



**CRAY-1® AND CRAY X-MP
COMPUTER SYSTEMS**

SOLID-STATE STORAGE DEVICE (SSD)
REFERENCE MANUAL

HR-0031

CRAY

RESEARCH, INC.

CRAY-1® AND CRAY X-MP COMPUTER SYSTEMS

**SOLID-STATE STORAGE DEVICE (SSD)
REFERENCE MANUAL**

HR-0031

Copyright© 1982 by CRAY RESEARCH, INC. This manual or parts thereof may not be reproduced in any form without permission of CRAY RESEARCH, INC.

Each time this manual is revised and reprinted, all changes issued against the previous version in the form of change packets are incorporated into the new version and the new version is assigned an alphabetic level. Between reprints, changes may be issued against the current version in the form of change packets. Each change packet is assigned a numeric designator, starting with 01 for the first change packet of each revision level.

Every page changed by a reprint or by a change packet has the revision level and change packet number in the lower righthand corner. Changes to part of a page are noted by a change bar along the margin of the page. A change bar in the margin opposite the page number indicates that the entire page is new; a dot in the same place indicates that information has been moved from one page to another, but has not otherwise changed.

Requests for copies of Cray Research, Inc. publications and comments about these publications should be directed to:

CRAY RESEARCH, INC.,
1440 Northland Drive,
Mendota Heights, Minnesota 55120

<u>Revision</u>	<u>Description</u>
	December, 1982 - Original printing.

PREFACE

This publication describes the Cray Research, Inc., Solid-state Storage Device (SSD) for Cray Computer Systems. This publication contains operation and programming information for the SSD and assumes the reader has a familiarity with digital computers.

Details of Cray Research, Inc., mainframes that interface with the SSD are given in the CRAY X-MP Series Mainframe Reference Manual, publication HR-0032, the CRAY-1 S Series Mainframe Reference Manual, publication HR-0029, and the CRAY-1 M Series Mainframe Reference Manual, publication HR-0064.

////////////////////////////////////

WARNING

This equipment generates, uses, and can radiate radio frequency energy and if not installed and used in accordance with the instructions manual, may cause interference to radio communications. It has been tested and found to comply with the limits for a Class A computing device pursuant to Subpart J of Part 15 of FCC Rules, which are designed to provide reasonable protection against such interference when operated in a commercial environment. Operation of this equipment in a residential area is likely to cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

////////////////////////////////////

CONTENTS

<u>PREFACE</u>	iii
1. <u>GENERAL INFORMATION</u>	1-1
INTRODUCTION	1-1
Conventions	1-1
SSD MEMORY PORTS	1-3
Port 0 - Internal Memory Refresh	1-4
Port 1 - Maintenance Port	1-4
Port 2 - CRAY-1 S and CRAY-1 M Port	1-4
Port 3 - CRAY X-MP Port	1-4
DATA PROTECTION	1-5
SSD MEMORY	1-5
Memory control and addressing	1-6
Memory cycle	1-6
2. <u>SSD AND CRAY X-MP OPERATION</u>	2-1
INTRODUCTION	2-1
SSD WRITE OPERATION	2-2
SSD READ OPERATION	2-2
SSD PROGRAMMING	2-3
SSD block address	2-3
Central Memory address	2-3
Transfer block length read/write mode	2-3
Status word error flags	2-4
Block length registers	2-4
Clear interface/abort sequence	2-4
Master clear after power on	2-5
3. <u>SSD AND CRAY-1 S OR CRAY-1 M OPERATION</u>	3-1
INTRODUCTION	3-1
HIGH SPEED CONTROLLER	3-1
HSC COMMAND CHANNEL	3-3
HSC command word	3-4
Read/write field	3-4
CPU loopback field	3-5
Clear 100 Mbytes per second channel error field	3-5
Diagnostic mode field	3-6

HSC command word (continued)	
Transfer length field	3-6
CPU address field	3-6
SSD address field	3-7
HSC STATUS CHANNEL	3-7
HSC status word	3-8
Source field	3-8
Status code field	3-9
SSD PROGRAMMING	3-13
Channel instructions	3-13
Channel master clear sequence	3-15
HSC programming sequence	3-15
HSC command sequence example	3-15
HSC Status Channel interrupt routine example	3-16

FIGURES

1-1	SSD chassis	1-2
1-2	SSD memory organization	1-6
2-1	CRAY X-MP Series Computer System with SSD	2-1
3-1	CRAY-1 S Series Computer System with SSD	3-2
3-2	Data and control paths between SSD and CRAY-1 S or CRAY-1 M mainframe	3-3
3-3	HSC command word format	3-5
3-4	Address words sent from HSC to CPU side of 100 Mbytes per second channel	3-6
3-5	Address words sent from HSC to SSD port interface	3-7
3-6	HSC status word format	3-8

TABLES

1-1	SSD operating characteristics	1-3
1-2	Memory cycles	1-7
3-1	Command channel error status	3-9
3-2	SSD read status codes	3-9
3-3	SSD write status codes	3-11
3-4	Channel instruction summary	3-14

INDEX

INTRODUCTION

The Cray Research, Inc., Solid-state Storage Device (SSD) is a device for temporary storage of datasets and can be configured with a CRAY X-MP, a CRAY-1 S, or a CRAY-1 M Computer System to significantly increase data transfer rates. The maximum data transfer burst rate is dependent upon the SSD memory size and configuration.

The SSD is housed in a stand-alone, 4-column chassis (figure 1-1). Three columns contain the SSD memory modules; one column contains channels and control logic. The power supplies and cooling system are similar to those used in a Cray mainframe. The bench at the base of each column houses the DC power supplies for that column, and a power distribution unit for the SSD supplies 400 Hz power from the Cray Computer System motor-generators to the SSD chassis power supplies. The SSD chassis is cooled by the system condensing unit. An auxiliary condensing unit may be required to support increased system heat load.

The SSD is located adjacent to a CRAY X-MP mainframe but can be 70 feet from a CRAY-1 S or a CRAY-1 M mainframe.

Operating characteristics of the SSD are summarized in table 1-1. SSD memory ports, data protection, and memory are described in this section.

CONVENTIONS

Unless otherwise indicated, numbers in this manual are decimal. Octal numbers are indicated with an 8 subscript.

Italicized lowercase letters, such as *jk*, indicate variable information.




Figure 1-1. Solid-state Storage Device

Table 1-1. SSD operating characteristics

Characteristic	Description
Word size	64 data bits and 8 check bits
Memory size	8,388,608 64-bit words 16,777,216 64-bit words 33,554,432 64-bit words
Number of 64-word blocks	65,536 for 8-million word memory 131,072 for 16-million word memory 262,144 for 32-million word memory
Memory Ports	Port 0 - Internal memory refresh Port 1 - Maintenance port Port 2 - CRAY-1 S or CRAY-1 M mainframe port Port 3 - CRAY X-MP mainframe port
Maximum data transfer burst rate for:	
Port 2	100 Mbytes per second for 8-million word, 16-million word, or 32-million word memory
Port 3	320 Mbytes per second for 8-million word memory 640 Mbytes per second for 16-million word memory 1250 Mbytes per second for 32-million word memory
Data protection	SECDED in memory SECDED on I/O interface

SSD MEMORY PORTS

The SSD has four memory ports. One is used internally; one is used for maintenance purposes; and two are used for data input and output. All four ports connect to SSD memory and are described in the following paragraphs. An Error Log passes port error information and memory error

information to a maintenance computer, and the maintenance computer program logs the error information for maintenance analysis.

