CRAY T90 SERIES I/O THEORY OF OPERATIONS

Figures

Figures (continued)

Tables

Overview

The I/O module is the interface between the CPU and channels and the outside world of the CRAY T90 series system. Each I/O module supports the following channels:

- 8 low-speed (LOSP) channels
- 8 high-speed (HISP) channels
- 4 very high-speed (VHISP) channels
- 4 LOSPX channels

The term *LOSPX* refers to either the error logger channel, the maintenance channel, the support channel, or the programmable interrupt/maintenance control unit interrupt (PINT/MCUI) channel. Quadrant position on the I/O module determines the type of a given LOSPX channel. LOSPX channel assignments are as follows: quadrant 0 is the maintenance channel; quadrant 1 is the support channel; quadrant 2 is the error logger channel; and quadrant 3 is the PINT/MCUI channel.

A fully configured CRAY T90 series system (a CRAY T932 system with four I/O modules) can support the following channels:

- 32 LOSP
- 32 HISP
- 16 VHISP
- 16 LOSPX

A CRAY T916 system can support the following channels:

- 16 LOSP
- 16 HISP
- 8 VHISP
- 8 LOSPX

A CRAY T94 system can support the following channels:

- 8 LOSP
- 8 HISP
- 4 VHISP
- 4 LOSPX

Logical channel numbers are as follows:

- LOSP: 100 through 177
- HISP: 400 through 437
- VHISP: 20 through 37
- support channel: 60 through 67

An I/O module is connected directly to a CPU module with electronic zero insertion force (EZIF) orthogonal interconnect module (OIM) connectors. OIMs connect modules that are on an orthogonal plane. Within a stack of 8 CPU modules (in a CRAY T916 or CRAY T932 system), CPUs 2, 3, 4, and 5 connect to the I/O module. In a CRAY T94 system, CPUs 0, 1, 2, and 3 connect to the I/O module.

External cabling starts from the I/O module via I/O connectors. Cables within the CRAY T90 series system are routed through the chassis and connect to the bulkhead where external cable connections reside.

Any CPU can access any I/O channel. There is not always a direct connection between the CPU making a request and the I/O module receiving it; therefore, I/O commands are always sent to the local shared module first. If the local shared module cannot access the I/O module, it sends the command to the remote shared module. The remote shared module sends the command to an I/O channel through the CPU module that lies between the shared module and the I/O module.

Channel Assignments

Table 1 lists all channel assignments for the CRAY T932 system. The column labeled "Quadrant" refers to the four quadrants on the I/O module. The column labeled "I/O Module" refers to the physical position of the module within the system. Note that HISP channel numbers are for reference only because the CPU cannot issue a HISP instruction.

Table 2 provides similar information for the CRAY T916 system, and Table 3 provides information for CRAY T94 systems.

Table 1. I/O Channel Assignments (CRAY T932 System)

I/O Module	CPU Physical Number	I/O Module Quadrant	VHISP Number	LOSP Number	HISP Number	Support Channel Number
A	2	0	20	100/101,102/103	400, 401	60/61
	3	1	21	104/105,106/107	402, 403	
	4	2	22	110/111,112/113	404, 405	
	5	3	23	114/115,116/117	406, 407	
E	12	0	24	120/121,122/123	410, 411	62/63
	13	1	25	124/125,126/127	412, 413	
	14	2	26	130/131,132/133	414, 415	
	15	3	27	134/135,136/137	416, 417	
I	22	0	30	140/141,142/143	420, 421	64/65
	23	1	31	144/145,146/147	422, 423	
	24	2	32	150/151,152/153	424, 425	
	25	3	33	154/155,156/157	426, 427	
M	32	0	34	160/161,162/163	430, 431	66/67
	33	1	35	164/165,166/167	432, 433	
	34	2	36	170/171,172/173	434, 435	
	35	3	37	174/175,176/177	436, 437	

Table 2. I/O Channel Assignments (CRAY T916 System)

I/O Module	CPU Physical Number	I/O Module Quadrant	VHISP Number	LOSP Number	HISP Number	Support Channel Number
E	2	0	20	100/101,102/103	400, 401	60/61
	3	1	21	104/105,106/107	402, 403	
	4	2	22	110/111,112/113	404, 405	
	5	3	23	114/115,116/117	406, 407	
I	12	0	24	120/121,122/123	410, 411	62/63
	13	1	25	124/125,126/127	412, 413	
	14	2	26	130/131,132/133	414, 415	
	15	3	27	134/135,136/137	416, 417	

Table 3. I/O Channel Assignments (CRAY T94 System)

I/O Module	CPU Physical Number	I/O Module Quadrant	VHISP Number	LOSP Number	HISP Number	Support Channel Number
A	0	0	20	100/101,102/103	400, 401	60/61
	1	1	21	104/105,106/107	402, 403	
	2	2	22	110/111,112/113	404, 405	
	3	3	23	114/115,116/117	406, 407	

I/O Module Description

The following paragraphs describe the I/O module's orientation within the entire system as well as its logical and physical layout. The options and connectors on the I/O module are also described.

I/O Module Orientation in the System

Figure 1 shows the position of the I/O modules within the system, which is divided into four quadrants: 0, 1, 2, and 3. An I/O module can reside in any or all quadrants. The I/O modules are designated A (quadrant 0), E (quadrant 1), I (quadrant 2), and M (quadrant 3).

Figure 1 illustrates a CRAY T932 system. The CRAY T916 system can contain two I/O modules, and the CRAY T94 system contains only one.

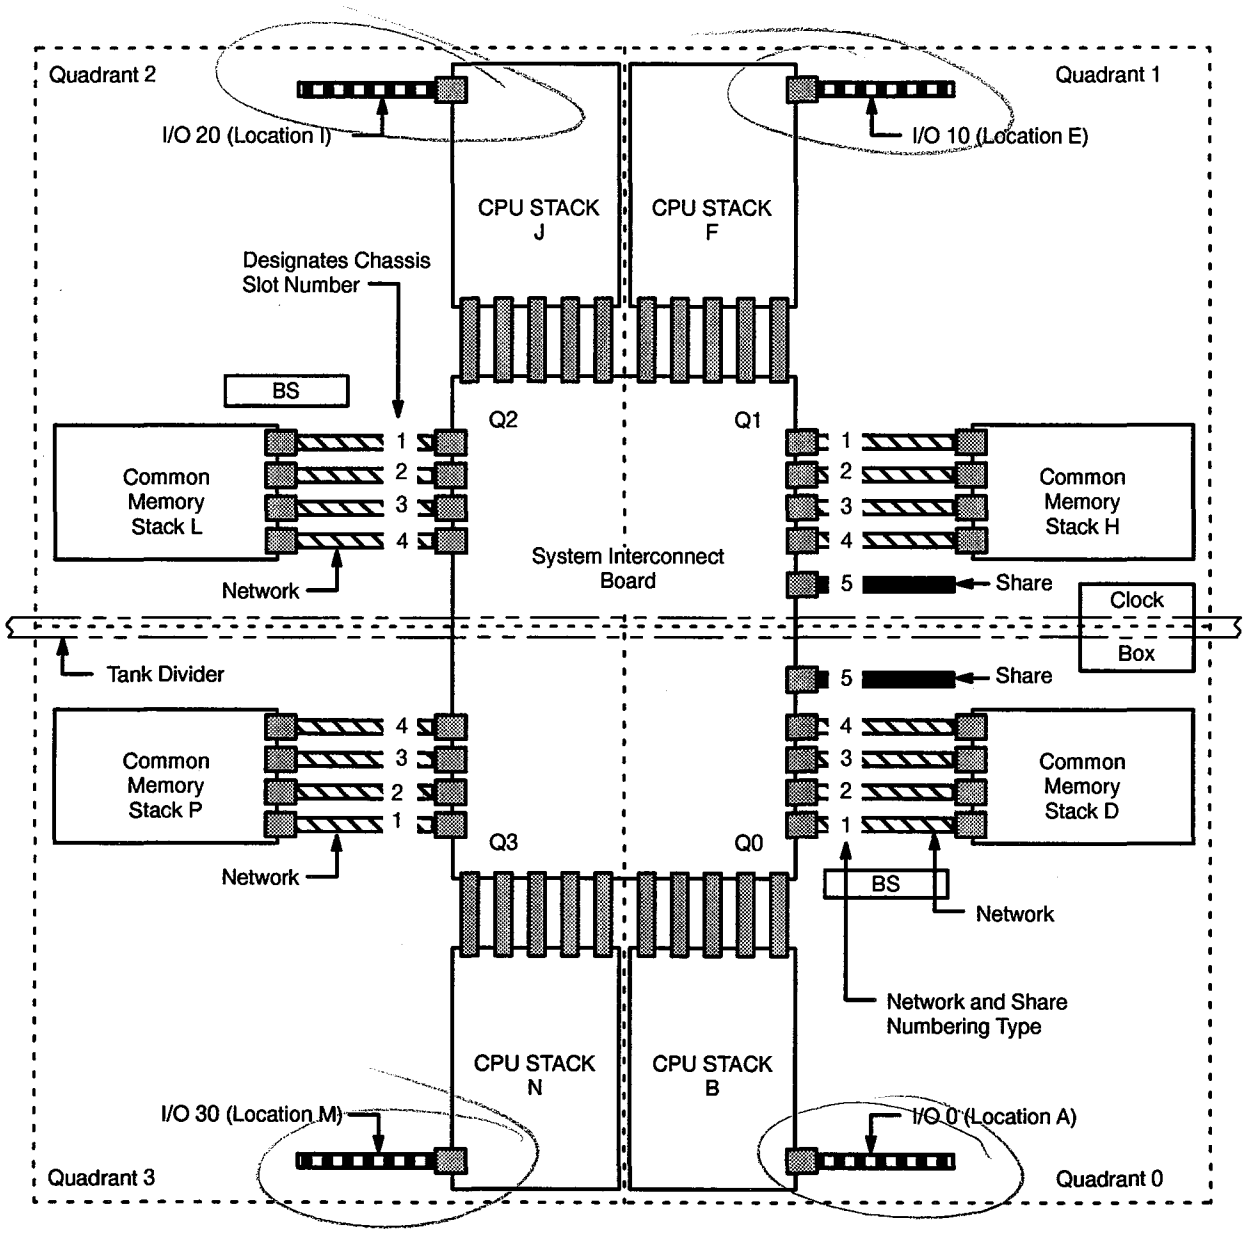

Figure 1. I/O Module Position within System

Logical Layout of Module

The I/O module is divided into four logical groups or quadrants. Each quadrant on the module can handle transfers for the following channels:

- 2 LOSP channels
- 2 HISP channels
- 1 VHISP channel
- 1 LOSPX channel

Each quadrant is logically connected to one of the four CPUs that make physical connection to the I/O module.

Figure 2 shows options and channel assignments per quadrant for the I/O modules in a CRAY T916 or CRAY T932 system. Software channel numbering is based on the position of the module within the system.

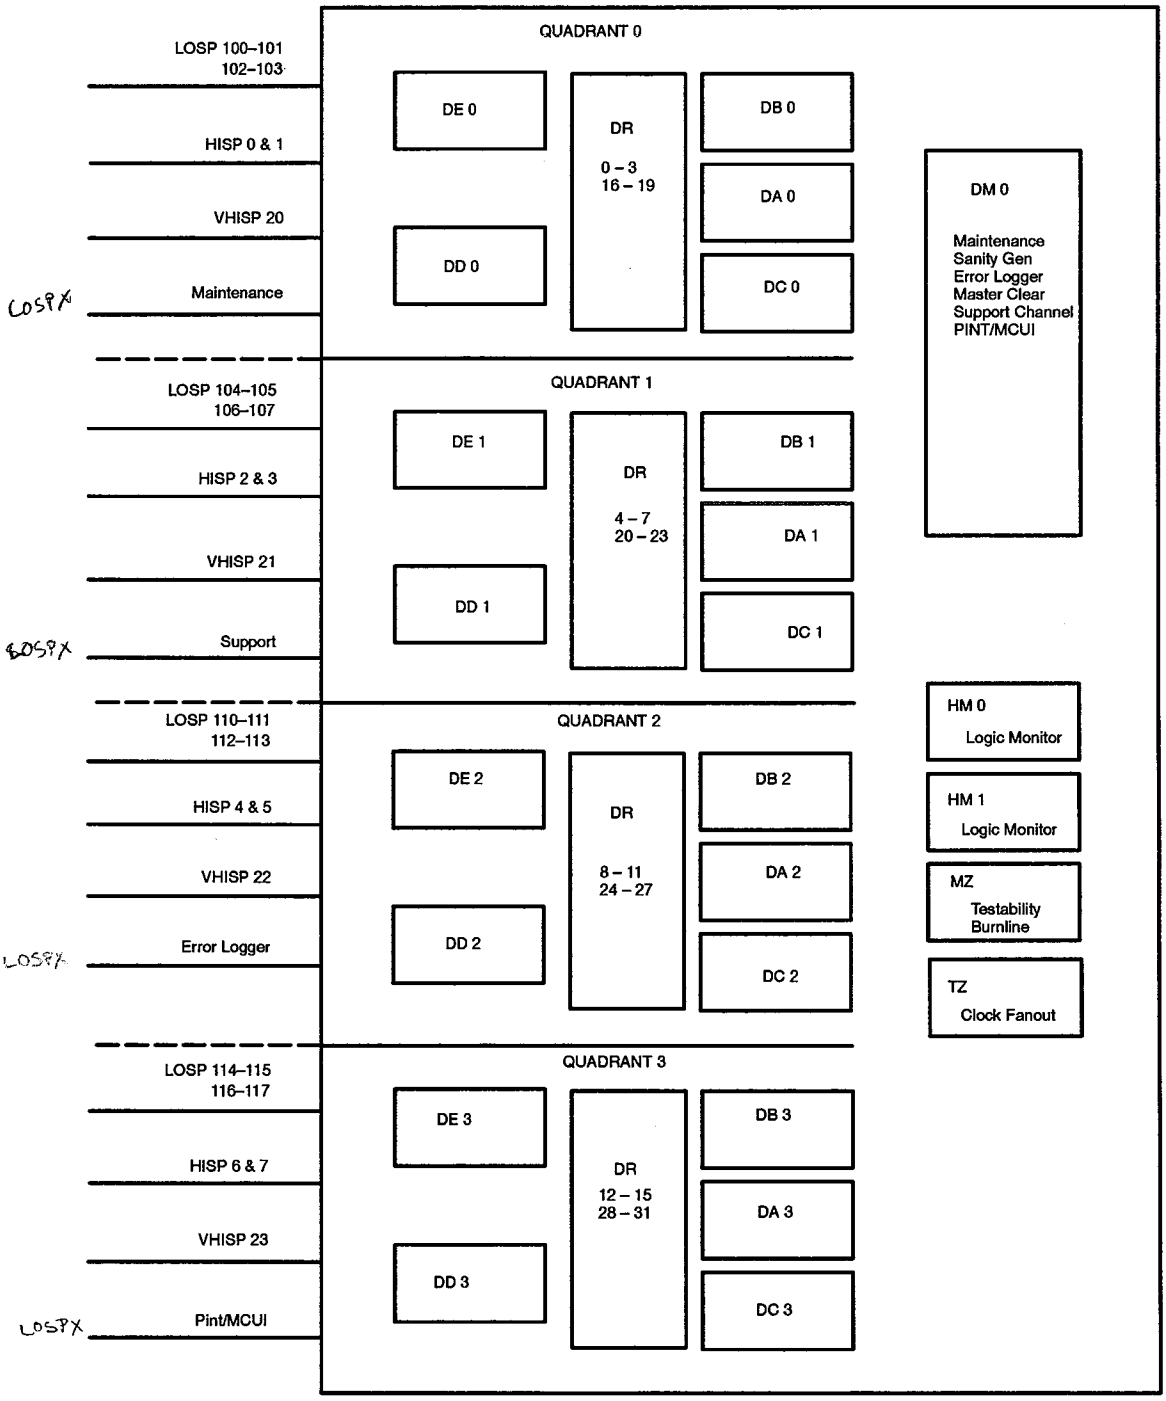

Figure 2. I/O Module Logical Layout

Physical Layout of Module

The I/O module is a single assembly consisting of one 8-layer printed circuit board (A board) laminated between two 52-layer printed circuit boards (boards 1 and 2). There are logic chips mounted on boards 1 and 2, as shown in Figure 3.

Board 1	52 layers
A board	8 layers
Board 2	52 layers

Figure 3. I/O Module Construction

There are 16 channel connectors, which are described later in this document. Terminators for VHISP control signals are also located on the module. Figure 4 is a diagram of the physical layout of board 1; Figure 5 shows board 2.

