ERROR LOGGER CHANNEL

Figures

Tables

Tables (continued)

Error Logger Channel Description

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 (refer to Figure 1). From an I/O bulkhead, the LOSP error logger channel connects to an IO module. From the IO module, the serial error logger channel connects to the rest of the modules in the system. The serial error logger channel is composed of two paths as described in the following subsection.

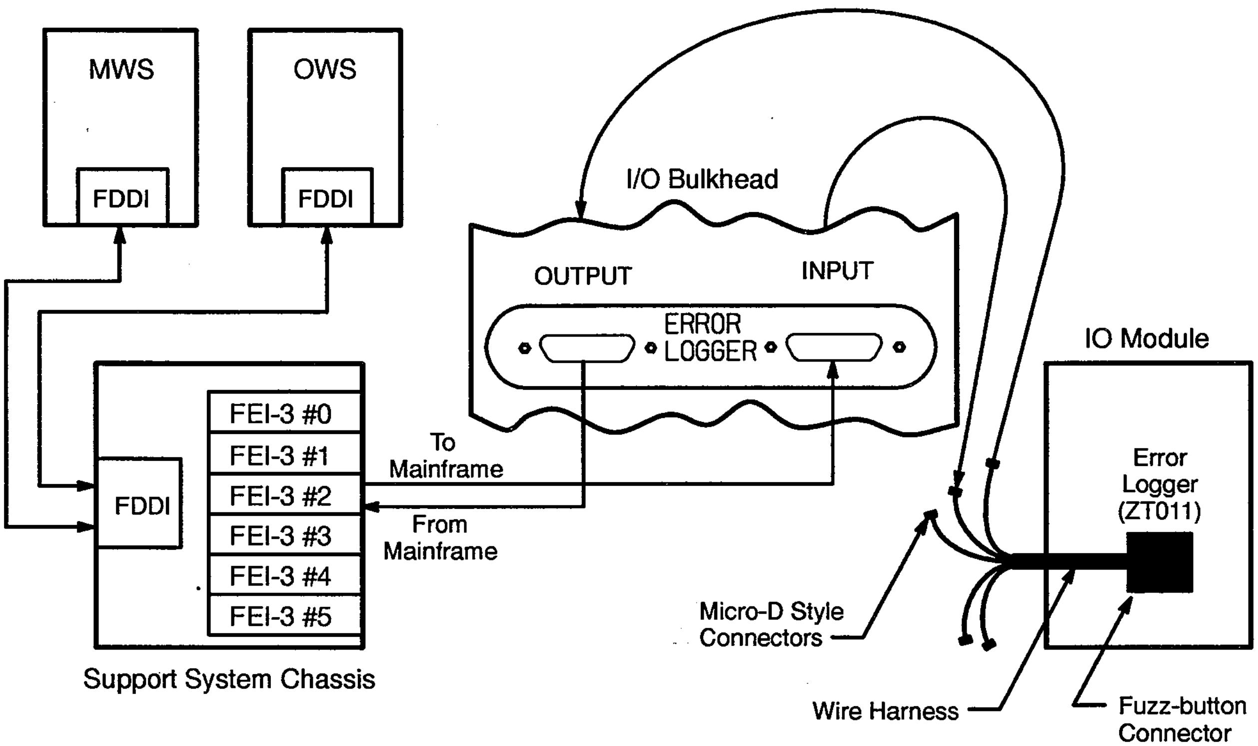

Figure 1. Error Logger Cabling Diagram

Error Logger Channel Paths

Errors are detected on either the IO, CP, or shared modules. Errors from IO modules (not including the master IO module) are sent to the shared module. Errors detected on the master IO module, the IO module that owns the master sanity code generator, are sent directly to the LOSP error channel interface (DD2 option). From the master IO module, errors are sent through the LOSP error logger channel to the I/O bulkhead.

Memory errors (detected on CP modules) and CPU errors are sent from the CP module to the local shared module. From the local shared module, errors are routed back through a CP module to the master IO module. The home port is the shared module port from which the shared module received sanity code from the IO module. The home port path is via one of the four physically connected CP modules.

Errors from the remote shared module (CRAY T932 systems) are forwarded to the local shared module. These errors are then sent from the local shared module through the CP module to the home port on the master IO module as shown in Figure 2.

NOTE: Detailed information that describes the error logger channel paths from IO, CP, and shared modules is provided in the subsection that describes errors detected on the corresponding type of module.

