAT)

- TC** Terminal Count; when set to 1, indicates that the transfer was completed by a terminal count condition (BLTC Register reached zero).
- EOT** External End Of Transfer. This bit is set to 1 when the transfer is externally terminated by the assertion of the $\overline{\text{EOT}}$ signal.
- OVR** Channel Overrun. Used only with double-buffered transfers. OVR is set to 1 when the present transfer is completed (BLTC = 0), but the parameters for the next transfer (address and block length) are not valid.
- CHAC** Channel Active. When set to 1, indicates that the channel is active (CHEN bit in register CNTL is 1 and BLTC > 0).

The TC, EOT and OVR bits are sticky. This means that, once set by the occurrence of the specific condition, they will remain set until explicitly cleared by software. These bits can be individually cleared by writing a value into the STAT register with the bit positions to be cleared set to 1. The CHAC bit continuously reflects the active or inactive status of the channel, and therefore, it is READ only.

IMSK—Interrupt Mask. This register is used to enable or disable interrupts for various conditions recorded in the STAT register. The format of IMSK is shown in *Figure 2-13*.




FIGURE 2-13. DMA Interrupt Mask Register (IMSK)

The DIP bit selects the DMA Interrupt Priority level in the Interrupt Control Unit. When a DIP is 0, the priority level is 6; when DIP is 1, the priority level is 14. Bits 0–2 and 4–6 are the Interrupt Mask Bits. An interrupt is enabled when the corresponding mask bit is set to 1.

DSTAT—DMA Debug Status Register. DSTAT is an image of the STAT register and provides an alternate means for accessing STAT. Writing a value into DSTAT stores the value itself into the STAT register, as opposed to only clearing bits as when writing directly into STAT. Setting a bit into DSTAT affects the associated interrupt as well as the CHEN bit in the CNTL register, according to the respective mask bit in IMSK.

ADC—Device Address Counter. Holds the current address of either the source data item or the destination location in the Addressed Device. If the AD bit in the MODE register is set to 1, ADC is incremented after each transfer cycle by the number of bytes transferred.

ADR—Device Address Register. Holds the starting address of either the source data block or the destination data area in the Addressed Device.

BLTC—Block Length Counter. Holds the current number of bytes to be transferred. BLTC is decremented after each transfer cycle by the number of bytes transferred.

BLTR—Block Length Register. Holds the number of bytes in the next block to be transferred. Writing a zero value into BLTR while the VLD and CHEN bits in CNTL are both set to 1 may cause unpredictable results.

Note 1: The ADR and BLTR registers are used to store the transfer parameters (i.e., address and block length) for the next data block to be transferred, for either auto-initialize or double-buffer modes of operation.

Note 2: The values programmed into ADC, ADR, BLTC and BLTR must be multiples of the bus width programmed in the mode register.

MODE—Mode Control Register. This register is used to specify the channel operating mode. The format of MODE is shown in *Figure 2-14*. All the reserved bits should be set to 0.




FIGURE 2-14. DMA Mode Control Register (MODE)

- OT** Operation Type.
 OT = 0 → Auto-Initialize Mode Disabled
 OT = 1 → Auto-Initialize Mode Enabled
- DIR** Transfer Direction. Specifies the direction of transfer between memory and peripheral device.
 DIR = 0 → Peripheral Device is Destination
 DIR = 1 → Peripheral Device is Source
- BW** Device Bus Width.
 BW = 00 → 8 Bits
 BW = 01 → 16 Bits
 BW = 10 → Reserved
 BW = 11 → Reserved
- AD** Device Address Control. Enables the incrementing of the device address after each transfer cycle.
 AD = 0 → Address Unchanged
 AD = 1 → Address Incremented
- BLT** Block Transfer Length. This 5-bit field is used by the bus fairness mechanism, and specifies the maximum number of bytes that can be transferred before the DMA channel relinquishes the bus. Refer to Section 3.4.2.5 for details.

CNTL—Channel Control Register. CNTL is used to synchronize the channel operation with the programming of the block transfer parameters. The format of CNTL is shown in *Figure 2-15*.




FIGURE 2-15. Channel Control Register (CNTL)

- CHEN** Channel Enable.
 CHEN = 0 → Channel Disabled
 CHEN = 1 → Channel Enabled
- VLD** Transfer Parameters Valid. Specifies whether the transfer parameters for the next block to be transferred are valid.
 VLD = 0 → Parameters Not Valid
 VLD = 1 → Parameters Valid

2.0 Architectural Description (Continued)

The CHEN bit is set to 0 in the following cases.

- Upon Reset
- Software clears it by writing to the CNTL register
- The EOT bit in STAT is set to 1
- The OVR bit in STAT is set to 1 and is unmasked

In the last two cases the CHEN bit is forced to 0 and cannot be set to 1 by software unless EOT is cleared and OVR is either cleared or masked by clearing the corresponding bit in IMSK.

2.1.7 Interrupt Control Unit (ICU) Registers

The interrupt control unit contains two memory-mapped registers: IVCT and ISRV. These registers are 8 bits and 16 bits wide respectively, and must be accessed by specifying operand lengths equal to their widths. Specifying different operand lengths may cause unpredictable results. The address map for IVCT and ISRV is shown in *Figure 2-16*.

Register Names	Register Addresses
IVCT	FFFFFFE0
ISRV	FFFFFFE04

FIGURE 2-16. ICU Registers Address Map

IVCT—Interrupt Vector Register. This is a read-only register containing the current interrupt request vector to be used by the CPU in accessing the interrupt dispatch table when the request is acknowledged. Bits 0–3 contain an encoded value representing the interrupt request priority level. The binary value 1111 represents the highest priority level, while the value 0000 indicates that no interrupt request is pending. *Figure 2-17* shows the format of IVCT.

7	6	5	4	3	2	1	0
0	0	0	1	V	V	V	V

FIGURE 2-17. Interrupt Vector Register (IVCT)

ISRV—Interrupt In-Service Register. ISRV is used by the ICU to keep track of the priority levels of interrupts currently being serviced.

A value of 1 in bit position i (where $1 \leq i \leq 15$), indicates that the i -th priority level is currently in-service. Bit position 0 is always forced to 0. Upon reset the ISRV register is cleared to 0. Note that a zero value in the ISRV register indicates that there are no in-service interrupts.

2.1.8 Timers Registers

Each of the three timers in the NS32CG160 is controlled by a set of four registers. These registers are all 16 bits wide, and are memory-mapped. Their address map is shown in *Figure 2-18*. Accesses to the timer registers must specify operand lengths of 16 bits otherwise undefined results may be obtained. Refer to Section 3.4.4 for details on the timer's operation.

	Register Names	Register Addresses
	Reserved	FFFFF800
	Reserved	FFFFF804
	Reserved	FFFFF808
	Reserved	FFFFF80C
Timer 0	TC	FFFFF810
	TRCA	FFFFF814
	TRCB	FFFFF818
	TCNTL	FFFFF81C
Timer 1	TC	FFFFF820
	TRCA	FFFFF824
	TRCB	FFFFF828
	TCNTL	FFFFF82C
Timer 2	TC	FFFFF830
	TRCA	FFFFF834
	TRCB	FFFFF838
	TCNTL	FFFFF83C

FIGURE 2-18. Timer Registers Address Map

TC—Counter. TC is a down-counter which, upon underflow, is reloaded with the contents of either TRCA or TRCB for modes 1 and 2, and with all 1's for mode 3.

TRCA, TRCB—Reload/Capture Registers A and B. These registers hold either the counter reload values or a snapshot of TC.

TCNTL—Timer Control Register. Used to control the operation of each timer, and to enable timer interrupts. The TCNTL format is shown in *Figure 2-19*.

31	12	11	10	9	8	7	5	4	0	
Res	WIS	WIA	WIB	PRC	TMC	TCS	IPFA	IENA	IPFB	IENB

FIGURE 2-19. Timer Control Register (TCNTL)

IENB Interrupt Enable Bit B. When set to 1, enables the interrupt from IPFB.

IPFB Interrupt Pending Flag B.

IENA Interrupt Enable Bit A. When set to 1, enables the interrupt from IPFA. If Mode 3 is selected, it enables the interrupt from TCS as well.

IPFA Interrupt Pending Flag A.

TCS Timer Control and Status. In modes 1 or 2, this bit is used to start and stop the timer. The timer starts when TCS is 1. In mode 3, TCS is the underflow interrupt pending flag.

TMC Timer Mode Control. This three-bit field selects the timer mode of operation. (See Section 3.4.4.)

PRC Prescaler Control. Used only in modes 1 and 3. PRC controls the frequency of the timer input clock (TCLK). When $PRC = 0$, $TCLK = CTTL/8$; when $PRC = 1$, $TCLK = CTTL/4096$.

2.0 Architectural Description (Continued)

WIB Write Inhibit B. When WIB is set to 1, writing into IPFB is inhibited.

WIA Write Inhibit A. When WIA is set to 1, writing into IPFA is inhibited.


WIS Write Inhibit S. When WIS is set to 1, writing into TCS is inhibited.

Note: The value read from WIB, WIA and WIS is undefined.

2.2 MEMORY ORGANIZATION


The 32CG160 implements full 32-bit addresses internally, of which only the least significant 24 bits are available externally. This allows the device to access up to 16 Mbytes of memory.

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




Byte at Address A

Two contiguous bytes are called a word. Except where noted, 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 address 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 NS32CG160 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 NS32CG160, 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 control registers of the on-chip peripherals. *Figure 2-20* shows the NS32CG160 address mapping.

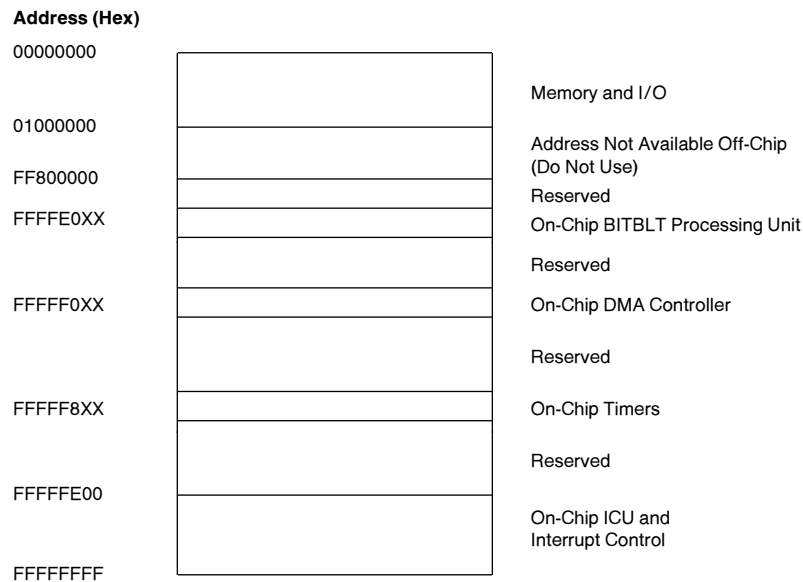


FIGURE 2-20. NS32CG160 Address Mapping

2.0 Architectural Description (Continued)

2.3 MODULAR SOFTWARE SUPPORT

The NS32CG160 provides special support for software modules and modular programs.

Each module in a NS32CG160 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 NS32CG160 uses a module table in memory and two registers in the CPU.

The Module Table is located within the first 64 kbytes 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:

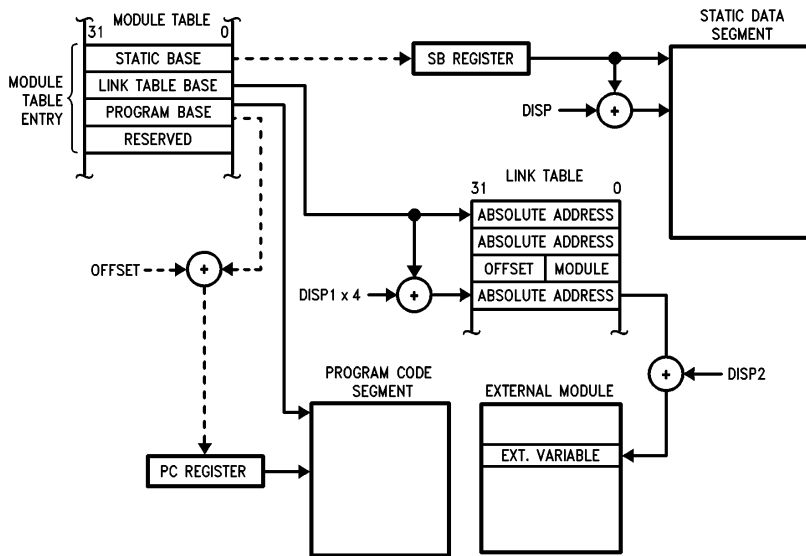
- The Static Base entry contains the address of the beginning of the module's static data segment.
- The Link Table Base points to the beginning of the module's Link Table.
- The Program Base is the address of the beginning of the code and constant data for the module.
- A fourth entry is currently unused 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 NS32CG160 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 2-21* shows a typical NS32CG160 run-time environment.



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

FIGURE 2-21. NS32CG160 Run-Time Environment.

TL/EE/10752-2

2.0 Architectural Description (Continued)

2.4 INSTRUCTION SET

2.4.1 General Instruction Format

Figure 2-22 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 2-24, 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).

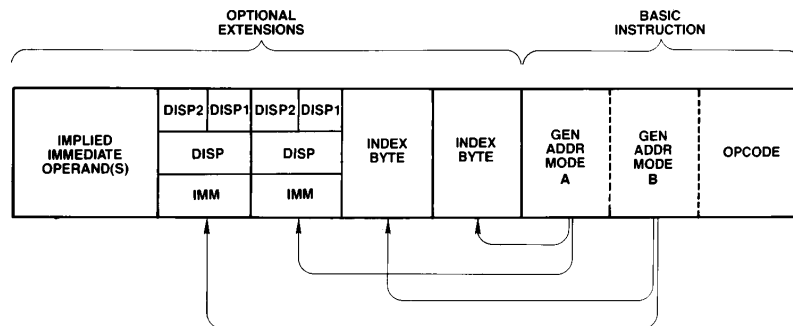


FIGURE 2-22. General Instruction Format

TL/EE/10752-3




FIGURE 2-23. Index Byte Format

TL/EE/10752-4

2.0 Architectural Description (Continued)


Byte Displacement: Range -64 to +63



Word Displacement: Range -8192 to +8191



Double Word Displacement:
Range (Entire Addressing Space)



TL/EE/10752-5

FIGURE 2-24. Displacement Encodings

2.4.2 Addressing Modes

The NS32CG160 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 NS32CG160 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.

NS32CG160 Address 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 registers. 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 2-4 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 warp or pitch are available on several of the graphics instructions.