PORT 0 - INTERNAL MEMORY REFRESH

Port 0 controls refreshing of SSD memory. Because SSD memory is refreshed within each memory chip, there is no actual data path associated with Port 0. Memory is refreshed every two milliseconds. The refresh operation uses much of the SSD control logic also used in other port operations.

PORT 1 - MAINTENANCE PORT

Port 1, the maintenance port, connects the SSD to a maintenance computer through a 6 Mbytes per second channel that has 16-bit asynchronous control logic.[†] Port 1 is programmable and executes maintenance commands received from the maintenance computer. Port 1 cannot be used as a data input or data output port.

PORT 2 - CRAY-1 S OR CRAY-1 M PORT

Port 2 connects the SSD with a CRAY-1 S or a CRAY-1 M mainframe by adapting the SSD to a Memory Channel that has a maximum data transfer rate of 100 Mbytes per second.^{††} The channel pair is connected to a CRAY-1 S or CRAY-1 M mainframe and has an input (write to SSD) and an output (read from SSD) channel. Section 3 of this manual describes the SSD and CRAY-1 S or CRAY-1 M mainframe operation.

PORT 3 - CRAY X-MP PORT

Port 3 connects the SSD directly to the CRAY X-MP mainframe using a channel that has maximum data transfer rate of 1250 Mbytes per second. Section 2 of this manual describes the SSD and CRAY X-MP operation.

[†] For a description of the 6 Mbytes per second (16-bit asynchronous) control channel, refer to a Cray mainframe reference manual.

^{††} For a description of the 100 Mbytes per second channel, refer to a Cray mainframe reference manual.

DATA PROTECTION

To protect data, single error correction/double error detection (SECDED) logic is used in SSD memory and on data channels to or from the SSD. SECDED logic used is similar to the SECDED logic used in a Cray mainframe Central Memory[†].

When data is written into SSD memory, a checkword is generated for the word to be checked and stored with that word. The checkword is an 8-bit Hamming^{††} code. When the word is read from SSD memory, the checkword and data word are processed to determine if any bits were altered. If no errors occurred, the word is passed to the mainframe.

If an error occurred, the 8 bits of the checkword are analyzed by the logic to find the number of altered bits. If only a single bit was altered, the correction logic resets that bit to the correct state and passes the corrected word out to the mainframe. The Error Log receives details of the error.

If more than a single bit was altered, the logic cannot correct the word. When a double error is detected, it is reported in the 1-word status that is transmitted to the Status Channel (CRAY-1 S or CRAY-1 M configuration) or an error flag is set in the A register (CRAY X-MP configuration). The Error Log receives details of the error.

If more than two bits are in error, results are unpredictable.

SECDED is also used on the data channel when writing data into the SSD. Errors that occur on the channel or in Port 2 or Port 3 logic are corrected and processed as described above. Therefore, there is data protection for write data before storage and data protection for read data after storage.

SSD MEMORY

The SSD has three memory size options: 8 million, 16 million, and 32 million 64-bit words. SSD memory is organized into two groups, Group 0 and Group 1 (figure 1-2) that are used as separate, parallel memories. Each group has up to four sections each divided into eight

[†] For a description of Central Memory, refer to a Cray mainframe reference manual.

^{††} Hamming, R. W., "Error Detection and Correcting Codes," Bell System Technical Journal, 29, No. 2, pp. 147-160 (April 1950).

banks. The sections and banks used are determined by the SSD memory size and the data transfer rate. All eight sections (64 banks) are used in an SSD with a 32-million word memory; four sections (32 banks) are used in an SSD with a 16-million word memory; and two sections (16 banks) are used in an SSD with an 8-million word memory.

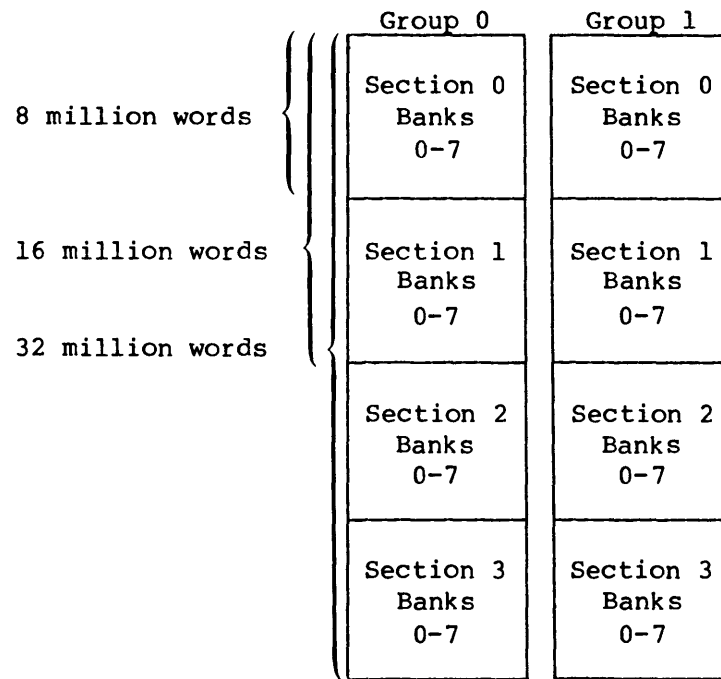


Figure 1-2. SSD memory organization

MEMORY CONTROL AND ADDRESSING

The SSD uses a 64-word block for all transfers into or out of memory. Individual words are not accessible by addressing. A user provides only a starting address for a 64-word block, and a full block is read or written. To read a particular word, the entire block is transferred to Central Memory, then the word is selected using software methods, similar to disk storage data handling methods. Control and addressing are done simultaneously for both memory groups. SSD memory control logic routes each word to the correct location. Once a 64-word block transfer begins, it continues to completion without interruption. Control logic is the same for all memory size options.

MEMORY CYCLE

Control timing of the SSD is based on how long it takes to access half the banks of a 32-million word memory and is the same for all memory size options. After a bank is accessed, there is a delay time (approximately 400 nanoseconds) before it can be accessed again. A bank is considered busy for almost a full memory cycle. Banks are paired as follows: 0 and 4, 1 and 5, 2 and 6, and 3 and 7. If bank n is in the memory cycle, then bank $n+4$ is also considered busy.

The number of words read or written in one memory cycle depends on the number of sections (size) of SSD memory. Therefore, the number of memory cycles needed to transfer a full 64-word block also depends on the size of memory.

A memory cycle for a 32-million word SSD is the transfer of 32 words from half the total number of memory banks, that is, from 32 of its 64 banks. The first memory cycle uses Banks 0, 1, 2, and 3 of all sections and both groups. The second memory cycle uses Banks 4, 5, 6, and 7 of all sections and both groups.