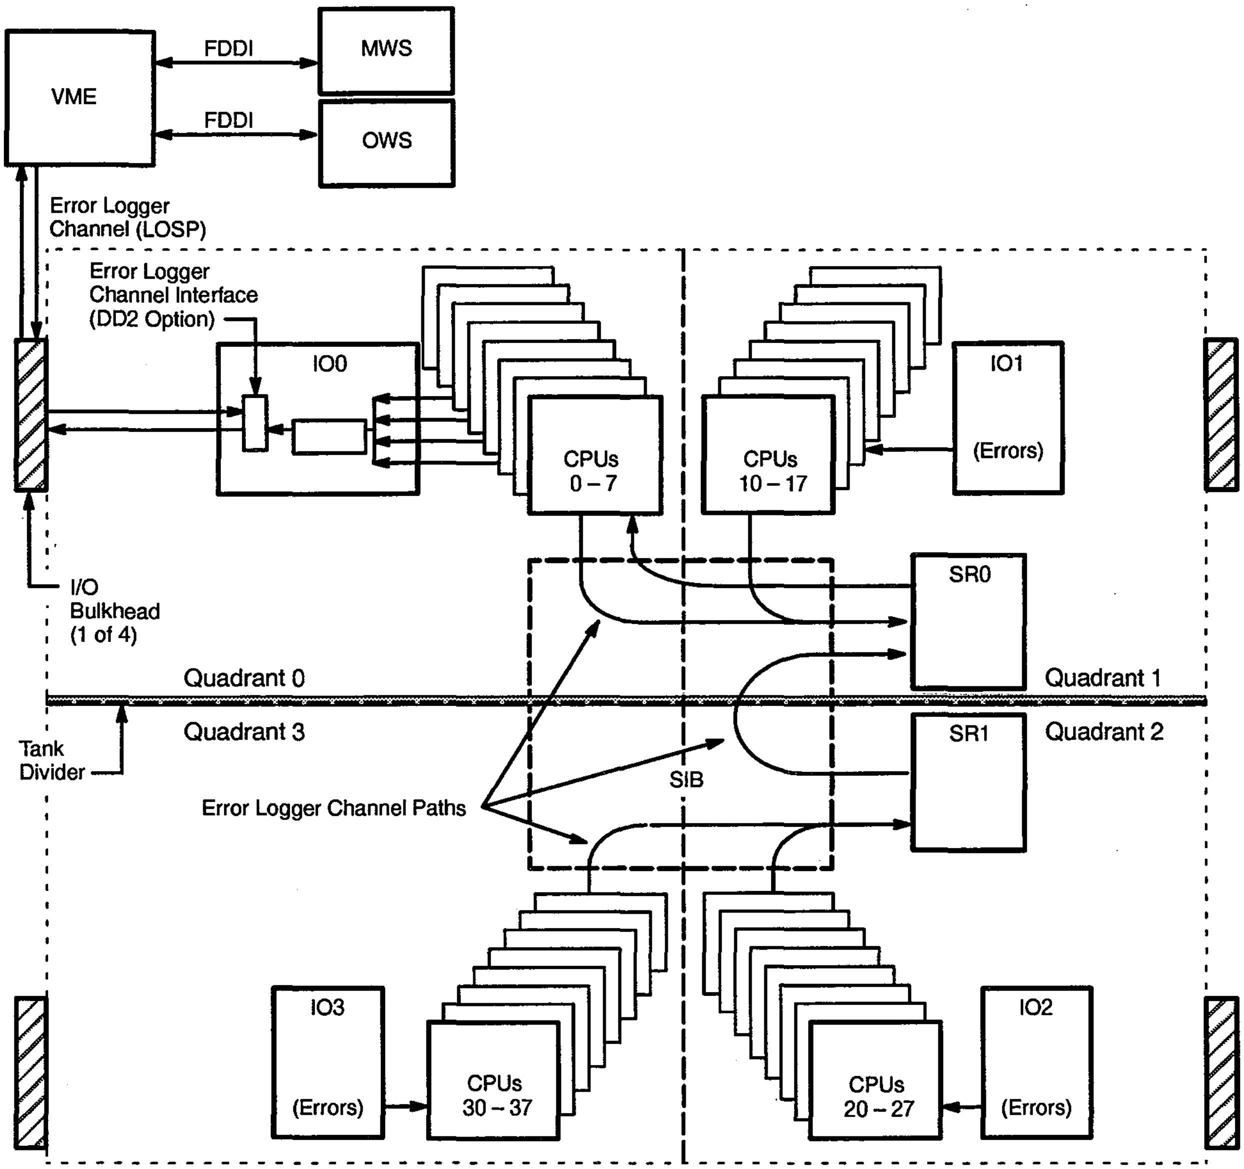

NOTE: IO module 0 is the master IO module in this example illustration.
IO module 2 could be the master.

Figure 2. Error Logger Channel Paths

Error Logger Data Words

All error logger messages or data words are 69 bits long as shown in Figure 3. Two types of error logger paths are used: Error Logger and Error Logger Acknowledge. These paths are grouped with maintenance and sanity code paths because they follow the sanity tree that was established when the mainframe was configured. Error messages are sent on Error Logger paths and acknowledgements are sent on Error Logger Acknowledge paths. The error logger paths flow in the opposite direction of the maintenance channel and sanity code paths.

The Error Logger path provides a path for reporting errors; it flows in the opposite direction of the sanity tree. The only functions of the Error Logger Acknowledge path are to acknowledge that one module has received an error from another module and to indicate that another error can be sent.

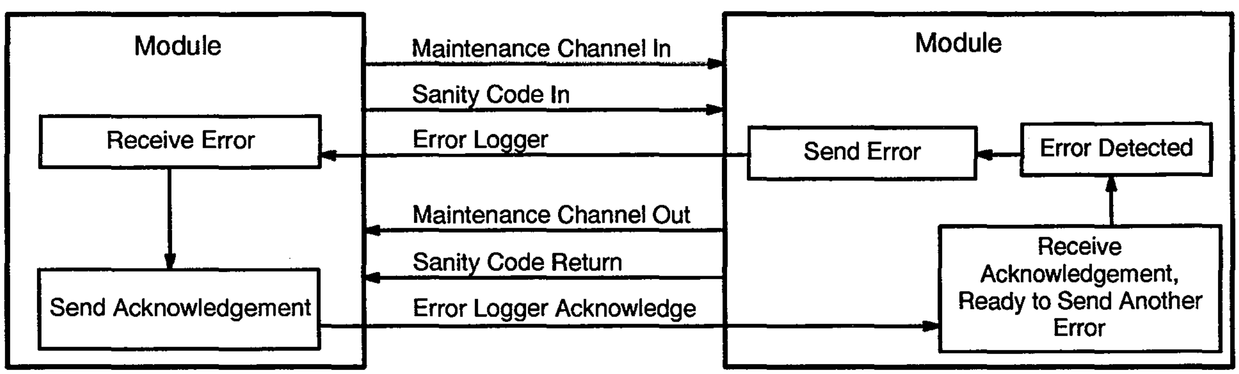

Error Data Word Sent on Error Logger Path

Bits 68	63	55	7 0
37 Header	Starting Error Code	Error Message Data (Data Fields are Different for Each Error Type)	Ending Error Code

Error Acknowledge Sent on Error Logger Acknowledge Path

Bits 68	63 0
37 Header	64 Zeroes

Figure 3. Error and Error Acknowledge 69-bit Words

Error Codes

The format of each type of error word differs. Each type of error word is identified by an 8-bit error code at the beginning and end of each error word as shown in Figure 4.

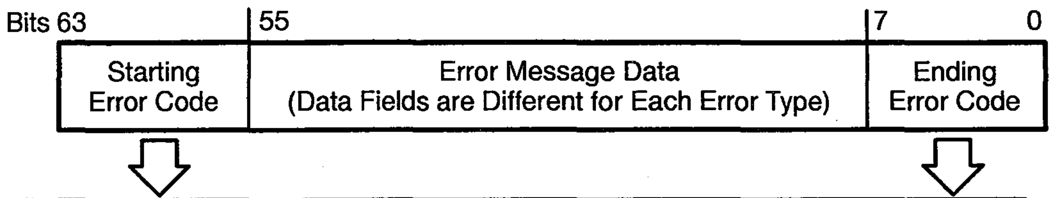

Error Type	Starting Error Code		Ending Error Code	
Common Memory	1 0 0 0 1	0 0 0	0 1 1 1 0	1 1 1
CPU HSR	1 0 0 0 1	0 0 1	0 1 1 1 0	1 1 0
Shared Module	1 0 0 0 1	0 1 1	0 1 1 1 0	1 0 0
IO Module				
Channel Errors	1 0 0 0 1	1 0 0	0 1 1 1 0	0 1 1
SECDED	1 0 0 0 1	1 0 1	0 1 1 1 0	0 1 0
Parity - HM	1 0 0 0 1	1 1 0	0 1 1 1 0	0 0 1

Figure 4. Error Codes

The error codes have two functions:

- To identify the type of error
- To verify parcel-to-data word synchronization

Identify Error Type

The last 3 bits (56 through 58) of the starting error code are used to identify different error types, as shown in Figure 4. Each error type has a unique error code.