2.0 Architectural Description (Continued)

TABLE 2-4. NS32CG160 Addressing Modes

ENCODING	MODE	ASSEMBLER SYNTAX	EFFECTIVE ADDRESS
Register			
00000	Register 0	R0 or F0	None: Operand is in the 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 Disp1 + Register. "SP" is either SP0 or SP1, as selected in PSR.
10001	Stack Memory Relative	disp2(disp1 (SP))	
10010	Static Memory Relative	disp2(disp1 (SB))	
Reserved			
10011	(Reserved for Future Use)		
Immediate			
10100	Immediate	Value	None: Operand is input from instruction queue.
Absolute			
10101	Absolute	@disp	Disp.
External			
10110	External	EXT (disp1) + disp2	Disp2 + Pointer; Pointer is found at Link Table Entry number Disp 1.
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.
<p>"Mode" and "n" are contained within the Index Byte. EA (mode) denotes the effective address generated using mode.</p>			

2.0 Architectural Description (Continued)

2.4.3 Instruction Set Summary

Table 2-5 presents a brief description of the NS32CG160 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 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 2-5. NS32CG160 Instruction Set Summary

MOVES

Format	Operation	Operands	Description
4	MOVi	gen,gen	Move a Value
2	MOVQi	short,gen	Extend and Move a Signed 4-Bit Constant
7	MOVMi	gen,gen,disp	Move Multiple: Disp Bytes (1 to 16)
7	MOVZBW	gen,gen	Move with Zero Extension
7	MOVZID	gen,gen	Move with Zero Extension
7	MOVXBW	gen,gen	Move with Sign Extension
7	MOVXID	gen,gen	Move with Sign Extension
4	ADDR	gen,gen	Move Effective Address

INTEGER ARITHMETIC

Format	Operation	Operands	Description
4	ADDi	gen,gen	Add
2	ADDQi	short,gen	Add Signed 4-Bit Constant
4	ADDCi	gen,gen	Add with Carry
4	SUBi	gen,gen	Subtract
4	SUBCi	gen,gen	Subtract with Carry (Borrow)
6	NEGi	gen,gen	Negate (2's Complement)
6	ABSi	gen,gen	Take Absolute Value
7	MULi	gen,gen	Multiply
7	QUOi	gen,gen	Divide, Rounding toward Zero
7	REMi	gen,gen	Remainder from QUO
7	DIVi	gen,gen	Divide, Rounding Down
7	MODi	gen,gen	Remainder from DIV (Modulus)
7	MELi	gen,gen	Multiply to Extend Integer
7	DELi	gen,gen	Divide Extended Integer

PACKED DECIMAL (BCD) ARITHMETIC

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

2.0 Architectural Description (Continued)

TABLE 2-5. NS32CG160 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 2-5. NS32CG160 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	Adjst 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 2-5. NS32CG160 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
6	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 NS32CG160 graphics support capabilities. Basic discussions on frame buffer addressing and BITBLT operations are also provided. More detailed information on the NS32CG160 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 2-25*. 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.