A 32-million word SSD requires two memory cycles to transfer a 64-word block. A 16-million word and an 8-million word SSD use the same memory cycle time as the 32-million word memory but transfer fewer words in each memory cycle. The 16-million word SSD transfers 16 words per memory cycle with words coming from Group 0 and Group 1, Section 0 and Section 1, Banks 0, 1, 2, and 3 and then Banks 4, 5, 6, and 7. The 8-million word SSD transfers 8 words per memory cycle with words coming from Group 0 and Group 1, Section 0, Banks 0, 1, 2, and 3 and then Banks 4, 5, 6, and 7.

Table 1-2 shows the number of words transferred per memory cycle and number of memory cycles needed to transfer a 64-word block for each SSD memory size.

Table 1-2. Memory cycles

Memory size (words)	Words per memory cycle	Memory cycles per 64-word block
32 million	32	2
16 million	16	4
8 million	8	8

INTRODUCTION

When configuring an SSD with a CRAY X-MP mainframe (figure 2-1), the SSD connects directly to the mainframe using a channel with a burst rate capability of 1250 Mbytes per second.

The SSD uses I/O channels 4, 5, 6, and 7 of the CRAY X-MP.[†] The SSD provides high speed data transfer to or from Central Memory under the mainframe's software control.




Figure 2-1. CRAY X-MP Series Computer System with SSD

[†] For additional I/O information, refer to the CRAY X-MP Series Mainframe Reference Manual, publication HR-0032.

Data is transferred between the SSD buffers and the SSD memory at 2 words per SSD clock period (CP). Data is transferred between the SSD buffers and the CRAY X-MP mainframe in 32-word blocks, 2 words per mainframe clock period. For a data transfer, each CPU memory port handles 16 words through one of two buffers. Each CPU memory port handles every other word to or from memory; that is, CPU 0 memory port handles the first CPU address and every other word thereafter. CPU 1 memory port handles every other word starting at CPU address+1.

SSD write and read operations for an SSD operating with a CRAY X-MP and SSD programming information are described in this section.

SSD WRITE OPERATION

Each CPU I/O port has a double buffer containing two 16-word buffers for transmitting data to and receiving data from the SSD. For an SSD write operation, data written to the SSD is first transferred from Central Memory through both CPU I/O ports (one in CPU 0 and one in CPU 1) and written into one of the 16-word buffers of each double buffer. When one 16-word buffer from each double buffer is full, they are paired for a 32-word transfer to the SSD. A check is made to determine if one of the four 32-word buffers in the SSD is available.

When a 32-word buffer is available, the data is sent to that SSD buffer in 16 CPU CPs, two words per CP. While the transfer is in process, the alternate 16-word buffer in each the two CPU I/O port double buffers proceeds to fill in the CPU. When these two 16-word buffers are full, they are paired for a 32-word transfer to the next available 32-word SSD buffer. When a 32-word buffer fills in the SSD, a request is sent to SSD memory, and data is written from the 32-word buffer to SSD memory two words per SSD CP. This operation continues until all data is written to the SSD; then an interrupt is sent to the CPU that initiated the SSD data transfer.

SSD READ OPERATION

For an SSD read operation, when any of the four 32-word buffers in the SSD is empty, a reference is made to read from the SSD memory to Central Memory. Data is then read from SSD memory to a 32-word buffer two words per SSD CP. If the SSD has a full 32-word buffer and each CPU I/O port double buffer has an empty 16-word buffer, the data is sent to the two 16-word buffers at a rate of two words per CPU CP. When that 16-word buffer pair is full, a reference is made from the two CPU I/O memory ports to Central Memory. This operation continues until the last 16-word buffer pair is written to Central Memory, then an interrupt is sent to the CPU that initiated the SSD data transfer.

SSD PROGRAMMING

The 1250 Mbyte channel is available to the software as I/O channel 7 of the CRAY X-MP. A sequence of instructions must be issued to this channel to start the read or write operation. Either a read operation can be active or a write operation can be active. The instruction sequence to start an operation is defined below with $(A_j)=7$ for each instruction.

<u>Octal code</u>	<u>CAL syntax</u>	<u>Description</u>
0012 <i>j</i> 0	CI, <i>A_j</i>	Clear interrupt
0010 <i>j</i> <i>k</i>	CA, <i>A_j</i> <i>A_k</i>	(<i>A_k</i>)=SSD block address
0010 <i>j</i> <i>k</i>	CA, <i>A_j</i> <i>A_k</i>	(<i>A_k</i>)=Central Memory address
0011 <i>j</i> <i>k</i>	CL, <i>A_j</i> <i>A_k</i>	(<i>A_k</i>)=Transfer block length/read write mode; transfer initiated.
-----	-----	Wait interrupt (channel #7)
033 <i>i</i> <i>j</i> 1	A <i>i</i> CE, <i>A_j</i>	Read error flags
0012 <i>j</i> 0	CI, <i>A_j</i>	Clear interrupt/clear interface

SSD BLOCK ADDRESS

All transfers to or from the SSD must be in blocks of 64 words. Therefore, the A register value must be the SSD absolute address divided by 100_8 .

CENTRAL MEMORY ADDRESS

The Central Memory address is used to transfer information into Central Memory. This address can start anywhere within the address field of the CPU.

TRANSFER BLOCK LENGTH READ/WRITE MODE

The transfer block length is the positive number of 64-word blocks to transfer. The maximum block length is 777777_8 . To specify the direction of the transfer, the high-order bit of the *A_k* register (instruction 011*j**k*) is used. The bit clear indicates a read operation (SSD to Central Memory transfer). The bit set indicates a write operation (Central Memory to SSD transfer).

STATUS WORD ERROR FLAGS

The low-order bits of the A_k register contain the remaining block count. The high-order bits are used for error flags and status bits as follows.

<u>Bits</u>	<u>Description</u>
2 ²³	Fatal error (bit 2 ²² , 2 ²¹ , or 2 ²⁰ set)
2 ²²	Central Memory double-bit error
2 ²¹	SSD memory double-bit error (during read); double-bit channel error (during write).
2 ²⁰	Block length error
2 ¹⁹	SSD memory single-bit error (during read); single-bit channel error (during write).
2 ¹⁸	Request in progress
2 ¹⁷⁻²⁰	Remaining block length

BLOCK LENGTH REGISTERS

Two block length registers connect a port on the SSD to an interface on the CRAY X-MP. One register is on the mainframe side of the interface; one register is on the SSD side. Instruction 033*i*jl reads only the mainframe side block length register.

NOTE

On an error, the remaining block length counter can be off by up to 5 blocks. This is because there are six 32-word buffers in the path for any one direction, and on some errors it takes time to shut down the channel. Therefore, if an attempt is made to use some of the data read, or to rewrite less than the entire transfer, a possible 5-block discrepancy should be taken into account.