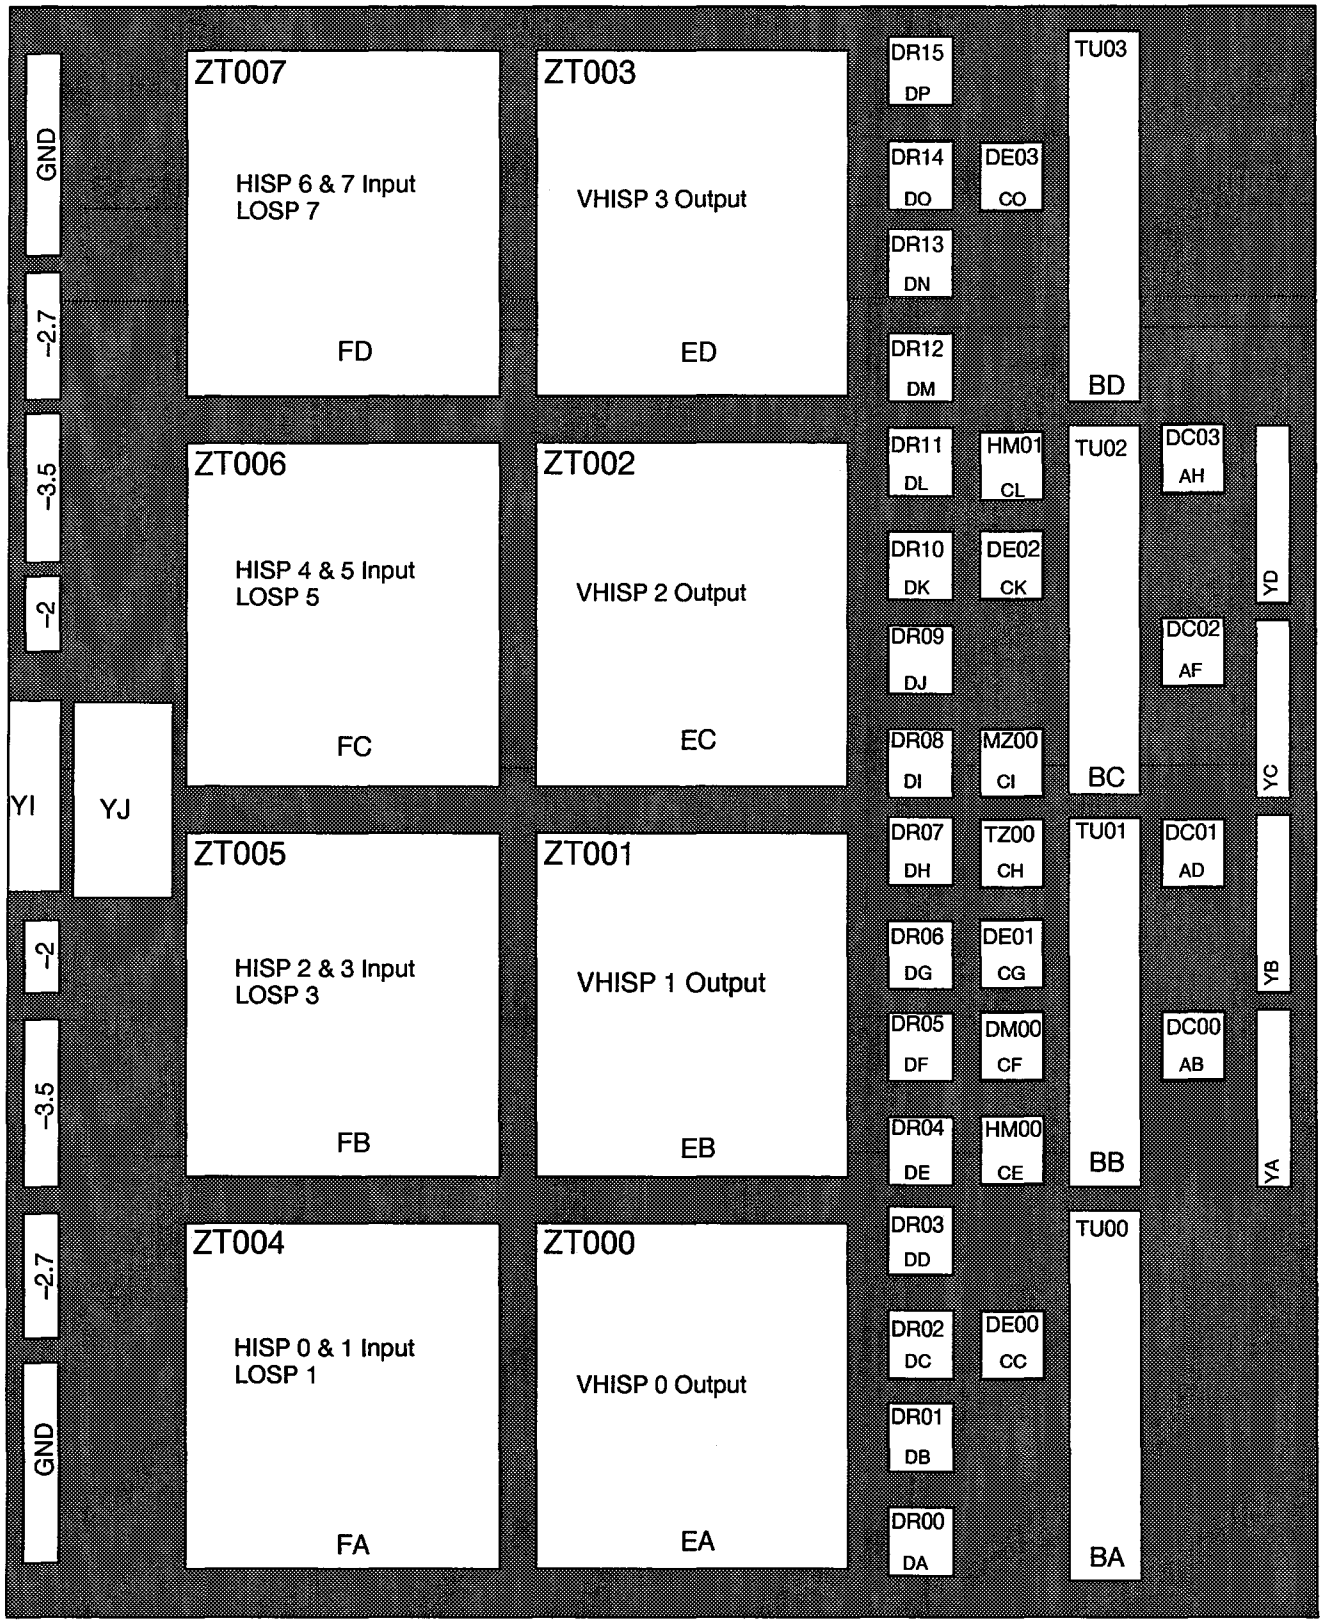

Figure 4. I/O Module Board 1

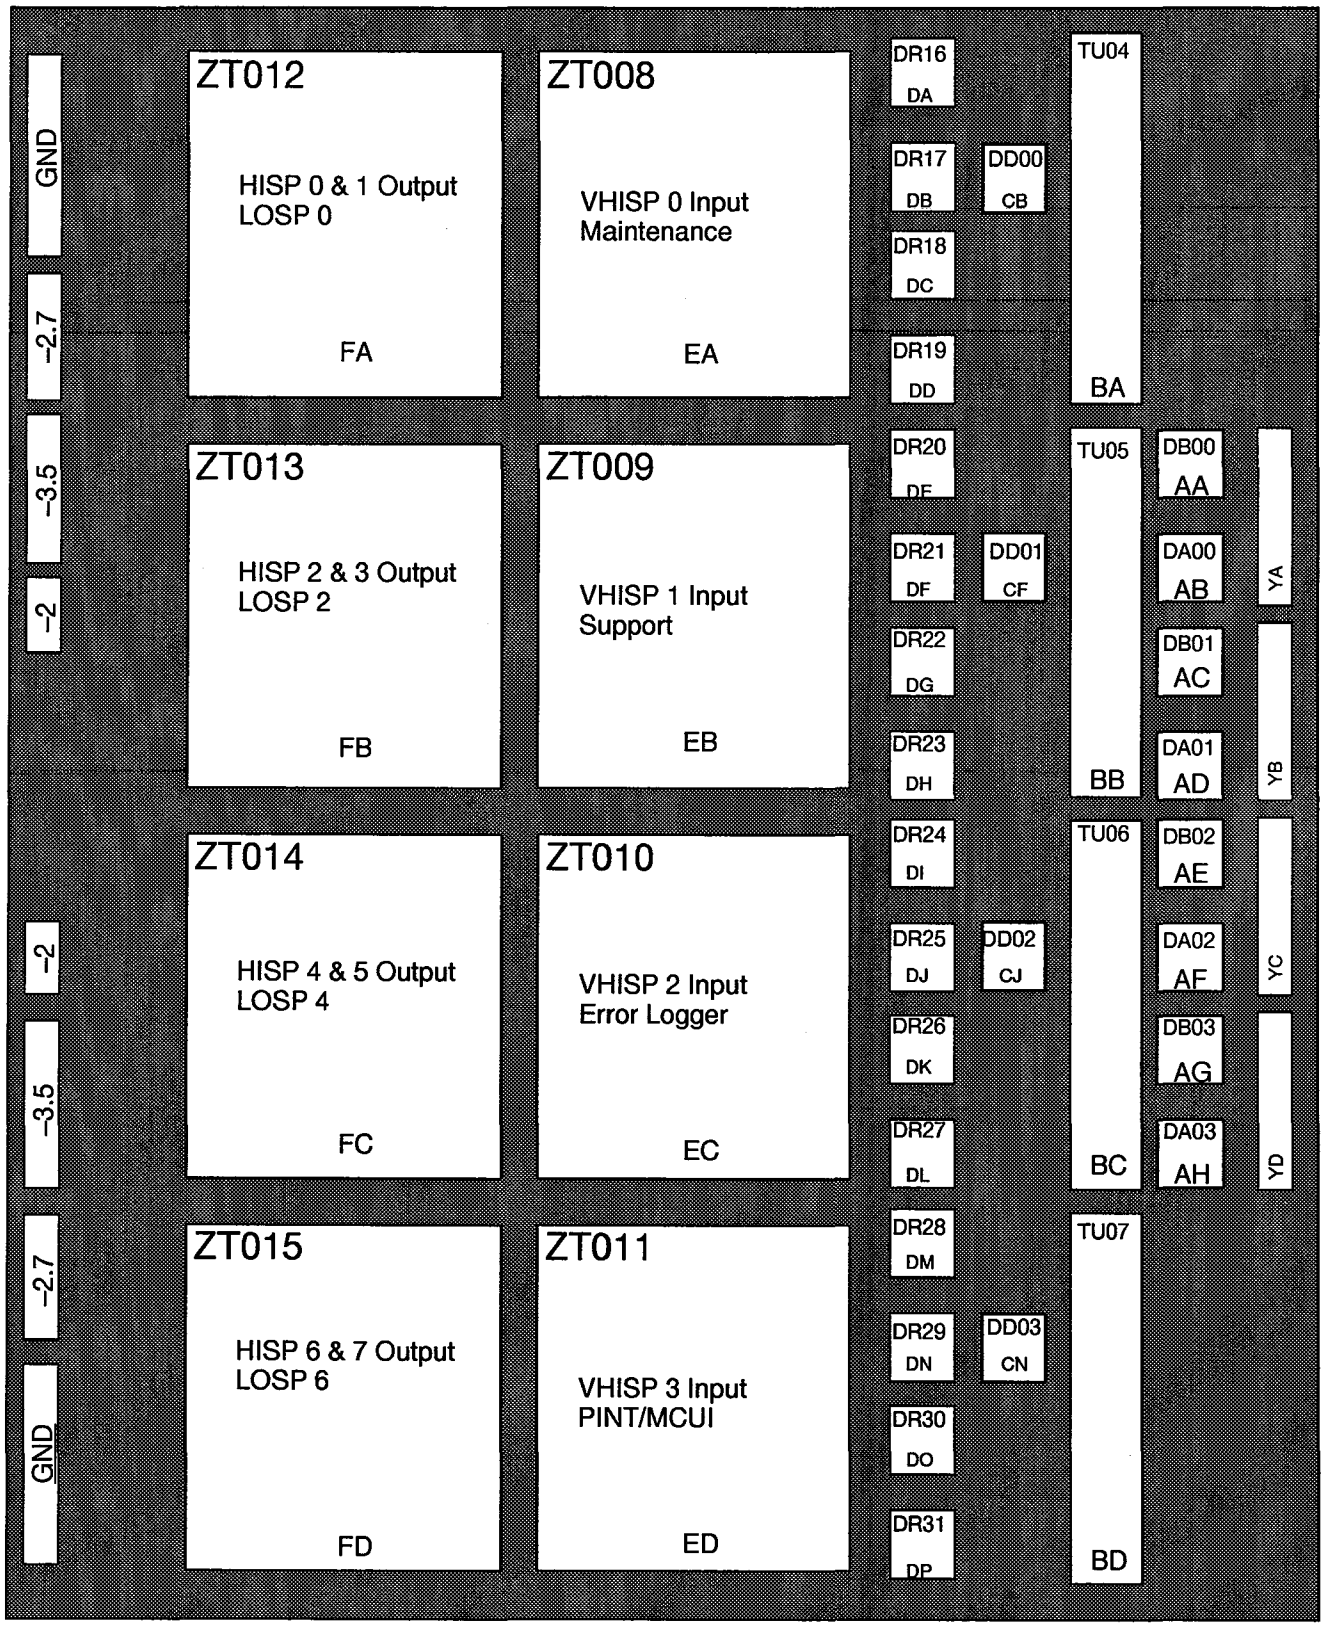

Figure 5. I/O Module Board 2

Options on the I/O Module

Each of the four logical quadrants contains the following options:

- DA option (1 each quadrant)
- DB option (1 each quadrant)
- DC option (1 each quadrant)
- DD option (1 each quadrant)
- DE option (1 each quadrant)
- DR option (8 each quadrant)

The I/O module also contains one DM option for maintenance channel distribution, and error logger channel and support channel routing to available CPUs, and sanity code generation. The following paragraphs briefly describe the function of the options found on the I/O module.

DA Option

This option is one of two data interface options used between the CPU and I/O channels. An I/O module contains four DA options. These options are labeled DA00, DA01, DA02, and DA03, which correspond to the four quadrants of the module: 0, 1, 2, and 3.

During a read operation (memory to I/O) the DA option receives 32 data bits and 4 check bits of each memory word. This option handles only the odd data bits and odd check bits (the DC option handles the even bits). Data is then passed to the DR options of that quadrant, based on the channel transfer.

During a write operation (I/O to memory), the DA option passes 32 data bits (bits 00 through 15 and 32 through 47) and all 8 check bits to the CPU. The DA option receives data from the DR options in the quadrant.

The DA option also verifies SECDED on the data it passes for either the read or write operation. On a LOSP write, the DA option generates check bits before the data passes to the CPU.

DB Option

Four DB options reside on the I/O module. The DB option is the interface between the CPU and the channel controller options (DD/DE) for the I/O module. As with the DA and DC options, each quadrant of the I/O module uses one DB option.

The DB option is used to pass reference request information to the CPU for either a read or a write operation. The reference request packet contains a read/write bit, memory group, SECDED control, reference designator, block length, and address.

The DB option receives reference request, interrupt request, and 033 response information from the DD and DE options for all channel types (LOSP, HISP, and VHISP). The DB communicates with the DA, DC, and DR options to send/receive data to and from the CPU.

DC Option

Four DC options reside on the I/O module. These options are labeled DC00, DC01, DC02, and DC03.

During a read operation (memory to I/O), the DC option receives 32 data bits and 4 check bits of each memory word. This option handles only the even data and check bits.

During a write operation (I/O to memory), the DC option passes 32 data bits (bits 16 through 31 and 48 through 63) to the CPU. The data is received from the DR options of the quadrant.

The DC option verifies SECDED on the data it passes for either the read or write operation, and it passes channel address information to the I/O module on LOSP and VHISP transfers.

DD Option

The DD option is the channel control option for the following channels: HISP0 out, HISP1 out, LOSP0, and LOSPX. The DD option sends LOSP control signals to, and receives HISP control signals from, the external devices. Each I/O module has four DD options, one in each quadrant.

The DD option contains the memory addressing registers for the channels indicated, including both current address registers and limit address registers. This option multiplexes its addresses and sends the resulting CPU address to the DB option.

The DD option sends LOSP, VHISP, and HISP control terms to, and receives them from, external devices.

DE Option

The DE option is the channel control option for the HISP0 in, HISP1 in, LOSP1, and VHISP channels. The DE option contains the memory addressing registers for the channels indicated, including both current address registers and limit address registers. This option multiplexes its addresses and sends the resulting CPU address to the DB option.

The DE option sends LOSP, VHISP, and HISP control terms to and receives them from external devices.

DR Option

The DR option is the data buffer option between the channels and external devices. There are 32 DR options on the I/O module. Each quadrant of the I/O module contains 8 DR options. Refer again to Figure 2, the I/O module logical layout block diagram, to see the grouping of options per quadrant.

During a memory to I/O transfer, each DR option handles 18 bits (16 data bits and 2 check bits) received from the DA and DC options. During an I/O to memory transfer, each DR option sends 8 data bits and 1 check bit to the DA and DC options.

The DR option contains two data buffers for read data and two buffers for write data. Each buffer has four ranks. Each rank holds 16 words of data, 16 data bits and 2 check bits. These buffers can be circular; in other words, while one rank is being emptied, another rank can be filling.

The DR option receives memory addressing for the HISP channels. It also receives memory addressing and block length for the VHISP channels.

Table 4 describes the HISP and VHISP data buffer option assignments. Note that the bit assignments for the DR options on board 2 (DR16-31) are swapped with those on board 1. Table 5 describes LOSP channel data buffer option assignments.

Table 4. HISP and VHISP Data Buffer Option Assignments

Channel	Option	Bits
HISP0 input	DR000	0–3, 16–19, 32–35, 48–51, C.B. 0, 4
HISP1 input	DR001	4–7, 20–23, 36–39, 52–55, C.B. 1, 5
	DR002	8–11, 24–27, 40–43, 56–59, C.B. 2, 6
VHISP0 output	DR003	12–15, 28–31, 44–47, 60–63, C.B. 3, 7

Table 4. HISP and VHISP Data Buffer Option Assignments (continued)

Channel	Option	Bits
HISP0 output	DR016	16–19, 0–3, 48–51, 32–35, C.B. 4, 0
HISP1 output	DR017	20–23,4–7, 52–55, 36–39, C.B. 5,1
	DR018	24–27, 8–11, 56–59,40–43, C.B. 6, 2
VHISP0 input	DR019	28–31,12–15, 60–63, 44–47, C.B. 7, 3
HISP2 input	DR004	0–3, 16–19, 32–35, 48–51, C.B. 0, 4
HISP3 input	DR005	4–7, 20–23, 36–39, 52–55, C.B. 1, 5
	DR006	8–11, 24–27, 40–43, 56–59, C.B. 2, 6
VHISP1 output	DR007	12–15, 28–31, 44–47, 60–63, C.B. 3, 7
HISP2 output	DR020	16–19, 0–3, 48–51, 32–35, C.B. 4, 0
HISP3 output	DR021	20–23,4–7, 52–55, 36–39, C.B. 5,1
	DR022	24–27, 8–11, 56–59,40–43, C.B. 6, 2
VHISP1 input	DR023	28–31,12–15, 60–63, 44–47, C.B. 7, 3
HISP4 input	DR008	0–3, 16–19, 32–35, 48–51, C.B. 0, 4
HISP5 input	DR009	4–7, 20–23, 36–39, 52–55, C.B. 1, 5
	DR010	8–11, 24–27, 40–43, 56–59, C.B. 2, 6
VHISP2 output	DR011	12–15, 28–31, 44–47, 60–63, C.B. 3, 7
HISP4 output	DR024	16–19, 0–3, 48–51, 32–35, C.B. 4, 0
HISP5 output	DR025	20–23,4–7, 52–55, 36–39, C.B. 5,1
	DR026	24–27, 8–11, 56–59,40–43, C.B. 6, 2
VHISP2 input	DR027	28–31,12–15, 60–63, 44–47, C.B. 7, 3
HISP6 input	DR012	0–3, 16–19, 32–35, 48–51, C.B. 0, 4
HISP7 input	DR013	4–7, 20–23, 36–39, 52–55, C.B. 1, 5
	DR014	8–11, 24–27, 40–43, 56–59, C.B. 2, 6
VHISP3 output	DR015	12–15, 28–31, 44–47, 60–63, C.B. 3, 7
HISP6 output	DR028	16–19, 0–3, 48–51, 32–35, C.B. 4, 0
HISP7 output	DR029	20–23,4–7, 52–55, 36–39, C.B. 5,1
	DR030	24–27, 8–11, 56–59,40–43, C.B. 6, 2
VHISP3 input	DR031	28–31,12–15, 60–63, 44–47, C.B. 7, 3

Table 5. LOSP Channel Data Buffer Option Assignments

Channel	Option	Bits
LOSP0	DR016	0–3, parity 0
	DR017	4–7, parity 1
	DR018	8–11, parity 2
	DR019	12–15, parity 3
LOSP1	DR000	0–3, parity 0
	DR001	4–7, parity 1
	DR002	8–11, parity 2
	DR003	12–15, parity 3
LOSP2	DR020	0–3, parity 0
	DR021	4–7, parity 1
	DR022	8–11, parity 2
	DR023	12–15, parity 3
LOSP3	DR004	0–3, parity 0
	DR005	4–7, parity 1
	DR006	8–11, parity 2
	DR007	12–15, parity 3
LOSP4	DR024	0–3, parity 0
	DR025	4–7, parity 1
	DR026	8–11, parity 2
	DR027	12–15, parity 3
LOSP5	DR008	0–3, parity 0
	DR009	4–7, parity 1
	DR010	8–11, parity 2
	DR011	12–15, parity 3
LOSP6	DR028	0–3, parity 0
	DR029	4–7, parity 1
	DR030	8–11, parity 2
	DR031	12–15, parity 3
LOSP7	DR012	0–3, parity 0
	DR013	4–7, parity 1
	DR014	8–11, parity 2
	DR015	12–15, parity 3

DM Option

The DM option is used for maintenance channel distribution and for error logger channel and support channel routing to available CPUs. The DM option also contains the sanity code generator. Each I/O module has one DM option.

Options on the CPU Module

On the CPU module, the HA, HB, and HC option types perform I/O functions.

HA Option

There are four HA options (HA0 through HA3) on the CPU module. The HA options contain the input and output data buffers for read and write references. HA0 and HA1 serve as the write data buffers, and HA2 and HA3 contain the read data buffers. The HA options also perform error correction. During a write operation, the HA options receive a 72-bit word from the HB and HC options, and they perform SECDED. The 8 check bits are then stripped off, and a new 12-bit checkbyte is generated for SBCDBD. The HA options then pass the data to the CH options, which steer the data to the CI options (the section drivers).

During a read operation, options HA2 and HA3 receive 64-bit data words from the CJ options and generate the checkbytes. The HA options send 72-bit words to the HB and HC options in two transfers.

HB and HC Options

There is one HB option, and one HC option on the CPU module. The HB option controls the interface to the I/O module module. During a memory write operation, the HB and HC options work in unison, each receiving 36 bits from the I/O module.

During a read operation, the HB and HC options each receive data, in two 36-bit transfers, from two HA options. The HB option receives the odd bits; the HC option receives the even bits. The HB and HC options send the data to the I/O module.

Figure 6 and Figure 7 illustrate CPU options (read and write).