$$ADDR = ORG + y * WARP + x$$

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 NS32CG160 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
ADD   R0, R1            ; +X = BIT OFFSET
SBITD R1, ORG           ; SET PIXEL
```

Instruction Sequence 2:

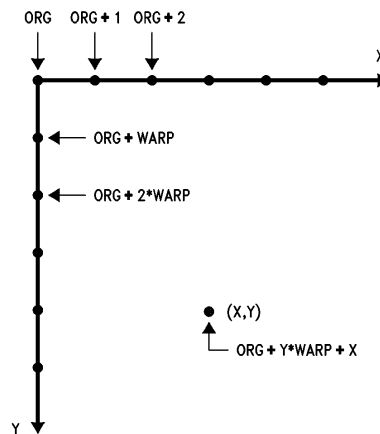
```
INDEXD R1, (WARP-1), R0 ; Y*WARP + X
SBITD  R1, ORG          ; SET PIXEL
```

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
ADD   (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/10752-6

FIGURE 2-25. 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 Raster-Op: 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.

2.5.2.1 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.

2.0 Architectural Description (Continued)

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 2-26*, 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 book.

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 2-26* 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 2-26. 32-Pixel by 32-Scan Line Frame Buffer

TL/EE/10752-7

2.0 Architectural Description (Continued)

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.

2.5.2.4 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 NS32CG160 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.

Figure 2-27(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 2-27(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 occurs 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 2-27(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.



TL/EE/10752-8

(a)



TL/EE/10752-9

(b)

FIGURE 2-27. Overlapping BITBLT Blocks

2.0 Architectural Description (Continued)

2.5.2.5 BITBLT Variations

The “classical” definition of BITBLT, as described in “Small talk-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” processor used in printing.

Texture op 1 Source op2 Destination → Destination

While the NS32CG160 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 NS32CG160 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 NS32CG160.

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 2-26* 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 2-28* and *Table 2-6* 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 2-28. BB Instructions Format

TABLE 2-6. “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.

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 2-29*.




FIGURE 2-29. BITWT Instruction Format

2.0 Architectural Description (Continued)

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.

This instruction performs an entire BITBLT operation in conjunction with the on-chip BPU as well as any external BPU. The BPU Control Registers should be loaded by the software before the instruction is executed. The NS32CG160 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.

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 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 2-30 shows the EXTBLT format. The bus activity for a simple BITBLT operation is shown in Figure 2-35.




FIGURE 2-30. 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 2-31.




FIGURE 2-31. 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 equaled the maximum run length value and the bit was not found. It is cleared otherwise. Figure 2-32 shows the TBITS instruction format.



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

FIGURE 2-32. TBITS Instruction Format

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.

2.0 Architectural Description (Continued)

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 1 kbyte 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 2-33* shows the SBITS instruction format.




FIGURE 2-33. 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° 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° rotation. A positive raster warp value leads to a 270° rotation. If the R3 value is = (space warp + 1 or - 1), then the result is a 45° line. If the R3 value is + 1 or - 1, a horizontal line results.

SBITS and SBITPS allow expansion on any 90° angle, giving portrait, landscape and mirror images from one font. *Figure 2-34* shows the SBITPS instruction format.




FIGURE 2-34. SBITPS Instruction Format

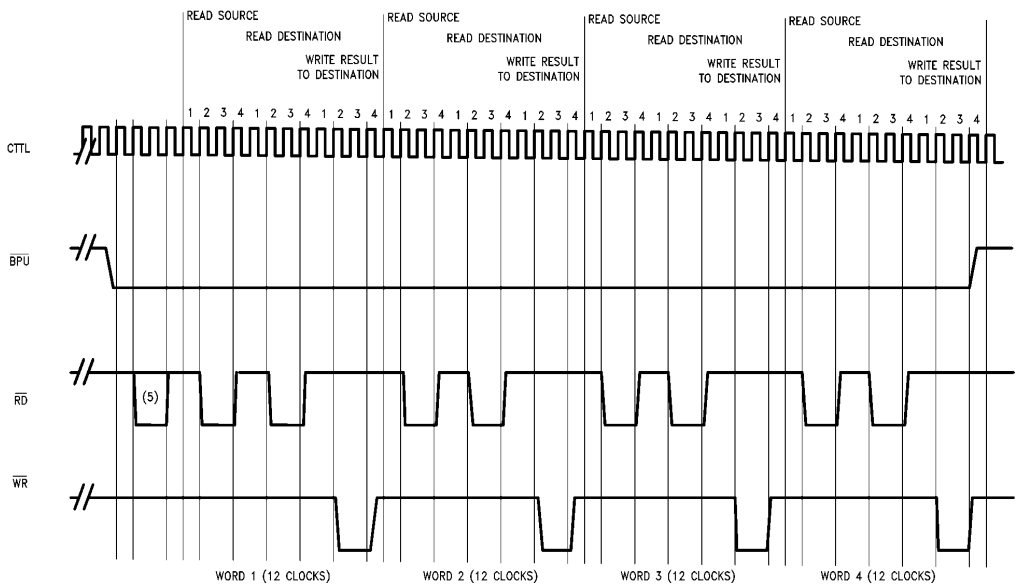
2.5.3.3.1 Magnifying Compressed Data

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.

2.0 Architectural Description (Continued)



TL/EE/10752-10

FIGURE 2-35. 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.

3.0 Functional Description

This chapter provides details on the functional characteristics of the NS32CG160 microprocessor.

The chapter is divided into five main sections:

Instruction Execution, Exception Processing, Debugging, On-Chip Peripherals and System Interface.

3.1 INSTRUCTION EXECUTION

To execute an instruction, the NS32CG160 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.




FIGURE 3-1. Operating States

TL/EE/10752-11

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.

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 (see note).

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.

Note: When the Direct-Exception mode is enabled, the CPU does not save the MOD Register contents nor does it read the module linkage information for the exception service procedure. Refer to Section 3.2 for details.

3.1.2 Instruction Endings

The NS32CG160 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.

3.1.2.1 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.

3.0 Functional Description (Continued)

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 T-flag must be copied to the saved P-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. A terminated instruction cannot be completed.

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 NS32CG160 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.

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).

3.0 Functional Description (Continued)

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

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.

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).




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 generated either internally, by the on-chip DMA Channels or Timers, or externally, by activating NMI or one or more of the IRO-3 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 optionally 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. Trap (TRC) always disabled. Maskable interrupts are also disabled if the exception is caused by an interrupt.
2. Vector Acquisition. A vector is either obtained from the on-chip interrupt control unit or is supplied internally by default.
3. Service Call. The CPU performs one of two sequences common to all exceptions to complete the acknowledge process and enter the appropriate service procedure. The selection between the two sequences depends on whether the Direct-Exception mode is disabled or enabled.

Direct-Exception Mode Disabled

The Direct-Exception mode is disabled while the DE bit in the CFG register is 0 (Section 2.1.4). In this case the CPU first pushes the saved PSR copy along with the contents of the MOD and PC registers on the interrupt stack. Then it reads the double-word entry from the Interrupt Dispatch table at address "INTBASE + vector × 4". See Figures 3-4 and 3-5. The CPU uses this entry to call the exception service procedure, interpreting the entry as an external procedure descriptor.

A new module number is loaded into the MOD register from the least-significant word of the descriptor, and the static-base pointer for the new module is read from memory and loaded into the SB register. Then the program-base pointer for the new module is read from memory and added to the most-significant word of the module descriptor, which is interpreted as an unsigned value. Finally, the result is loaded into the PC register.

3.0 Functional Description (Continued)

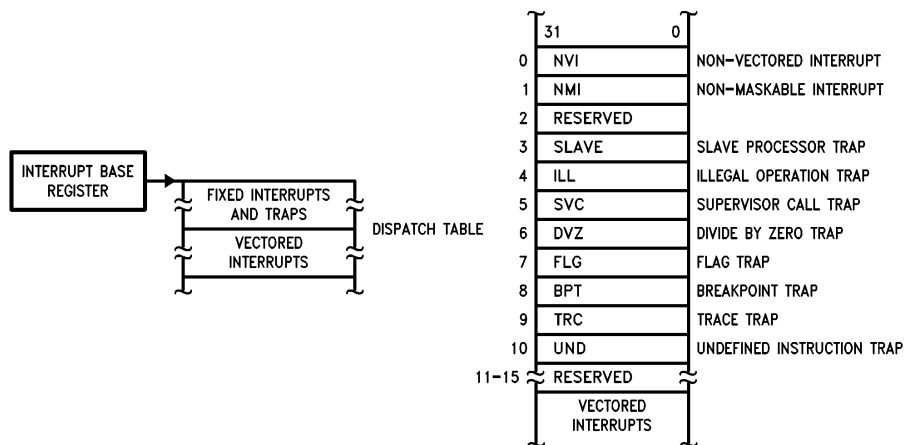


FIGURE 3-4. Interrupt Dispatch Table

TL/EE/10752-13

Direct-Exception Mode Enabled

The Direct-Exception mode is enabled when the DE bit in the CFG register is set to 1. In this case the CPU first pushes the saved PSR copy along with the contents of the PC register on the Interrupt Stack. The word stored on the Interrupt Stack between the saved PSR and PC register is reserved for future use; its contents are undefined. The CPU then reads the double-word entry from the Interrupt Dispatch Table at address "INTBASE + vector × 4". The CPU uses this entry to call the exception service procedure, interpreting the entry as an absolute address that is simply loaded into the PC register. *Figure 3-6* provides a pictorial of the acknowledge sequence. It is to be noted that while the direct-exception mode is enabled, the CPU can respond more quickly to interrupts and other exceptions because fewer memory references are required to process an exception. The MOD and SB registers, however, are not initialized before the CPU transfers control to the service procedure. Consequently, the service procedure is restricted from executing any instructions, such as CXP, that use the contents of the MOD or SB registers in effective address calculations.

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 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 the on-chip ICU as well as any external interrupt control logic 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 Program Counter (PC) and the Processor Status Register from the interrupt stack. If the Direct-Exception mode is disabled, they also restore the MOD and SB register contents. *Figure 3-7* and *3-8* show the RETT and RETI instruction flows when the Direct-Exception mode is disabled.

3.0 Functional Description (Continued)




TL/EE/10752-14




TL/EE/10752-15

FIGURE 3-5. Exception Acknowledge Sequence:
Direct-Exception Mode Disabled.

3.0 Functional Description (Continued)



TL/EE/10752-16



TL/EE/10752-17

**FIGURE 3-6. Exception Acknowledge Sequence:
Direct-Exception Mode Enabled.**

3.0 Functional Description (Continued)

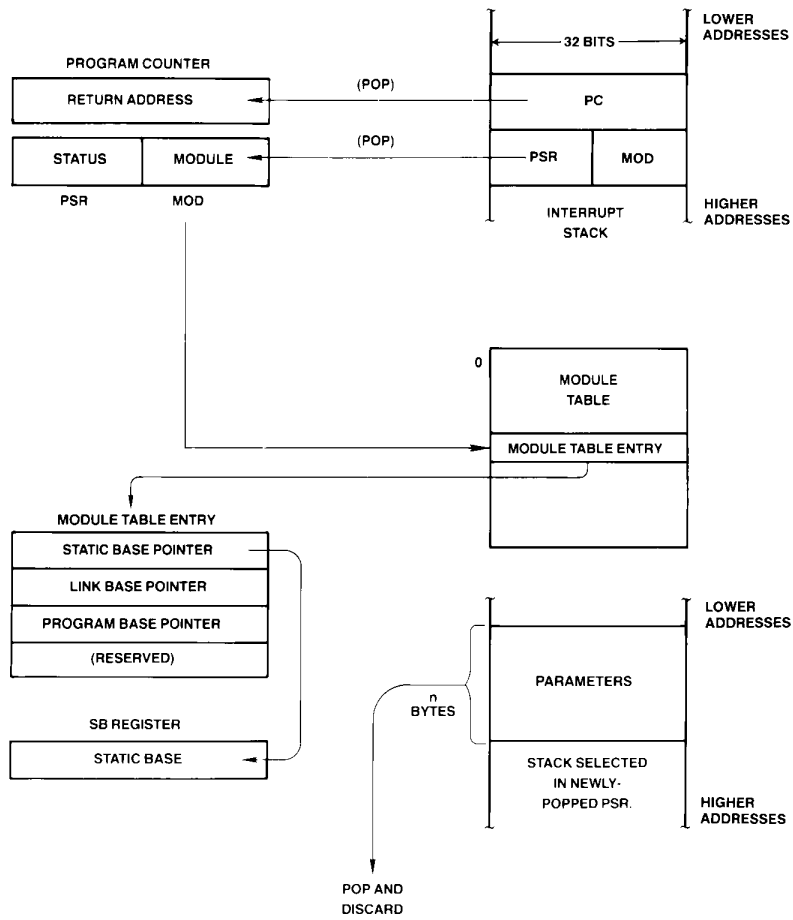


FIGURE 3-7. Return from Trap (RETT n) Instruction Flow: Direct-Exception Mode Disabled.

TL/EE/10752-18

3.0 Functional Description (Continued)

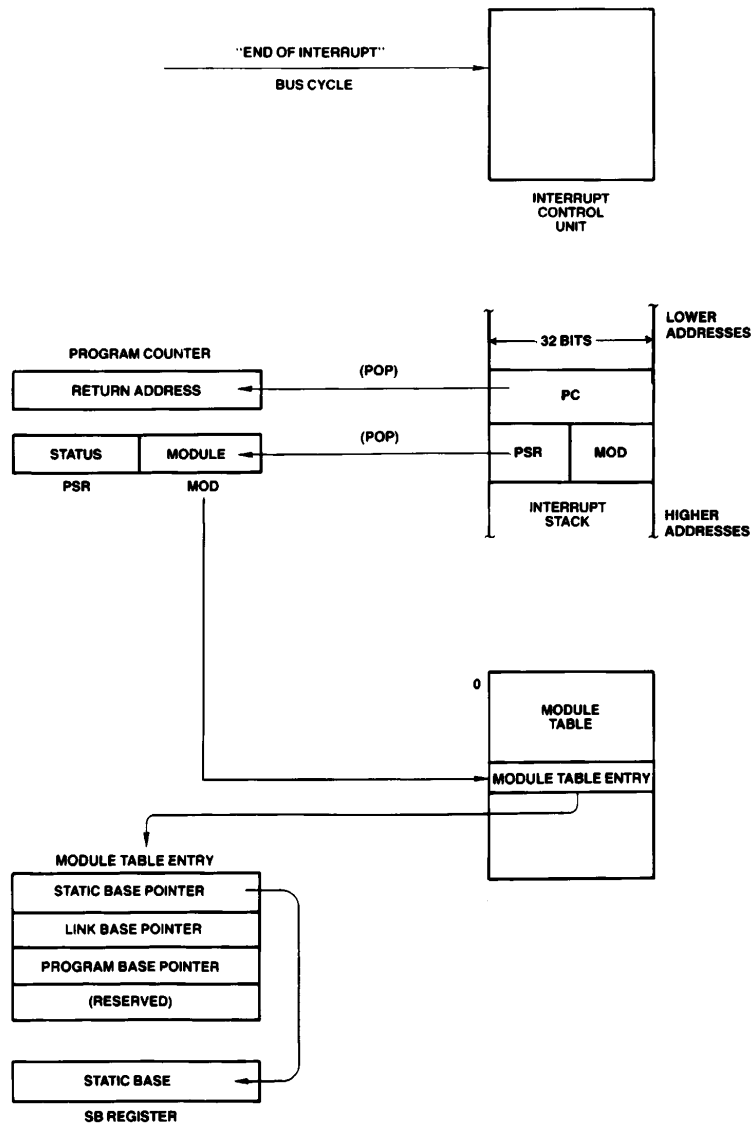


FIGURE 3-8. Return from Interrupt (RETI) Instruction Flow: Direct-Exception Mode Disabled.

TL/EE/10752-19

3.0 Functional Description (Continued)

3.2.3 Maskable Interrupts

Maskable interrupt requests are generated either externally through the $\overline{IR}0-3$ pins or internally by the DMA controller or the timers. These requests are enabled to generate an interrupt only while the I-bit in the PSR register is set to 1. The I-bit is automatically cleared during service of a maskable interrupt or NMI, and is restored to its original setting upon return from the interrupt service routine via the RETT or RETI instruction.

Maskable interrupts can be configured through the I-bit in the CFG register to be either non-vectored (CFG bit I = 0) or vectored (CFG bit I = 1).

If the non-vectored mode is selected, a default vector value of zero is always used. For the vectored mode instead, the on-chip Interrupt Control Unit will provide the CPU with a vector value. This vector value 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 on-chip ICU that it may re-prioritize any interrupt requests still pending.

3.2.4 Non-Maskable Interrupt

The Non-Maskable Interrupt is triggered whenever a falling edge is detected on the \overline{NMI} pin. The CPU performs an "Interrupt Acknowledge" bus cycle from Address $FFFFFF00_{16}$ 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 NS32CG160 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 3-9 shows an exception processing flowchart.

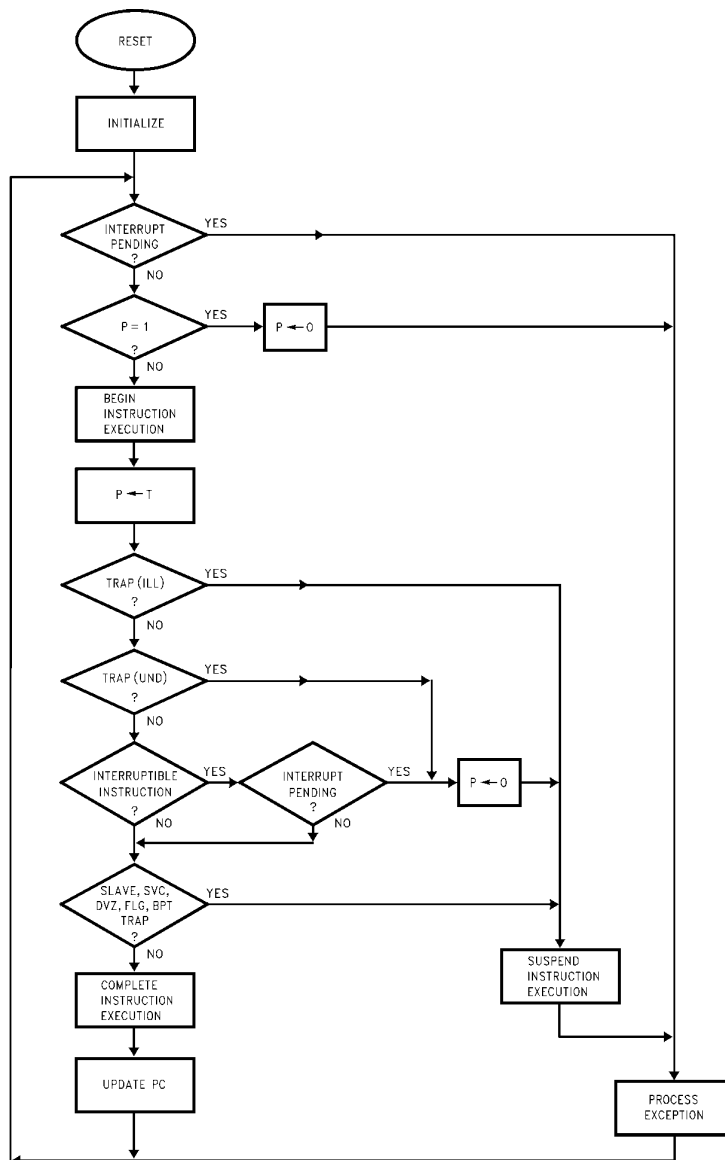
Before executing an instruction, the CPU checks for pending interrupts, or Trap (TRC). The CPU responds to any pending interrupt requests; nonmaskable interrupts are recognized with higher priority than maskable interrupts. If no interrupts are pending, then the CPU checks the P-flag in the PSR to determine 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/10752-20

FIGURE 3-9. Exception Processing Flowchart

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 3-10*.

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{NMI} pin receives a falling edge, or an interrupt event is signaled either through $\overline{IR0-3}$ or by the on-chip peripherals and PSR I-bit is 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, TBITPS), at the next interruptible point during its execution. The graphics instructions are interruptible.

3.0 Functional Description (Continued)

1. If an interruptible 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, T, U, S, P and I.
3. If the interrupt is Non-Maskable:
 - a. Read a byte from address FFFFFFF0₁₆, applying Status Code 0100 (Interrupt Acknowledge). Discard the byte read.
 - b. Set "Vector" to 1.
 - c. Go to Step 6.
4. If the interrupt is Non-Vectored:
 - a. Read a byte from address FFFFFFFE0₁₆, applying Status Code 0100 (Interrupt Acknowledge). Discard the byte read.
 - b. Set "Vector" to 0.
 - c. Go to Step 6.
5. Here the interrupt is Vectored.
 - a. Do a dummy read cycle from address FFFFFFFE0₁₆, applying status code 0100 (Interrupt Acknowledge), and discarding the byte read. This is to notify external circuitry that the interrupt is being acknowledged.
 - b. Read vector byte from the IVCT register of the on-chip Interrupt Control Unit.
6. Perform Service (Vector, Return Address), *Figure 3-10*.

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 3-10*.

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 T, U and S.
3. Set "Vector" to 9.
4. Set "Return Address" to the address of the next instruction.
5. Perform Service (Vector, Return Address), *Figure 3-10*.

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

Service (Vector, Return Address):

1. Push the PSR copy onto the Interrupt Stack as a 16-bit value.
2. Read 32-bit Interrupt Dispatch Table (IDT) entry at address "INTBASE + vector × 4".
3. If Direct-Exception mode is selected, then go to Step 10.
4. Move the LS word of the IDT entry (Module Field) into the temporary MOD register.
5. Read the Program Base pointer from memory address "MOD + 8", and add to it the M.S. word of the IDT entry (Offset Field), placing the result in the Program Counter.
6. Read the new Static Base pointer from the memory address contained in MOD, placing it into the SB Register.
7. Push MOD Register into the Interrupt Stack as a 16-bit value.
8. Copy temporary MOD register into MOD Register.
9. Go to Step 11.
10. Place IDT entry in the Program Counter.
11. Push the Return Address onto the Interrupt Stack as a 32-bit quantity.
12. Flush queue: Non-sequentially fetch first instruction of Exception Service Routine.

FIGURE 3-10. Service Sequence

3.0 Functional Description (Continued)

3.3 DEBUGGING SUPPORT

The NS32CG160 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 NS32CG160 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 NS32CG160 instruction will not be re-entered, due to a

single-step trap. Both the NMI and maskable 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 NS32CG160 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 NS32CG160 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 NS32CG160 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 ON-CHIP PERIPHERALS

Four types of on-chip peripherals are provided in the NS32CG160: A BITBLT processing unit, a DMA controller, an interrupt control unit and timers. Details on the operation of these peripherals are provided in the following sections.

3.4.1 BITBLT Processing Unit (BPU)

The on-chip BPU is designed to efficiently execute BITBLT operations, and is used by the EXTBLT instruction.

The BPU control registers as well as the general purpose registers must be loaded with the appropriate data before the EXTBLT instruction is executed. Refer to sections 2.1.5 and 2.5, or to the NS32CG16 Printer/Display Processor Programmer's Reference Supplement for details.

For color applications, where multiple bit-planes are required, external BPUs, like the National DP8511, can be used in conjunction with the on-chip BPU. In this case each BPU controls one bit-plane.

The software can enable or disable the on-chip BPU or the external BPUs individually by programming the appropriate bits in the various BPU's control registers.


3.0 Functional Description (Continued)

TABLE 3-3. Barrel Shifter Truth Table

SN Field	c0	c1	c2	c3	c4	c5	c6	c7	c8	c9	c10	c11	c12	c13	c14	c15
0	a0	a1	a2	a3	a4	a5	a6	a7	a8	a9	a10	a11	a12	a13	a14	a15
1	a1	a2	a3	a4	a5	a6	a7	a8	a9	a10	a11	a12	a13	a14	a15	b0
2	a2	a3	a4	a5	a6	a7	a8	a9	a10	a11	a12	a13	a14	a15	b0	b1
3	a3	a4	a5	a6	a7	a8	a9	a10	a11	a12	a13	a14	a15	b0	b1	b2
4	a4	a5	a6	a7	a8	a9	a10	a11	a12	a13	a14	a15	b0	b1	b2	b3
5	a5	a6	a7	a8	a9	a10	a11	a12	a13	a14	a15	b0	b1	b2	b3	b4
6	a6	a7	a8	a9	a10	a11	a12	a13	a14	a15	b0	b1	b2	b3	b4	b5
7	a7	a8	a9	a10	a11	a12	a13	a14	a15	b0	b1	b2	b3	b4	b5	b6
8	a8	a9	a10	a11	a12	a13	a14	a15	b0	b1	b2	b3	b4	b5	b7	b7
9	a9	a10	a11	a12	a13	a14	a15	b0	b1	b2	b3	b4	b5	b6	b7	b8
10	a10	a11	a12	a13	a14	a15	b0	b1	b2	b3	b4	b5	b6	b7	b8	b9
11	a11	a12	a13	a14	a15	b0	b1	b2	b3	b4	b5	b6	b7	b8	b9	b10
12	a12	a13	a14	a15	b0	b1	b2	b3	b4	b5	b6	b7	b8	b9	b10	b11
13	a13	a14	a15	b0	b1	b2	b3	b4	b5	b6	b7	b8	b9	b10	b11	b12
14	a14	a15	b0	b1	b2	b3	b4	b5	b6	b7	b8	b9	b10	b11	b12	b13
15	a15	b0	b1	b2	b3	b4	b5	b6	b7	b8	b9	b10	b11	b12	b13	b14

3.4.1.1 BPU Operation

A block diagram of the on-chip BPU is shown in *Figure 3-11*. The BPU performs the fundamental operations required for raster graphics applications: shifting, masking and bitwise logic operations. These operations are carried out in the BPU via two major functional blocks: the Barrel Shifter and the BITBLT Logic Unit.



TL/EE/10752-21

FIGURE 3-11. BPU Block Diagram

The function of the Barrel shifter is to align the source data to the destination data. Bit alignment may cross word boundaries. The Barrel Shifter is implemented as a 32-to-16-bit multiplexer. The shifting operation is controlled by the SN field in the BBSC register. Let A = a0 – a15 and B = b0 – b15 represent the least significant and most significant 16-bit input words to the Barrel Shifter, and C = c0 – c15 the 16-bit output word. An n-shift operation on A and B is denoted by concatenating the least significant n bits of word B to the most significant 16-n bits of word A. A pictorial representation of the Barrel Shifter operation is shown in Table 3-3.

A multiplexer precedes the Barrel Shifter to allow the input words to be swapped if necessary.

The BIS bit in the BBSC register controls the routing of the two 16-bit input words, one in the Barrel Input Latch BIL and the other in the Source Data Input Latch SDIL, to the least significant and most significant word positions in the Barrel Shifter. When BIS is 0 the word in SDIL is routed to the least significant position with BIL data going to the most significant position. The opposite takes place when BIS is 1.

The BITBLT Logic Unit is responsible for the Mask and Bitwise logical operations performed on the source and destination data.

To execute a BITBLT operation, a source data word is first fetched from memory and stored into the Source data Input Latch SDIL. A second source data word is subsequently fetched and also stored into SDIL while the first data word is simultaneously transferred into the Barrel Input Latch BIL. These two data words are loaded into the Barrel Shifter which produces an aligned 16-bit output word. This word is then routed to the BITBLT Logic Unit, where it is masked and logically combined with a destination data word previously fetched and stored into the Destination Data Input Latch DDIL. The result is then routed to the output port, so it can be written back to memory.

3.4.2 DMA Controller

The on-chip DMA Controller provides 2 channels for transferring blocks of data between memory and I/O devices with minimal CPU intervention. Source or destination address as well as block size and type of operation are set up in advance by programming the appropriate control register. Actual transfers are handled by the DMA channels in response to external transfer requests. Upon receiving a transfer request from an I/O device, the DMA Controller performs the following operations:

1. Acquires control of the bus.
2. Acknowledges the requesting I/O device by asserting the appropriate DAK signal.
3. Starts executing data transfer cycles according to the values stored into the control registers of the channel being serviced.
4. Terminates the data transfer operation whenever one of the following events occurs:
 - The specified number of bytes has been transferred

3.0 Functional Description (Continued)

- The \overline{EOT} signal is activated during a data transfer cycle
- The software writes into the appropriate control registers

Each channel can perform data transfers in flyby (direct) mode. In addition, multiple transfer operations can be chained together by programming the appropriate control registers. Detailed descriptions of the different modes of operation are provided in the following sections.

3.4.2.1 DMA Data Transfers

In a DMA data transfer each data item is transferred using a single bus cycle without reading the data into the DMA controller. This mode offers a fast transfer rate, but the source and destination bus widths must be the same. Data transfers cannot occur between two memory elements. One of the elements must always be an I/O device selected by the channel's DAK signal. This device is referred to as the implied I/O device. The other element can be either memory or another I/O device, and is referred to as the addressed I/O device.

Since only one address is required, this address is taken from the corresponding ADC counter. The DMA Channel generates either a read or a write bus cycle according to the setting of the DIR bit in the MODE register. When the DIR bit is 0, a read bus cycle from the addressed device is performed and the data is written to the implied I/O device. When the DIR bit is 1, a write bus cycle to the addressed device is performed, and the data is read from the implied I/O device.

The number of bytes transferred in each cycle is taken from the BW field in the MODE register.

After the data element has been transferred, the BLTC counter is decremented by the number of bytes transferred. If the addressed device is memory, the ADC counter is also incremented by the same amount.

3.4.2.2 Single Transfer Operation

This mode provides the simplest way to accomplish a single block transfer operation. The block transfer address and byte count should be first written into the corresponding ADC and BLTC counters, and the OT bit in the MODE register should be programmed for non auto-initialize mode. When the CHEN bit in the CNTL register is set to 1 and VLD bit is set to 0, the channel becomes active and responds to external transfer requests.

When the BLTC counter reaches 0, the transfer operation terminates, the TC and OVR bits in the STAT register are set to 1, the CHAC bit is set to 0 and, if OVR is unmasked, the CHEN bit in the CNTL register is forced to 0. An interrupt can also be generated at the end of the transfer by setting the appropriate bits in the IMASK register.

3.4.2.3 Double-Buffer Operation

This mode allows the software to set up the next block transfer specification while the current block transfer is in progress. The operation is initialized by writing the block transfer address and byte count into the ADC and BLTC counters, and programming the OT bit in the MODE register for non auto-initialize mode. When the CHEN bit in the CNTL register is set to 1, the channel becomes active and responds to external transfer requests.

While the current block transfer is in progress, the software can write the address and byte count for the next block into the ADR and BLTR registers, and then set the VLD bit in the CNTL register to 1.

When the BLTC counter reaches 0, the TC bit is set to 1 and the DMA channel checks the value of the VLD bit. If it is 1, the channel copies the ADR and BLTR values into ADC and BLTC, and becomes ready to start the next block transfer. If the VLD bit is 0, the channel sets the OVR bit in the STAT register to 1, clears the CHAC bit and, if OVR is unmasked, it forces the CHEN bit to 0.

3.4.2.4 Auto-Initialize Operation

This mode allows the DMA controller to continuously fill the same memory area without software intervention. The operation is initialized by writing the block address and byte count into the ADC and BLTC counters as well as the ADR and BLTR registers, and programming the OT bit in the MODE register for auto-initialize mode. When the CHEN bit in the CNTL register is set to 1, the channel becomes active and responds to external requests.

When the BLTC counter reaches 0, the TC bit in the STAT register is set to 1, the contents of the ADR and BLTR registers are copied to the ADC and BLTC counters, and the operation is repeated.

3.4.2.5 Bus Arbitration

Whenever a DMA channel needs to perform a data transfer, it first needs to acquire control of the bus. Bus arbitration is performed according to a fixed priority scheme. An external HOLD request has the highest priority, followed by channel 0, channel 1 and, finally the CPU at the lowest priority. Once the bus is granted in response to a channel request and no higher priority request is pending, the channel can use the bus for a certain number of back-to-back transfers before it is forced to release it. This is controlled by a Bus Fairness mechanism whose purpose is to prevent bus monopolization from a DMA channel. The maximum number of back-to-back transfers can be programmed through the BLT field in the MODE register. When the programmed number of transfer is reached, the channel will release the bus for at least one clock cycle, so that it can be granted to the CPU. Table 3-4 shows the maximum number of back-to-back transfers for different values of the BLT field.