CLEAR INTERFACE/ABORT SEQUENCE

The interface on the CRAY X-MP connects to an SSD port and is cleared after every transfer. This operation takes approximately 15 CPs. Clearing the interface is initiated on the first Clear Interrupt instruction (0012*j*0) after a transfer. The entire interface is cleared, allowing the next transfer to occur. The error flags are cleared in the process. Under normal operating conditions (no errors or

aborts), the next operation can begin immediately following the clear interrupt instruction. A pending operation is held until the clear finishes.

An abort can be attempted by issuing a Clear Interface instruction (0012*j*0) on a system that is operational. If an abort is attempted or if errors are encountered during the transfer, the Request in Progress status bit (2^{18}) must be checked after the Clear Interrupt instruction (0012*j*0) and before the new operation executes the Set SSD Address instruction (0010*j**k*). The status bit must be clear before starting the next operation. Because a pending reference to the SSD must complete before a new transfer can start, an abort attempt can vary in time depending on refresh in the SSD and other I/O activity within the SSD. Since the software directly controls the Central Memory address, the address should not be changed before any pending references are completed.

NOTE

The address register on the channels is directly available to instruction 0010*j**k*. Since a new operation can be initiated and stacked when clearing the interface (due to instruction 0012*j*0), there cannot be any SSD I/O references in progress due to aborts or errors when the stacked instruction 0010*j**k* is issued. The Request in Progress flag (bit 2^{18}) checks for SSD I/O references in progress and should be clear before an instruction 0010*j**k* is issued.

MASTER CLEAR AFTER POWER ON

Clearing the interface on the CRAY X-MP initially clears all controls associated with the SSD port. The following instruction clears the interface:

<u>Octal code</u>	<u>CAL syntax</u>	<u>Description</u>
0012 <i>j</i> 0	CI, A <i>j</i> (A <i>j</i>)=7	Clear interface

The Request in Progress status bit (2^{18}) should be checked before the first transfer. The Request in Progress bit should never be set for more than 40 CPs.

SSD AND CRAY-1 S OR CRAY-1 M OPERATION

3

INTRODUCTION

The SSD uses the 100 Mbytes per second channel modules on a CRAY-1 S or on a CRAY-1 M mainframe that are otherwise used for transmitting data between an I/O Subsystem and CRAY-1 mainframe. In addition to these modules, the mainframe controls the operations of the channel to the SSD through a special high speed controller (HSC), which is a group of modules installed in the CRAY-1 S or CRAY-1 M mainframe and cabled to both ends of the 100 Mbytes per second channel link. The mainframe sends transfer commands to and receives status information from the HSC using a modified I/O channel pair. This channel pair is referred to as the HSC command channel (output channel) and the HSC status channel (input channel). The HSC sends the appropriate control signals to both ends of the 100 Mbytes per second channel link to control the transfer of commands and status information.

Figure 3-1 shows a CRAY-1 S Series Computer System configured with an SSD. Special use of the the HSC and I/O channel pair and SSD programming information are described in the following paragraphs.

HIGH SPEED CONTROLLER

The high speed controller (HSC) expands the capabilities of a CRAY-1 S or a CRAY-1 M 100 Mbytes per second channel by allowing the central processing unit (CPU) to originate data transfers. Without the HSC, the CPU 100 Mbytes per second channel can only respond to channel requests from other devices (such as the I/O Subsystem).

The HSC is comprised of two logic modules residing in the mainframe. The HSC decodes commands received from the CPU, returns status to the CPU, initiates 100 Mbytes per second channel transfers, detects error conditions, and forces diagnostic conditions. The only HSC connection to the CPU and SSD 100 Mbytes per second channel is via the channel's control cables; the data cables are connected directly between the CPU and SSD. Figure 3-2 shows the data and control paths between a CRAY-1 S or CRAY-1 M mainframe and SSD.

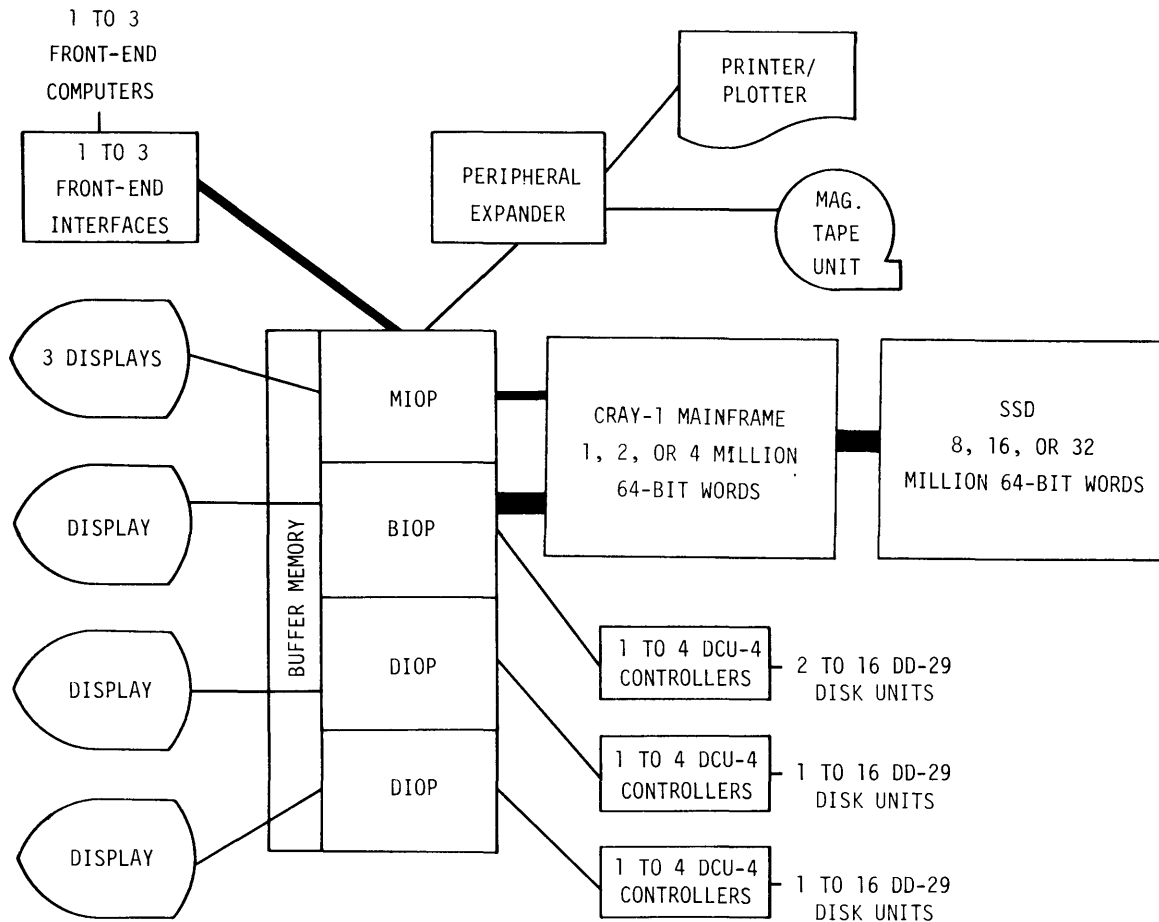


Figure 3-1. CRAY-1 S Series Computer System with SSD

Three control signals (Data Ready, Last Word flag, and Transmit Data) control the flow of data in each direction between a CRAY-1 S or CRAY-1 M mainframe and an SSD. The control signals pass through the HSC unaltered (except under certain diagnostic conditions). The remaining channel signals all originate or terminate at the HSC and are used primarily to initiate data transfers and detect any 100 Mbytes per second channel errors. Once a data transfer is initiated (that is, immediately after the address words are transferred), the data transfer is directed by the three control signals controlling the data flow.