Verify Parcel Synchronization

The 8-bit patterns of the starting and ending error codes are used to verify that data channel parcels 0 through 3 are synchronized (in sync) with data word transfers. The parcels are counted from left to right. The starting error code (8 bits, first parcel) is compared to the ending error code (8 bits, fourth parcel) to indicate that the proper 4 parcels are being transferred as a word.

The 8 bits of the ending error code (0 through 7) are inverted from the starting error code bits (56 through 63). If the starting and ending codes do not compare, the error data parcels are not in sync with the data word transfer (this means that the data word being transferred probably has parcels from more than one error word).

Error Types

The following error types are reported through the error logger channel. Each of these error types is described in detail later in this document.

- Common memory errors (detected on CP modules)
- CPU high-speed register (HSR) parity errors
  - Vector registers
  - B and T registers
  - Instruction buffers
  - Logic monitor (HM options) test-point buffer parity errors
  - Data cache
- IO module errors
  - SECDED - I/O read/write data to/from memory
  - HISP and VHISP channel errors
    (these errors are currently not reported to the error logger)
  - Logic monitor (HM options) test-point buffer parity errors
- Shared module errors
  - SR option errors
  - Logic monitor (HM options) test-point buffer parity errors

CP Module Errors

The following error types are detected on CP modules and reported to the error logger channel. Each error type is described in detail after the following error logger routing subsection.

- Common memory errors
- I/O write errors
- CPU HSR parity errors

Error Logger Routing from CP Modules

Common memory errors and HSR parity errors are reported to the HG option on the CP module. These error messages, along with error messages from IO or shared modules, are reported using one of two error logger paths out from the HG. The error logger channel that is used depends on whether the HG is owned by the IO module or the shared module, as shown in Figure 5. That is, whether sanity code was initially received from the IO module (IQA) or the shared module (IRA). The CP module receives sanity from an IO module initially only in a test environment, if there is no shared module available.

When a full system is operating normally, CP modules always initially receive sanity code from a shared module. The IQA input provides a secondary sanity code path for use in a STCO environment when there is no shared module. The IQA input can also be used in the field to run an IO module connected to a CP module and to isolate the shared module.

I/O write errors are also detected on the CPU module. These errors do not follow the regular error logger path. The HA0 and HA1 options on the CPU perform SECDED on the write data as it is received from the IO module. Any error information is sent directly to the IO module on the I/O command path. This path is shared with the I/O commands. The CPU can report an error every clock period if there are no conflicts with the I/O commands. The I/O commands have a higher priority for use of the path. The only event that prevents the CPU from immediately reporting errors to the IO module is if a command is being sent from the CPU to the IO module. This path is illustrated in Figure 15 on page 28.

A physical CPU number is returned with all CP module errors to identify the reporting CPU. This physical CPU number is assigned at configuration time with the System Configuration Environment (SCE) tool.

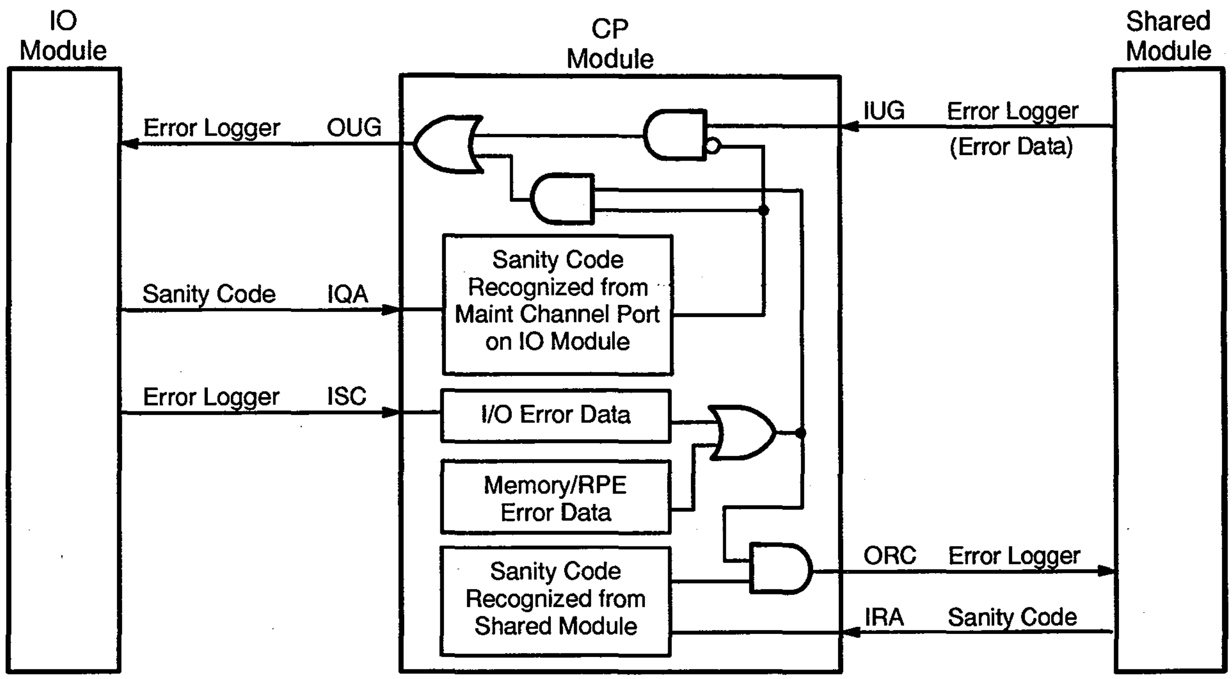

Figure 5. Error Logger Channel Routing from a CP Module

Common Memory Errors

In CRAY T90 series systems, common memory is divided into 8 sections. Memory errors are reported to the following three areas:

- To the error logger channel
- To interrupt flags in the exchange package (Figure 7)
- To CPU status registers 4, 5, and 6

Memory error correction is done on the CJ options on the CP modules (CJ0 through CJ7, which correspond to memory sections 0 through 7) as shown in Figure 6.