TABLE 3-4. Maximum Number of Back-to-Back DMA Transfers

BLT Field	Max. Number of Transfers	
	Byte Transfers	Word Transfers
00000	Unlimited	Unlimited
00001	1	1
00010	2	1
00100	4	2
01000	8	4
10000	16	8

Note: The values shown for the BLT field are the only ones allowed. Specifying a different value may cause unpredictable results.

3.4.3 Interrupt Control Unit (ICU)

The on-chip Interrupt Control Unit (ICU) manages up to 15 levels of prioritized interrupt requests. Requests can be generated either externally or internally. External requests are binary encoded as a 4-bit value, and are input to the ICU through the $\overline{IR0}$ – $\overline{IR3}$ pins. Internal requests are generated by the on-chip DMA controller and timers. Table 3-5 shows the possible interrupt sources and related priority levels.

3.0 Functional Description (Continued)

The ICU keeps track of the interrupt priority levels currently in-service, and signals to the CPU only interrupt requests whose priority level is higher than the level of the highest priority interrupt currently being serviced. In addition, the ICU monitors the system bus and responds to Interrupt-Acknowledge and End-of-Interrupt bus cycles, by providing vector values to the CPU and updating the appropriate bits in the ISRV register.

Note: The Series 32000 interrupt handling specification remains unchanged, except for the elimination of cascaded maskable interrupt requests and their associated bus cycles. In particular, the vector numbers are always positive, in the range 1₁₆ through 1F₁₆. The CPU interpretation and handling of the PSR I-bit and CFG I-bit remains unchanged. From the CPU standpoint, the on-chip ICU can be regarded as an independent module.

TABLE 3-5. Interrupt Sources and Priority Levels

Priority Level	Interrupt Source
INT15 (Highest)	External Only
INT14	External or DMA (DIP Bit in IMSK is 1)
INT13	External or Timer 0 IPFA or Capture Mode Underflow
INT12	External or Timer 0 IPFB
INT11	External Only
INT10	External or Timer 1 IPFA or Capture Mode Underflow
INT9	External or Timer 1 IPFB
INT8	External Only
INT7	External Only
INT6	External or DMA (DIP Bit in IMSK is 0)
INT5	External or Timer 2 IPFA or Capture Mode Underflow
INT4	External Only
INT3	External or Timer 2 IPFB
INT2	External Only
INT1 (Lowest)	External Only
—	No Interrupt

3.4.3.1 Interrupt-Acknowledge Processing

When an Interrupt acknowledge cycle is performed, the bit in the ISRV register corresponding to the priority level of the current interrupt request is set to 1. This is specified in the least significant 4 bits of the IVCT register and represents the higher of the encoded priority value from the $\overline{IR}0$ – $\overline{IR}3$ pins and the highest priority of any pending internal request. During the acknowledge cycle a special bus cycle is also executed. (Section 3.5.5.5).

Note that the IVCT register is not latched by the ICU even after the interrupt request is acknowledged by the CPU. This allows the software to examine the priority level of the current request by reading the IVCT register.

For proper ICU operation, the priority level of an interrupt request must not be decreased unless one of the following conditions is met.

1. The CPU performed an Interrupt-Acknowledge bus cycle
2. Maskable interrupts are currently disabled (PSR I-flag is 0)
3. A higher or same priority level interrupt is currently in-service

The first condition can be used by a requesting I/O device to determine when it is appropriate to remove the interrupt request. The other conditions could be used by the software to remove an interrupt by accessing the I/O device's control registers.

3.4.3.2 End-Of-Interrupt Processing

In response to an End-of-Interrupt bus cycle, the ICU clears the bit in the ISRV register corresponding to the highest priority interrupt currently in-service. The CPU is assumed to have returned to the next lower priority interrupt service routine.

In addition, the ICU returns the IVCT register value on the data bus. This value is not related to the priority level of the terminated interrupt routine that executed the RETI instruction. The return IVCT value is ignored by the CPU.

Note that it is also possible to clear bits in the ISRV explicitly by software, for example, inside an interrupt service routine in order to reenables interrupts at the same or lower priority levels. In this case, either the RETT instruction should be used to terminate that interrupt service routine, or else the corresponding ISRV bit should be set to 1 again before executing the RETI instruction.

3.4.4 Timers

The NS32CG160 provides three on-chip timer blocks. Since these blocks are identical, the descriptions that follow are equally applicable to any one of them. Each timer block consists of a 16-bit counter TC, a control register TCNTL, and two support registers TRCA and TRCB. Two external signals TXA and TXB are also provided for each block to handle all the interactions with external logic. Each timer can operate in one of three modes: Processor Independent, External Event Counter, and Input Capture. Table 3-6 shows the TMC field encodings for the different modes. Details on the operation of each mode are given in the following sections.

3.4.4.1 Processor Independent Mode (Mode 1)

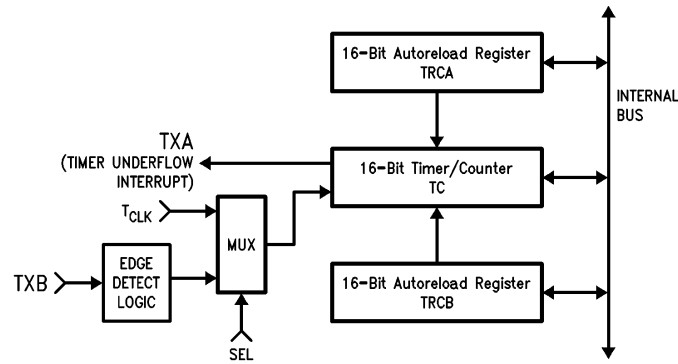
This mode can be used to generate an output signal with minimal software intervention. The software only needs to define the ON and OFF times for the waveform to be generated. Once started, the timer will generate a periodic waveform without further intervention, except when the parameters need to be updated.

In this mode the timer counts down at the Tclk rate (Section 2.1.8). Upon the occurrence of every underflow the timer is alternately reloaded with the contents of the support registers TRCA and TRCB. The first timer underflow causes a reload from the TRCA register. Reloads from subsequent underflows alternate from the two support registers, starting with TRCB. Each underflow toggles the TXA output pin.

Timer underflows are alternately latched into the IPFA and IPFB flags. The software is responsible for resetting these flags. Two enable bits, IENA and IENB, allow timer underflow interrupts to be enabled or disabled. Setting the IENA bit will cause an interrupt when a timer underflow causes a reload from TRCA, while setting IENB will cause an interrupt when the reload is from TRCB.

The IENA and IENB bits give the user the flexibility to enable or disable interrupts on either or both edges of the timer output waveform.

3.0 Functional Description (Continued)



TL/EE/10752-22

Note 1: Tclk = CCTL/8 if TCNT PRC-bit = 0
Tclk = CCTL/4096 if TCNT PRC-bit = 1

Note 2: Specifying a counter value of 0 yields a count of 2^{16} .

FIGURE 3-12. Timer Block Diagram for the Processor Independent and External Event Counter Modes

TABLE 3-6. Timer Modes

TMC Field	Timer Mode	Interrupt A Source	Interrupt B Source	Timer Counts On
000	IDLE			
001	MODE 1 (Processor Independent)	Autoreload TRCA	Autoreload TRCB	Tclk
010	MODE 2 (External Event Counter)	Autoreload TRCA	Autoreload TRCB	TXB Rising Edge
100	MODE 3 (Input Capture) Captures on: TXA Rising Edge TXB Rising Edge	TXA Rising Edge or Timer Underflow	TXB Rising Edge	Tclk

Note: The values shown for the TMC field (Section 2.1.8) are the only ones allowed. Specifying a different value may cause unpredictable results.

3.4.4.2 External Event Counter Mode (Mode 2)

This mode is similar to the processor independent mode. The only difference is that the timer is clocked by the rising edge of the signal applied on the TXB input pin.

3.4.4.3 Input Capture Mode (Mode 3)

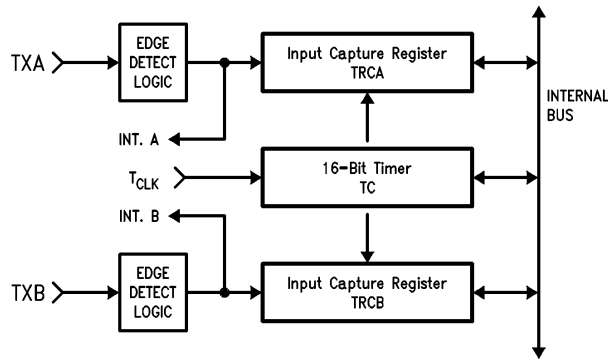
This mode allows to perform precise measurements of external frequencies and to time external events.

The timer constantly runs at the Tclk rate. The registers TRCA and TRCB act as capture registers, and are controlled by external signals applied on the TXA and TXB pins. The timer value gets copied into the corresponding register when a trigger event is signaled on either TXA or TXB. A trigger event is specified as a rising edge of the input signal.

Trigger events can be programmed to generate interrupts. The occurrence of a trigger event on either TXA or TXB will be latched into IPFA or IPFB respectively. Interrupts are controlled by the setting of IENA and IENB.

Timer underflows can also generate interrupts. Since the underflow interrupt pending flag TCS has a different function in the other timer modes, the software should always reset it when the Input Capture Mode is selected. Timer underflow interrupts are also enabled by the IENA bit. Therefore, when IENA is set to 1 and an interrupt occurs, both IPFA and TCS should be checked to determine the origin of the interrupt.

3.0 Functional Description (Continued)



TL/EE/10752-23

Note: When the Timer is idle TXA is in input state.

FIGURE 3-13. Timer Block Diagram for the Input Capture Mode

3.5 SYSTEM INTERFACE

This section provides general information on the NS32CG160 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 NS32CG160 requires a single 5V power supply, applied on 7 pins. The logic voltage pin V_{CC} supplies the power to the on-chip logic. The buffer voltage pins V_{CCB1-6} supply the power to the on-chip output drivers.

Grounding connections are made on 7 pins. The Logic Ground Pin GNDL provides the ground connection to the on-chip logic. The buffer ground pins GND $B1-6$ are the ground pins for the on-chip output drivers.

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 μ F ceramic capacitors can be used for this purpose. They should attach to V_{CC} , GND pairs as close as possible to the NS32CG160.

A 1.0 μ F tantalum capacitor should also be connected between V_{CC} and ground.

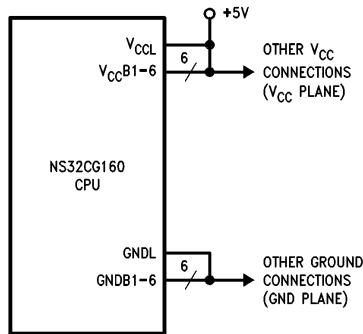
During prototype using wire-wrap or similar methods, the capacitors should be soldered directly to the power pins of the NS32CG160 socket, or as close as possible, with very short leads.

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 NS32CG160. 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 NS32CG160.

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

3.0 Functional Description (Continued)



TL/EE/10752-24


FIGURE 3-14. Power and Ground Connections

3.5.2 Clocking

The NS32CG160 provides an internal oscillator that interacts with an external clock source through two signals; OSCIN and OSCOUT.


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 3-7 provides the crystal characteristics and the values of the R, C, and L components, including stray capacitance, required for various frequencies.



TL/EE/10752-25

FIGURE 3-15. Crystal Interconnections—30 MHz



TL/EE/10752-26

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

TABLE 3-7. 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 NS32CG160 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 signal. 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	130 mA
0	1	2	12.5 MHz	65 mA
1	0	4	6.25 MHz	33 mA
1	1	8	3.13 MHz	17 mA

3.0 Functional Description (Continued)

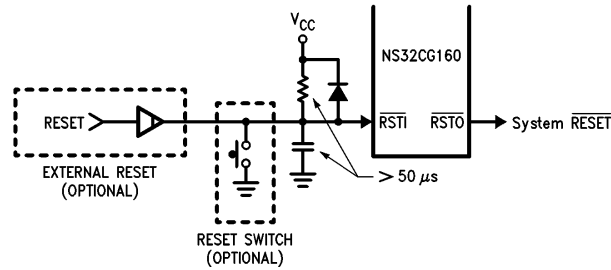


FIGURE 3-17. Recommended Reset Connections


TL/EE/10752-27

3.5.4 Resetting

The \overline{RSTI} input pin is used to reset the NS32CG160. 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.

On application of power, \overline{RSTI} must be held low for at least $50 \mu s$ 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/10752-28

FIGURE 3-18. Power-On Reset Requirements



TL/EE/10752-29

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{TSC} , \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 timer signals $\overline{TXA0-TXA2}$ are set to input mode.

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. \overline{NMI} is enabled to allow Non-Maskable Interrupts. The following conditions are present after reset due to the PSR being reset to 0:

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{NMI} is enabled. Maskable interrupts are disabled.

\overline{BPU} is inactive high.

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

Note that vector/non-vector interrupts have not been selected. While interrupts are disabled, a SETCFG [I] instruction must be executed to enable vectored interrupts. If non-vector 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, full CPU clock rate, and with direct exception mode enabled, execute a SETCFG [DE, F, I] instruction. To declare non-vector interrupts, no FPU, full CPU clock rate, and with direct exception mode disabled, execute a SETCFG [] instruction.