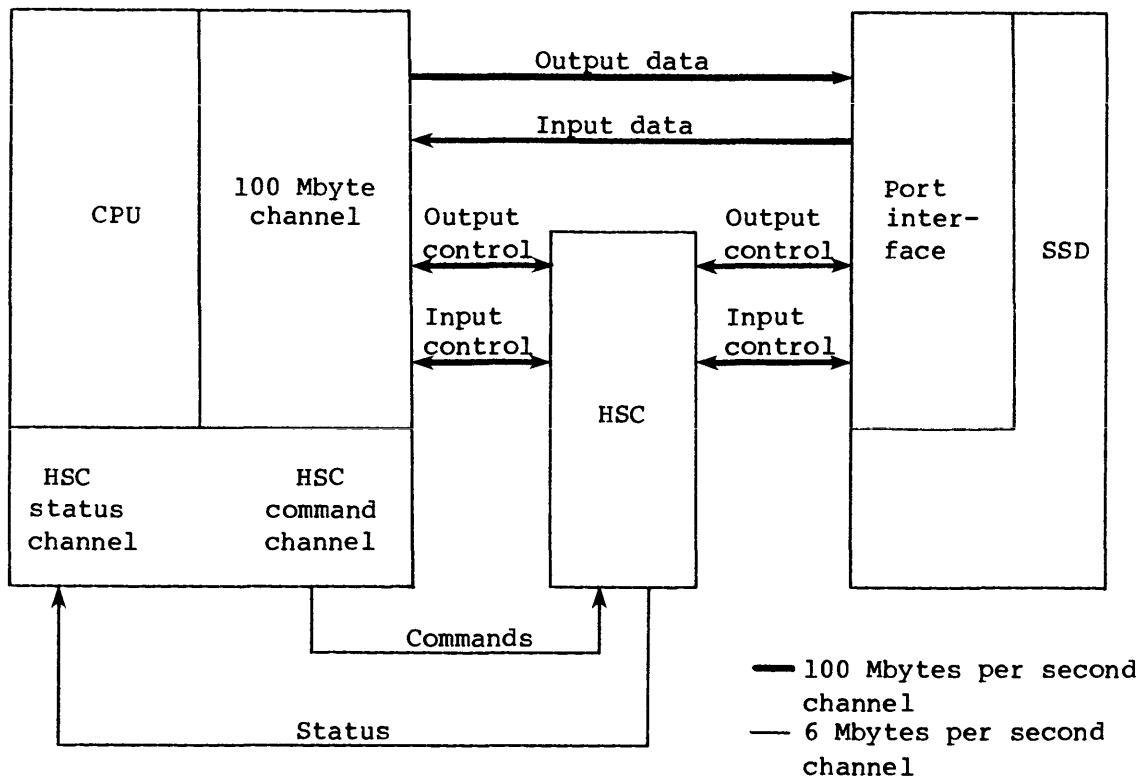


Figure 3-2. Data and control paths between SSD and CRAY-1 S or CRAY-1 M mainframe

HSC COMMAND CHANNEL

To initiate an SSD data transfer, a single 64-bit command word (see subsection on HSC command word) is sent to the HSC from a special mainframe CPU output channel designated the HSC command channel. The HSC command channel occupies one CPU output channel. Because the HSC command channel does not have an interrupt capability, it can be assigned to any output channel without regard to interrupt priority. The HSC command channel differs from a normal CRAY-1 S or CRAY-1 M CPU output channel in the following respects.

- Interrupt request: The HSC command channel cannot generate an interrupt request.
- Channel Error flag: Since no interrupt capability exists on the HSC command channel,[†] the function of the Channel Error flag is redefined to allow testing for

[†] Interrupt capability does exist on the HSC status channel.

Channel Error flag: completion of a command transfer. The Channel Error flag is renamed Command Channel Busy flag and is set whenever the HSC command channel is active. The flag remains set until the Command Channel Disconnect is transmitted or until a Clear Channel Interrupt instruction (0012*jx*) is executed. An I/O master clear from the Maintenance Control Unit, or equivalent, also clears the Command Channel Busy flag. The state of the Command Channel Busy flag is sensed with a normal Read Channel Error flag instruction (033*ijl*). The usual error information conveyed by the Channel Error flag (that is, resume while inactive or resume during memory reference) does not apply to the HSC command channel.

Channel limit address: Since all transfers over the HSC command channel consist of a single CPU word, the Channel Limit (CL) register is not implemented. A Set CL instruction (0011*jk*) for the HSC command channel is executed as a pass. In addition, the Channel Address (CA) register is not incremented after the channel requests a memory reference.

HSC COMMAND WORD

The format of the HSC command word is shown in figure 3-3; the fields are described below.

Read/write field

When set, the 1-bit read/write field indicates that the command pertains to the SSD write data path (write from Central Memory to SSD). When this bit is clear, the command affects the SSD read data path (read to Central Memory from SSD). If the Clear 100 Mbyte Channel Error bit is zero and the diagnostic mode field is not 05, 10, or 15,[†] an SSD data transfer is initiated. One SSD read operation and one SSD write operation can be in progress simultaneously.

[†] Mode 05	Transmit Clear Channel pulse to SSD output side of the 100 Mbytes per second channel without deactivating HSC or CPU input side of the channel.
Mode 10	Transmit Last Word flag to CPU input side of the 100 Mbytes per second channel (only if transfer is in progress)
Mode 15	Transmit Clear Channel pulse to CPU input side of the 100 Mbytes per second channel without deactivating HSC or SSD output side of the channel.