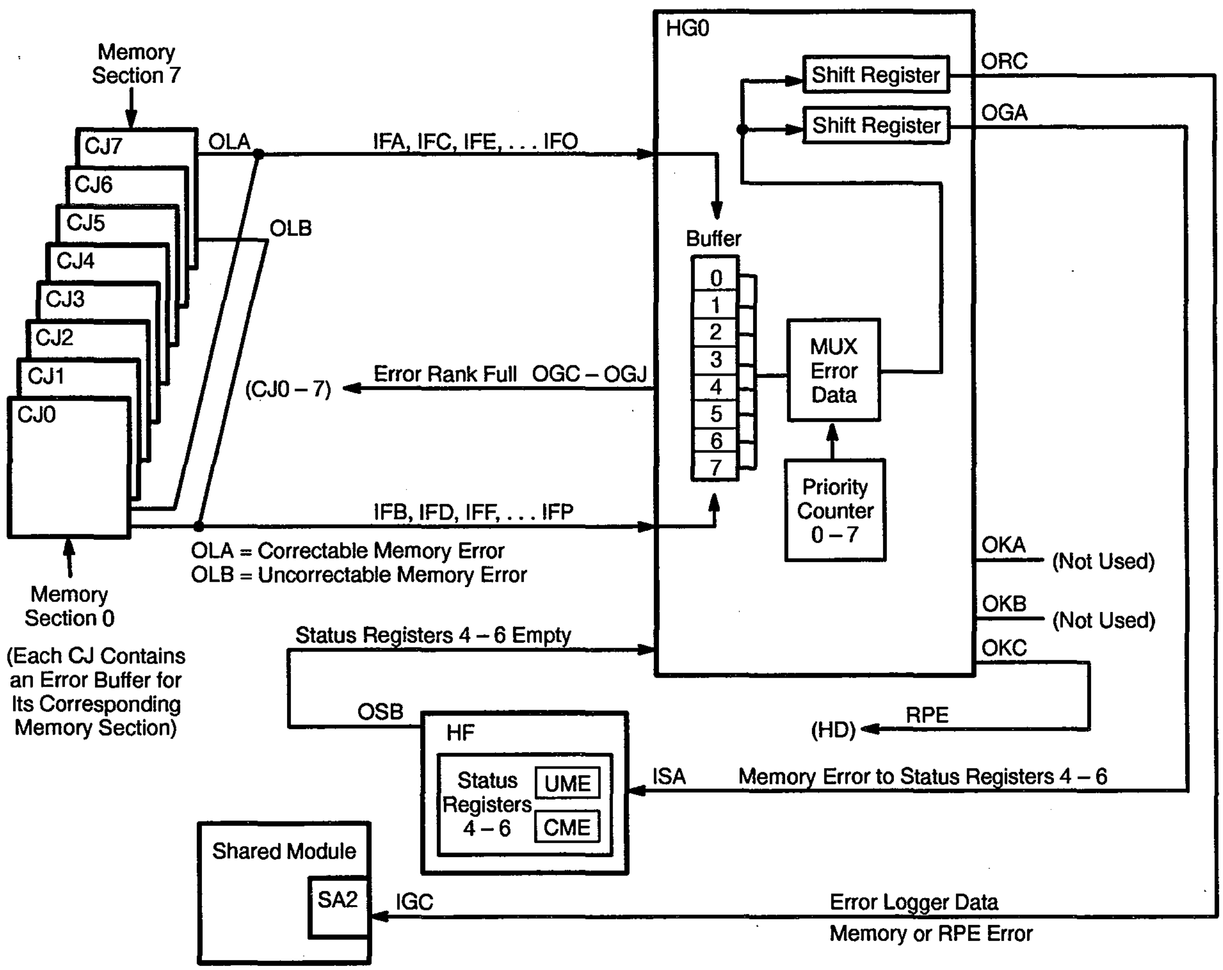

Figure 6. Memory Error Reporting to Status Registers and the Error Logger Channel

Each CJ option has an error rank, or buffer, that follows the error-reporting priority as shown in Figure 10; one buffer holds errors and the other shifts or sends errors to the HG option.

The CJ options send single-byte correction/double-byte detection (SBCDBD) error data to the HG option [the CJs send 38 bits (bits 55 through 18 as shown in Figure 8) serially to the HG]. Of the two paths from each CJ option, one path is used for uncorrectable memory errors and the other path is used for correctable memory errors. The HG option receives SBCDBD errors on 16 inputs (2 from each of the 8 CJ options) and stores them in a buffer as shown in Figure 6.

If the buffer on the HG is full, the HG sends an Error Rank Full signal to the CJ option whose errors have filled both ranks (two errors from the same memory section). The Error Rank Full signal is dropped after the HG reports the error; the CJ option then sends any remaining errors. Each CJ option also has two error buffers.

Errors are loaded into two shift registers in the HG option: one for error logger channel reporting (ORC) and the other for reporting to the status registers (OGA). The HG has a priority counter that cycles sequentially through each section and stops when it detects an error from one of the memory sections. After the error from that memory section is reported, the counter continues and cycles to the next section; the counter does not reset to start with section 0.

Interrupt Flags

The HD option receives error information from the CJ options for setting the memory error correctable (MEC) and the memory error uncorrectable (MEU) flags. Refer to Figure 7. The HD option does not have enough pins to receive correctable and uncorrectable error flags from all 8 CJ options. Therefore, the error signals (OMB to IMA and OMC to IMB) are daisy chained through the CJs and then sent to the HD.

Although errors may be lost on the CJ option and never reported to the HG option, interrupt flags should still be set on the HD option. Therefore, the CJ option reports correctable and uncorrectable memory errors to the HD option. The OKA (correctable memory error) and OKB (uncorrectable memory error) pins on the HG option are not used.