3.5.5 Bus Cycles

The NS32CG160 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 devices registers.
5. To transfer information to or from a Slave Processor.

3.0 Functional Description (Continued)

3.5.5.1 Bus Status

The NS32CG160 CPU presents four bits of Bus Status information on pins ST0–ST3. The various combinations on these pins indicate why the CPU is performing a bus cycle, or, if it is idle on the bus, then why it is idle.

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— Reserved.

0011— The bus is idle because the CPU is waiting for a Slave Processor to complete an instruction.

0100— Interrupt Acknowledge.
The CPU is performing a Read cycle to acknowledge an interrupt request.

0101— Reserved.

0110— End of Interrupt.
The CPU is performing a Read cycle to indicate that it is executing a Return from Interrupt (RET) instruction at the completion of an interrupt's service procedure.

0111— Reserved.

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 signaled 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.

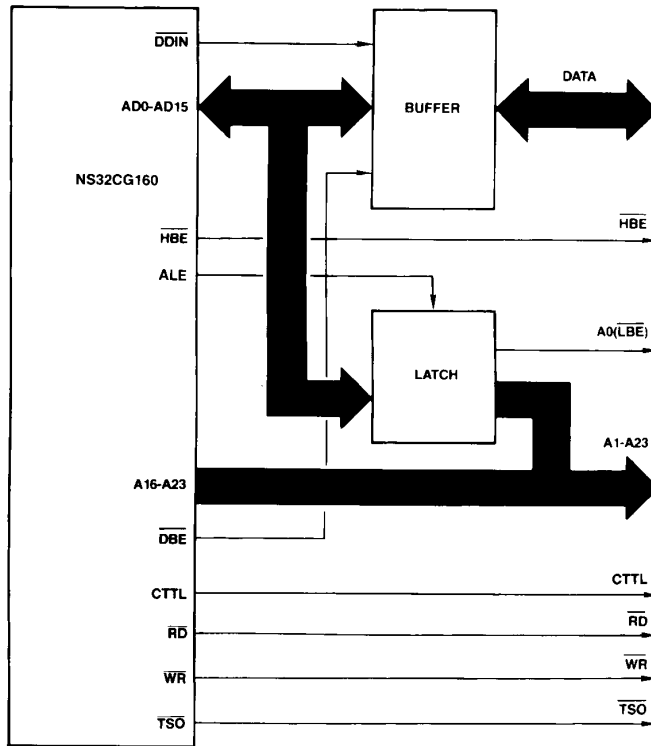
The cases shown assume that the selected memory or peripheral device is capable of communicating with the CPU at full speed. If not, then cycle extension may be requested through \overline{CWAIT} and/or $WAIT1-2$.

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{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{ADS} to generate the strobe for the address latches. See *Figure 3-20*. During this time also the status signals \overline{DDIN} , indicating the direction of the transfer, and \overline{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.

3.0 Functional Description (Continued)



TL/EE/10752-30

FIGURE 3-20. Bus Connections

3.0 Functional Description (Continued)

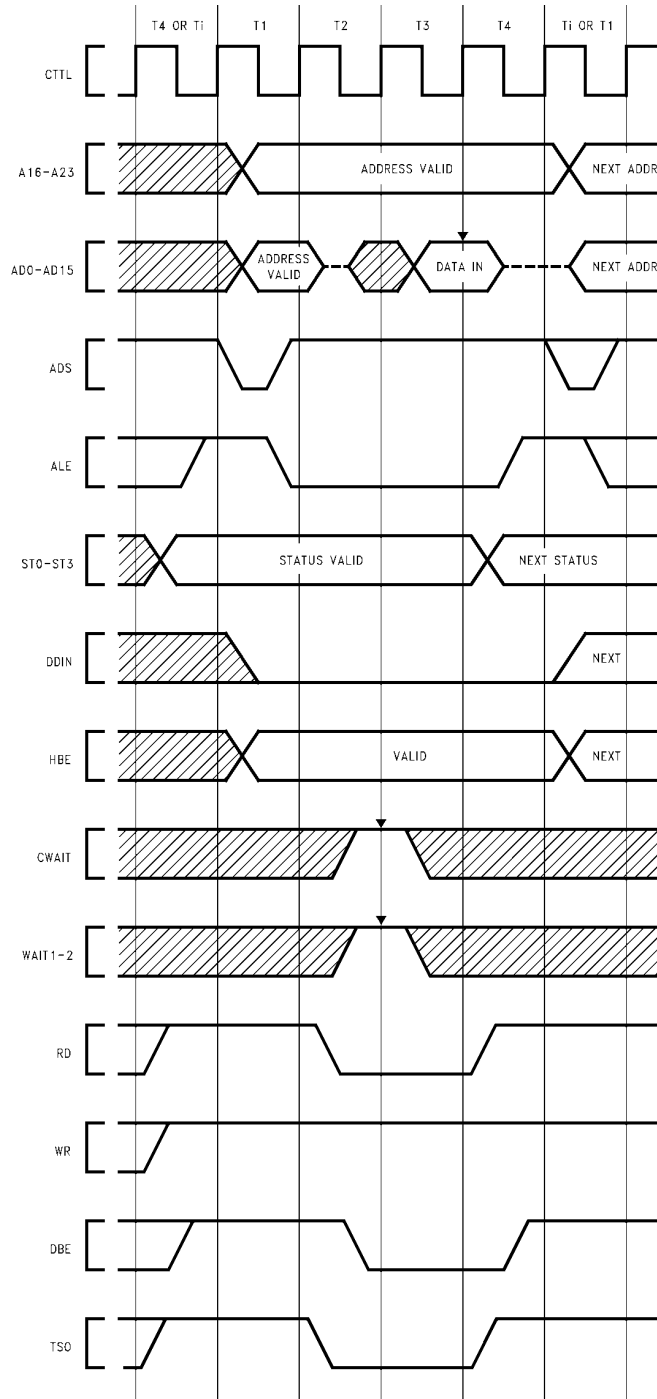


FIGURE 3-21. Read Cycle Timing

TL/EE/10752-31

3.0 Functional Description (Continued)

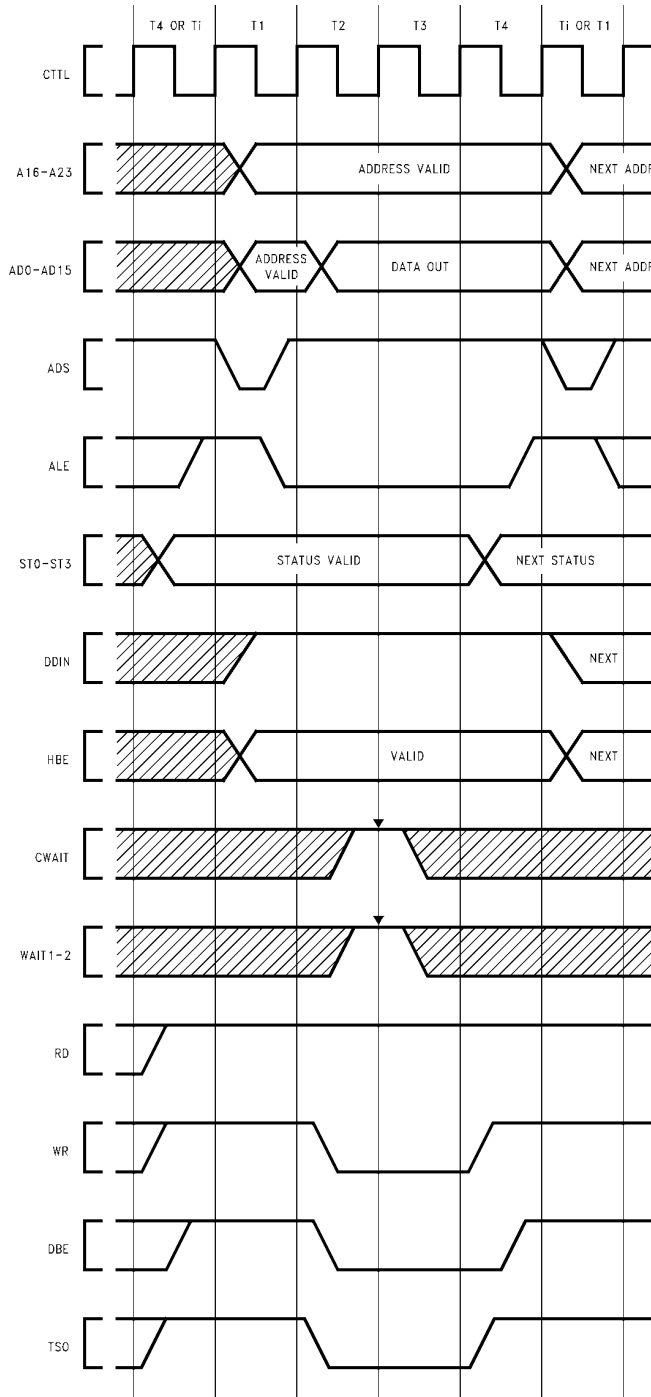


FIGURE 3-22. Write Cycle Timing

TL/EE/10752-32

3.0 Functional Description (Continued)

At this time the signals \overline{TSO} (Timing State Output), \overline{DBE} (Data Buffer Enable) and either \overline{RD} (Read Strobe) or \overline{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{CWAIT} and $\overline{WAIT1-2}$ 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{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{RD} or \overline{WR} , and \overline{TSO} signals go inactive, and on the falling edge of CTTL, \overline{DBE} goes inactive, having provided for necessary data hold times. Data during Write cycles remains valid from the CPU throughout T4. Note that the Bus Status lines (ST0-ST3) change at the beginning of T4, anticipating the following bus cycle (if any).

3.5.5.3 Cycle Extension

To allow sufficient access time for any speed of memory or peripheral device, the NS32CG160 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 $\overline{WAIT1-2}$ and \overline{CWAIT} input signals.

At the end of state T2, on the rising edge of CTTL, $\overline{WAIT1-2}$ and \overline{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 additional 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 $\overline{WAIT1-2}$ inputs are only sampled by the CPU at the end of state T2. They are ignored at all other times.

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

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

\overline{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 $\overline{WAIT1-2}$ and \overline{CWAIT} inputs.

1. Start bus cycle
2. Sample $\overline{WAIT1-2}$ and \overline{CWAIT} at the end of state T2.
3. If the $\overline{WAIT1-2}$ inputs are both inactive, then go to step 6.
4. Insert the number of wait states selected by $\overline{WAIT1-2}$.
5. Sample \overline{CWAIT} again.
6. If \overline{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 $\overline{WAIT2}$, and one is due to \overline{CWAIT} .


3.5.5.4 Instruction Fetch Cycles

Instructions for the NS32CG160 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.

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 3-9).

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.

3.0 Functional Description (Continued)



TL/EE/10752-33

FIGURE 3-23. Cycle Extension of a Read Cycle

3.0 Functional Description (Continued)

3.5.5.5 Special Bus Cycles

Special bus cycles are performed during CPU accesses to the top area of the address space. These cycles may be used by external logic to track CPU activities involving the on-chip I/O devices as well as to monitor maskable interrupt acknowledges and returns.

A special bus cycle starts with the assertion of the special output signal $\overline{\text{IAS}}$. The ALE signal stays high during the en-

tire cycle, and the signals $\overline{\text{ADS}}$, $\overline{\text{TSO}}$, $\overline{\text{DBE}}$, $\overline{\text{RD}}$ and $\overline{\text{WR}}$ are not activated. $\overline{\text{CWAIT}}$ and $\overline{\text{WAIT1-2}}$ are ignored, so the cycle is always performed with no wait states. The CPU drives the data bus with the same data that is being written into the on-chip peripherals registers during special write cycles, and ignores the data placed on the data bus, during special read cycles. *Figure 3-24* shows the timing for special read and write cycles.




FIGURE 3-24. Special Bus Cycle Timing

TL/EE/10752-34

3.0 Functional Description (Continued)

TABLE 3-8. Interrupt Sequences

Cycle	Status	Address	\overline{DDIN}	\overline{HBE}	A0	Internal Data Bus	
						Byte 1	Byte 0
A. Non-Maskable Interrupt Control Sequences							
Interrupt Acknowledge							
1	0100	FFFFFF00 ₁₆	0	1	0	X	X
Interrupt Return							
None: Performed through Return from Trap (RETT) instruction.							
B. Non-Vectored Interrupt Control Sequences							
Interrupt Acknowledge							
1	0100	FFFFFE00 ₁₆	0	1	0	X	X
Interrupt Return							
1	0110	FFFFFE00 ₁₆	0	1	0	X	X
C. Vectored Interrupt Sequences							
Interrupt Acknowledge							
1	0100	FFFFFE00 ₁₆	0	1	0	X	Vector: Range: 11 ₁₆ –1F ₁₆
Interrupt Return							
1	0110	FFFFFE00 ₁₆	0	1	0	X	IVCT Register Content

Note 1: The bus cycles for maskable interrupts are special read cycles.

Note 2: The Bus cycle for the non-maskable interrupt acknowledge is a normal read cycle. Since the top 8 bits of the address are not available externally, this cycle can be detected by decoding the ST0–3 signals.

3.5.5.6 BPU Bus Cycles


Whenever the EXTBLT instruction is executed, the NS32CG160 performs read and write bus cycles to access the on-chip BPU as well as any off-chip BPUs. These bus cycles are identical to normal memory access cycles, except that the \overline{BPU} output signal is asserted.

The \overline{BPU} output is used to signal a BPU memory transaction to both the on-chip and off-chip BPUs.

In response, the BPUs will access their respective memory planes using the address issued by the CPU, and pass the data to their data paths. The number of wait states should be determined by the slowest BPU interface.

Figure 3-25 shows a BPU read bus cycle.

3.0 Functional Description (Continued)



TL/EE/10752-35

FIGURE 3-25. BPU Read Cycle

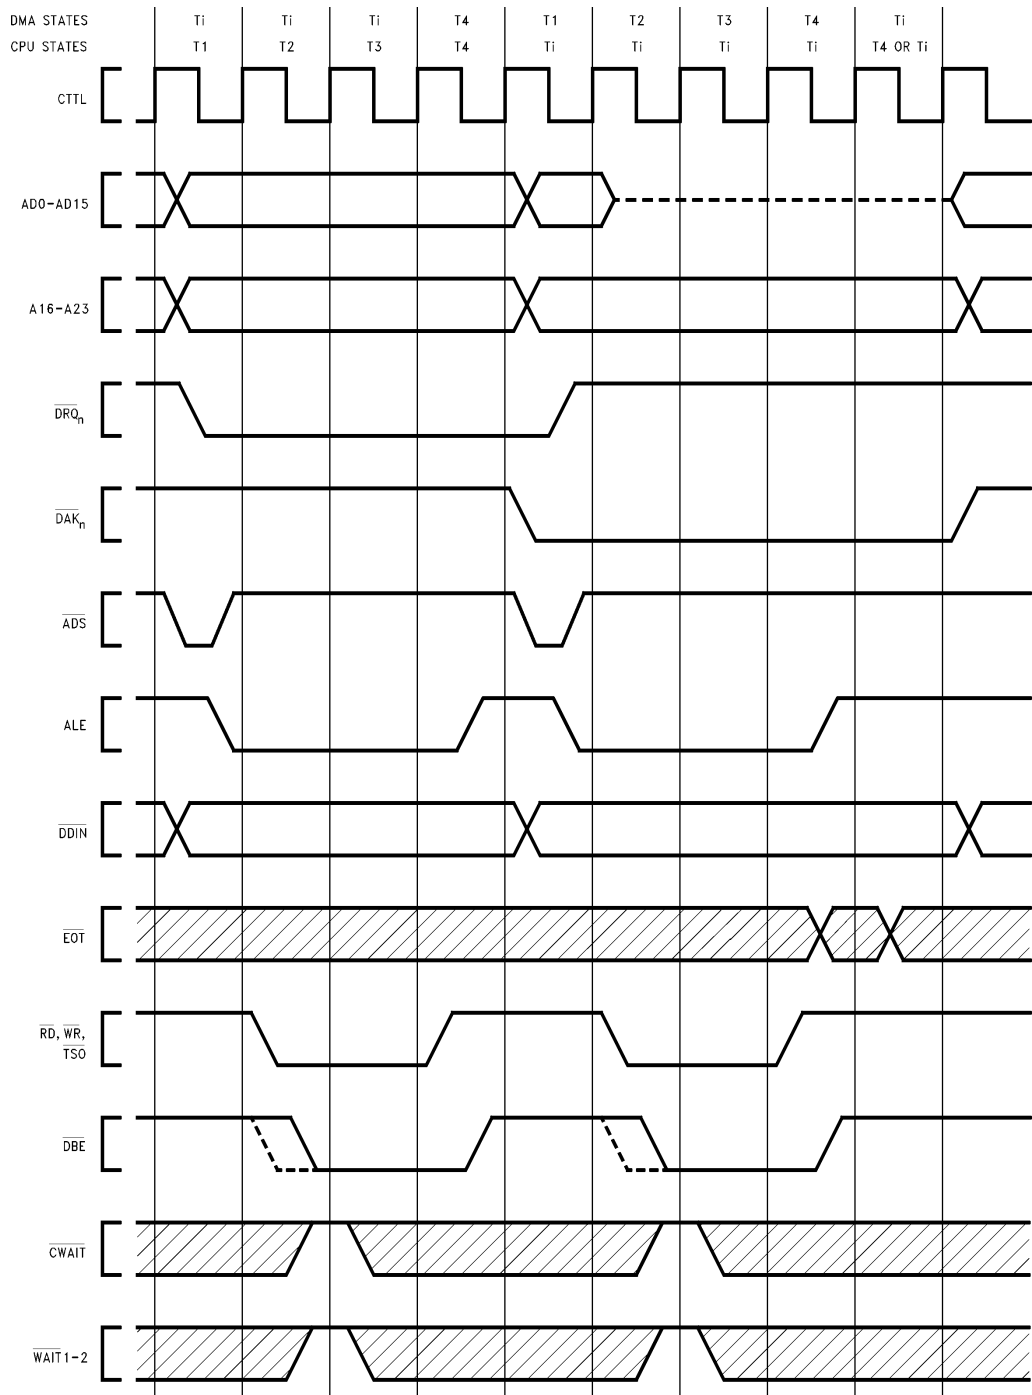
3.5.5.7 DMA Controller Bus Cycles

DMA bus cycles are executed by the NS32CG160 whenever a data transfer is performed by one of the on-chip DMA channels. These cycles have the same timing as normal read and write cycles. However, the data is transferred directly between a peripheral device and memory or another peripheral device. The data bus is not driven during DMA cycles. The DDIN signal indicates whether a read or a write

to the addressed device is being performed. The other bus control signals are handled as in normal bus cycles with the exception of ST0-3 which are undefined. Either $\overline{DAK0}$ or $\overline{DAK1}$ is asserted, depending on which DMA channel is performing the transfer. *Figure 3-26* shows the timing diagram of DMA bus cycles.

Two idle clock cycles are always inserted between consecutive DMA bus cycles.

3.0 Functional Description (Continued)



TL/EE/10752-36

FIGURE 3-26. DMA Bus Cycle Timing

3.0 Functional Description (Continued)

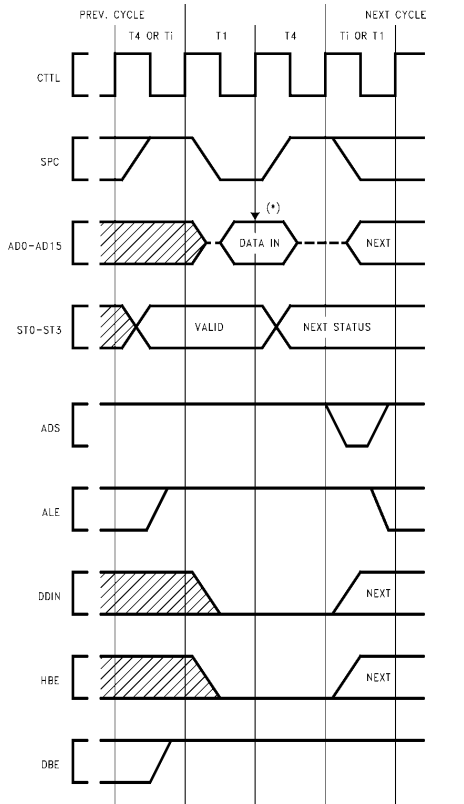
3.5.5.8 Slave Processor Bus Cycles

A Slave Processor bus cycle always takes exactly two clock cycles, labeled T1 and T4 (see *Figures 3-27* and *3-28*). 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 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 3-29 shows the NS32CG160 and FPU connection diagram.



Note: CPU samples Data Bus here. TL/EE/10752-37
FIGURE 3-27. Slave Processor Read Cycle



*Note: Slave Processor samples Data Bus here. TL/EE/10752-38
FIGURE 3-28. Slave Processor Write Cycle


3.5.5.9 Data Access Sequences

The 24-bit address provided by the NS32CG160 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 NS32CG160 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 NS32CG160 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 3-30*.


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 3-9 gives the state of A0 and \overline{HBE} for each category.

3.0 Functional Description (Continued)



TL/EE/10752-74

FIGURE 3-29. NS32CG160 and FPU Interconnections



TL/EE/10752-39

FIGURE 3-30. Memory Interface

TABLE 3-9. 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 3-10 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 3-10. Data Access Sequences

Cycle	Type	Address	HBE	A0	High Bus	Low Bus					
A. Odd Word Access Sequence											
					Byte 1	Byte 0	← A				
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	← A		
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	← A		
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	← A
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	← A
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)

3.5.5.10 Bus Access Control

The NS32CG160 CPU has the capability of relinquishing its control of the bus upon request from an external DMA controller or another CPU. This capability is implemented by means of the $\overline{\text{HOLD}}$ (Hold Request) and $\overline{\text{HLDA}}$ (Hold Acknowledge) pins. By asserting $\overline{\text{HOLD}}$ low, an external device requests access to the bus. On receipt of $\overline{\text{HLDA}}$ 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., $\overline{\text{TSC}}$, $\overline{\text{DBE}}$, $\overline{\text{RD}}$ 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 $\overline{\text{HLDA}}$ 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 3-31* shows the timing sequence when the CPU is idle. In this case, the CPU grants the bus during the immediately following clock cycle. *Figure 3-32* 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 External DMA cycles the $\overline{\text{WAIT}}1-2$ 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/10752-40

FIGURE 3-31. $\overline{\text{HOLD}}$ Timing, Bus Initially Idle

3.0 Functional Description (Continued)



TL/EE/10752-41

FIGURE 3-32. $\overline{\text{HOLD}}$ Timing, Bus Initially Not Idle

3.0 Functional Description (Continued)

3.5.5.11 External Interrupt Requests

Five signals are provided by the NS32CG160 to externally request interrupts. $\overline{IR}0-3$ and \overline{NMI} are for maskable and non-maskable interrupts respectively. These signals are sampled on the rising edge of CTTL.

$\overline{IR}0-3$ are level sensitive and, once asserted, they must be kept asserted until acknowledged. They can be asynchronous to CTTL, since the NS32CG160 internally synchronizes them. Nevertheless, if $\overline{IR}0-3$ meet the required setup and hold times, then they are recognized deterministically. The on-chip synchronization circuitry compares the values of $\overline{IR}0-3$ sampled in two consecutive CTTL edges. An interrupt request that is held constant for two consecutive edges is considered valid. The sampled value of $\overline{IR}0-3$ indi-

cates an external interrupt request at the encoded priority. When $\overline{IR}0-3$ are all high, then no external interrupt is requested. When they are all low, then a level-15 request is generated.

\overline{NMI} is edge sensitive; a high-to-low transition is detected by the CPU and stored in an internal latch, so that there is no need to keep this signal asserted until the request is acknowledged. \overline{NMI} can be asserted asynchronously to CTTL, but it should be at least 2 clock cycles wide in order to be recognized.

If \overline{NMI} meets the specified setup and hold times, it will be recognized on the rising edge of CTTL deterministically.

Refer to *Figure 4-17* for more details on the timing of the above signals.