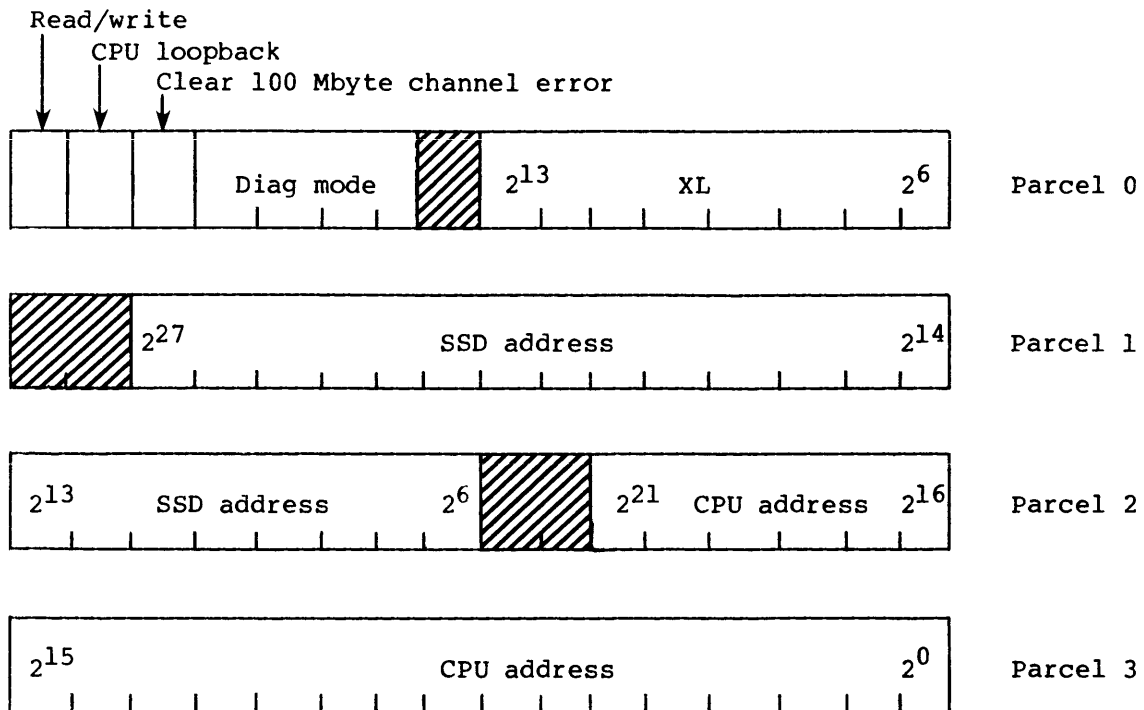


Figure 3-3. HSC command word format

CPU loopback field

The 1-bit CPU loopback field is used for maintenance purposes only. The bit is set only when the CPU output and input sides of the 100 Mbytes per second channel are connected together through the HSC for checking either the HSC or the 100 Mbytes per second channel without using other equipment (that is, the I/O Subsystem or the SSD).

Clear 100 Mbytes per second channel error field

When a 100 Mbytes per second channel error occurs, the channel with the error should be cleared before using it for a subsequent data transfer. Setting this command bit causes a 100-nanosecond Clear Channel signal to be transmitted to the appropriate SSD and CPU sides of the 100 Mbytes per second channel and resets the read or write logic in the HSC. The Read/write bit indicates which data path between the CPU and SSD will be cleared. For example, if an error occurs during a SSD write data transfer, the error is reported via the HSC status channel. A subsequent HSC command with the Read/write bit set and the Clear 100 Mbyte Channel Error bit set results in a Clear Channel signal sent over both the SSD input side and the CPU output side of the 100 Mbytes per second channel. The command also resets the HSC write logic. The clear operation does not result in a status word input. Any SSD read operation in progress at the time is not affected.

When the Clear 100 Mbyte Channel Error bit is set, all other command fields except the Read/write bit are ignored. In addition, the Clear Error command is ignored by the HSC if the error code is zero.

Diagnostic mode field

This 4-bit field is for maintenance purposes only. During normal operation, this field must contain all zeros.

Transfer length field

The 8-bit transfer length (XL) field designates the number of 64-word blocks of data to be transferred over the 100 Mbytes per second channel. From 1 to 256 blocks of data can be transferred with a single request. A field of all zeros in the command word indicates the maximum transfer length of 256 blocks (16,384 words). Figures 3-4 and 3-5 illustrate the mapping of address parameters into the three address words transmitted from the HSC to the CPU side of the 100 Mbytes per second channel and to the SSD port interface.




Figure 3-4. Address words sent from HSC to CPU side of 100 Mbytes per second channel

CPU address field

The 22-bit CPU address field contains the starting CPU address for the indicated transfer. Transfers to or from Central Memory can begin at any word address.




Figure 3-5. Address words sent from HSC to SSD port interface

The 22-bit SSD address field contains the starting SSD address divided by 100_8 .

Upon completing an SSD data transfer or upon detecting an error condition, a single 64-bit status word (figure 3-6) is transmitted from the HSC to a special mainframe CPU input channel designated the HSC status channel. An interrupt request is generated upon receiving the status word (see subsection on HSC status word). The HSC status channel occupies one CPU input channel and differs from a normal CRAY-1 S or CRAY-1 M CPU input channel in the following respects.

Programmable master clear: and to pass the master clear pulse along to both the CPU and SSD sides of the 100 Mbytes per second channel. The HSC status channel is also reset in the process. The command channel is not affected by this operation.
(continued)

HSC STATUS WORD

The format of the status word received from the HSC is shown in figure 3-6, and the fields are described below.




Figure 3-6. HSC status word format

Source field

The 2-bit source field indicates the operation pertaining to the accompanying status code. The source field is coded as follows:

<u>Code</u>	<u>Function</u>
00	Indicates an error was detected on the HSC command channel
01	Pertains to the most recently executed SSD read operation
10	Pertains to the most recently executed SSD write operation
11	Undefined

Status code field

The interpretation of the 4-bit status code field depends upon the source of the status word as indicated in the source field description above. Tables 3-1 through 3-3 list the possible status codes and their descriptions for command channel errors, read status, and write status, respectively.

Table 3-1. Command channel error status; source field = 00₂

Status code (binary)	Name	Description
00x1	Command parity error	A parity error was detected by the HSC in the HSC command word; command rejected.
001x	Command channel error	An error in the HSC command channel protocol was detected by the HSC; command rejected.

Table 3-2. SSD read status codes; source field = 01₂

Status code (octal)	Name	Description
00	Transfer complete	SSD read data transfer completed normally
01	Function error	A command to initiate a new SSD read transfer or a command to clear a read error was received before the previous read operation completed or before it was terminated by an error.
02	SSD active error	The SSD output side of the 100 Mbytes per second channel went inactive before sending Last Word flag while the CPU side was still active, or the SSD port enable switch is in the disable position.

Table 3-2. SSD read status codes; source field = 01₂ (continued)

Status code (octal)	Name	Description
03	SSD Transmit Address timeout	An SSD read transfer was attempted but the SSD output side of the 100 Mbytes per second channel failed to respond with a Transmit Address signal within three milliseconds, or the SSD channel enable switch is in the disable position.
04	SSD address error	The SSD output side of the 100 Mbytes per second channel received greater than or less than three Address Ready pulses, or there was a parity error in one of the address words transferred to the SSD.
05	SSD data error	The SSD detected a multiple-bit error upon reading memory.
06		Not used
07	SSD Data Ready timeout	The SSD output side of the 100 Mbytes per second channel is active but has not transmitted any data for three milliseconds.
10		Not used
11		Not used
12	CPU active error	The CPU input side of the 100 Mbytes per second channel went inactive while the SSD side of the channel was still active.
13	CPU Transmit Address timeout	An SSD read transfer was attempted but the CPU input side of the 100 Mbytes per second channel failed to respond with a Transmit Address within three milliseconds.
14	CPU address error	The CPU input side of the 100 Mbytes per second channel received greater than or less than three Address Ready pulses, or there was a parity error in one of the address word transfers to the CPU.