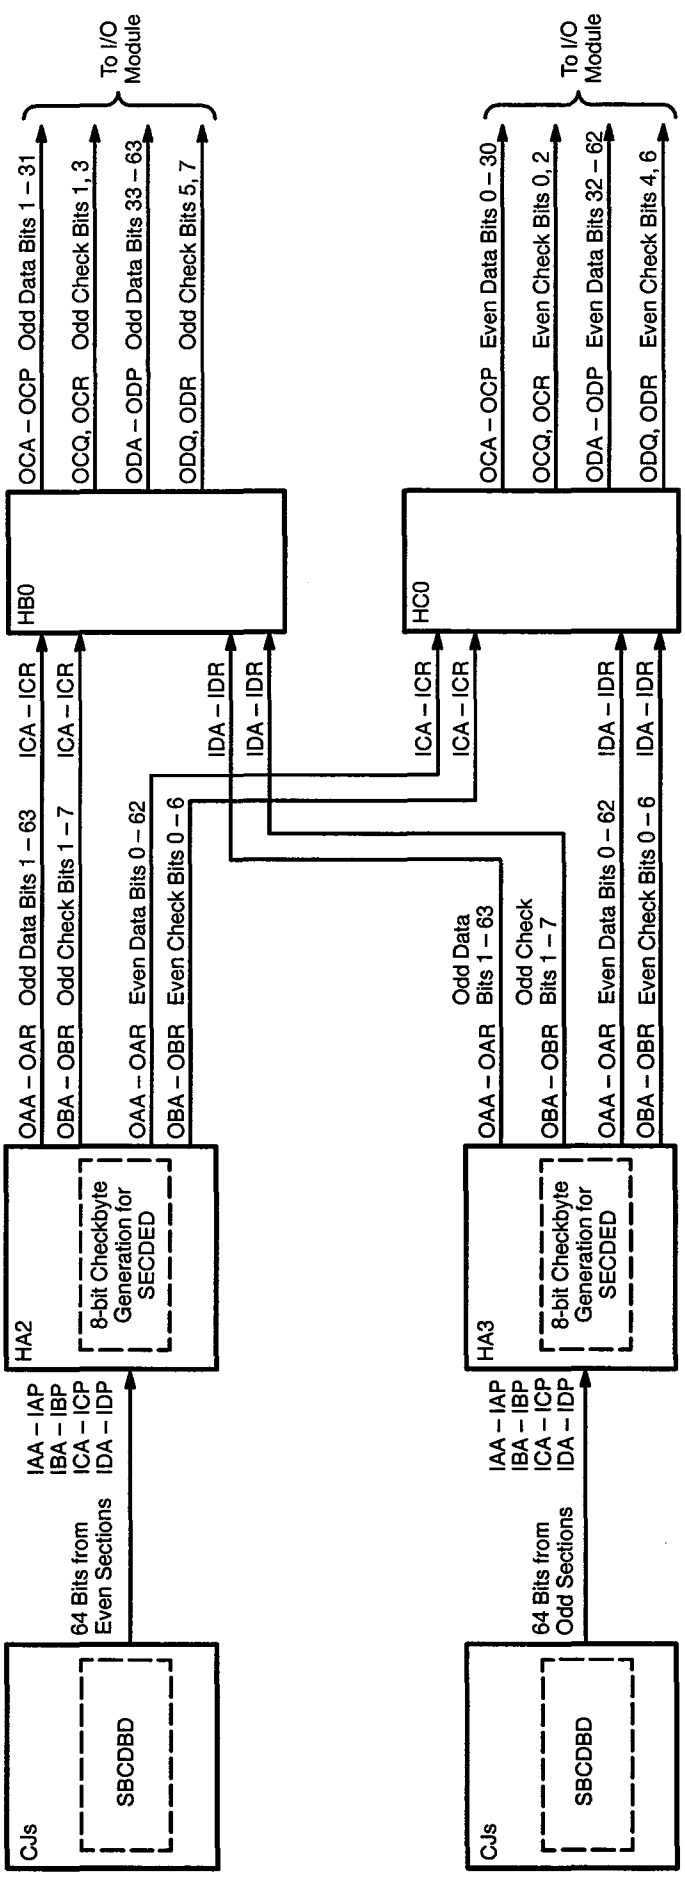

NOTE: The terms OAA, OAC, OAE, etc. go to the HB0 option.
The terms OAB, OAD, OAF, etc. go to the HC0 option.
The terms OBA, OBC, OBE, etc. go to the HB0 option.
The terms OBB, OBD, OBF, etc. go to the HC0 option.

Figure 6. CPU Options – Read (Memory to I/O)

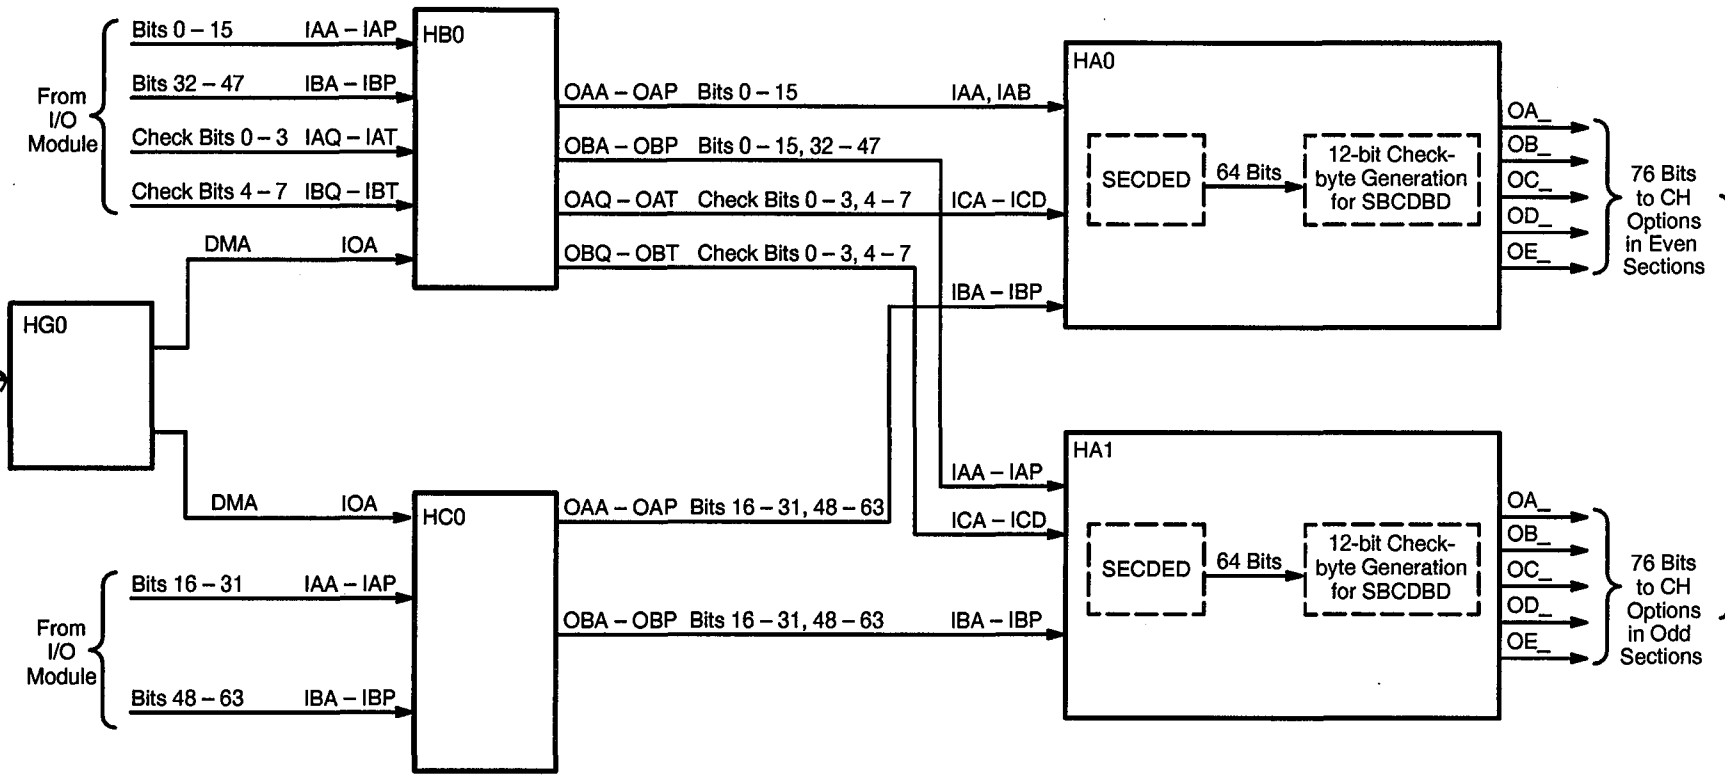

Figure 7. CPU Options – Write (I/O to Memory)

Connectors on the I/O Module

There are three basic types of connections on the I/O module. The first type is the I/O module-to-CPU module connection, second is the I/O module-to-bulkhead connection, and third is the maintenance connection. All three types of connections and their respective connectors are described in the following paragraphs.

I/O Module to CPU Module

The I/O module connects to the CPU modules through four OIM connectors. There are four connection pads on the I/O module, which are labeled YA, YB, YC, and YD. Each connection pad connects to the YF pad of a CPU module. All CPU-I/O commands and I/O data pass through these connections.

I/O Module to Bulkhead

The external channels physically connect to the I/O module through an I/O connector harness. The I/O connector harness makes the connection between the I/O module and the bulkhead. The connectors on the I/O module side of the harness are sometimes referred to as *fuzz-button* connectors.

The I/O module contains 16 I/O connectors: eight on board 1 and eight on board 2. There are five types of I/O connectors; each type has a unique part number. Types of connectors differ only in how they branch into the connections at the bulkhead. For example, one I/O connector may branch into ten 51-pin connectors, and another may branch into four 100-pin connectors. Refer again to Figure 3 and Figure 4 for more detailed information about the placement and characteristics of these connectors.

Each I/O module is wired to one bulkhead. In a CRAY T932 system that has four I/O modules, each module is wired to one of four bulkheads. External cabling starts from the bulkhead.

Maintenance Connector

The maintenance connector is located on the power regulator of the module. The maintenance connector is used to send/receive the following signals to/from the module:

- Boundary scan signals – TCLK, TM, TDI, TDO, and RCLCK
- OIM and SIM connector interlock signals
- Continuity line signals
- Thermal diode to output module temperature status

Channel Functions

This section covers the functions of all the I/O channels in CRAY T90 series systems. It includes information on each type of LOSPX channel, as well as a description of the LOSP, HISP, and VHISP channels.

LOSPX Channel

The term *LOSPX* refers to an error logger channel, a maintenance channel, a support channel, or a PINT/MCUI channel. Each I/O module supports up to four LOSPX channels. Quadrant position on the I/O module indicates the type of LOSPX channel. The support channel can operate in either LOSP or MISP mode. The following paragraphs describe each of the four types of LOSPX channels.

Error Logger Channel

The error logger channel in CRAY T90 series systems uses a standard 6-Mbyte/s LOSP channel. The LOSP portion of the error logger channel runs from an FEI-3 board in the system support VME chassis to one of the I/O bulkheads. From an I/O bulkhead, the LOSP error logger channel runs through a boundary scan module (normal operation) to an I/O module. From the I/O module, the serial error logger channel connects to the rest of the modules in the system. For more information, refer to the document entitled *Error Logger Channel.*

Maintenance Channel

The maintenance channel provides a channel for performing the following CRAY T90 series system functions:

- Configuring the system
- Master clearing CPUs
- Master clearing memory
- Initializing the system
- Enabling and disabling sections of the system
- Monitoring system activity
- Selecting test points
- Initializing I/O resources
- Sending and controlling sanity codes

Each half of a CRAY T932 system can have a separate and independent maintenance channel, if the system is logically and physically divided into two separate systems. Not all the maintenance ports on the I/O modules in a CRAY T932 system are used. If one of the maintenance ports fails, the maintenance channel can be recabled to a different I/O module; however, this recabling changes the system configuration.

The maintenance channel originates as a LOSP channel that is routed from the support system VME chassis through the boundary scan module and then to a maintenance port interface on one of the I/O modules.

Refer to the document entitled *Maintenance Channel* for more information.

PINT/MCUI

The PINT/MCUI (programmable interrupt and maintenance control unit interrupt) channel provides the selective real-time interrupt to any logical CPU or group of CPUs. The actual CPU or group of CPUs is encoded in the 16 parallel lines that make up the PINT channel, as shown in Figure 8.

The MCU interrupt has the ability to interrupt all CPUs at once. The CRAY T90 series system can target the interrupt to any combination of the four possible partitions. All CPUs are capable of receiving MCUI and real-time interrupts.

The PINT/MCUI channel enters the I/O module via the bulkhead. It originates from the PINT quarter board within the IOS. The 16 bits enter the DD3 option. From that option, the data is sent serially to the DM

option. The DM option directs the data to one of four DD/DE options. The DE option reassembles the data back into 16-bit packets (interrupt request packet) and sends it to the shared module via the attached CPU.

Figure 8 shows the bit assignments in the 16-bit packet that comes across the PINT/MCUI channel from the IOS-E PINT quarter board. Bits 0 - 7 identify the CPU number, which can be either a physical or logical CPU number, as selected by bit 15. Bits 8 through 11 select the I/O group or groups that contain the CPUs to be interrupted. Bit 12 selects either a PINT or MCU interrupt.

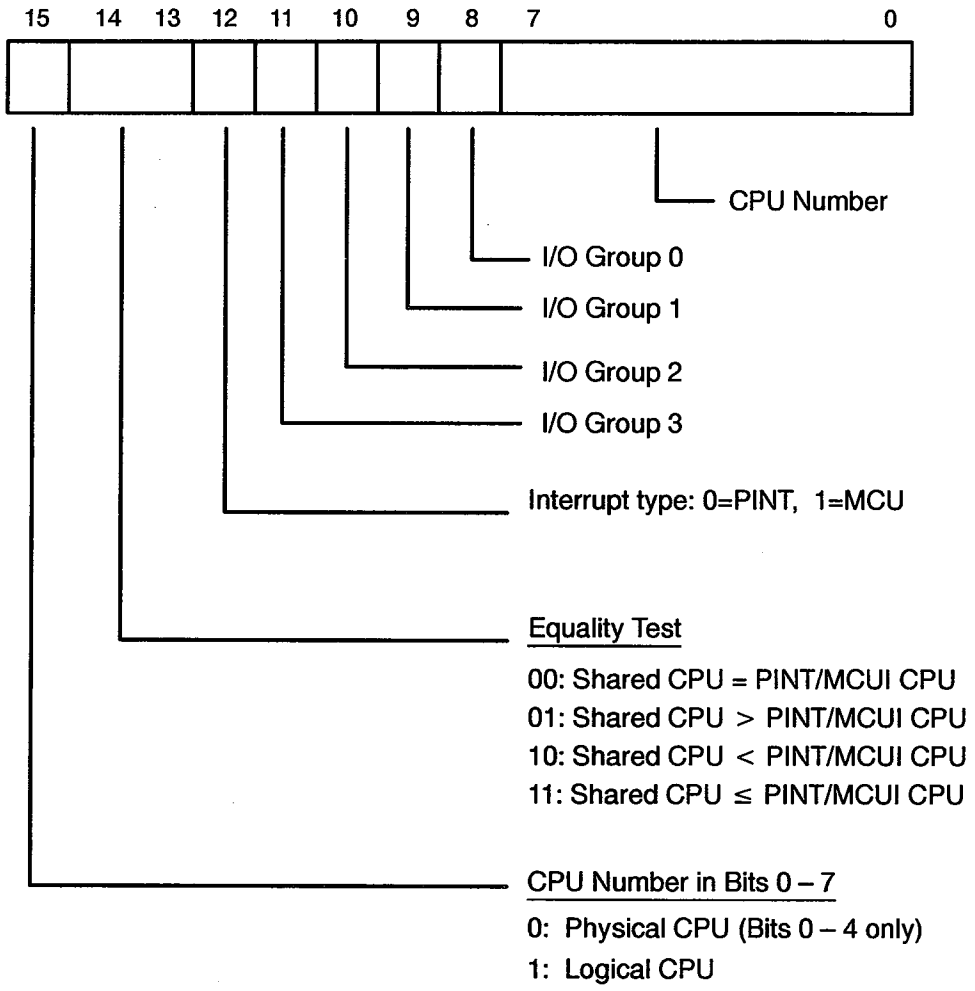

Figure 8. PINT/MCUI Channel Format

The equality test bits determine exactly which CPUs are interrupted. For example, bits 0 through 7 and bit 15 select physical CPU 5. If the equality test bits are 00, then physical CPU 5 gets the interrupt. If the equality test bits are 01, all CPUs numbered 6 and greater get the interrupt. If the

equality test bits are 10, all CPUs numbered 4 and less get the interrupt. If the equality test bits are 11, all CPUs numbered 5 and less get the interrupt. This option provides the capability to interrupt all CPUs, which is done by setting the CPU number to the greatest number of CPUs in the system and setting the equality test bits to 11.

Support Channel

Each I/O module contains eight LOSP channels, which are connected to IOS clusters. The support channel is the ninth LOSP that originates from the OWS. It provides a TCP/IP connection for system calls and can operate in either LOSP or MISP mode.

The support channel enters the I/O module on the DD1 option. From there, the data is sent serially to the DM option. The DM option sends the serial data to one of four DA/DC options; each DA/DC option corresponds to one physically connected CPU.

LOSP Channel

The function of the LOSP channel is to connect the IOS to the system's CPU. The LOSP channel path is 20 bits wide (16 data bits and 4 parity bits). Odd parity is performed on the channel. All LOSP transfers are word transfers, with four 16-bit parcels making one word. During a LOSP input sequence, a parity error causes all transfers to memory to be equal to zero following the word with the parity error. A LOSP-in parity error sets the Error flag, but does not cause an interrupt. A LOSP channel interrupt is caused by a Disconnect signal from the connected device, (CA = CL) or a fatal error.

The LOSP and support channels on CRAY T90 series systems can operate in either LOSP mode or MISP mode. MISP mode is selected with a software command and is described in greater detail later in this section.

The control and data paths and options associated with the LOSP channel input and output sequences are illustrated in Figure 9 and Figure 10.