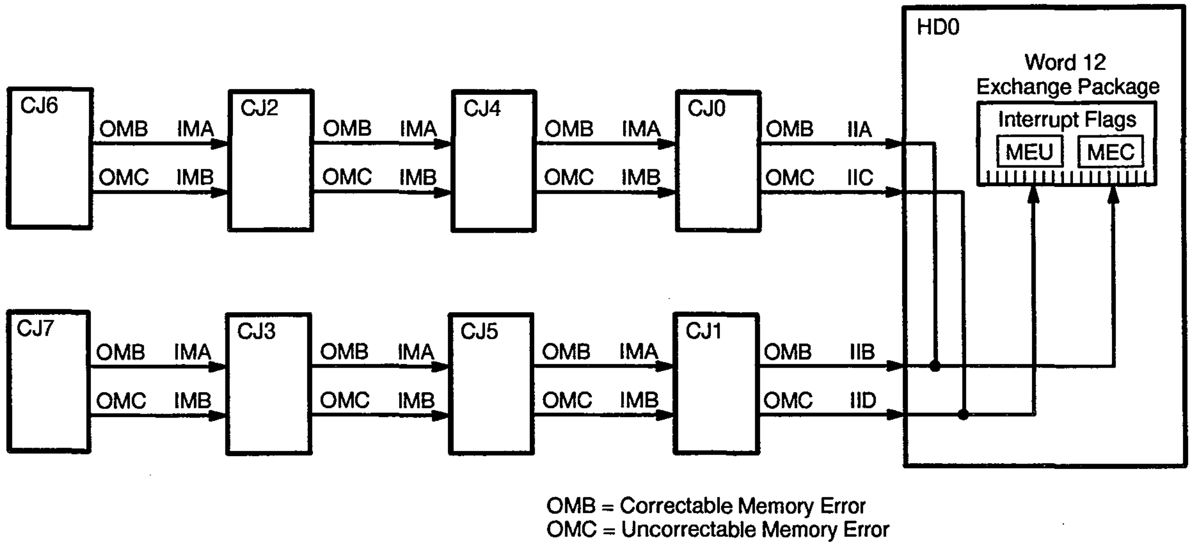

Figure 7. Setting Memory Interrupt Flags

Status Register Flags

Status register 4 contains the correctable and uncorrectable memory error flags. Bits 32 through 45 of status register 4 define the destination code associated with the error. Status register 5 bits 32 through 43 contain the syndrome code of the error. Status register 6 bits 32 through 44 contain the error address for a CPU memory error.

Common Memory Error Words

Figure 8 shows the format of common memory error words; the text following the figure describes the fields within the word.

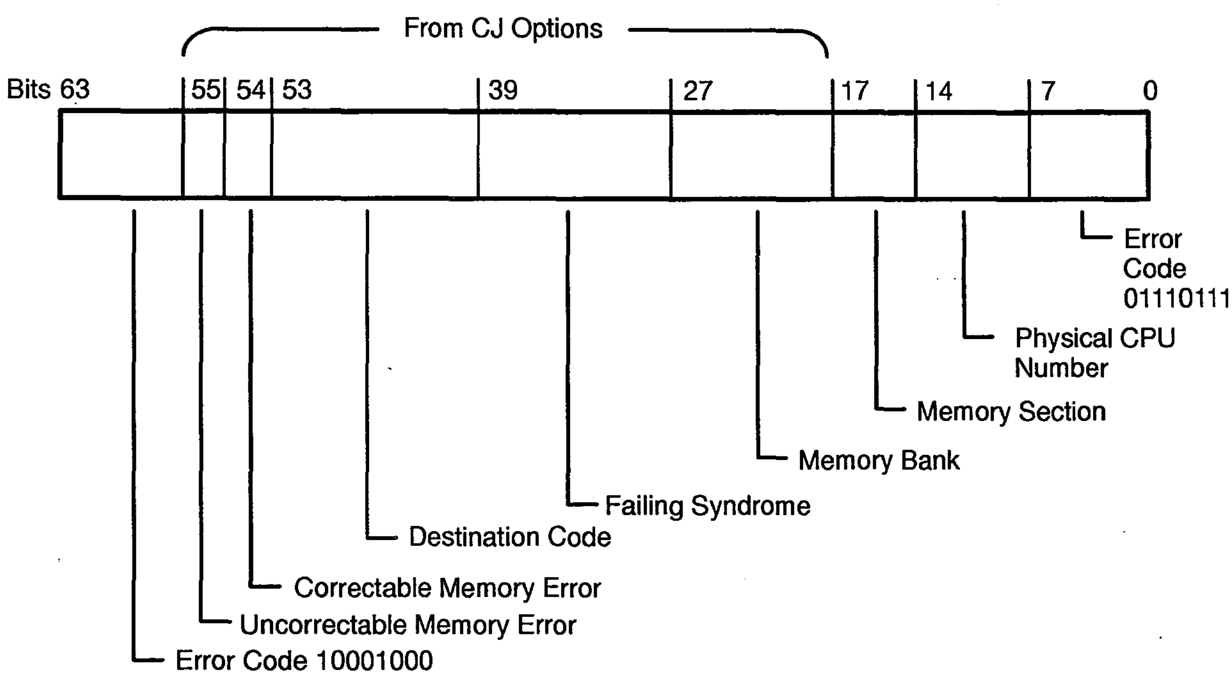

Figure 8. Common Memory Error Word Format

Physical CPU Number. This is the physical CPU number assigned at configuration time.

Memory Section. This is the memory section number in which the failing data word occurred. This is the physical memory section number; it is not affected by any memory degradation options. Table 1 lists the memory sections and corresponding mainframe memory stacks.

Table 1. Memory Section-to-Memory Stack Map

Memory Section	CRAY T932 Physical Stack	CRAY T916 Physical Stack	CRAY T94 Stack/Module
0,1	H	L	C/02
2,3	L	H	C/04
4,5	P	H	C/03
6,7	D	L	C/05

Memory Bank. This number indicates the memory subsection (for CRAY T932 systems) and bank number of the failing data word. These are the physical memory subsection and bank numbers; they are not affected by any configuration or memory degradation options. Figure 9 shows the bit assignments for the memory bank field.