Table 3-2. SSD read status codes; source field = 01₂ (continued)

Status code (octal)	Name	Description
15	CPU data/block length error	The CPU input side of the 100 Mbytes per second channel detected a multiple-bit error in the data received from the SSD, or the last word of data was received and the block length count was not zero, or the last word of data was not received and the block length count was zero, or the CPU side received data when its data buffers were full.
16	CPU Last Word timeout	The last word of data was sent from the SSD output side of the 100 Mbytes per second channel but the CPU side failed to go inactive within three milliseconds.
17	CPU Transmit Data timeout	The Transmit Data signal from the CPU input side of the 100 Mbytes per second channel was not received by the HSC for three milliseconds while active.

Table 3-3. SSD write status codes; source field = 10₂

Status code (octal)	Name	Description
00	Transfer complete	SSD write data transfer completed normally.
01	Function error	A command to initiate a new SSD write transfer or a command to clear a write error was received before the previous write operation completed or before it was terminated by an error.
02	SSD active error	The SSD input side of the 100 Mbytes per second channel went inactive while the CPU side was still active, or the SSD port enable switch is in the disable position.

Table 3-3. SSD write status codes; source field = 10₂ (continued)

Status code (octal)	Name	Description
03	SSD Transmit Address timeout	An SSD write transfer was attempted but the SSD input side of the 100 Mbytes per second channel failed to respond with a Transmit Address signal within three milliseconds, or the SSD channel enable switch is in the disable position.
04	SSD address error	The SSD input side of the 100 Mbyte channel received greater than or less than three Address Ready pulses, or there was a parity error in one of the address words transferred to the SSD.
05	SSD data/block length error	The SSD detected a multiple-bit error while writing data received by the SSD input side of the 100 Mbytes per second channel, or the last word of data was received by the SSD side and the block length count was not zero, or the last word of data was not received and the block length count was zero, or the SSD side received data when its data buffers were full.
06	SSD Last Word timeout	The last word of data was sent from the CPU output side of the 100 Mbytes per second channel but the SSD side failed to go inactive within three milliseconds.
07	SSD Transmit Data timeout	The Transmit Data signal from the SSD input side of the 100 Mbytes per second channel was not received by the HSC for three milliseconds while active.
10		Not used
11		Not used
12	CPU active error	The CPU output side of the 100 Mbyte channel went inactive before sending Last Word flag while the SSD side was still active.

Table 3-3. SSD write status codes; source field = 10₂ (continued)

Status code (octal)	Name	Description
13	CPU Transmit Address timeout	An SSD write transfer was attempted but the CPU output side of the 100 Mbytes per second channel failed to respond with a Transmit Address within three milliseconds.
14	CPU address error	The CPU output side of the 100 Mbyte channel received greater than or less than three Address Ready pulses, or there was a parity error in one of the address word transfers to the CPU.
15	CPU data error	The CPU detected a multiple-bit error upon reading memory.
16		Not used
17	CPU Data Ready timeout	The CPU output side of the 100 Mbyte channel is active but has not transmitted any data for three milliseconds.

SSD PROGRAMMING

The following programming notes apply when configuring an SSD with a CRAY-1 S or a CRAY-1 M Series Computer System.

CHANNEL INSTRUCTIONS

Table 3-4 lists the I/O instructions implemented on the CPU. Each instruction's use is summarized for normal I/O channel pairs and for the HSC command/status channel pair.

Table 3-4. Channel instruction summary

I/O instruction	Normal usage	Command/status channel usage
$\dagger 0010jk$	Set CA of channel (A_j) to (A_k); activate channel; interrupt when transfer is complete.	Set CA of channel (A_j) to (A_k), set CL to (A_k)+1, activate channel. HSC status channel interrupts when status word transfer is done. HSC command channel never interrupts.
$\dagger 0011jk$	Set CL of channel (A_j) to (A_k).	Master clear HSC via HSC status channel (A_j). Ignored by command channel.
$\dagger 0012jx$	Clear Interrupt flag and Error flag on channel (A_j) and deactivate channel.	Clear Interrupt flag and Error flag on HSC status channel (A_j) and deactivate channel. Clear HSC Command Channel Busy flag on HSC command channel (A_j) and deactivate channel.
$033i0x$	Channel number of highest priority interrupt request to A_i .	Channel number of highest priority interrupt request to A_i .
$033ij0$	CA of channel (A_j) to A_i .	CA of channel (A_j) to A_i .
$033ij1$	Error flag of channel (A_j) to A_i .	Error flag of HSC status channel (A_j) to A_i . Channel Busy flag of HSC command channel (A_j) to A_i .

 \dagger Privileged to Monitor Mode

CHANNEL MASTER CLEAR SEQUENCE

The HSC command/status channel pair can be master cleared with the following instruction sequence.

<u>Octal code</u>	<u>CAL syntax</u>	<u>Description</u>
022 <i>ijk</i>	A1 ICH	HSC status channel number
022 <i>ijk</i>	A2 OCH	HSC command channel number
0011 <i>jk</i>	CL,A1 Ak	Master Clear HSC status channel, HSC, and 100 Mbytes per second channel
0012 <i>jk</i>	CI,A2	Deactivate HSC command channel

After executing this sequence, neither the HSC command channel nor the HSC status channel should be activated for 8 CPs. This sequence clears the HSC command channel, the HSC status channel, the HSC, and the SSD and CPU control channels of the 100 Mbytes per second channel connected to the HSC.

This operation differs from transmitting a Clear 100 Mbyte Channel Error command. In this operation, the master clear sequence is all encompassing, whereas the Clear 100 Mbyte Channel Error command only clears the selected 100 Mbytes per second channel data path, that is, one input channel, one output channel, and a portion of the HSC logic.

HSC PROGRAMMING SEQUENCE

The HSC provides only one I/O interrupt request (generated by the HSC status channel) for each data transfer. Two program sequences are involved in communicating with the HSC. First, a command must be sent to the HSC. Subsequently, the HSC status channel interrupt request must be processed upon completing the data transfer.

HSC command sequence example

An HSC command sequence example is listed below. Upon activating the command channel, the channel remains active for approximately 48 CPs, assuming no Central Memory conflicts. If the HSC command channel is active for a substantially longer period of time, the channel can be assumed to be hung. For this reason, the timeout check in the sample sequence checks the channel's active time.

<u>Octal code</u>	<u>CAL syntax</u>		<u>Description</u>
022206	CMDSEQ A2	6	Timeout value; approximately 1 usec.
0221 <i>jk</i>	A1	OCH	HSC command channel number
033011	CHKCMD A0	CE,A1	Sample Command Channel Busy flag
011 <i>ijkm</i>	JAZ	OUTCMD	Jump if channel not busy
031020	A0	A2-1	Else, wait for not busy
031220	A2	A2-1	
011 <i>ijkm</i>	JAN	CHKCMD	Loop if timeout not expired
006 <i>ijkm</i>	J	ERR	Exit; timeout error.
0202 <i>jk</i>	OUTCMD A2	CMD	Address of command word
001012	CA,A1	A2	Set CA and activate HSC command channel
006 <i>ijkm</i>	J	EXIT	

HSC status channel interrupt routine example

An example of an HSC status channel interrupt routine is listed below. This example assumes the CPU responded to the HSC status channel interrupt request with an exchange jump into the interrupt handler. If it is preferable to remain in Monitor Mode after executing the command sequence, the HSC Status Channel address register can be monitored to indicate completion of the transfer. An SSD transfer of maximum block length (256 blocks or 16,384 words) completes in approximately 1.25 milliseconds.