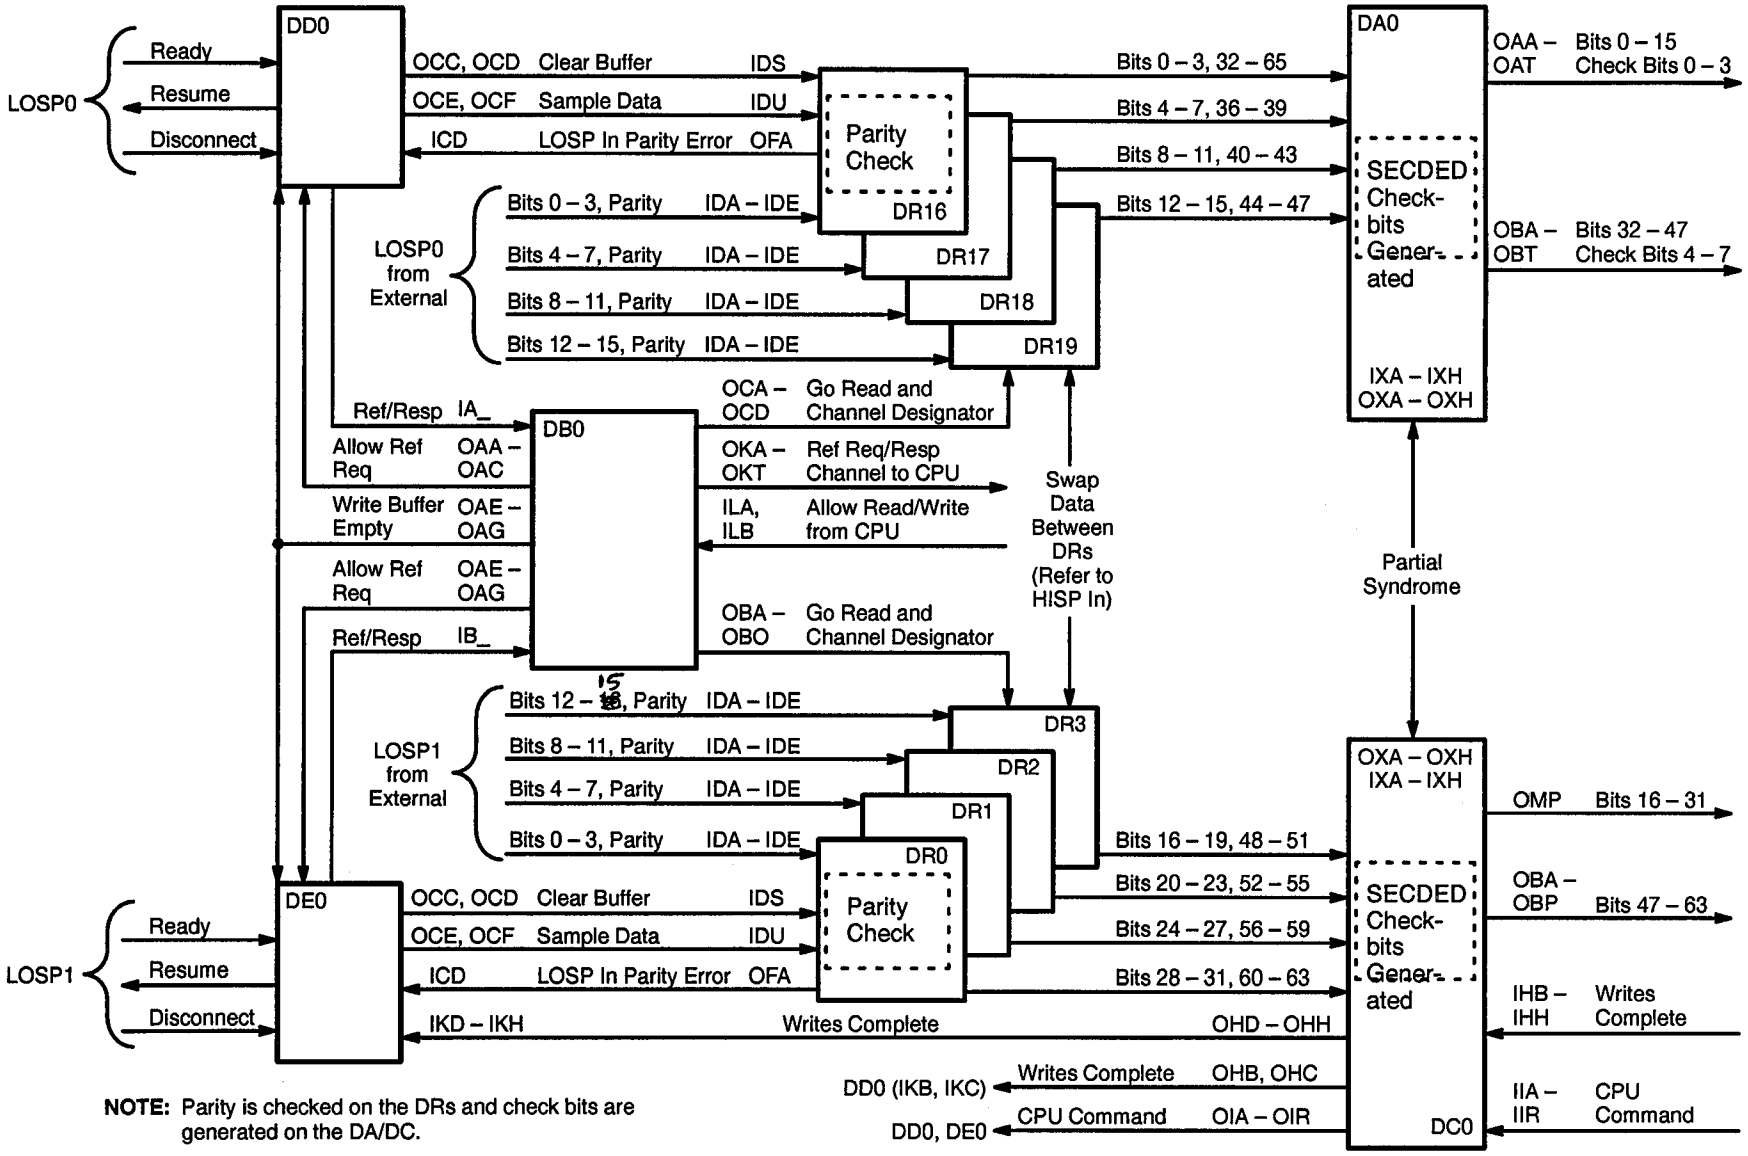

Figure 9. LOSP0 and 1 Input Control and Data

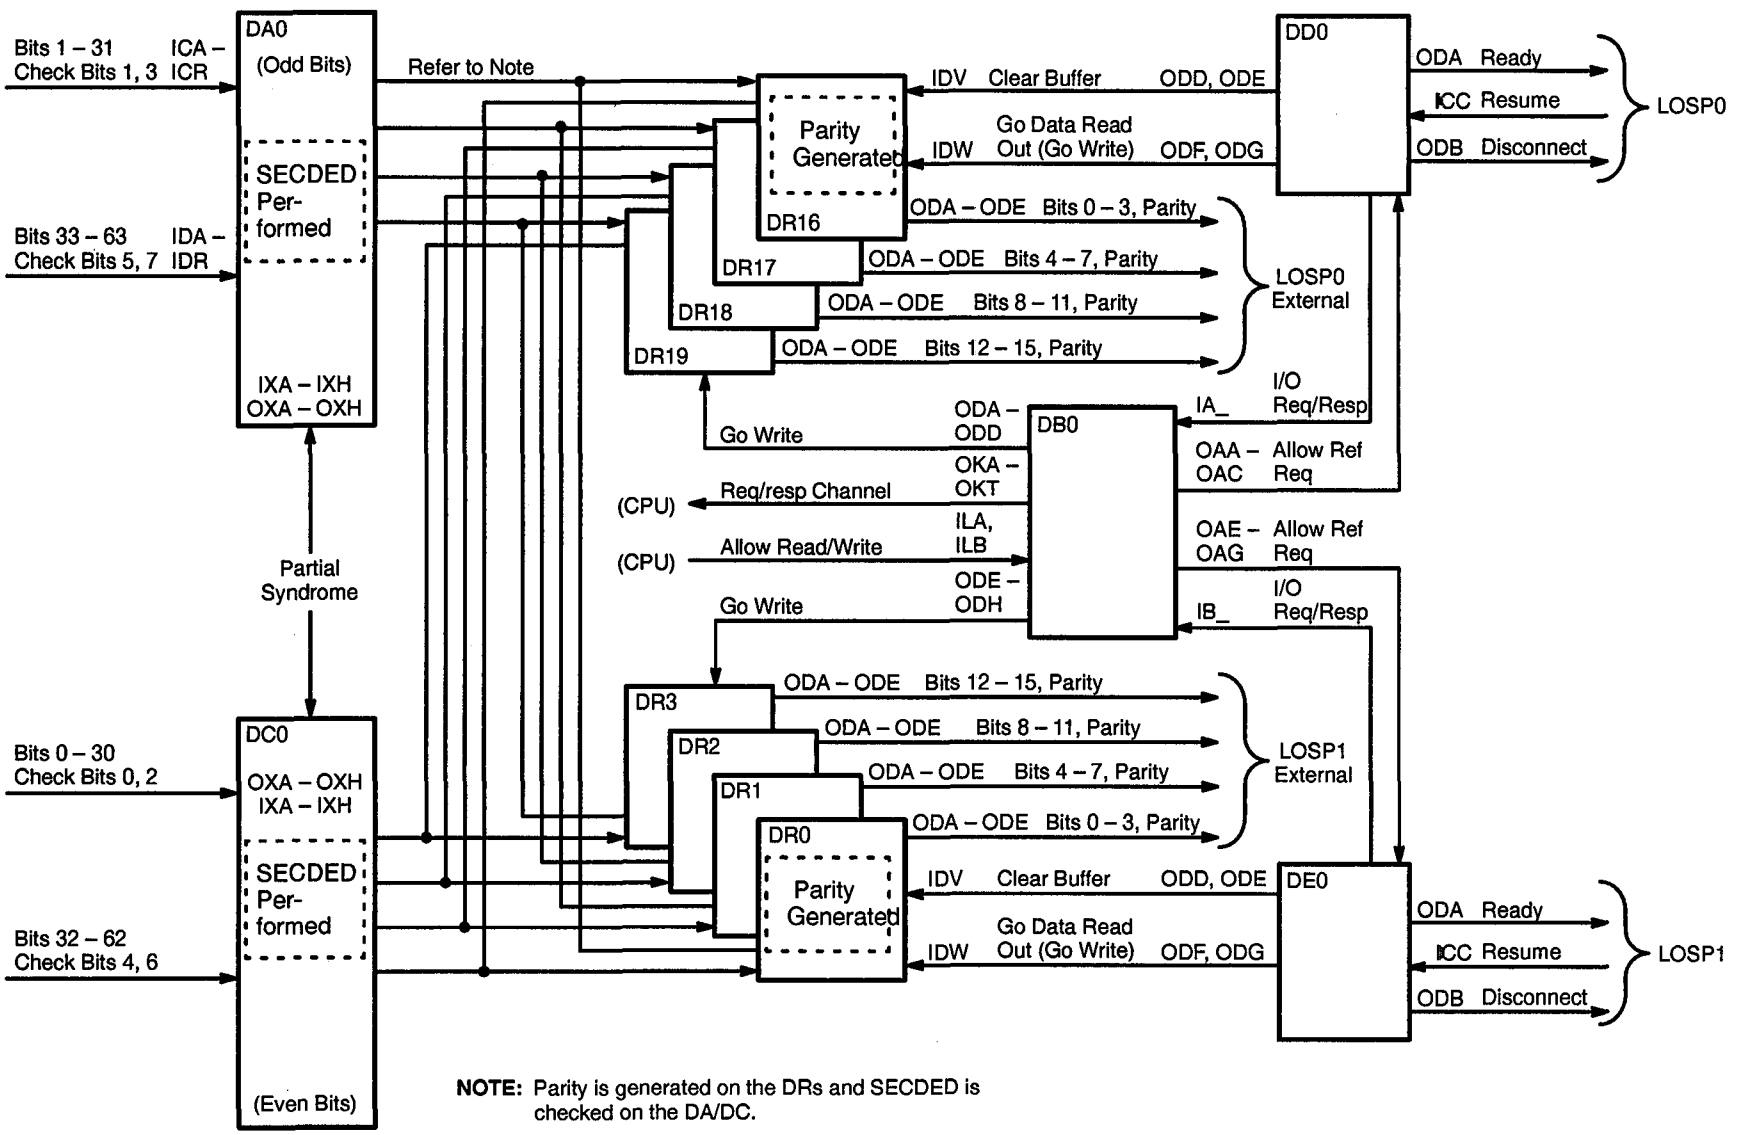

NOTE: Parity is generated on the DRs and SECDED is checked on the DA/DC.

Figure 10. LOSP0 and 1 Output Control and Data

LOSP Protocol

The LOSP channel operates at 6 Mbytes/s and is similar to the LOSP channels in the CRAY Y-MP and CRAY C90 computer systems. The data transfer is controlled by two signals: Ready and Resume. The Ready signal is sent with the data; the Resume signal is sent by the receiving device. This protocol is used for both the input and output channels.

Ready Signal

The Ready signal is sent with each parcel of data. This signal indicates that the data on the data line is valid. The Ready signal is a 50-ns pulse.

Data Lines 2^0 through 2^{15}

The data lines transmit 16-bit parcels of data between the transmitting and receiving devices. The channel samples data after the leading edge of the Ready signal is received. Data integrity is verified by 4 parity bits transmitted with the data.

Parity Bits

Parity bit signals 0 through 3 accompany each 16-bit parcel of data. Each of the four parity bit signals is assigned to one 4-bit group of data bits. The parity bits are set or cleared to give odd parity to the resulting 5-bit group (4 data, 1 parity).

Resume Signal

The Resume signal is sent from the receiving device to the sending device to indicate that data has been received and that the receiving device is ready for another 16 bits of data. The Resume signal is a 50-ns pulse.

Disconnect Signal

The Disconnect signal is sent from the sending device to indicate the end of the transmission. The Disconnect signal deactivates the channel. When a channel is disconnected in the mainframe, an interrupt is generated (if the channel is enabled for interrupts).

LOSP Channel Instructions

Any CPU can request any LOSP channel to perform a read or write operation. The following list shows the sequence of the instructions issued.

- 0012*j*0 - (A*j*) is the channel number. The instruction issued clears any interrupts: clear error, clear done, and enable.
- 0011*jk* - (A*j*) is the channel number, and (A*k*) is the channel limit address (CL).
- 0010*jk* - (A*j*) is the channel number, and (A*k*) is the channel current address (CA). This instruction activates the channel.

Table 6 lists and explains all of the LOSP channel instructions for both read and write operations.

Table 6. LOSP Channel Instructions

Machine Instruction	CAL Syntax	Input LOSP	Output LOSP
0010*jk*	CA,A*j* A*k*	Set 35-bit CA Start channel	Set 35-bit CA Start channel
0011*jk*	CL, A*j* A*k*	Set 32-bit CL	Set 32-bit CL
0012*j*0	CI,A*j*	Clear interrupt Clear error Clear done Enable interrupt	Clear interrupt Clear error Clear done Clear master clear Enable interrupt
0012*j*1	MC,A*j*	Clear interrupt Clear error Clear done Clear ready held Enable interrupt	Clear interrupt Clear error Clear Done Set master clear Enable interrupt
0012*j*2	DI,A*j*	Disable interrupt	Disable interrupt
0012*j*3	EI,A*j*	Enable interrupt	Enable interrupt
033*i*00	A*i* CI	Transmit interrupting channel number to Ai	Transmit interrupting channel number to Ai
033*ij*0	A*i* CA,A*j*	Read CA	Read CA
033*ij*1	A*i* CE,A*j*	Read status & errors Bits 0 –28: 0 Bit 29: parity error Bit 30: error Bit 31: not done	Read status & errors Bits 0 – 29: 0 Bit 30: error Bit 31: not done

I/O Commands

The DC option receives, from the CPU, four 16-bit packets containing either the limit address register contents or the current address register contents along with the channel number, CPU number, and command bits 0 through 7. A Go I/O Command signal is sent with the first packet. These packets of data are routed through the DC option to both the DD option (for LOSP0) or DE option (for LOSP1). Table 7 describes bit information for each of the four packets.

Because the DD and DE options handle both input and output for a LOSP channel, there are two channel limit registers and two channel current address registers per option.

Command codes interpreted by the DD or DE option are as follows:

- Command 11: enter CA
- Command 12: enter CL
- Command 14, 15, 16, and 17: clear interrupt
- Command 20: read current address
- Command 21: read error status

Table 7. DC Option I/O Command Packets

Bit Position	1st Packet	2nd Packet	3rd Packet	4th Packet
Bit 00	A*k* data 00	A*k* data 16	XX	Command bit 00
Bit 01	A*k* data 01	A*k* data 17	XX	Command bit 01
Bit 02	A*k* data 02	A*k* data 18	XX	Command bit 02
Bit 03	A*k* data 03	A*k* data 19	XX	Command bit 03
Bit 04	A*k* data 04	A*k* data 20	XX	Command bit 04
Bit 05	A*k* data 05	A*k* data 21	XX	Command bit 05
Bit 06	A*k* data 06	A*k* data 22	XX	Command bit 06
Bit 07	A*k* data 07	A*k* data 23	XX	Command bit 07
Bit 08	A*k* data 08	A*k* data 24	Channel number 00	CPU number bit 00
Bit 09	A*k* data 09	A*k* data 25	Channel number 01	CPU number bit 01
Bit 10	A*k* data 10	A*k* data 26	Channel number 02	CPU number bit 02
Bit 11	A*k* data 11	A*k* data 27	Channel number 03	CPU number bit 03
Bit 12	A*k* data 12	A*k* data 28	Channel number 04	CPU number bit 04
Bit 13	A*k* data 13	A*k* data 29	Channel number 05	CPU number bit 05
Bit 14	A*k* data 14	A*k* data 30	Channel number 06	CPU number bit 06
Bit 15	A*k* data 15	A*k* data 31	Channel number 07	CPU number bit 07

50-Mbyte MISP Protocol

The LOSP channel and the support channel can run in mid-speed (MISP) mode. A MISP channel operates at 50 Mbytes/s. MISP protocol, described below, is similar to the 6-Mbyte/s LOSP channel protocol.

Ready Signal

The Ready signal is a pulse 18.75 ns long. The Ready signal indicates to the channel that data transmission will follow on the data lines. Each ready pulse corresponds to a parcel of data.

Data Lines 2^0 through 2^{15}

The data lines transmit 16-bit parcels of data between the transmitting and receiving devices. The channel typically samples data 37.5 ns after the leading edge of the appropriate Ready signal is received. Typically, data is sent 18.75 ns after the leading edge of the Ready signal pulse and remains valid for 37.5 ns.

Parity Bits

Parity bit signals 0 through 3 accompany each parcel of data. Each of the 4 parity bit signals is assigned to one 4-bit group of data bits. During a LOSP-out sequence, the parity is generated prior to sending the 16 bits of data to the external device.

Resume Signal

The Resume signal is a pulse 18.75 ns long. The receiver sends the Resume signal to the data transmitter, indicating that parcel 0 of the transmitted word has been received and that the receiver is ready for the next full-word (64-bit) transfer. Each Resume signal corresponds to 4 parcels of data and 4 Ready signals. The final Resume signal of a transfer is an exception, however; it is issued in response to a Disconnect signal and can be sent only after the final word transfer is successfully stored.

Disconnect Signal

Like the Ready and Resume signals, the Disconnect signal is a pulse that is 18.75 ns long. The data transmitter sends a Disconnect signal to the receiver, which indicates that the transmission is complete. A Disconnect signal can be sent only after the last Resume signal for the last word of data is received by the transmitting device.

Disconnecting Resume Signal

The Disconnecting Resume signal reflects an error or interrupt condition in the receiving device. This signal disables channel communications between a receiver and a transmitter. It acts as a Master Clear signal to both the transmitter and receiver. The Disconnecting Resume signal uses the same wire as the Resume signal and is differentiated by the width of the signal. The Disconnecting Resume signal has a 56.25-ns pulse instead of the 18.75-ns pulse of a normal Resume signal.

HISP Channel

Each I/O module contains eight HISP channels. These channels are distributed across the four logical quadrants on an I/O module. The HISP channels are referred to as HISP0 and HISP1 in quadrant 0, HISP2 and HISP3 in quadrant 1, HISP4 and HISP5 in quadrant 2, and HISP6 and HISP7 in quadrant 3.

A HISP channel operates at 200 Mbytes/s. A block of data in a HISP transfer consists of sixteen 72-bit words. Memory address and block length are transferred in five 12-bit packets from the master device to the mainframe. The block length register contains 16 bits; each bit represents a transfer of sixteen 72-bit words. The high-speed address register contains 32 bits. Table 8 lists the options that receive address and block length bits from the master device.

The mainframe operates as a slave device for a HISP transfer. There are no instructions that the mainframe issues that activate a HISP channel transfer; all channel transfers start from a master device such as the IOS-E.

Table 8. HISP Address and Block Length Options

Channel	Address Bits	Block Length Bits
HISP0 input	DR000	DR001
HISP0 output	DR018	DR019
HISP1 input	DR002	DR003
HISP1 output	DR016	DR017
HISP2 input	DR004	DR005
HISP2 output	DR022	DR023
HISP3 input	DR006	DR007
HISP3 output	DR020	DR021
HISP4 input	DR008	DR009
HISP4 output	DR026	DR027
HISP5 input	DR010	DR011
HISP5 output	DR024	DR025
HISP6 input	DR012	DR013
HISP6 output	DR030	DR031
HISP7 input	DR014	DR015
HISP7 output	DR028	DR029

HISP Input Sequence

The IOS-E initiates a HISP transfer. The following exchange of control and data signals occurs between the IOS-E and the mainframe during a write operation. Figure 11 illustrates the exchange of control during the sequence.

1. Clear Channel: The Clear Channel signal clears the mainframe and enables it to receive data.

2. Transmit Address: The mainframe activates the Transmit Address signal when it can accept the address and block length. The mainframe deactivates the signal after it has received five Address Ready signals.

3. Address Ready: The Address Ready signal informs the mainframe that the IOS-E is sending the address and block length.

4. Address and Block Length: The IOS-E sends five packets containing the address and block length. Each transfer accompanies an Address Ready signal. The five transfers consist of the following bits:

   - First transfer: address bits 0 through 3 and 20 through 23; block length bits 0 through 3
   - Second transfer: address bits 4 through 7 and 24 through 27; block length bits 4 through 7
   - Third transfer: address bits 8 through 11 and 28 through 31; block length bits 8 through 11
   - Fourth transfer: address bits 12 through 15 and 32 through 35; block length bits 12 through 15
   - Fifth transfer: address bits 16 through 19 and 36 through 39; block length bits 16 through 19

5. Transmit Data: The mainframe sends a Transmit Data signal when it can accept 16 words of data from the IOS-E.

6. Data Ready: A Data Ready signal is sent to the mainframe with each 72-bit word of data from the IOS-E.

7. Steps 5 and 6 repeat until the last word of data is transferred.

8. Last Word flag: The Last Word Flag signal is sent from the IOS-E when the last word of data is sent. The Last Word Flag accompanies the Data Ready signal on the last word.

9. Done flag: The Done flag is set within the IOS-E to indicate that block length is zero and no errors have occurred.

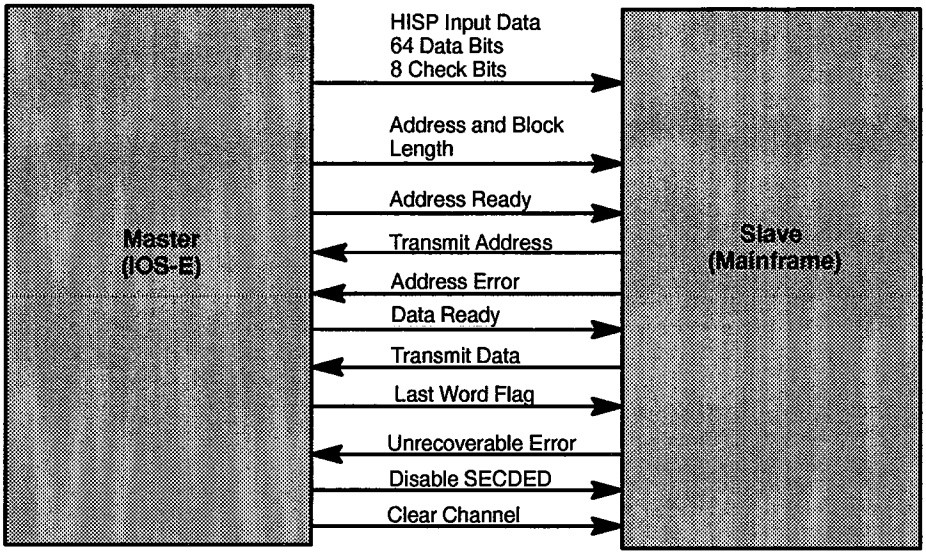

Figure 11. HISP Input Sequence

Figure 12 is a block diagram of the data flow through the options during an input sequence. Figure 13 illustrates the control signal, address, and block length flow through the options on the I/O module during an input sequence.