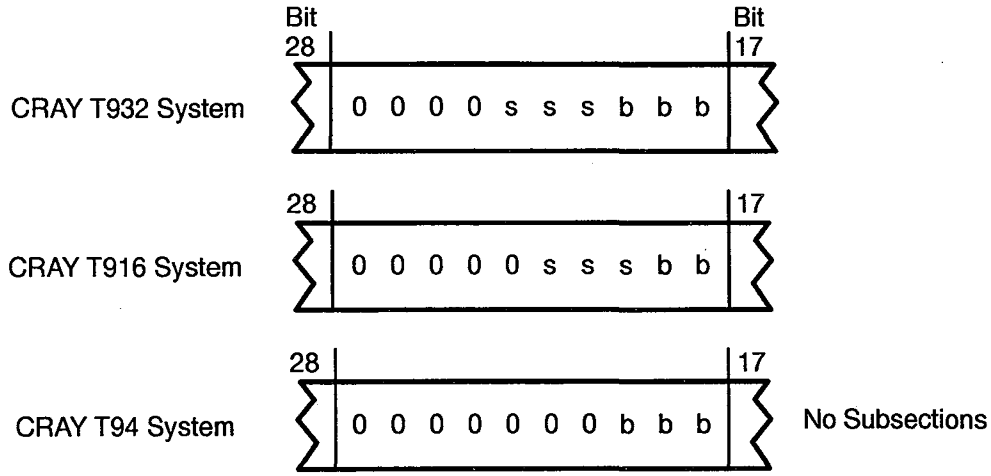

s = subsection; b = bank

NOTES: On CRAY T932 systems, subsection numbers correspond to the module number of the stack that is defined in the memory section field of the error word (refer to Table 1).

Figure 9. Memory Bank Field Bit Assignments

Failing Syndrome. If the error is correctable, this syndrome number is used by the error correction algorithm to determine the failing memory chip. [You can flaw a bad memory chip and configure the system to use one of the spare memory chips by issuing special direct memory access (DMA) commands from the maintenance channel.] The syndrome is not required for most maintenance actions because the memory bank indicated in the error word (bits 27 through 18) identifies the replaceable component.

Destination Code. This is the final destination of the requested word as described in Table 2.

Table 2. Destination Codes

53	52	51	50										40	Description	
1	1	1	-	e	e	e	e	e	e	e	e	e	e	Cache	Word
1	1	0	r	r	r	-	e	e	e	e	e	e	e	Vector	Number, Element
1	0	1	r	r	r	0	-	-	-	-	-	-	-	S Register	Number
1	0	1	r	r	r	1	-	-	-	-	-	-	-	A Register	Number
1	0	0	-	-	-	0	-	e	e	e	e	e	e	T Register	Number
1	0	0	-	-	-	1	-	e	e	e	e	e	e	B Register	Number
0	1	1	r	r	r	-	-	e	e	e	e	e	e	Instruction	Buffer, Word
0	1	0	r	r	r	-	-	e	e	e	e	e	e	I/O	Buffer, Word
0	0	1	-	-	-	-	-	-	e	e	e	e	e	Exchange	Word

(Header "Destination Code Bits" spans bits 53–40.)

e = element, r = register number, and - equals unused bit position returned as 0.

Correctable Memory Error. When this bit is set to 1, it indicates a correctable memory error.

Uncorrectable Memory Error. When this bit set to 1, it indicates an uncorrectable memory error.

NOTE: A correctable memory error can overwrite an uncorrectable memory error in a CJ option. When this happens, both the correctable and uncorrectable bits of the error word are set.

Memory Error-reporting Priority

Memory error-reporting logic is set to function as defined in the following bulleted list and as shown in Figure 10:

- Error logic attempts to report errors from each section of memory to CPU status registers 4, 5, and 6 and to the error logger channel.

  NOTE: An error in a status register can prevent further error updating. All three status registers must be read to enable another error to be stored.

- Each of the 8 memory sections has a 1-word error buffer that is located on the CP module. Each shared module has a buffer for each of the 16 CP modules to which it directly connects. Both of these types of error buffers are written to and read from in a circular fashion.

- Each memory section has two error buffers located in a CJ option. The first buffer feeds the error logger. The error in the second buffer is held until the first buffer is empty; then the error in the second buffer is sent to the first buffer. If an error from the same section of memory occurs when the second buffer is full, the error in the second buffer is overwritten with the new error, as long as the new error is of a different type (correctable or uncorrectable) than the errors in both the first and second buffer.

  NOTE: The uncorrectable error flag (MEU) remains set if a correctable error overwrites an uncorrectable error.

- If a CPU is waiting to report an error to the error logger (because another CPU is reporting an error) and to the status registers and already has two errors stacked on the CJ, any new errors from the same section may be lost. If another memory error is detected from a different section of memory, it is saved because an error buffer exists for each section of memory. The error logic tries to report errors to status registers 4, 5, and 6 and to the error logger channel. Once an error is reported to either the status registers or the error logger channel, then that error can be overwritten by a new error.

Action	First Buffer	Second Buffer	New Error
New Error is Lost.	Correctable	Correctable	Correctable
New Error Overwrites Second Buffer.	Correctable	Correctable	Uncorrectable
New Error is Lost.	Correctable	Uncorrectable	Uncorrectable or Correctable
New Error is Lost.	Uncorrectable	Correctable	Uncorrectable or Correctable
New Error Overwrites Second Buffer. MEC and MEU Flags are Set.	Uncorrectable	Uncorrectable	Correctable
New Error is Lost.	Uncorrectable	Uncorrectable	Uncorrectable

NOTE: All error data recorded in these buffers and in the CPU status registers is assigned a physical CPU number.

Figure 10. Memory Error Buffering Rules

CPU HSR Parity Errors

The five types of HSR options used on the CP module each have internal parity checking logic to verify the integrity of data stored in the registers. Table 3 lists the HSR options.

Table 3. CPU HSR Option Types and Functions

Option Type	Function
VR	Vector registers
CH	Data cache
IC	Instruction buffers
BT	B and T registers
HM	Test-point buffers and logic

Merging of Parity Errors

Parity is checked when data is read from one of five register types. The OA options provide a location for all parity errors to merge with all HSR options on CP modules as shown in Figure 11. Each HSR option sends a parity error signal to one of six OA options. The OA selects one of eight parity error inputs based on a fixed priority scheme.

Each OA that receives a parity error generates a 4-bit code (OOA through OOD) composed of a 3-bit Parity Error Code and a 1-bit Valid Error Code. The 4-bit codes are sent to OA0 for the last level of the parity error merge. OA0 then selects one of the six 4-bit codes based on a fixed priority scheme and generates a 6-bit Parity Code (OOE through OOJ). The Parity Code is routed through CH options (2 bits each from options CH12, CH13, and CH14) before it is sent to the HG option. (The HG option provides error logger functions for the CP module.)