FIGURE 3-33. Interrupt Request Encoding Logic

TL/EE/10752-42

3.0 Functional Description (Continued)

3.5.5.12 Instruction Status

In addition to the four bits of Bus Cycle status (ST0–3), the NS32CG160 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.

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

U/ \bar{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 bus cycle except for a DMA cycle. See the Timing Specifications in Section 4.

$\bar{I}L\bar{O}$ (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. $\bar{I}L\bar{O}$ is guaranteed to be active during the operand accesses performed by the interlocked instructions.

Note: The acknowledge of $\bar{H}O\bar{L}D$ and $\bar{D}R\bar{Q}0-1$ is on a cycle by cycle basis. Therefore, it is possible to have either $\bar{H}L\bar{D}A$ or $\bar{D}A\bar{K}0-1$ active when an interlock operation is in progress. In this case, $\bar{I}L\bar{O}$ remains low and the interlocked instruction continues only after the DMA cycles are completed.

4.0 Device Specifications

4.1 NS32CG160 PIN DESCRIPTIONS

The following is a brief description of all NS32CG160 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 $\bar{H}O\bar{L}D$ acknowledge.

4.1.1 Supplies

V_{CCL} Logic Power.

+5V Positive Supply for On-Chip Logic.

V_{CCB1-6} Buffers Power.

+5V Positive Supplies for On-Chip Output Buffers

GN_{DL} Logic Ground.

Ground Reference for On-Chip Logic.

GN_{DB1-6} Buffers Ground.

Ground Reference for On-Chip Output Buffers.

4.1.2 Input Signals

$\bar{R}ST\bar{I}$ 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.

$\bar{H}O\bar{L}D$

Hold Request.

When active, causes the CPU to release the bus for external DMA or multiprocessing purposes. See Section 3.5.5.10.

Note: If the $\bar{H}O\bar{L}D$ signal is generated asynchronously, its set up and hold times may be violated. In this case, it is recommended to synchronize it with CTT \bar{L} to minimize the possibility of metastable states.

The CPU provides only one synchronization stage to minimize the $\bar{H}L\bar{D}A$ latency. This is to avoid speed degradations in cases of heavy $\bar{H}O\bar{L}D$ activity (i.e., DMA controller cycles interleaved with CPU cycles).

$\bar{I}R0-3$

Interrupt Request.

Low levels on these signals request a prioritized interrupt (Section 3.5.5.11).

$\bar{N}M\bar{I}$

Non-Maskable Interrupt.

A High-to-Low transition on this signal requests a non-maskable interrupt.

Note: $\bar{I}N\bar{T}$ and $\bar{N}M\bar{I}$ are true asynchronous inputs. Therefore, no external synchronizing circuit is needed

$\bar{C}W\bar{A}I\bar{T}$

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.

$\bar{W}A\bar{I}T1-2$ Two-Bit Wait State Inputs.

These inputs, collectively called $\bar{W}A\bar{I}T1-2$, allow from zero to three wait states to be specified. They are binary weighted. See Section 3.5.5.3.

Note: During an external DMA cycle, $\bar{W}A\bar{I}T1-2$ should be kept inactive unless they are also monitored by the DMA Controller. Wait states, in this case, should be generated through $\bar{C}W\bar{A}I\bar{T}$.

$\bar{E}O\bar{T}$

DMA End of Transfer.

When asserted in the T4 state of any DMA bus cycle causes the DMA channel currently in control of the bus to terminate the data transfer.

$\bar{D}R\bar{Q}0-1$

DMA Requests.

When active, these signals request DMA service from channels 0 and 1. $\bar{D}R\bar{Q}0-1$ are sampled on each rising edge of CTT \bar{L} . Once asserted, these signals should not be deasserted until the request has been acknowledged. To avoid multiple DMA cycles, they should be deasserted before the end of each cycle.

Note: If $\bar{D}R\bar{Q}0-1$ are generated asynchronously, the set up and hold times may be violated. In this case it is recommended to synchronize them with the rising edge of CTT \bar{L} to minimize the possibility of metastable states.

TXB0-2 Timer Trigger Signals.

A low-to-high transition on any of these pins will signal a trigger event to the corresponding timer. (Section 3.4.4).

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.

4.0 Device Specifications (Continued)

$\overline{\text{HBE}}$	<p>*High Byte Enable. Status signal used to enable data transfers on the most significant byte of the data bus.</p>	$\overline{\text{RSTO}}$	<p>Reset Output. This signal becomes active when $\overline{\text{RSTI}}$ is low, initiating a system reset.</p>
ST0–3	<p>Status. Bus cycle status code; ST0 is the least significant. Encodings are: 0000— Idle: CPU Inactive on Bus. 0001— Idle: WAIT Instruction. 0010— (Reserved.) 0011— Idle: Waiting for Slave. 0100— Interrupt Acknowledge. 0101— Reserved. 0110— End of Interrupt. 0111— Reserved. 1000— Sequential Instruction Fetch. 1001— Non-Sequential Instruction Fetch. 1010— Data Transfer. 1011— Read Read-Modify-Write Operand. 1100— Read for Effective Address. 1101— Transfer Slave Operand. 1110— Read Slave Status Word. 1111— Broadcast Slave ID.</p>	$\overline{\text{RD}}$	<p>Read Strobe. Activated during CPU or DMA read cycles to enable reading of data from memory or peripherals. See Section 3.5.5.2.</p>
$\text{U}/\overline{\text{S}}$	<p>User/Supervisor. User or Supervisor Mode status. High indicates User Mode; low indicates Supervisor Mode.</p>	$\overline{\text{WR}}$	<p>Write Strobe. Activated during CPU or DMA write cycles to enable writing of data to memory or peripherals.</p>
$\overline{\text{ILO}}$	<p>Interlocked Operation. When active, indicates that an interlocked operation is being executed.</p>	$\overline{\text{TSO}}$	<p>Timing State Output. The falling edge of $\overline{\text{TSO}}$ identifies the beginning of state T2 of a bus cycle. The rising edge identifies the beginning of state T4.</p>
$\overline{\text{HLDA}}$	<p>Hold Acknowledge. Activated by the CPU in response to the $\overline{\text{HOLD}}$ input to indicate that the CPU has released the bus.</p>	$\overline{\text{DBE}}$	<p>Data Buffers Enable. Used to control external data buffers. It is active when the data buffers are to be enabled.</p>
$\overline{\text{PFS}}$	<p>Program Flow Status. A pulse on this signal indicates the beginning of execution of an instruction.</p>	OSCOUT	<p>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.</p>
$\overline{\text{BPU}}$	<p>BPU Cycle. This signal is activated during a bus cycle to enable an external BITBLT processing unit. The EXTBLT instruction activates this signal. Note: $\overline{\text{BPU}}$ is low (Active) only during bus cycles involving pre-fetching instructions and execution of EXTBLT operands. It is recommended that $\overline{\text{BPU}}$, $\overline{\text{ADS}}$ and status lines (ST0–ST3) be used to qualify BPU bus cycles. If DMA cycles can occur during the execution of EXTBLT, then the $\overline{\text{HLDA}}$ and $\overline{\text{DAK0-1}}$ signals should be used to further qualify BPU cycles. $\overline{\text{BPU}}$ may become active during T4 of a non-BPU bus cycle, and may become inactive during T4 of a BPU bus cycle. $\overline{\text{BPU}}$ must be qualified by $\overline{\text{ADS}}$ and status lines (ST0–ST3) to be used as an external gating signal.</p>	$\overline{\text{DAK0-1}}$	<p>DMA Acknowledge Signals. Activated in response to DMA requests to notify external devices that the corresponding request has been acknowledged.</p>
		$\overline{\text{IAS}}$	<p>Special Cycle Address Strobe. Signals the beginning of a special bus cycle.</p>
		CTTL1–2	<p>System Clock. Output clock for bus timing. CTTL1 and CTTL2 must be externally connected together.</p>
		ALE	<p>Address Latch Enable. Active high signal that can be used to control external address latches.</p>

4.0 Device Specifications (Continued)

4.1.4 Input-Output Signals

AD0-15 *Address/Data Bus.

Multiplexed Address/Data Information. Bit 0 is the least significant bit of each.

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 3.5.5.8.

DDIN *Data Direction.

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


$\overline{\text{ADS}}$

*Address Strobe.

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

TXA0-2 Timer Control Signals.

These signals are used to either output the timer waveforms or to signal trigger events (Section 3.4.4).



TL/EE/10752-75

Bottom View

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)		180	250	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/10752-44

FIGURE 4-2. Output Signals Specification Standard


Abbreviations:

L.E.— Leading Edge

R.E.— Rising Edge

T.E.— Trailing Edge

F.E.— Falling Edge



TL/EE/10752-45

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, NS32CG160-15, NS32CG160-20, NS32CG160-25

- 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	NS32CG160-15		NS32CG160-20		NS32CG160-25		Units
				Min	Max	Min	Max	Min	Max	
t _{CTP}	4-16	CTTL Clock Period	R.E., CTTL to Next R.E., CTTL	66	1000	50	1000	40	1000	ns
t _{CTH}	4-16	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}	4-16	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}	4-16	CTTL Rise Time	0.8V to 2.0V on R.E., CTTL		6		5		4	ns
t _{CTf}	4-16	CTTL Fall Time	2.0V to 0.8V on F.E., CTTL		6		5		4	ns
t _{xCTd}	4-16	OSCIN to CTTL Delay	4.2V on R.E., OSCIN to R.E., CTTL		35		29		25	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}	4-8	AD0-AD15 Floating	After R.E., CTTL Ti		14		13		12	ns
t _{AHf}	4-8	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	ADS Signal Active	After R.E., CTTL T1		14		13		12	ns
t _{ADSi}	4-4	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	ADS Pulse Width	At 0.8V (Both Edges)	20		15		10		ns
t _{ADsf}	4-8	ADS Floating	After R.E., CTTL Ti		14		13		12	ns
t _{ALADSS}	4-4	AD0-AD15 Setup	Before 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}	4-8	HBE Signal Floating	After R.E., CTTL Ti		14		13		12	ns
t _{DDINv}	4-4	DDIN Signal Valid	After R.E., CTTL T1		14		13		12	ns
t _{DDINh}	4-4	DDIN Signal Hold	After R.E., CTTL Next T1 or Ti	0		0		0		ns
t _{DDINf}	4-8	DDIN Floating	After R.E., CTTL Ti		14		13		12	ns
t _{SPCa}	4-11	SPC Output Active	After R.E., CTTL T1		14		13		12	ns
t _{SPCia}	4-11	SPC Output Inactive	After R.E., CTTL T4		14		13		12	ns
t _{HLDa}	4-8	HLD Signal Active	After R.E., CTTL Ti		14		13		12	ns
t _{HLDaia}	4-9	HLD Signal Inactive	After R.E., CTTL Ti		14		13		12	ns
t _{STv}	4-4	Status ST0-ST3 Valid (before T1, See Note 1)	After R.E., CTTL T4		14		13		12	ns
t _{STh}	4-4	Status ST0-ST3 Hold	After R.E., CTTL T4	0		0		0		ns

4.0 Device Specifications (Continued)

4.4.2 Timing Tables (Continued)

4.4.2.1 Output Signals: Internal Propagation Delays, NS32CG160-15, NS32CG160-20, NS32CG160-25 (Continued)

Symbol	Figure	Description	Reference/ Conditions	NS32CG160-15		NS32CG160-20		NS32CG160-25		Units
				Min	Max	Min	Max	Min	Max	
t _{BPUv}	4-4	\overline{BPU} Signal Valid	After R.E., CTTL T4 or T _i		14		13		12	ns
t _{BPUh}	4-4	\overline{BPU} Signal Hold	After R.E., CTTL T4 or T _i	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}	4-14	\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}	4-14	\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} -6 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
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 _{IASw}	4-6	\overline{IAS} Pulse Width	At 0.8V (Both Edges)	20		15		10		ns
t _{ALIASs}	4-6	AD0-AD15 Setup	Before \overline{IAS} T.E.	10		10		10		ns
t _{DAKa}	4-7	\overline{DAK}_n Signals Active	After R.E., CTTL T _i		14		13		12	ns
t _{DAKia}	4-7	\overline{DAK}_n Signal Inactive	After R.E., CTTL		14		13		12	ns
t _{TXAd}	4-19	TXA _n Outputs Delay	After R.E., CTTL		14		13		12	ns
t _{TXAh}	4-19	TXA _n Outputs Hold	After R.E., CTTL	0		0		0		ns
t _{TXAf}	4-20	TXA _n Signals Floating (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 _{TXAnf}	4-20	TXA _n Signals Not Floating (Note 4)	After R.E., CTTL	0.5 t _{CTP} -6 ns		0.5 t _{CTP} -6 ns		0.5 t _{CTP} -6 ns		
t _{ILOa}	4-15	\overline{ILO} Signal Active	After R.E., CTTL		14		13		12	ns
t _{ILOia}	4-15	\overline{ILO} Signal Inactive	After R.E., CTTL		14		13		12	ns
t _{RSTOa}	4-22	\overline{RSTO} Signal Active	After R.E., CTTL		14		13		12	ns
t _{RSTOia}	4-22	\overline{RSTO} Signal Inactive	After R.E., CTTL		14		13		12	ns
t _{RTOI}	4-22	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: "... T_i, 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.

4.0 Device Specifications (Continued)

4.4.2 Timing Tables (Continued)

4.4.2.2 Input Signal Requirements: NS32CG160-15, NS32CG160-20 and NS32CG160-25

Symbol	Figure	Description	Reference/ Conditions	NS32CG160-15		NS32CG160-20		NS32CG160-25		Units
				Min	Max	Min	Max	Min	Max	
t _{Xp}	4-16	OSCIN Clock Period	R.E., OSCIN to Next R.E, OSCIN	33	500	25	500	20	500	ns
t _{Xh}	4-16	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}	4-16	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, 4-12	Data In Setup	Before R.E., CTTL T4	15		14		10		ns
t _{Dih}	4-4, 4-12	Data In Hold (Note 1)	After R.E., CTTL T4	2		2		2		ns