<u>Octal code</u>	<u>CAL syntax</u>		<u>Description</u>
033200	A2	CI	Read channel number of interrupt request
030002	A0	A2	
010 <i>ijkm</i>	JAZ	EXIT	Exit if not I/O interrupt
0221 <i>jk</i>	A1	ICH	HSC status channel number
031021	A0	A2-A1	
011 <i>ijkm</i>	JAN	EXIT	Exit if interrupt not from HSC
1207 <i>jk</i>	S7	STAT,0	Get status word
033011	A0	CE,A1	Read Status Channel Error flag
011 <i>ijkm</i>	JAN	ERR	Exit; channel error.
001210	CI,A1		Clear status channel interrupt request
0202 <i>jk</i>	A2	STAT	Address of status buffer
001012	CA,A1	A2	Reactivate HSC status channel
			Process status word.

INDEX

INDEX

- 100 Mbytes per second channel, 1-4, 3-1
 - programming, 3-3
- 1250 Mbytes per second channel, 1-4, 2-1
 - programming, 2-3
- Address words, 3-6
- Block length registers, 2-4
- Buffers, 2-2
- Central Memory, 2-2
- Central Memory address, 2-3
- Channel Error flag, 3-4
- Channel instruction summary, 3-14
- Channel instructions, 3-13
- Channel limit address, 3-4, 3-7
- Channel master clear sequence, 3-15
- Channels
 - 100 Mbytes per second channel, 1-4, 3-6
 - 1250 Mbytes per second channel, 1-4, 2-1
 - I/O, 2-1
- Clear 100 Mbytes per second channel error field, 3-5
- Clear interface/abort sequence, 2-4
- Command channel, 3-1, 3-3
 - command sequence example, 3-15
 - differences from CPU output channel, 3-3
 - error status, 3-9
 - HSC command word, 3-4
- Command Channel Busy flag, 3-4
- Command Word, 3-4
 - clear 100 Mbytes per second channel error field, 3-5
 - CPU address field, 3-6
 - CPU loopback field, 3-5
 - diagnostic mode field, 3-6
 - format, 3-5
 - read/write field, 3-4
 - SSD address field, 3-7
 - transfer length field, 3-6
- Configurations, 2-1, 3-2
- Control paths between SSD and CRAY-1 M mainframe, 3-3
- Control paths between SSD and CRAY-1 S mainframe, 3-3
- Control timing, 1-7
- Conventions, 1-1
- CPU address field, 3-6
- CPU loopback field, 3-5
- CRAY X-MP port, 1-4
- CRAY X-MP Series Computer System with SSD,
 - 2-1
 - CRAY-1 M port, 1-4
 - CRAY-1 S port, 1-4
 - CRAY-1 S Series Computer System with SSD, 3-2
- Data paths between SSD and CRAY-1 M mainframe, 3-3
- Data paths between SSD and CRAY-1 S mainframe, 3-3
- Data protection, 1-3, 1-5
- Data transfer, 2-2
 - 64-word block, 1-6, 2-3
 - burst rate, 1-3, 2-1
- Diagnostic mode field, 3-6
- Error correction/error detection, 1-5
- Error Log, 1-3, 1-5
- High speed controller (HSC), 3-1
 - command sequence example, 3-15
 - programming sequence, 3-15
 - status channel interrupt routine example, 3-16
- HSC, see High speed controller
- I/O channels, 2-1
- I/O port, 2-2
- Instructions, 2-3, 3-14
- Internal memory refresh port, 1-4
- Interrupt request, 3-3
- Maintenance port, 1-4
- Master clear after power on, 2-5
- Memory Channel, 1-4
- Port interface, 3-7
- Ports, 1-3
 - Port 0 - Internal Memory Refresh, 1-4
 - Port 1 - Maintenance Port, 1-4
 - Port 2 - CRAY-1 S and CRAY-1 M Port, 1-4
 - Port 3 - CRAY X-MP Port, 1-4
- Programmable master clear, 3-7

Read/write field, 3-4
Read operation, 2-2
Read status codes, 3-9

SECEDED, see Error correction/error detection

Source field, 3-8

SSD

- chassis, 1-2
- description, 1-1
- memory, 1-5, 2-2
 - addressing, 1-6
 - control, 1-6
 - cycle, 1-7
 - organization, 1-6
 - ports, 1-3, 2-2
 - size, 1-3, 1-5, 1-7
- operating characteristics, 1-3
- programming, 2-3, 3-13
 - block length registers, 2-4
 - Central Memory address, 2-3
 - channel instructions, 3-13
 - channel master clear sequence, 3-15
 - clear interface/abort sequence, 2-4
 - HSC command sequence example, 3-15
 - HSC programming sequence, 3-15
 - HSC status channel interrupt routine example, 3-16
 - master clear after power on, 2-5
 - SSD block address, 2-3
 - status word error flags, 2-4
 - transfer block length read/write mode, 2-3
- read operation, 2-2
- read status codes, 3-9
- write operation, 2-2
- write status codes, 3-11

SSD address field, 3-7

SSD and CRAY X-MP operation, 2-1

SSD and CRAY-1 M operation, 3-1

SSD and CRAY-1 S operation, 3-1

SSD block address, 2-3

SSD read status codes, 3-9

SSD write status codes, 3-11

Status channel, 1-5, 3-1, 3-7

- differences from CPU input channel, 3-7

- HSC status word, 3-8

- interrupt routine example, 3-16

Status code field, 3-9

Status word, 3-8

- format, 3-8

- source field, 3-8

- status code field, 3-9

Status word error flags, 2-4

Transfer block length read/write mode, 2-3

Transfer length field, 3-6

Write operation, 2-2

Write status codes, 3-11

READERS COMMENT FORM

Solid-state Storage Device (SSD) Reference Manual

HR-0031

Your comments help us to improve the quality and usefulness of our publications. Please use the space provided below to share with us your comments. When possible, please give specific page and paragraph references.

NAME _____

JOB TITLE _____

FIRM _____

ADDRESS _____

CITY _____ STATE _____ ZIP _____



CUT ALONG THIS LINE

FOLD



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



BUSINESS REPLY CARD

FIRST CLASS PERMIT NO. 6184 ST. PAUL, MN

POSTAGE WILL BE PAID BY ADDRESSEE

CRAY
RESEARCH, INC.

Attention:
PUBLICATIONS

1440 Northland Drive
Mendota Heights, MN 55120
U.S.A.

FOLD

STAPLE

Cray Research, Inc.
Publications Department
1440 Northland Drive
Mendota Heights, MN 55120
612-452-6650
TLX 298444