The HG option latches the 6-bit register parity error (RPE) code. When the HG receives a Status Register 7 Empty signal from the HF option, it sends the RPE to status register 7 in the HF option. The HG option also sends an RPE Interrupt signal to the HD option to set the interrupt flag in word 12 of the exchange package.

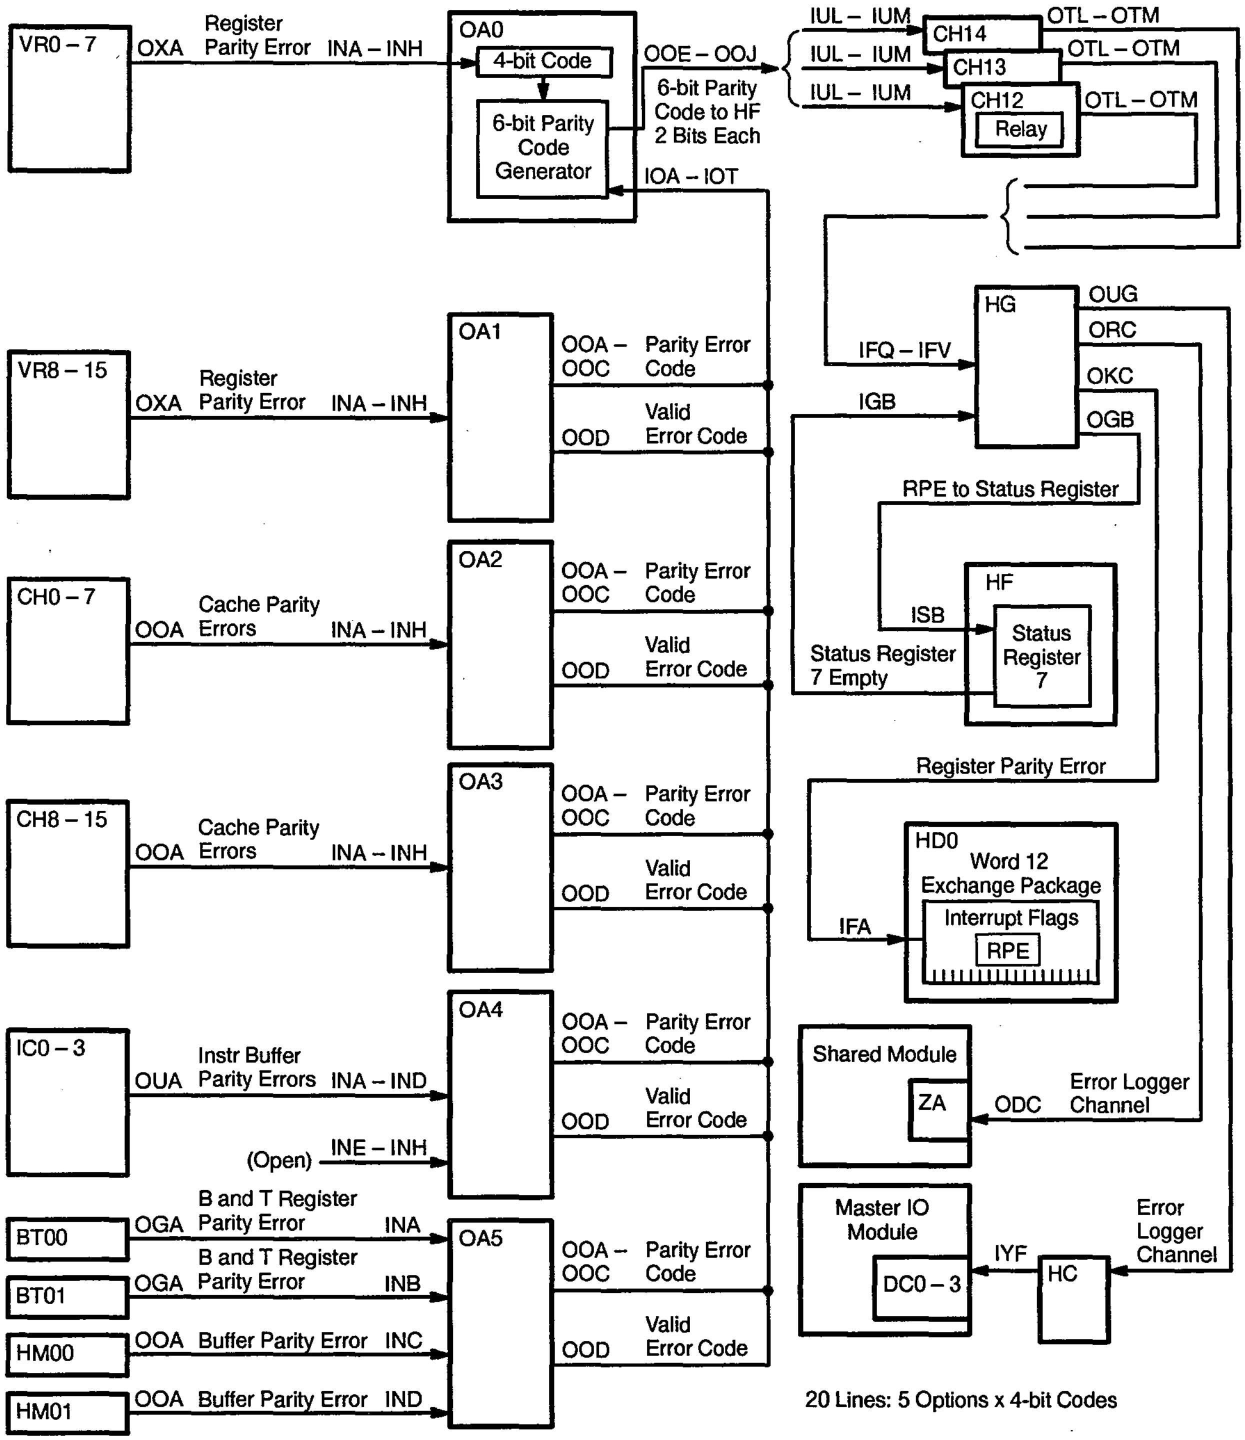

Figure 11. HSR Parity Error Merge

CPU HSR Error Words

Figure 12 illustrates the format of CPU HSR error words; the text following the figure describes the fields within the word.