t _{CWs}	4-4, 4-5	$\overline{C}WAIT$ Signal Setup	Before R.E., CTTL T3 or T3(w)	22		18		10		ns
t _{CWh}	4-4, 4-5	$\overline{C}WAIT$ Signal Hold	After R.E., CTTL T3 or T3(w)	2		2		2		ns
t _{Ws}	4-4, 4-5	$\overline{W}AIT_n$ Signals Setup	Before R.E., CTTL T3 or T3(w)	22		21		20		ns
t _{Wh}	4-4, 4-5	$\overline{W}AIT_n$ Signals Hold	After R.E., CTTL T3 or T3(w)	2		2		2		ns
t _{HLDs}	4-8, 4-9	$\overline{H}OLD$ Setup Time	Before R.E., CTTL T2 or Ti	16		15		14		ns
t _{HLDh}	4-8, 4-9	$\overline{H}OLD$ Hold Time	After R.E., CTTL Ti	2		2		2		ns
t _{PWR}	4-21	Power Stable to $\overline{R}STI$ R.E. (Note 2)	After V _{CC} Reaches 4.5V	50		40		30		μs
t _{RSTw}	4-22	$\overline{R}STI$ Pulse Width	At 0.8V (Both Edges)	64		64		64		t _{CTp}
t _{DRQs}	4-7	$\overline{D}RQ_n$ Setup Time	Before R.E., CTTL	16		15		14		ns
t _{DRQh}	4-7	$\overline{D}RQ_n$ Hold Time	After R.E., CTTL	2		2		2		ns
t _{EOTs}	4-7	$\overline{E}OT$ Setup Time	Before R.E., CTTL T4	16		15		14		ns
t _{EOTh}	4-7	$\overline{E}OT$ Hold Time	After R.E., CTTL T4	2		2		2		ns
t _{IRs}	4-17	$\overline{I}R_n$ Setup Time	Before R.E., CTTL	16		15		14		ns
t _{IRh}	4-17	$\overline{I}R_n$ Hold Time	After R.E., CTTL	2		2		2		ns
t _{NMI} s	4-17	$\overline{N}MI$ Setup Time	Before F.E., CTTL	15		14		12		ns
t _{NMI} h	4-17	$\overline{N}MI$ Hold Time	After F.E., CTTL	2		2		2		ns
t _{TXABs}	4-18	TXA _n and TXB _n Signals Setup	Before F.E., CTTL	15		14		12		ns
t _{TXABh}	4-18	TXA _n and TXB _n Signals Hold	After F.E., CTTL	2		2		2		ns
t _{SPCd}	4-13	SPC Pulse Delay from Slave (Note 2)	After F.E., CTTL T4	2		2		2		t _{CTp}
t _{SPCs}	4-13	$\overline{S}PC$ Input Setup	Before R.E., CTTL	22		21		20		ns
t _{SPCh}	4-13	$\overline{S}PC$ Hold Time	After R.E., CTTL	2		2		2		ns
t _{ADs}	4-10	$\overline{A}DS$ Input Setup	Before F.E., CTTL	15	t _{CTp} -3 ns	14	t _{CTp} -3 ns	12	t _{CTp} -3 ns	
t _{ADSh}	4-10	$\overline{A}DS$ Input Hold (Note 3)	After F.E., CTTL T1	2		2		2		ns
t _{DDINs}	4-10	$\overline{D}DIN$ Input Setup	Before F.E., CTTL	15		14		12		ns
t _{DDINh}	4-10	$\overline{D}DIN$ Input Hold	After R.E., CTTL T4	2		2		2		ns

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

Note 2: Not tested, guaranteed by design.

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

4.0 Device Specifications (Continued)

4.4.3 Timing Diagrams

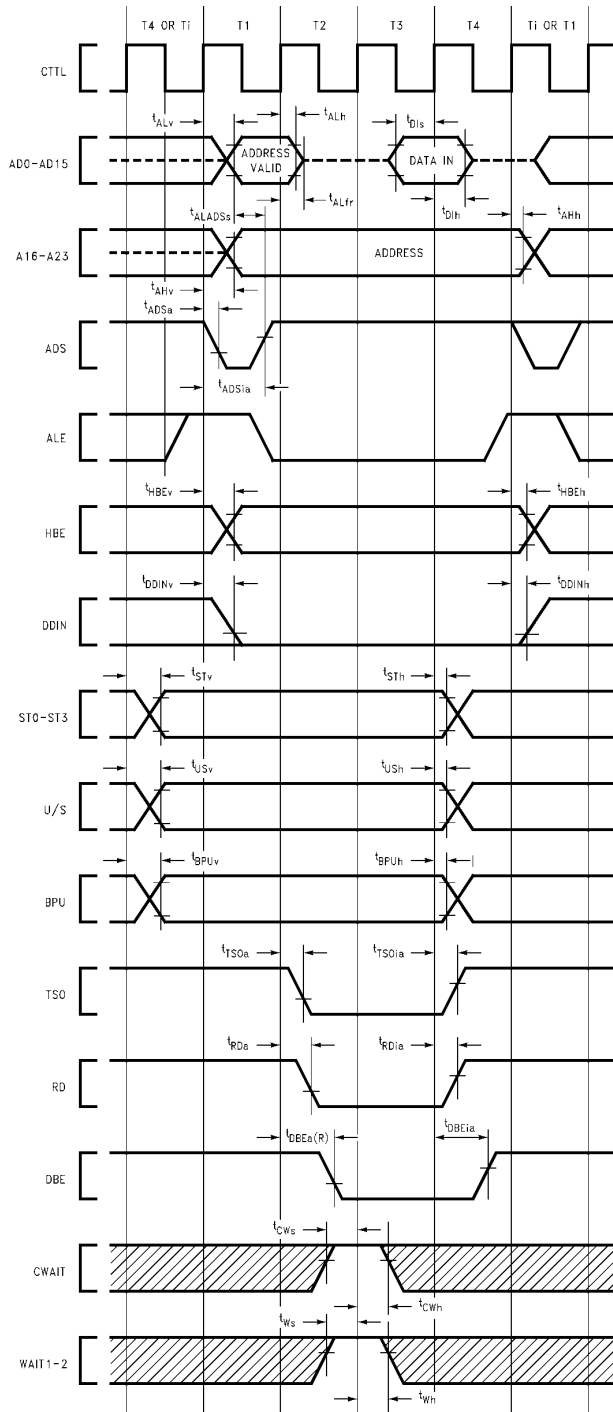


FIGURE 4-4. Read Cycle

TL/EE/10752-46

4.0 Device Specifications (Continued)

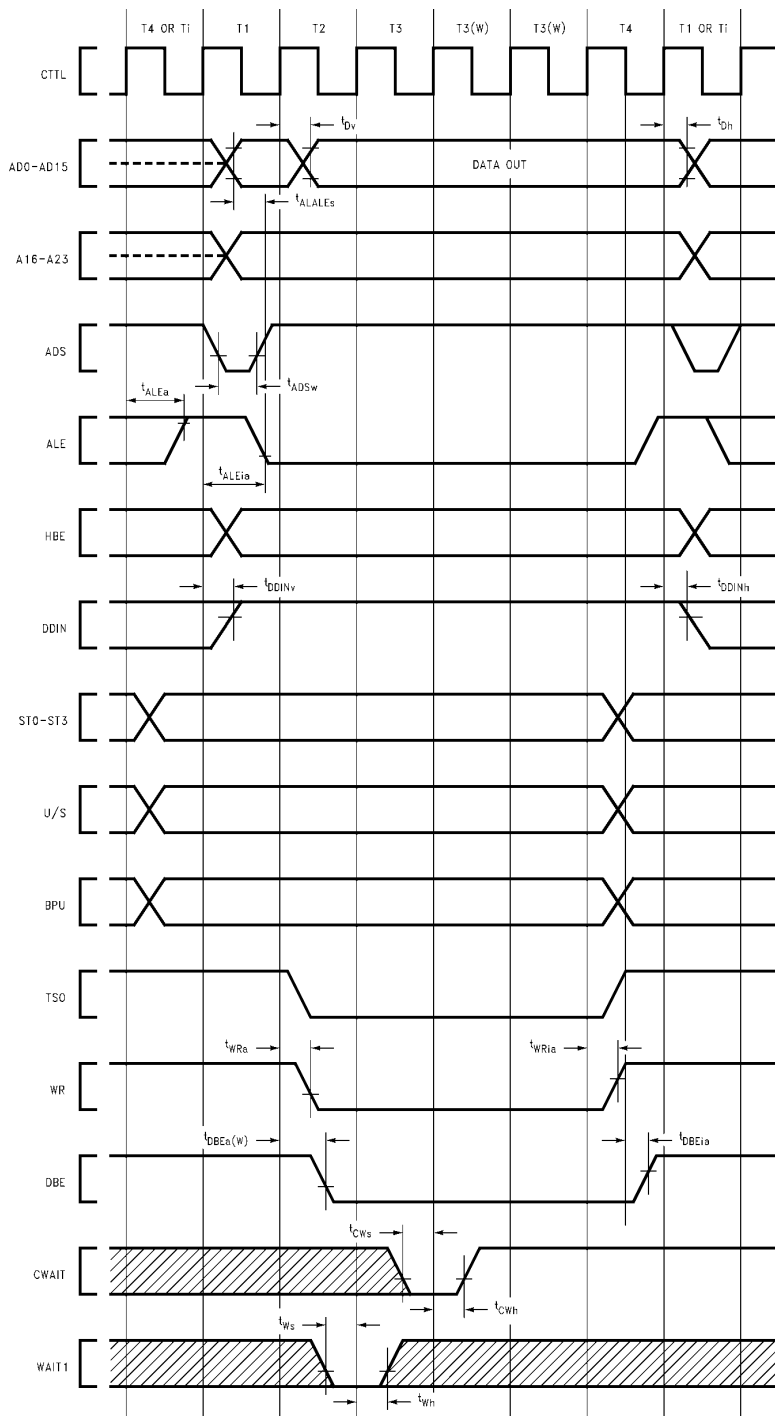


FIGURE 4-5. Write Cycle

TL/EE/10752-47

4.0 Device Specifications (Continued)




FIGURE 4-6. Special Bus Cycle

TL/EE/10752-48

4.0 Device Specifications (Continued)





FIGURE 4-7. On-Chip DMA Bus Cycle

TL/EE/10752-49

4.0 Device Specifications (Continued)




TL/EE/10752-50

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

FIGURE 4-8. $\overline{\text{HOLD}}$ Acknowledge Timing (Bus Initially Not Idle)

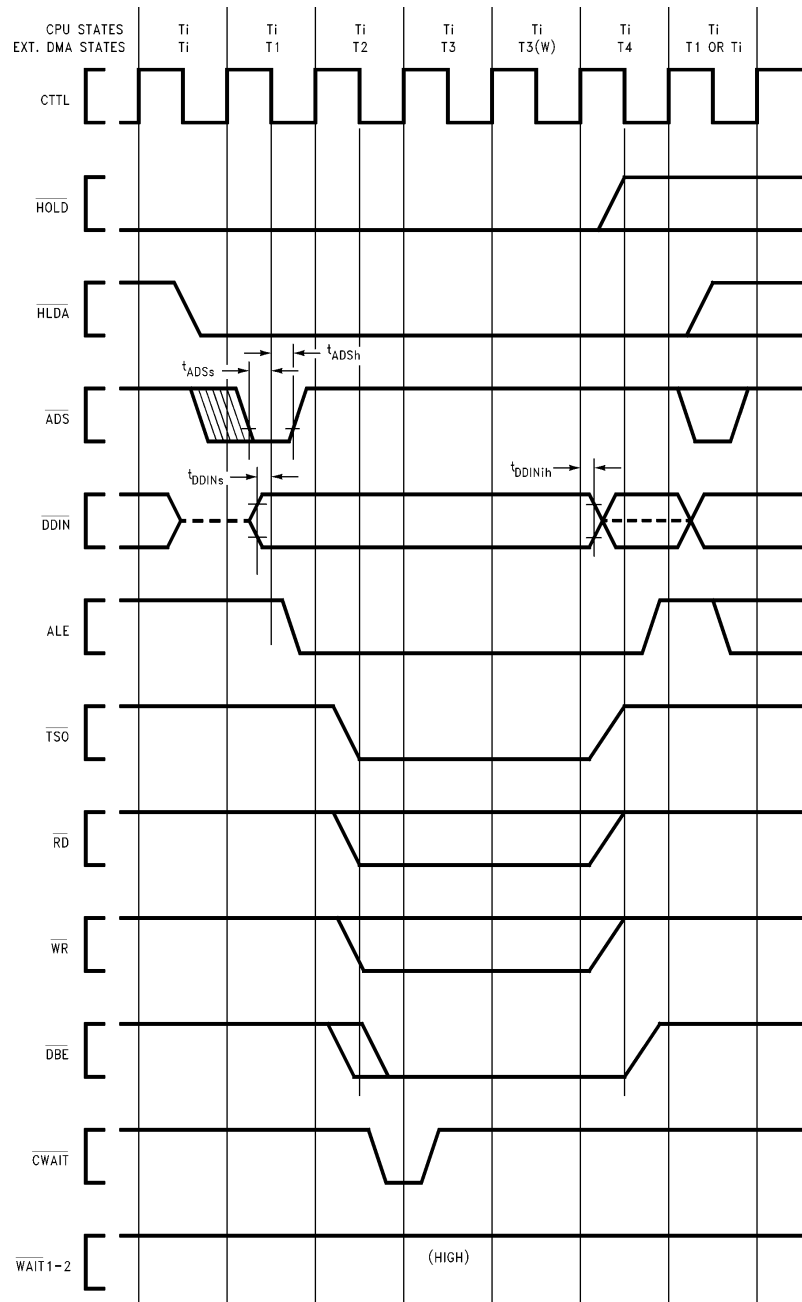
4.0 Device Specifications (Continued)



TL/EE/10752-51

FIGURE 4-9. HOLD Timing (Bus Initially Idle)

4.0 Device Specifications (Continued)



TL/EE/10752-76

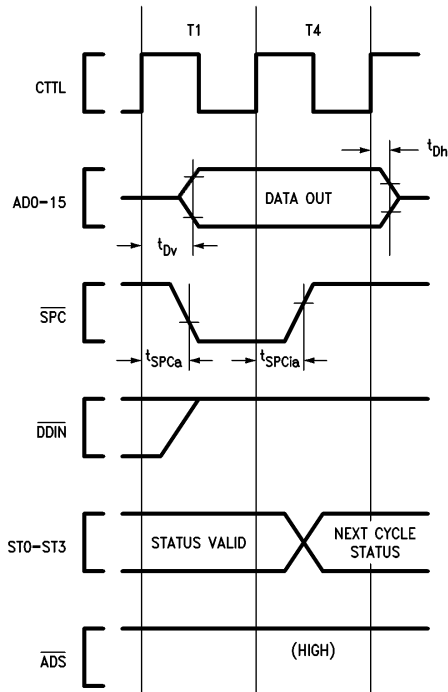
Note 1: \overline{ADS} must be deactivated before state T4 of the external DMA controller cycle.

Note 2: During an external DMA cycle $\overline{WAIT1-2}$ must be kept inactive unless they are monitored by the DMA Controller. An external DMA cycle is similar to a CPU cycle. The NS32CG160 generates \overline{TSO} , \overline{RD} , \overline{WR} , \overline{ALE} and \overline{DBE} . The external DMA controller drives the address/data lines \overline{HBE} , \overline{ADS} and \overline{DDIN} .

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

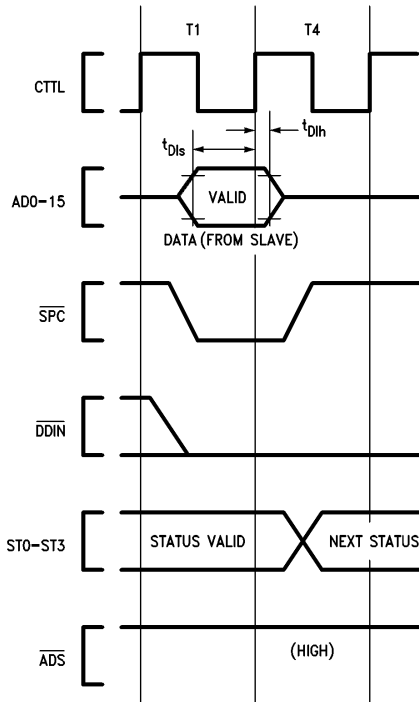
FIGURE 4-10. External DMA Controller Bus Cycle

4.0 Device Specifications (Continued)




TL/EE/10752-53

FIGURE 4-11. Slave Processor Write Timing



TL/EE/10752-54

FIGURE 4-12. Slave Processor Read Timing




TL/EE/10752-55

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 4-13. SPC Timing

4.0 Device Specifications (Continued)



TL/EE/10752-56


FIGURE 4-14. PFS Signal Timing



TL/EE/10752-57

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


FIGURE 4-15. \overline{ILO} Signal Timing



TL/EE/10752-58


FIGURE 4-16. Clock Waveforms

4.0 Device Specifications (Continued)



TL/EE/10752-59


FIGURE 4-17. External Interrupt Requests



TL/EE/10752-60

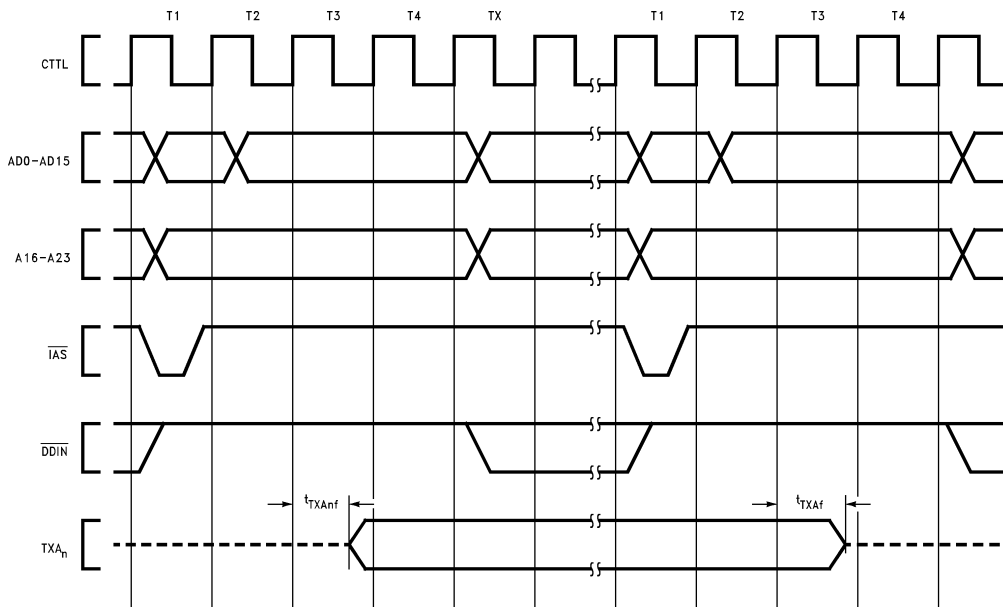
Note: Once TXAn and TXBn are changed, they must stay high or low for at least 4 CCTL cycles.

FIGURE 4-18. TXAn and TXBn Input Timing



TL/EE/10752-61

FIGURE 4-19. TXAn Output Timing



TL/EE/10752-62

FIGURE 4-20. TXAn Float/Non-Float Timing

4.0 Device Specifications (Continued)





FIGURE 4-21. Power-On Reset

TL/EE/10752-63



TL/EE/10752-64

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 4-22. 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 = Lower: 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 Scnd.

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:

DE	C	M	F	I
----	---	---	---	---




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
Scnd	—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

*Bit 23 can be either 0 or 1 for the SETCFG instruction. It is always 0 in all the other cases.




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




Format 8

EXT	—000	INDEX	—100
CVTP	—001	FFS	—101
INS	—010		
CHECK	—011		


Trap (UND) on 110 and 111

TL/EE/10752-65



Format 9

MOVif	—000	ROUND	—100
LFSR	—001	TRUNC	—101
MOVLf	—010	SFSR	—110
MOVFL	—011	FLOOR	—111




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 NS3381 is used.



Format 13


Trap (UND) Always




TL/EE/10752-68

Appendix A: Instruction Formats (Continued)


Format 14
Trap (UND) Always




Format 15
Trap (UND) Always




Format 16
Trap (UND) Always



Format 17
Trap (UND) Always




Format 18
Trap (UND) Always




Format 19
Trap (UND) Always


Implied Immediate Encodings:



Register Mask, Appended to SAVE, ENTER



Register Mask, Appended to RESTORE, EXIT



Offset/Length Modifier Appended to INSS, EXTS

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 NS32CG160.

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 2 3	Access Class Write Access Class Read 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 (Continued)

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 $\langle \rangle$ 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:

$\langle I \rangle$ Immediate

$\langle R \rangle$ CPU Register

$\langle M \rangle$ Memory

$\langle F \rangle$ FPU Register, either 32 or 64 Bits

$\langle x \rangle$ Any Addressing Mode

$\langle ab \rangle$ a and b represent the addressing modes of operand 1 and 2 respectively. Both a and b can be any addressing mode (e.g., $\langle MR \rangle$ 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 $\langle xM \rangle$ case as given in the Notes column.

Only the #TEA1, #TEA2, #TOPI and TCY columns have values assigned for the CMPW instruction. Therefore, they are the 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 $\langle xM \rangle$ case. This means that effective address times have to be calculated for both operands. (For the $\langle MR \rangle$ 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 μs , 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 rrmw 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:

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

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

$$(\text{TOPD} = 7 \text{ (Operand Aligned, No Wait)})$$

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

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

$$T_f = 4$$

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

$$\text{TPR} = 0$$

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	1 1	— —	— —	— —	2 2	— —	9 8	SCR < 0 SCR > 0
ACBi	1 1 — —	— — — —	— — — —	— — — —	— — — —	2 2 — —	— — — —	16 15%20 18 17%22	<M> no branch <M> branch <R> no branch <R> branch
ADDi	1 1 —	1 — —	— — —	— — —	— — —	3 1 —	— — —	3 4 4	<xM> <MR> <RR>
ADDCi	1 1 —	1 — —	— — —	— — —	— — —	3 1 —	— — —	3 4 4	<xM> <MR> <RR>
ADDPi	1 1	1 1	— —	— —	— —	3 3	— —	16 18	No Carry Carry
ADDQi	— —	1 —	— —	— —	— —	2 —	— —	6 4	<M> <R>
ADDR	1 1	1 —	— —	— —	1 —	— —	— —	2 3	<xM> <xR>
ADJSPi	1	—	—	—	—	1	—	6	
ANDi	1 1 —	1 — —	— — —	— — —	— — —	3 1 —	— — —	3 4 4	<xM> <MR> <RR>
ASHi	1	1	1	—	—	2	—	14 → 45	
Bcond	— —	— —	— —	— —	— —	— —	— —	7 6%10	no branch branch
BiCi	1 1 —	1 — —	— — —	— — —	— — —	3 1 —	— — —	3 4 4	<xM> <MR> <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	modular direct
	—	—	—	1	2	—	—	40	
BR	—	—	—	—	—	—	—	4%10	
BSR	—	—	—	—	1	—	—	6%16	
CASEi	1	—	—	—	—	1	—	4%9	
CBITi	1	1	2	—	—	1	—	15	<xM>
	1	—	—	—	—	1	—	7	<xR>
CBITii	1	1	2	—	—	1	—	15	<xM>
	1	—	—	—	—	1	—	7	<xR>
CHECKi	1	1	—	—	—	3	—	7	high low ok
	1	1	—	—	—	3	—	10	
	1	1	—	—	—	3	—	11	
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 field in register
	1	1	—	—	—	1	—	17 → 51	
EXTSi	1	1	—	—	1	1	—	26 → 36	
FFSi	1	1	2	—	—	1	24	24 → 28	
FLAG	—	—	—	4	3	—	—	6	no trap trap
	—	—	—	—	—	—	—	44	
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 field in register
	1	—	—	—	—	1	—	28 → 96	
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> <MR> <RR>
	1	—	—	—	—	1	—	3	
	—	—	—	—	—	—	—	3	
MOVMi	1	1	—	—	—	2 * n	—	3 * n + 20	n = # of elements in block
MOVQi	1	—	—	—	—	1	—	2	<M> <R>
	—	—	—	—	—	—	—	3	
MOVSB, W	—	—	—	—	—	2 * n	—	14 * n + 59	n = # elements no options B, W and/or U option in effect
	—	—	—	—	—	2 * n	—	24 * n + 54	
MOVSD	—	—	—	—	—	2 * n	—	10 * n + 59	n = # of elements no options B, W and/or U option in effect
	—	—	—	—	—	2 * n	—	24 * n + 54	
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	
MULB, W	1	1	—	—	—	3	—	9	
MULD	1	1	—	—	—	3	—	12	
NEGi	1	1	—	—	—	2	—	5	
NOP	—	—	—	—	—	—	—	3	
NOTi	1	1	—	—	—	2	—	5	
ORi	1	1	—	—	—	3	—	3	<xM> <MR> <RR>
	1	—	—	—	—	1	—	4	
	—	—	—	—	—	—	—	4	
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	—	—	—	2	2	—	—	64	modular
	—	—	—	1	2	—	—	59	direct
RETT	—	—	—	2	2	—	—	45	modular
	—	—	—	1	1	—	—	41	direct
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	modular
	—	—	—	1	2	—	—	40	direct
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,	—	—	—	—	—	—	17	8	<FF>
SUBf,	1	—	f	—	—	f	(14 → 17) + 3f	0	<MF>
MULf,	—	—	—	—	—	f	24 + f	0	<IF>
DIVf	—	1	2f	—	—	2f	(25 → 29) + 6f	0	<FM>
	—	1	2f	—	—	3f	(27 → 30) + 3f	0	<IM>
	1	1	3f	—	—	3f	(13 → 19) + 9f	0	<MM>
MOVf,	—	—	—	—	—	—	17	6	<FF>
ABSf,	1	—	f	—	—	f	(14 → 17) + 3f	0	<MF>
NEGf	—	—	—	—	—	f	24 + f	0	<IF>
	—	—	f	—	—	f	23 + 3f	6 + TEA2	<FM>
	—	—	f	—	—	2f	33 + f	TEA2 - 2 - f	<IM>
	1	—	2f	—	—	2f	(20 → 23) + 6f	TEA2 - 3	<MM>
MOVFL	—	—	—	—	—	—	17	8	<FF>
	1	—	1	—	—	1	17 → 20	0	<MF>
	—	—	—	—	—	1	25	0	<IF>
	—	—	2	—	—	2	35	6 + TEA2	<FM>
	—	—	2	—	—	3	43	TEA2 - 3	<IM>
	1	—	3	—	—	3	35 → 38	TEA2 - 3	<MM>
MOVLF	—	—	—	—	—	—	16	8	<FF>
	1	—	2	—	—	2	20 → 23	0	<MF>
	—	—	—	—	—	2	26	0	<IF>
	—	—	1	—	—	1	32	TEA2 + 6	<FM>
	—	—	1	—	—	3	42	TEA2 - 4	<IM>
	1	—	3	—	—	3	35 → 38	TEA2 - 3	<MM>
TRUNCfi,	—	—	—	—	1	—	20	9	<FR>
FLOORfi,	1	—	f	—	1	f	(17 → 20) + 3f	0	<MR>
ROUNDfi	—	—	—	—	1	f	25 + f	0	<IR>
	—	—	—	1	1	—	20	TEA2 + 6	<FM>
	—	—	—	1	1	f	26 + f	TEA2 - 2	<IM>
	1	—	f	1	1	f	(16 → 19) + 4f	TEA2 - 2 - f	<MM>
MOVif	—	—	—	—	1	—	25 - f	0	<RF>
	1	—	—	1	1	—	18	0	<MF>
	—	—	—	—	1	—	26	0	<IF>
	—	1	f	—	1	f	20 + 4f	0	<RM>
	—	1	f	—	1	f	22 + 5f	0	<IM>
	1	1	f	1	1	f	(10 → 13) + 5f	0	<MM>
CMPf	—	—	—	—	—	—	23	13	<FF>
	1	—	f	—	—	f	(20 → 23) + 3f	7	<MF>
	—	—	—	—	—	f	31 + f	7	<IF>
	—	1	f	—	—	f	(27 → 30) + 3f	0	<FM>
	—	1	f	—	—	2f	29	0	<IM>
	1	1	2f	—	—	2f	(15 → 21) + 6f	0	<MM>
	—	—	—	—	—	f	37 + f	0	<FI>
	1	—	f	—	—	2f	(21 → 29) + 8f	0	<MI>
	—	—	—	—	—	2f	35 + 2f	0	<II>
SFSR	—	—	—	—	—	1	19	7	<R>
	1	—	1	—	—	1	20	TEA1 + 4	<M>
LFSR	—	—	—	—	—	1	23	0	<R>
	1	—	1	—	—	1	18 → 21	0	<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	$16 + \text{Twaitrds} + \text{Twaitrdd} + \text{Twaitwr}$ $28 + \text{Twaitblt}$	Shift = 0 Shift = 1 → 8
EXTBLT	$35 + (19 + (12 + (\text{Twaitrds} + \text{Twaitrdd} + \text{Twaitwr})) * \text{width}) * \text{height}$ $35 + (13 + (12 + (\text{Twaitrds} + \text{Twaitrdd} + \text{Twaitwr})) * \text{width}) * \text{height}$	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	$39 + (2 * \text{Twaitrdd} + 2 * \text{Twaitwr} + 2 * \text{Twaitrds})$ $42 + (2 * \text{Twaitrdd} + 2 * \text{Twaitrds})$	$R2 \leq 25$ $R2 > 25$
SBITP	$8 + (34 * R2) + ((\text{Twaitrdd} + \text{Twaitwr}) * R2)$	

**NS32CG160-15/NS32CG160-20/NS32CG160-25
32-Bit Integrated System Processor**

Physical Dimensions inches (millimeters)

Lit. # 114301



84-Pin Plastic Leaded Chip Carrier (V)
Order Number NS32CG160V-15, NS32CG160V-20 or NS32CG160V-25
NS Package Number V84A

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: onjwge@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.

This datasheet has been download from:

www.datasheetcatalog.com

Datasheets for electronics components.