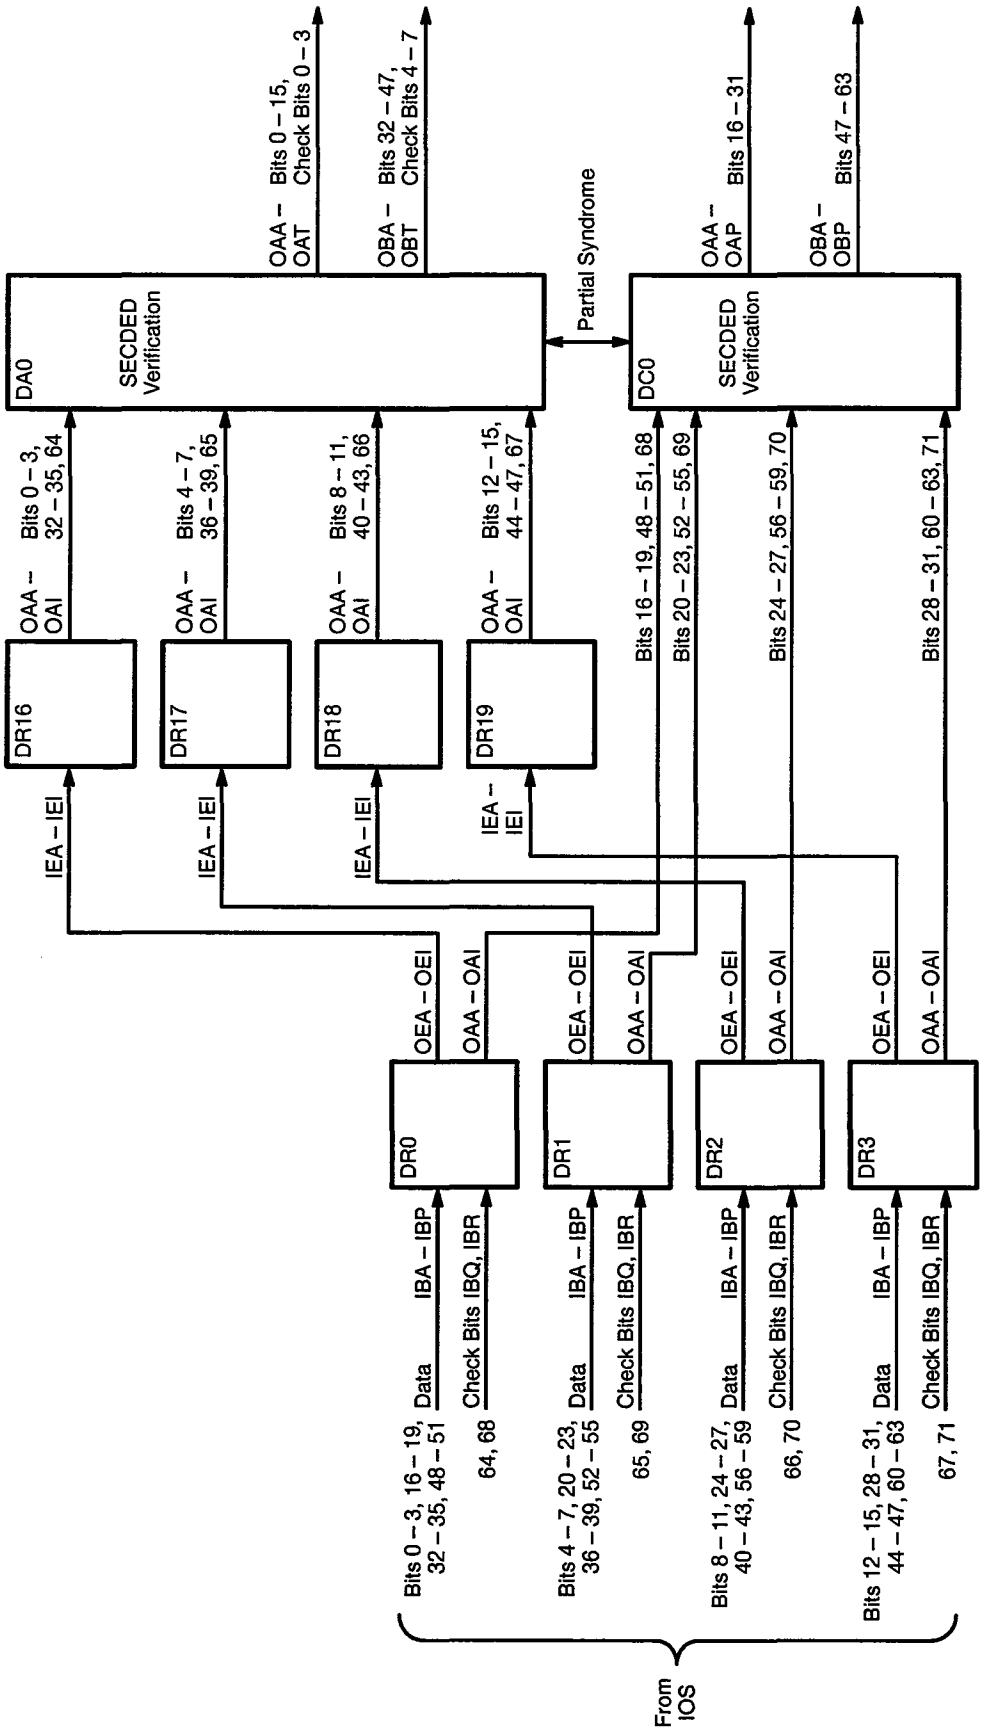

Figure 12. HISP0 Input Data

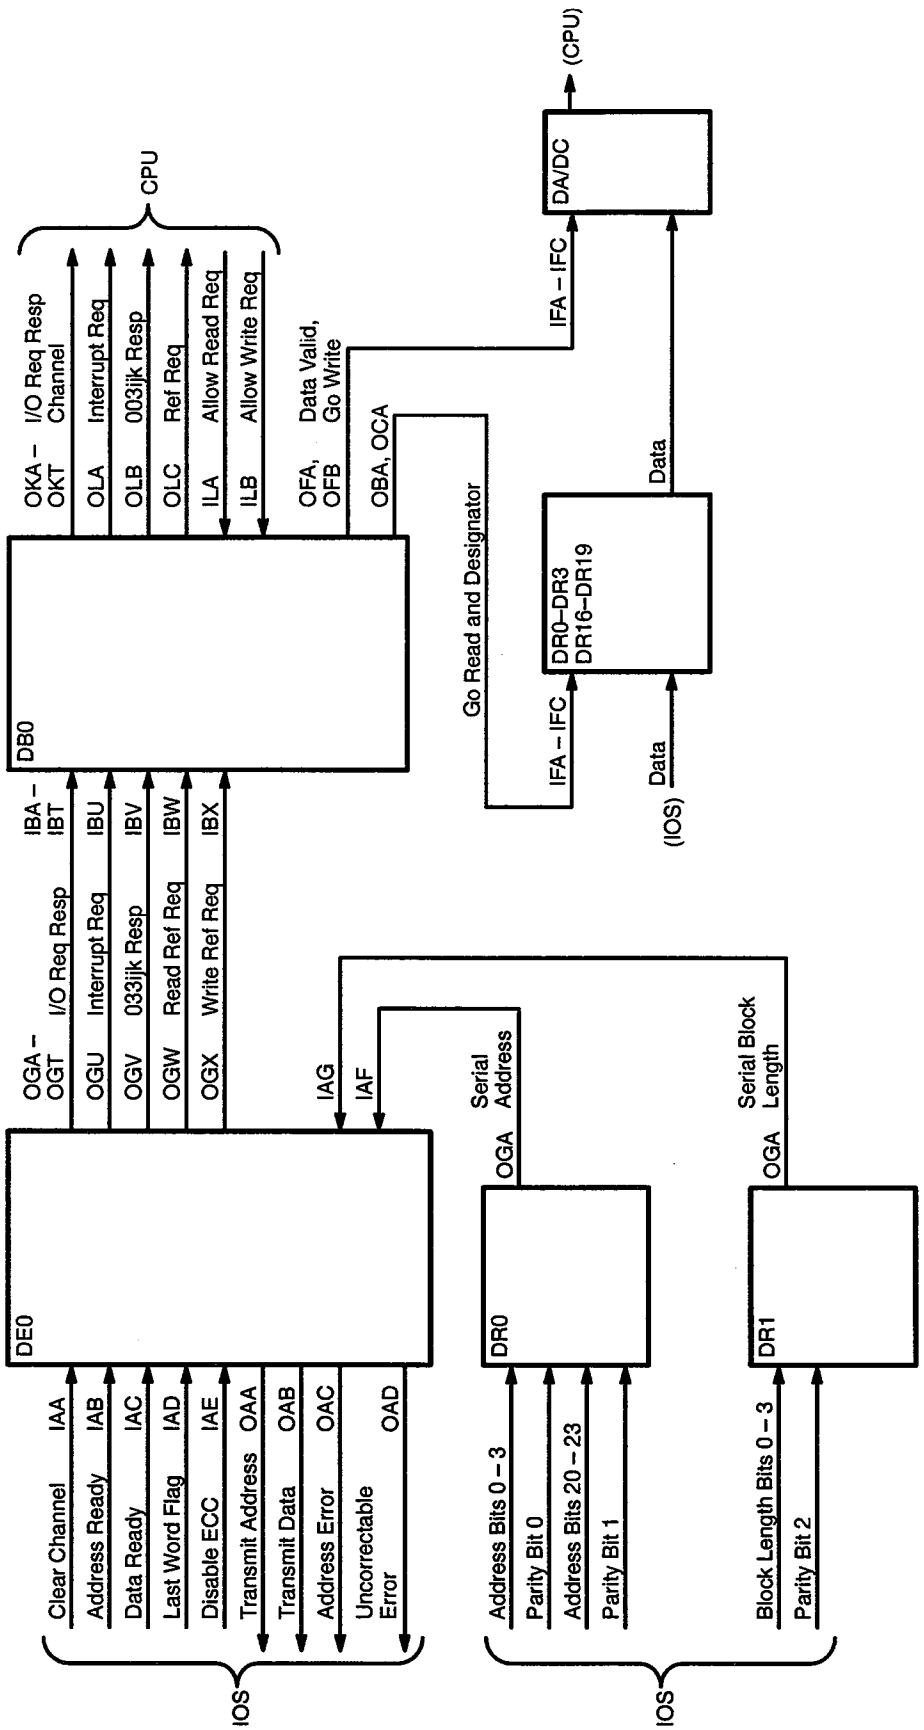

Figure 13. HISP0 Input Control and Address

HISP Output Sequence

The IOS-E initiates a HISP transfer. The following exchange of control and data signals occurs between the IOS-E and the mainframe during a read operation. Figure 14 illustrates the exchange of control that takes place during an output sequence.

1. Clear Channel: The Clear Channel signal clears the mainframe and enables it to receive data.

2. Transmit Address: The mainframe activates the Transmit Address signal when it can accept the address and block length. The mainframe deactivates the signal after it has received five Address Ready signals.

3. Address Ready: The Address Ready signal informs the mainframe that the IOS-E is sending the address and block length.

4. Address and Block Length: The IOS-E sends five packets containing the address and block length. Each transfer accompanies an Address Ready signal. The five transfers consist of the following bits:

   - First transfer: address bits 0 through 3 and 20 through 23; block length bits 0 through 3
   - Second transfer: address bits 4 through 7 and 24 through 27; block length bits 4 through 7
   - Third transfer: address bits 8 through 11 and 28 through 31; block length bits 8 through 11
   - Fourth transfer: address bits 12 through 15 and 32 through 35; block length bits 12 through 15
   - Fifth transfer: address bits 16 through 19 and 36 through 39; block length bits 16 through 19

5. Transmit Data: The IOS-E sends a Transmit Data signal when it can accept 16 words of data from the mainframe.

6. Data Ready: A Data Ready signal is sent to the IOS-E with each 72-bit word of data from the mainframe.

7. Steps 5 and 6 repeat until the last word of data is transferred.

8. Last Word Flag: The Last Word Flag signal is sent to the IOS-E when the last word of data is sent. The Last Word Flag accompanies the Data Ready signal on the last word.

9. Done flag: The Done flag is set within the IOS-E to indicate that block length is 0 and no errors have occurred.

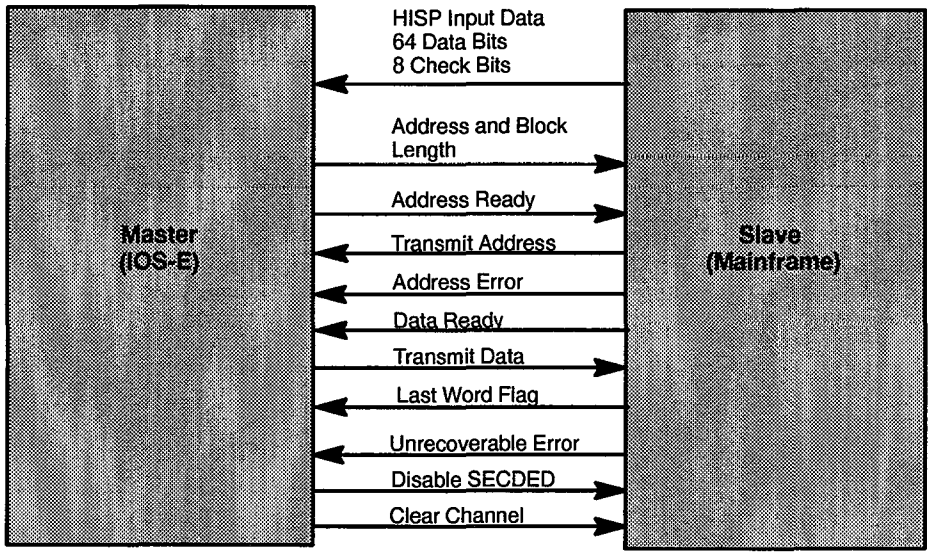

Figure 14. HISP Output Sequence

Figure 15 illustrates the control signal, address, and block length flow through the options on the I/O module during an output sequence. Figure 16 is a block diagram of the data flow through the options during an output sequence.

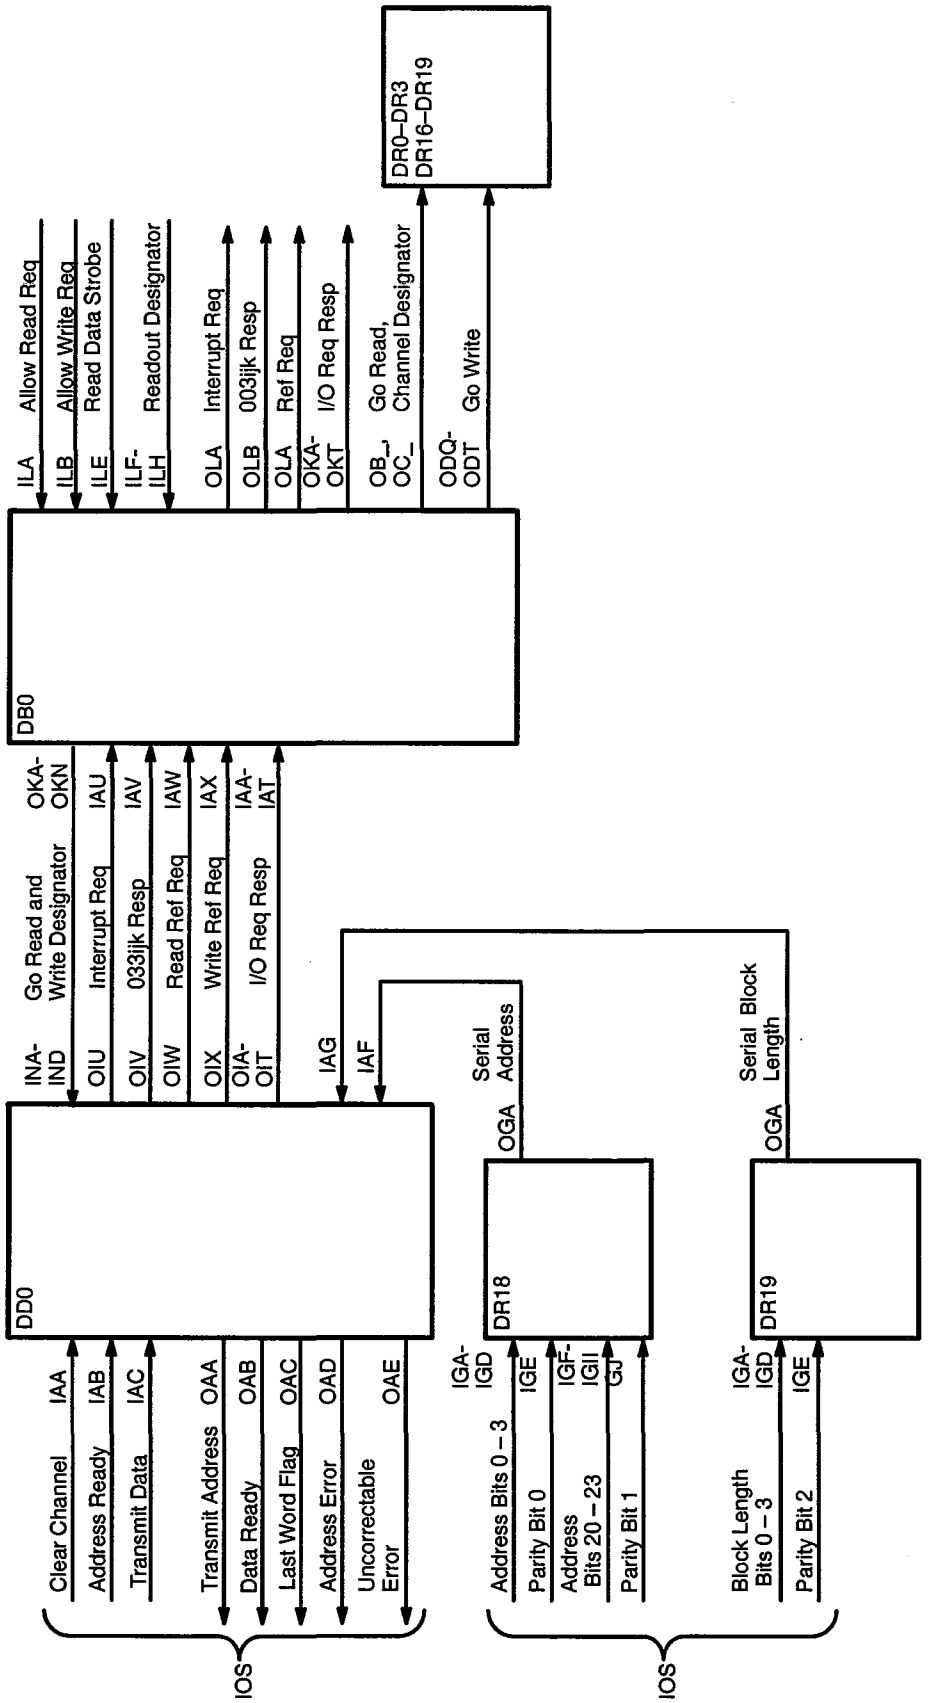

Figure 15. HISP0 Output Control and Address

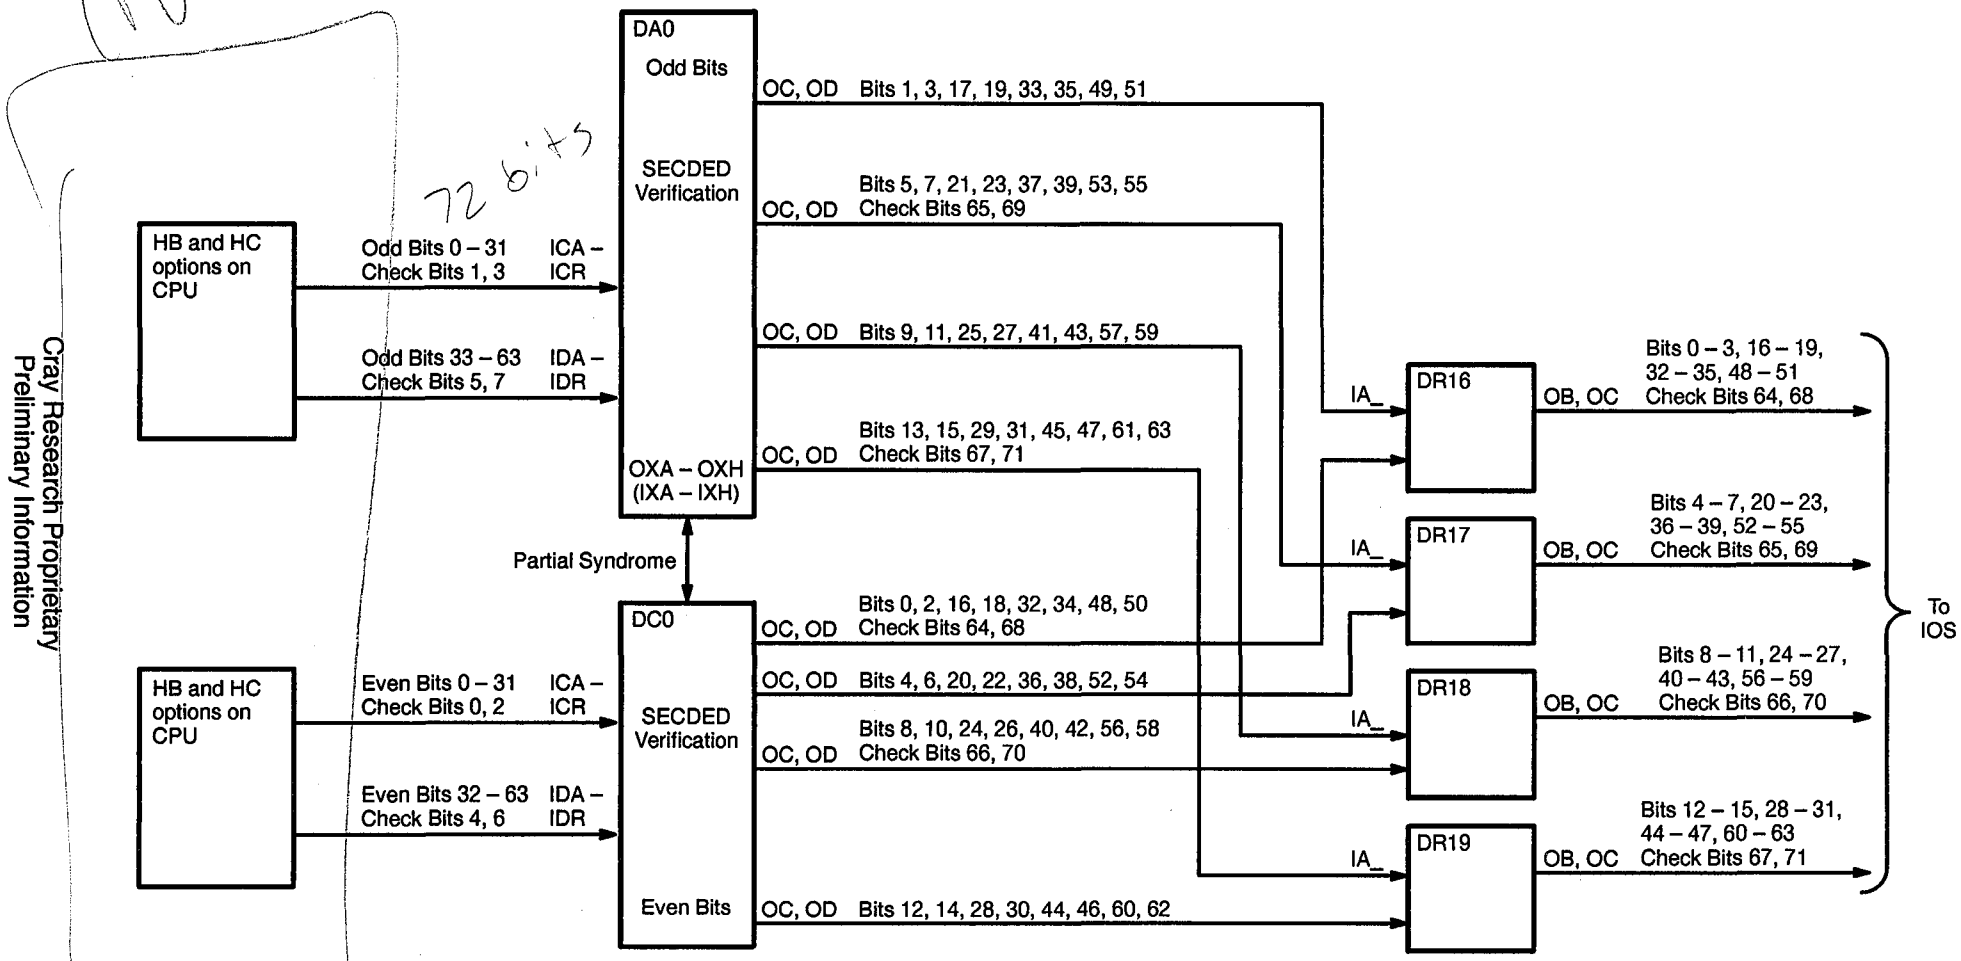

Figure 16. HISP0 Output Data

HISP Errors

There are certain error conditions that can occur during HISP channel activity. Some errors terminate the transfer, others do not. Table 9 lists and describes the HISP errors.

Table 9. HISP Errors

Error	Description
Address channel parity error	A parity error was detected on the DR option receiving memory address and/or block length. This error terminates the transfer.
Extra Address Ready signal	An Address Ready signal was detected after the fifth address/block length transfer. This error terminates the transfer.
Address Ready signal without Transmit Address signal	An Address Ready signal was detected but the mainframe had not sent the Transmit Address signal. This error terminates the transfer.
Missing Address Ready signal	A Data Ready signal was detected between the 1st and 5th address/block length transfer. This error terminates the transfer.
SECDED error detected on read data from CPU or write to CPU	For a read operation, an error occurred on the DA/DC options while data was en route from the CPU to the DR options. For a write operation, an error occurred on the DA/DC options while data was en route to the CPU.
SECDED error on CPU write buffers	Unrecoverable memory error signal.
SBCDBD error on CPU Read	Unrecoverable memory error signal.
Clear Channel signal and reference pending	If a Clear Channel signal arrives, and there is a pending write reference to the CPU, an error condition occurs to ensure the channel does not try a new transfer until the reference is completed. This error terminates the transfer.
Clear Channel signal and no HISP writes complete signal	If a Clear Channel signal arrives, and there is a pending write reference to the CPU, an error condition occurs to ensure that the channel does not try a new transfer until the reference is complete. This error terminates the transfer.
Last Word Flag signal and Block Length $\neq$ 00	Last Word Flag signal was detected but the block length register did not equal 0. This error terminates the transfer.
No Last Word Flag signal and Block Length = 00	The Block Length register equals 0 but no Last Word Flag signal is detected. This error terminates the transfer.
Last Word Flag signal and Block Length = 00 and another Data Ready	Last Word Flag signal received and the Block Length register equals 0 and another Data Ready signal is detected. This error terminates the transfer.

VHISP Channel

Each I/O module contains four VHISP channels. A fully configured CRAY T90 series system can include 16 VHISP channels. Each VHISP channel is identified by a specific channel number.

A VHISP transfer is under the direct control of the mainframe. The SSD-E is the slave device; the mainframe is the master device. Any CPU can request either a VHISP read from the SSD-E or write to the SSD-E using any VHISP channel.

VHISP Instructions

Table 10 lists the VHISP instructions.

Table 10. VHISP Instructions

Machine Instruction	Cal Syntax	Function
0010*jk*	CA,A*j* A*k*	Set CA (first sends 32-bit SSD address, second sends 32-bit CPU address)
0011*jk*	CL,A*j* A*k*	Set BL and start channel Bits 0 – 23: BL Bit 31: transfer direction Bit 24 and 25: reserved
0012*j*0	CI,A*j*	Clear interrupt Clear channel sequence Clear error status Clear done Enable interrupt
0012*j*2	DI,A*j*	Disable interrupt
0012*j*3	EI,A*j*	Enable interrupt
033*i*00	A*i* CI	Transmit interrupting channel number to Ai
033*ij*0	A*i* CA,A*j*	Read CA
033*ij*1	A*i* CE,A*j*	Read status and errors Bits 0 – 23: remaining BL Bits 24 – 41: 0 Bit 42: transfer in progress Bit 43: BL error Bit 44: double-bit error in SSD Bit 45: double-bit error in mainframe Bit 46: fatal error Bit 47: not done Bits 48 – 62: 0 Bit 63: not done

The 001000 instruction is a pass (no-operation) instruction. The 0010*jk* (*jk*≠0) and 0011*jk* instructions initiate a channel transfer. For a VHISP channel, the 0010*jk* must be issued twice, followed by instruction 0011*jk*. The first 0010*jk* instruction sets the starting block address in the SSD-E. The second 0010*jk* instruction sets the channel address (CA) in the mainframe. The 0011*jk* instruction sets the block length (BL) and starts the transfer.

VHISP Addressing

When the VHISP channel is online, the mainframe supplies the SSD-E address and block length for read or write operations. The address and block length are sent from the DR options. The DR options listed in the table refer only to quadrant 0 of the I/O module.

The mainframe sends 32 address bits that are loaded into the address register of the SSD-E. The mainframe also sends 32 block length bits to the SSD-E. The SSD-E defines a block of data to be 32 words, whereas the mainframe defines a block of data as 64 words.

VHISP Input Sequence

The following exchange of control and data takes place during a read to the mainframe. The mainframe (master) initiates the read operation.

1. Clear Channel: The mainframe sends the Clear Channel signal to the SSD-E to clear and initiate the VHISP channel.
2. Clear Acknowledge: Once the SSD-E has cleared the channel, it returns the Clear Acknowledge signal.
3. Address and Block Length: The mainframe sends address and block length bits to the SSD-E.
4. Transmit Enable: The mainframe sends this signal to the SSD-E to indicate that it has an empty buffer ready to be filled.
5. Read Data Ready: The SSD-E sends this signal once for every 16 144-bit words.
6. Data: Data leaves the SSD-E 144 bits at a time (72 bits for even word and 72 bits for odd word).
7. Steps 4, 5, and 6 continue until the last word of the last block is sent by the SSD-E.

Figure 17 illustrates the exchange of control that takes place during a VHISP input sequence.

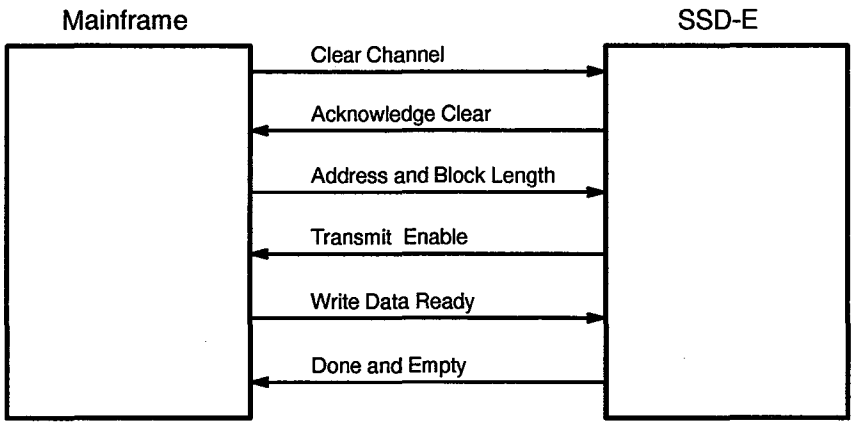

Figure 17. VHISP Input Control Signals

VHISP Input Data Path

The DR options receive input data from the SSD-E. Both an even and an odd word (72 bits each) are received. The DR options for quadrant 0 of the I/O module are DR16 through DR19. The even word is received on IBA through IBR and the odd word is received on ICA through ICR. Figure 18 shows the flow of data and control during a VHISP input sequence.

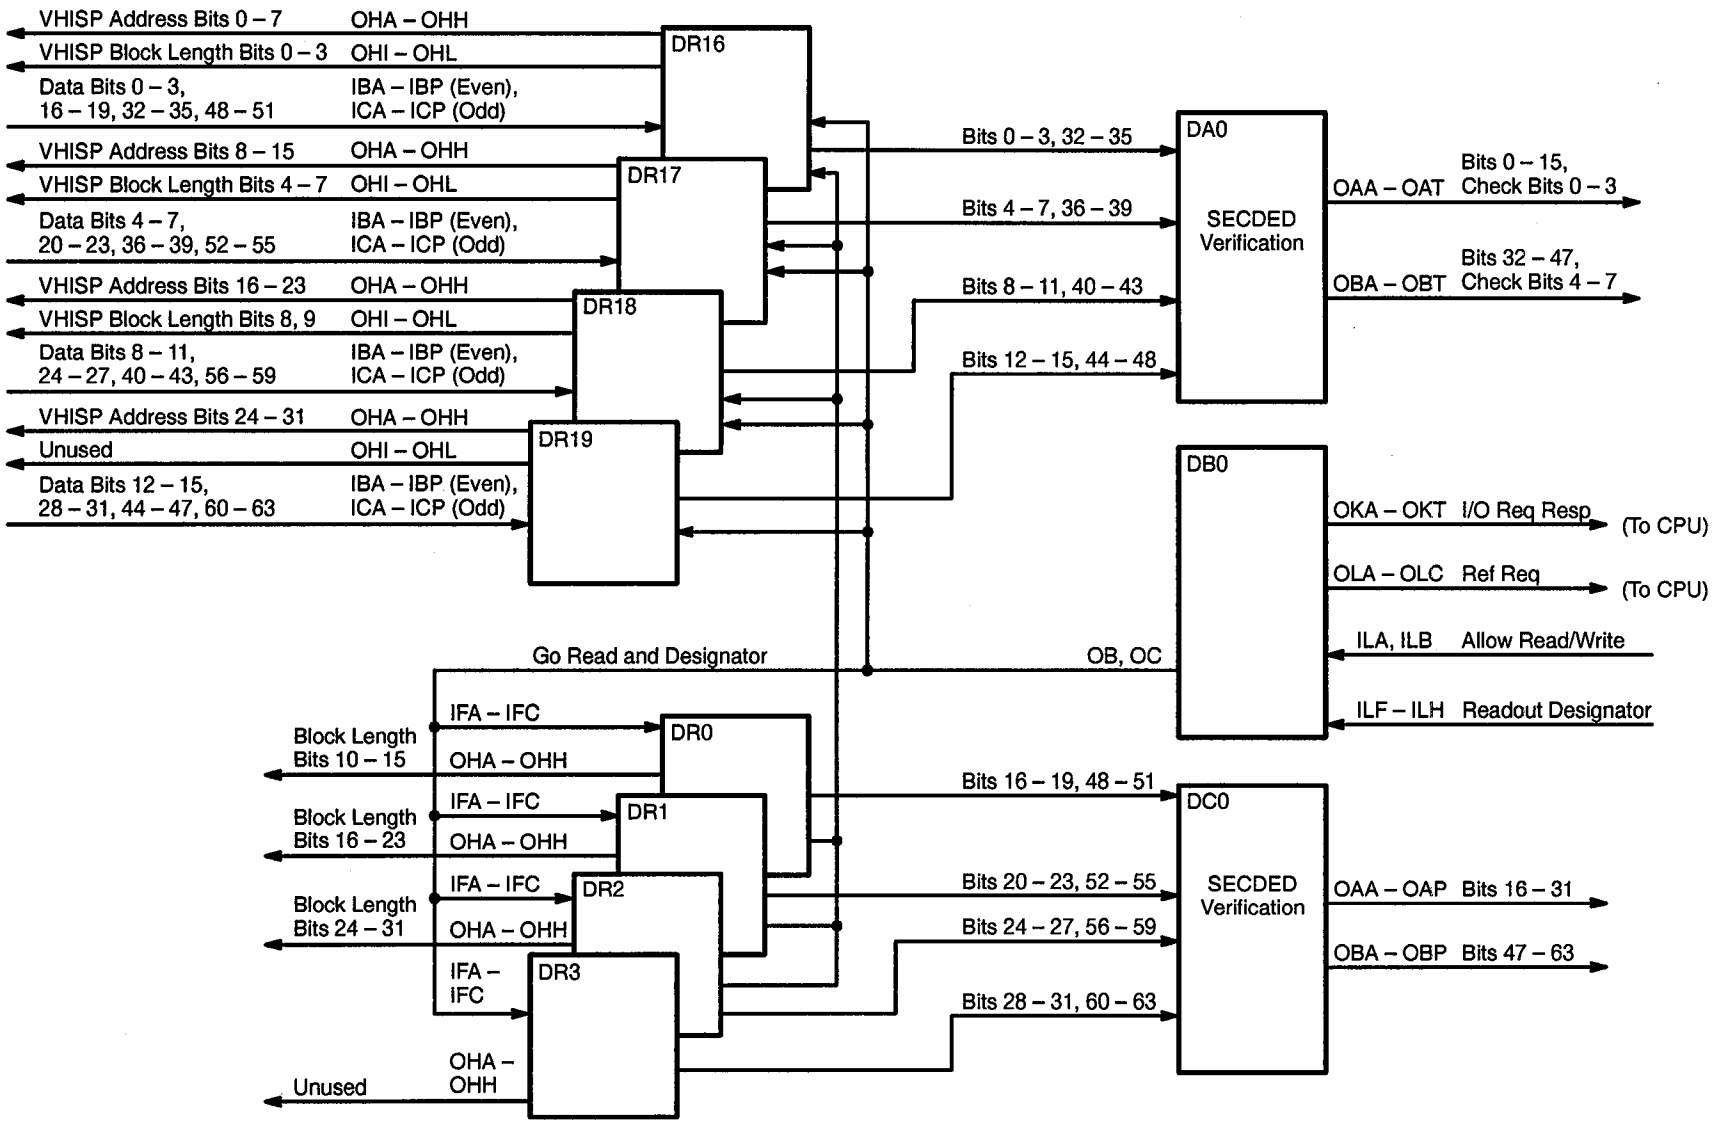

Figure 18. VHISP0 Input Control and Data

VHISP Output Sequence

The following exchange of control and data takes place during a read from mainframe memory. The mainframe (master) initiates the read operation.

1. Clear Channel: The mainframe sends the Clear Channel signal to the SSD-E to clear and initiate the VHISP channel.
2. Acknowledge Clear : Once the SSD-E has cleared the channel, it returns the Clear Acknowledge signal.
3. Address and Block Length: The mainframe sends address and block length bits to the SSD-E.
4. Transmit Enable: The mainframe sends this signal to the SSD-E to indicate that it has an empty buffer ready to be filled.
5. Read Data Ready: The SSD-E sends this signal once for every 16 144-bit words.
6. Data: Data leaves the SSD-E 144 bits at a time (72 bits for even word and 72 bits for odd word).
7. Steps 4, 5, and 6 continue until the last word of the last block is sent by the SSD-E.

Figure 19 illustrates the exchange of control that takes place during a VHISP read function.

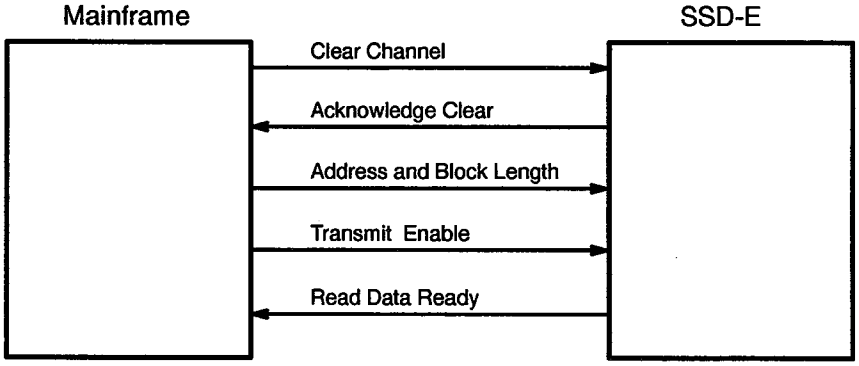

Figure 19. VHISP Output Control Signals

Figure 20 shows the flow of data and address/block length during a VHISP output sequence.

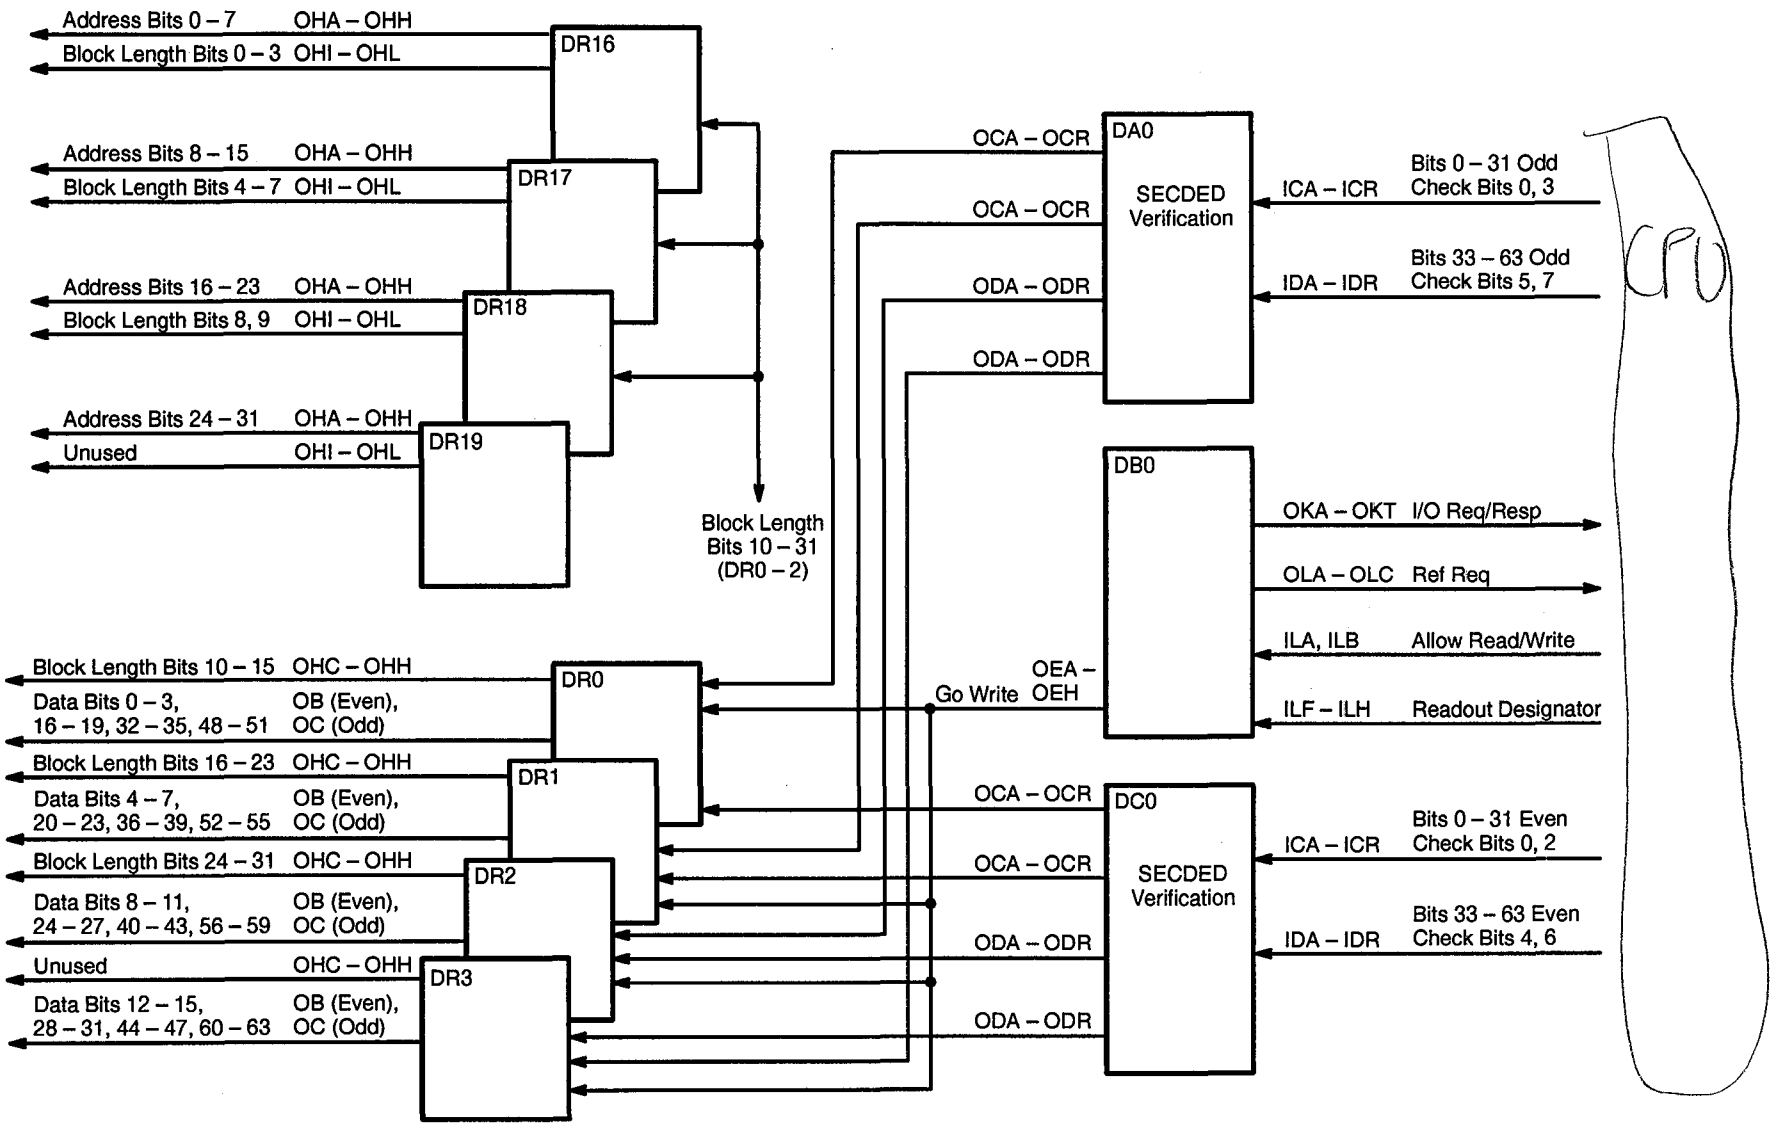

Figure 20. VHISP0 Output Control and Data

VHISP Errors

There are a number of conditions that can occur during a VHISP reference that cause VHISP channel errors. Some errors terminate the channel transfer; others do not. Table 11 describes the VHISP channel errors.

Table 11. VHISP Channel Errors

Error	Description
No Last Word and block length = 0	During transfer from mainframe to SSD, the SSD does not receive Last Word Flag signal, but the block length is 0. This error terminates the transfer.
Done and Empty and block length ≠0	During transfer from SSD to mainframe, the mainframe receives the Done and Empty signal, but block length is not zero. This error terminates the transfer.
Done signal is set, Done and Empty is not set	During a transfer from SSD to mainframe, the Done flag is set, but the mainframe does not receive the Done and Empty signal. This error terminates the transfer.
Done and Empty is set, but no Last Word flag	During a transfer from mainframe to SSD, Done and Empty flag is set, but mainframe has not sent the Last Word Flag signal. This error terminates the transfer.
Fatal Error	Caused by an OR condition of all above errors or an outstanding reference. This error terminates the transfer.
SECDED error on SSD during a read or write	The SSD detected a data error. If it is a single bit, it is corrected. If a double- or multiple-bit error has occurred, the transfer is terminated.
SECDED error on mainframe during a read or write.	The CPU detected a data error. If it is a single bit, it is corrected. If a double- or multiple-bit error has occurred, the transfer is terminated.
SBCDED error during CPU read	The CPU detected a data error. If it is a single byte, it is corrected. If a double- or multiple-byte error has occurred, the transfer is terminated.

Shared Module

All I/O channel request information follows a path from the issuing CPU, to the local shared module, and finally to the CPU that is connected to the I/O module that services the requested channel, as shown in Figure 21. If the local shared module determines that the requested channel is on the other physical half of the machine, it forwards the channel request to the remote shared module. The remote shared module then sends the information to the appropriate CPU module on that half of the machine.

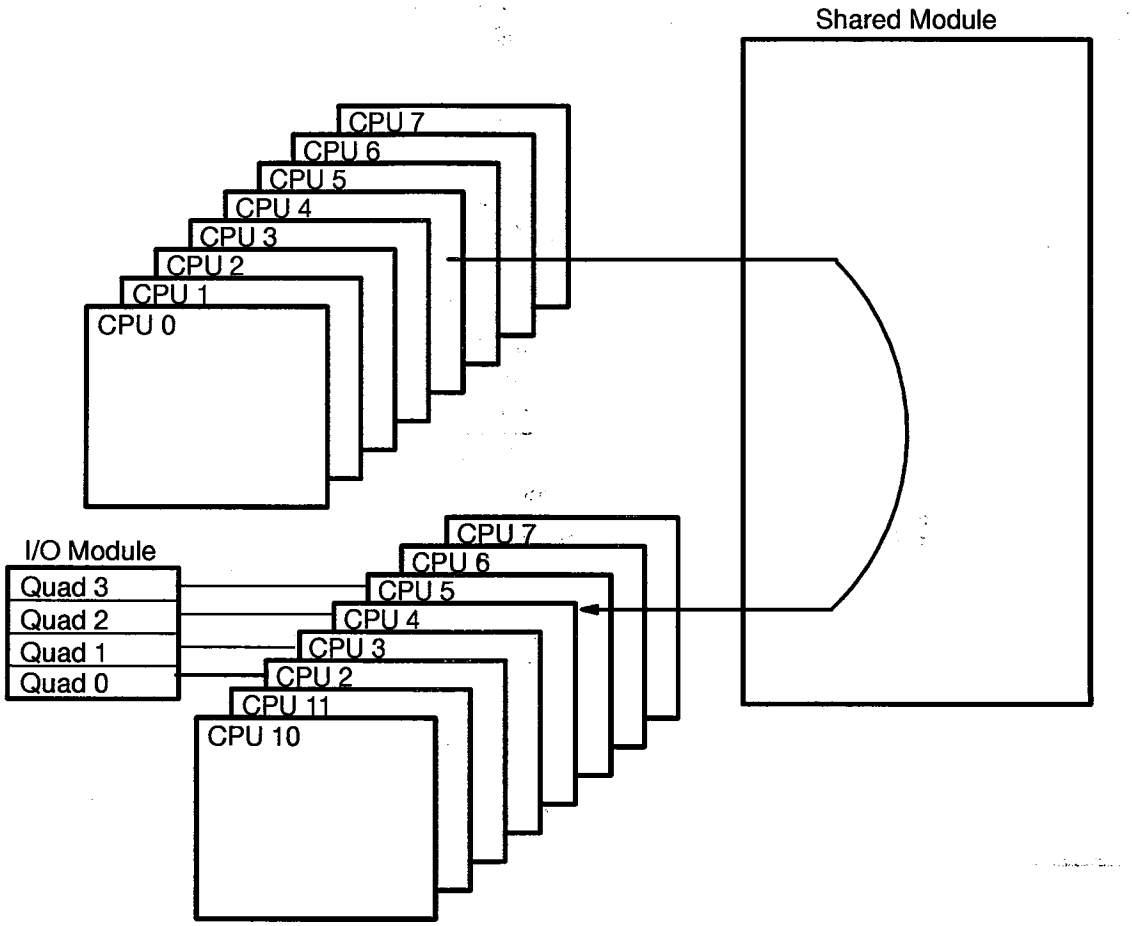

Figure 21. Channel Request Information Flow

Acronyms

The engineering documentation uses several acronyms identifying signals and commands used in the shared module. The following list defines these acronyms. Not all of these are used for an I/O instruction.

- CCMD - CPU shared command
- CMD - shared command
- SPTR - steering pointer
- RPTR - register pointer
- SCPU - source CPU
- DCPU - destination CPU
- EDCPU - effective destination CPU (from channel number)
- ECLN - effective CLN (from CLN+partition)
- MM/UM - monitor mode/user mode
- T&S - test and set
- DLI - deadlock interrupt

Shared Module Ports

The shared module is divided into ports, and the ports are then grouped. There are 16 ports on each shared module, one for each of the 16 CPUs that connect to it. The ports are numbered in decimal. Each port contains both an SA and an SB option. SA0 and SB0 make up the port for CPU0, SA1 and SB1 make up the port for CPU1, and so on. The ports are then divided into four groups for the passing of information and data. Group N contains ports 8, 0, 9, and 1; Group 1 contains ports 10, 2, 11, and 3; Group 2 contains ports 12, 4, 13, and 5; and finally Group 3 contains ports 14, 6, 15, and 7.

Figure 22 shows the grouping of the ports within the shared module. Group 0 is expanded in the diagram to show data and command flow from each port within the group to the other ports within that group.

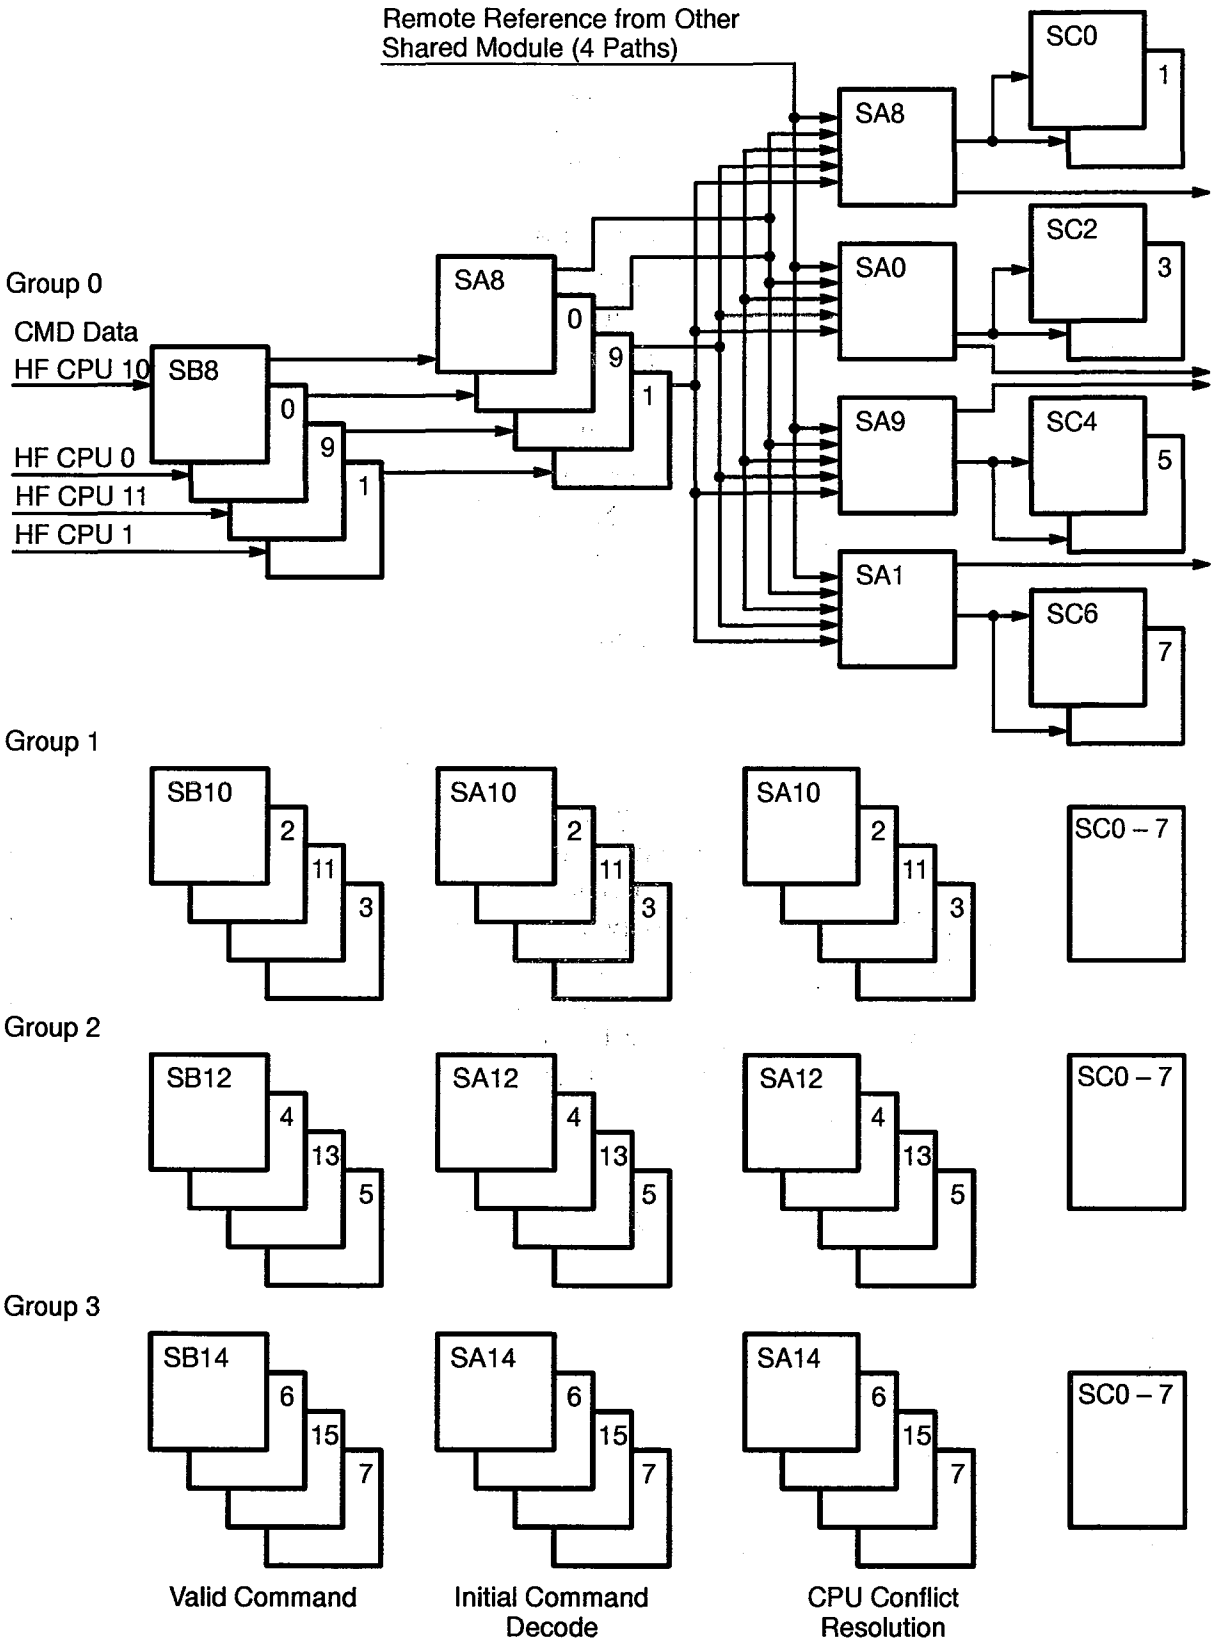

Figure 22. Shared Module Port/Group Designation

Shared Module Options

The shared module options are described below.

SA Option

There are 16 SA options, numbered SA00 through SA15, on the shared module. The SA option is one of two options that make up the 16 ports located on the module. This option decodes the channel number requested during an I/O instruction and generates the effective destination CPU (EDCPU) code and the steering pointer (SPTR) code.

In addition to collecting data and shared command codes, the SA option is designed to handle conflict resolution between CPU requests for CLN activity, I/O interrupts, remote references, and illegal source pointers.

SB Option

There are 16 SB options located on the shared module. These are numbered SB00 through SB15. The SB option is the second of two options that make up the 16 ports located on the module. The SB option is used to pass data and command information to the CPU that will service the I/O channel request.

The SB option receives the initial command code for an I/O request and, once the code is validated, sends the information to the SA option.

SC Option

There are 8 SC options used within the shared module, SC0 - SC7. The SC option performs functions related to channel interrupts and the real-time clock (RTC). During an I/O request, the SC option routes data to the SD option.

Channel interrupt and the SIE scoreboard are held in SC00. SC00 resolves RTC broadcasts with the remote shared module.

SD Option

There are 16 SD options on the shared module, numbered SD00 through SD16. The SD option routes data and command information to the CPU port that services the I/O request. Data and command information is routed to the correct port via the steering pointer code generated on the SA option.

Shared Command Codes

The shared module routes information with command codes. When the CPU decodes an I/O instruction, it generates a command code specific to that instruction. Table 12 lists the command codes generated for each I/O instruction.

Codes 50, 51, and 52 are shared module responses to a CPU. Codes 65, 66 and 67 are used internally on the shared module. Command 64 is generated on the CPU and sent to the shared module.

Table 12. Shared Command Codes

Command Code	Instruction	Description
c11	0010*jk*	Set channel address register
c12	0011*jk*	Set channel limit register
c14	0012*j*0	Clear I/O interrupt Clear error flag Clear device MC (output channel only) Enable channel interrupt
c15	0012*j*1	Clear I/O interrupt Clear error flag Set device MC (output channel only) Clear ready held (input channel only) Enable channel interrupts
c16	0012*j*2	Disable channel interrupt
c17	0012*j*3	Enable channel interrupts
c20	033*i*00 033*ij*0	Read lowest interrupting channel to Ai Read channel address from channel to Ai
c21	033*ij*1	Read error from channel
c50		Route I/O request from source CPU (SCPU) to effective destination CPU (EDCPU) where channel is held
c51	c20, c21	Same as 50 except start destination SB 033 timer
c52	c14, c15	Same as 50 except start destination SB 033 timer
c60	IO Respond	Route to SCPU in RPTR to respond to 033 request
c64	Set IO interrupt	I/O interrupt to shared scoreboard SC0
c65	Clear IO interrupt	Clear I/O interrupt channel in partition
c66	PINT	Broadcast to all CPU PINT
c67	MCU	Broadcast to all CPU MCU

Shared Module I/O Sequence

The following pages describe the sequence that occurs during a typical I/O reference. The HF option in the issuing CPU generates the command code from the I/O instruction. The HF option passes the command code, along with A*j* and A*k*, to an SB option on the shared module. If the SB option has received a Valid CPU Command signal from the SA option (indicating the presence of sanity), then the command code, A*j*, and A*k* are sent to the SA. This flow is illustrated in Figure 23.

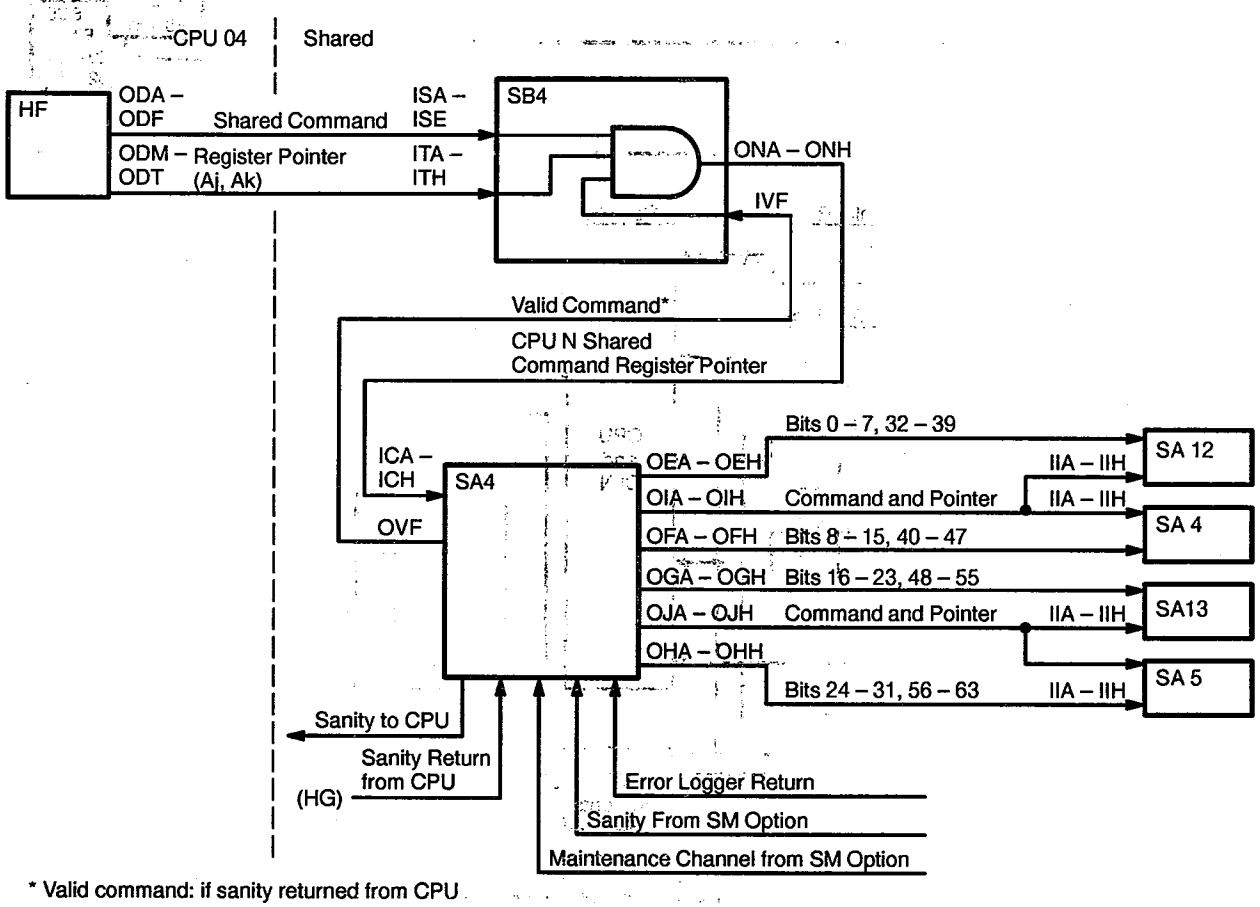

Figure 23. CPU-to-Shared Module Initial Data Flow

The SA options resolve any conflicts among the CPUs, I/O, and remote shared access within the group. The SA options can hold three reference requests. Figure 24 illustrates the conflict resolution in the SA option.

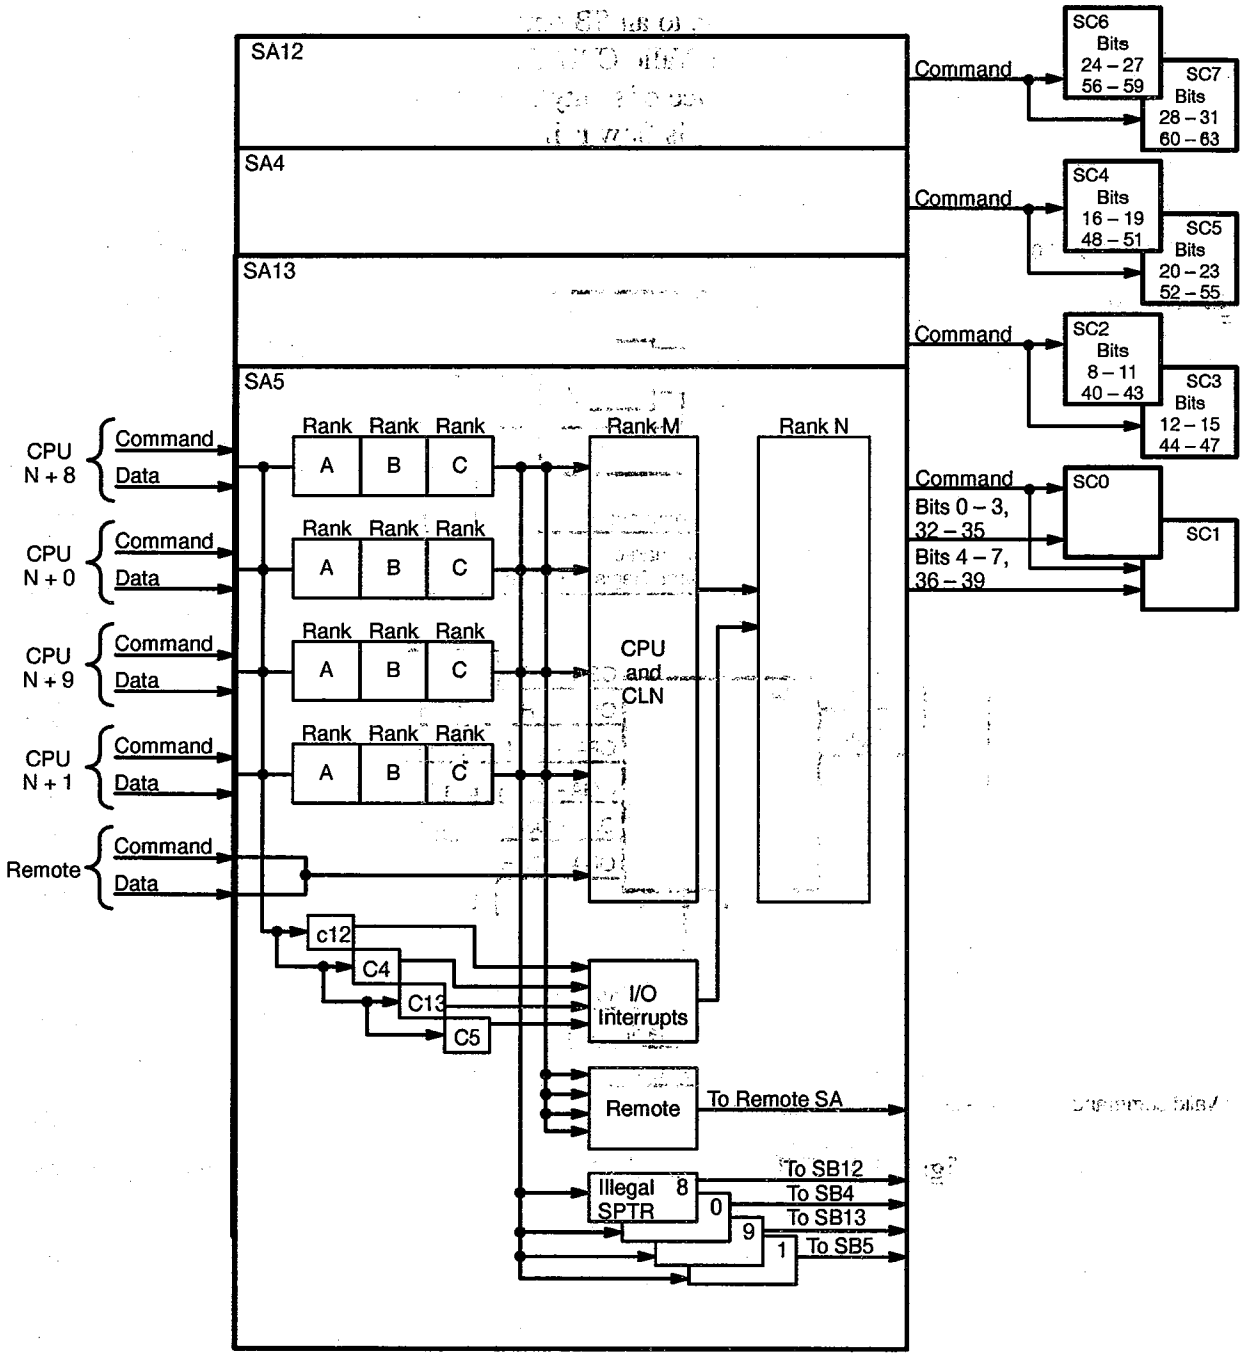

Figure 24. SA Option Conflict Resolution

The SA option also decodes the channel number and generates the effective destination CPU (EDCPU) code. This code is used to create a steering pointer (SPTR). The steering pointer allows the shared module to pass the data to the correct port for the CPU that services the I/O request.

The SC options send the data to the SD options, which steer the information to the appropriate port (SB options) for the destination CPU. The SB options send the information to the destination CPU in two 32-bit transfers. Figure 25 illustrates the format of data in this transfer.

CPU # (SCPU) 8 bits	Group # 2 bits	Command 6 bits	A*j* 8 bits	A*k* 40 bits

Figure 25. Data Format for Shared Module-to-CPU Transfers

On the CPU, the HF option reassembles the information into four 16-bit packets and sends it to the HC option, along with the Go I/O Command signal. The HC option is basically a pass-through option that passes the data to the DC option in the I/O module.

I/O Sequence

This last section details the events that occur during an I/O sequence. The example describes the system configuration, the operating system status, all signal transfers, data movement, and finally, all conditions required for the sequence to occur.

For the example, assume that the following configuration is in place: A CRAY T916 chassis with 16 CPUs, one shared module, two I/O modules, and the IOS configuration that is shown in Figure 26.

A user job is running in CPU 14. The job requests a file. At that point, the job swaps out, and the kernel swaps into CPU 14. The operating system knows that the requested file resides on the disk drive. The kernel issues instructions to initiate the LOSP channel. The 0012*j*0 instruction clears the channel, with A*j* designating the channel number. The next instruction issued is the 0011*jk*, which sets the channel limit, with A*j* identifying the channel number and A*k* containing the channel limit address. The 0010*jk* instruction is issued next. Again, A*j* contains the channel number and A*k* contains the current channel address. This instruction also starts the LOSP channel.

With each LOSP channel instruction that is issued, A*j*, A*k*, and an I/O command are sent to the local shared module. In this example, there is only one shared module. The shared module attaches an I/O group number and a requesting CPU number to the information received from the CPU.

The shared module determines which I/O module (via a CPU) to send the information to. In this case, the requested file resides on disk drive *x*, which is supported by IOS cluster 1, which is connected to I/O module I, quad 1. Because CPU 3 has access to quad 1 of the I/O module, the shared module sends the I/O information through that CPU to quad 1. The LOSP commands are sent through the LOSP channel to cluster 1, which initiates the disk transfer.

When the data is waiting in cluster 1's buffer board, the cluster sets up the HISP transfer. The data is sent through the HISP channel, through CPU 3, and into the designated memory location. The I/O logic on the CPU module is independent. CPU 3 can pass I/O data to or from memory while the CPU processes the user job without affecting the job running in the CPU.

When the HISP transfer is complete, the LOSP sends an interrupt through CPU 3 to the shared module. When an I/O channel generates an interrupt, that fact is recorded in an interrupt scoreboard on the shared module. There is one scoreboard per logical group, with a maximum of four groups. When a group has one or more interrupts pending, a *token* is placed into a ring that has a node at each CPU. At each node, the token is compared with the CPU's group number and process state. If the examined node does not fit the requirements, the token is advanced to the next node. When a CPU matches the group number and is enabled to service I/O interrupts, then that CPU handles the interrupt. In this example, CPU 7 matches the token and accepts the interrupt.

The kernel enters that CPU and notifies the operating system that the requested data is ready in memory. The operating system wakes up the original process and puts it into a queue. When the job is at the top of the queue, it rolls into the CPU and continues processing.

The following conditions must be met for this example to occur:

- The requesting CPU (CPU 14) and the I/O handling CPU (CPU 3) must have memory partitioned the same way.
- The I/O handling CPU (CPU 3) must be enabled to pass I/O.
- The requesting CPU (CPU 14) must be enabled to originate I/O.
- The requesting CPU (CPU 14) must be in the same I/O group as LOSP channel 104/105.
- The requesting CPU (CPU 14) must have shared access.

These conditions are actually parameters that are set up under the System Configuration Environment (SCE).

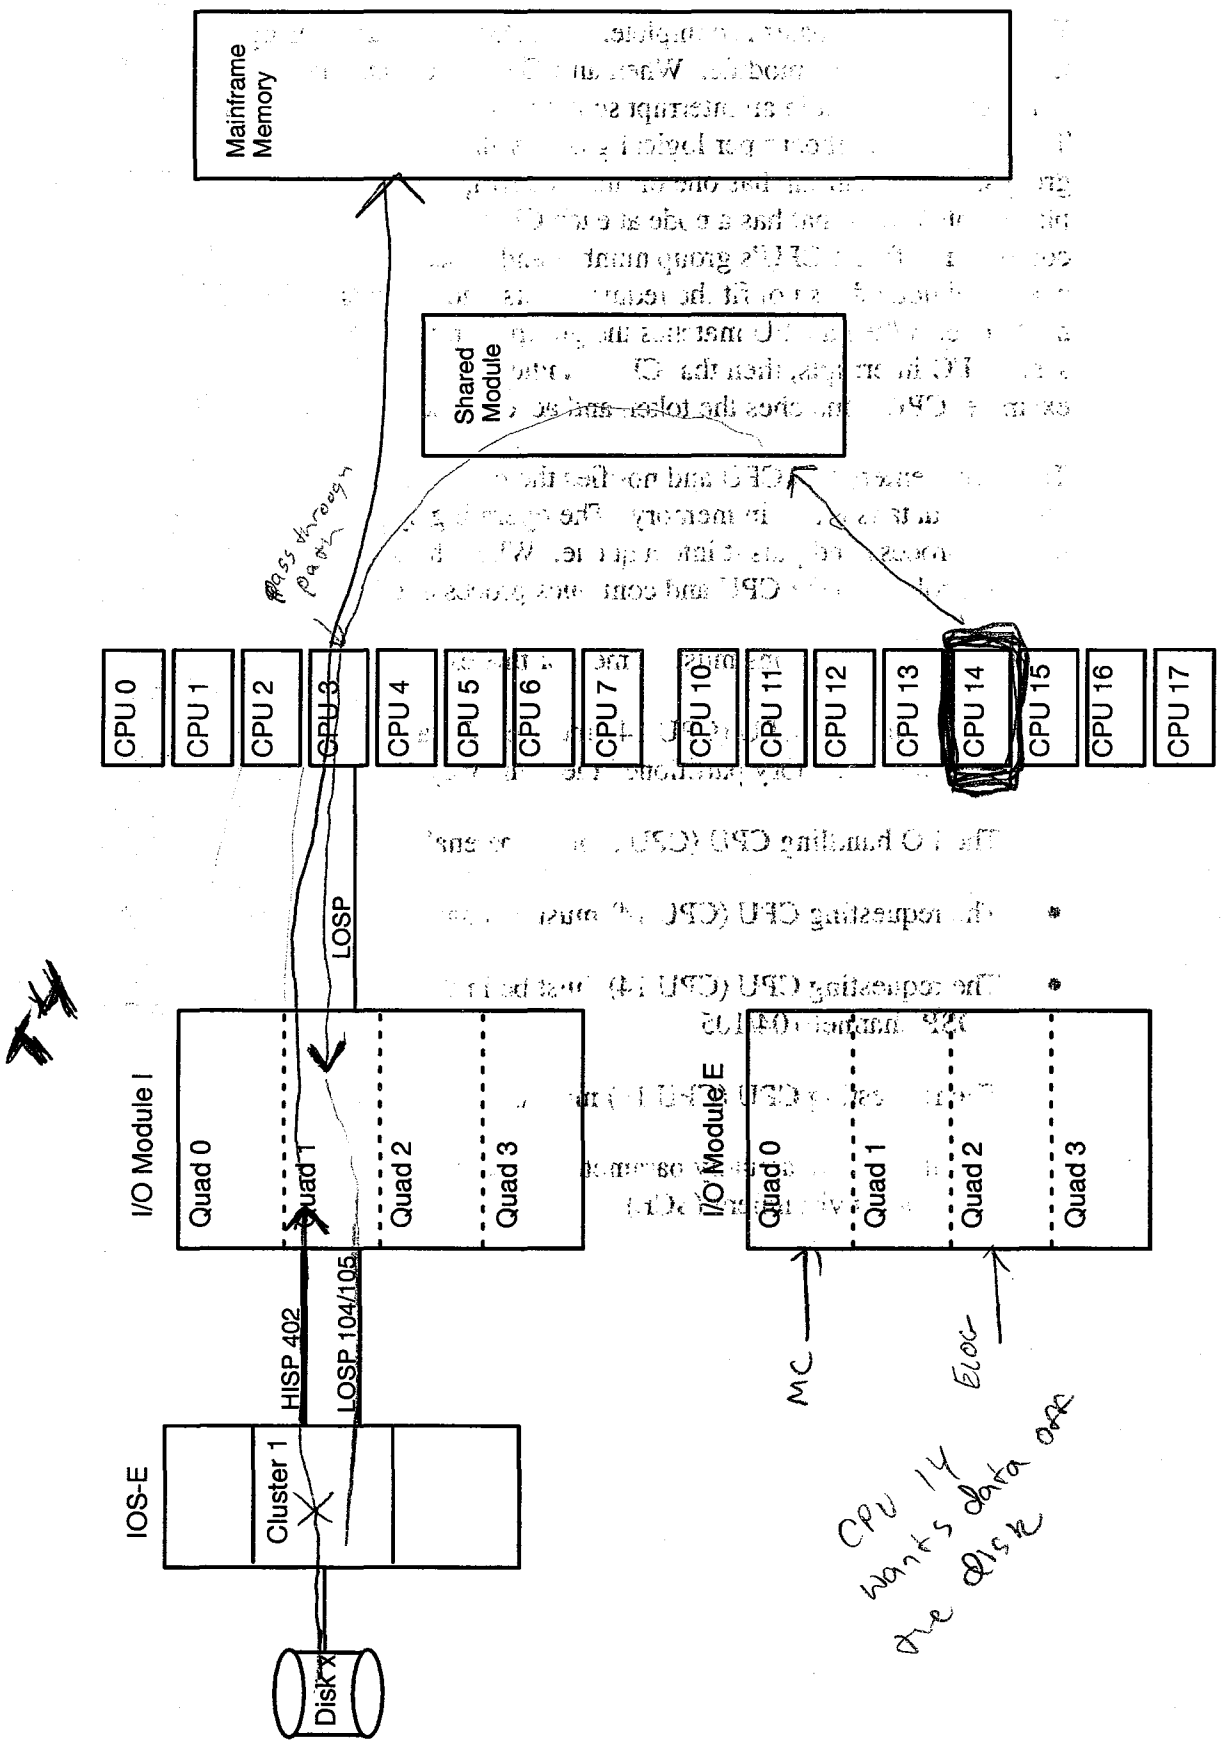

Figure 26. Example System Configuration

Reader Comment Form

Title: **CRAY T90™ Series**
I/O Theory of Operations
Preliminary Information

Number: **HTM-xxx-0**
December 1994

Your feedback on this publication will help us provide better documentation in the future. Please take a moment to answer the few questions below.

For what purpose did you primarily use this document?

_____Troubleshooting
_____Tutorial or introduction
_____Reference information
_____Classroom use
_____Other - please explain _____

Using a scale from 1 (poor) to 10 (excellent), please rate this document on the following criteria and explain your ratings:

_____Accuracy _____
_____Organization _____
_____Readability _____
_____Physical qualities (binding, printing, page layout) _____
_____Amount of diagrams and photos _____
_____Quality of diagrams and photos _____

Completeness (Check one)

_____Too much information _____

_____Too little information _____

_____Just the right amount of information

Your comments help Hardware Publications and Training improve the quality and usefulness of your publications. Please use the space provided below to share your comments with us. When possible, please give specific page and paragraph references. We will respond to your comments in writing within 48 hours.

NAME _____
JOB TITLE _____
FIRM _____
ADDRESS _____
CITY _____ STATE _____ ZIP _____
DATE _____

[or attach your business card]

CRAY RESEARCH, INC.

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.

Attn: Hardware Publications and Training
890 Industrial Boulevard
Chippewa Falls, WI 54729

Fold

STAPLE