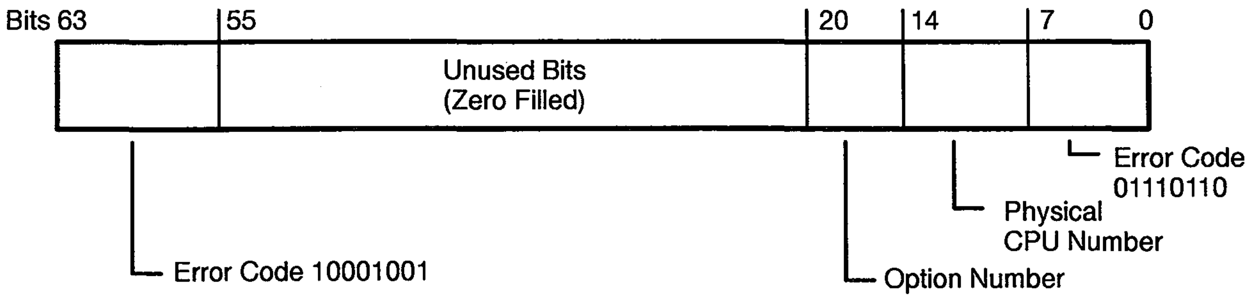

Figure 12. CPU High-speed Register Error Word Format

Physical CPU Number. The physical CPU number assigned at configuration time.

Option Number. All HSR options implement parity check logic on a byte basis. Some options support several data bytes, while other options split the data byte. For example, 1 byte may consist of data bits 0 through 3, and 32 through 35 for 1 parity bit. Table 4 lists 5-bit HSR option types and function codes.

Table 4. HSR Option Number Codes

Option	Code	Option	Code	Option	Code
VR000	001 000	CH000	011 000	IC000	101 000
VR001	001 001	CH001	011 001	IC001	101 001
VR002	001 010	CH002	011 010	IC002	101 010
VR003	001 011	CH003	011 011	IC003	101 011
VR004	001 100	CH004	011 100	-	-
VR005	001 101	CH005	011 101	-	-
VR006	001 110	CH006	011 110	-	-
VR007	001 111	CH007	011 111	-	-
VR008	010 000	CH008	100 000	BT000	110 000
VR009	010 001	CH009	100 001	BT001	110 001
VR010	010 010	CH010	100 010	HM000	110 010
VR011	010 011	CH011	100 011	HM001	110 011
VR012	010 100	CH012	100 100	-	-
VR013	010 101	CH013	100 101	-	-
VR014	010 110	CH014	100 110	-	-
VR015	010 111	CH015	100 111	-	-

IO Module Errors

The following error types are reported to the error logger channel from IO modules. Each of the error types is described in detail after the following error logger routing subsection.

- SECDED errors on I/O data
- I/O channel errors (currently not reported to the error logger)
- Logic monitor (HM options) buffer parity errors (for test-point data)

Error Logger Routing from an IO Module

If an IO module is the master (refer to Figure 13), errors from the IO module are routed directly out the LOSP error logger channel (DD2 interface). If the IO module is not the master, errors are sent to the shared module through the CP module that owns the home port to the IO module. Error data words are transferred serially between all modules in the system until they arrive at the DD2 option. The DD2 option sends errors through the error logger channel to the support system using 16-bit × 4 LOSP transfers.

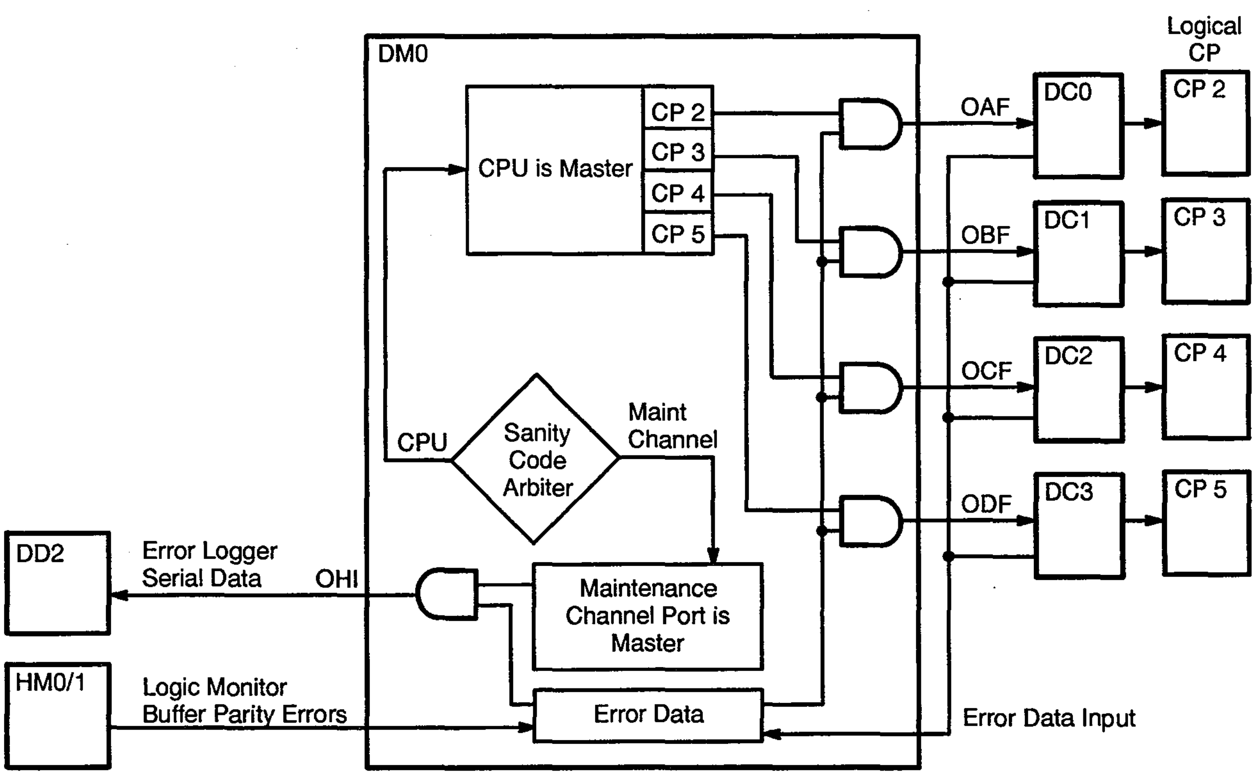

Figure 13. Error Logger Channel Routing from the IO Module

DM Option

The DM option on an IO module performs a 5-way port arbitration to route error data words. The five ports include four CP modules (physical CPs 2, 3, 4, and 5 on CRAY T916 and CRAY T932 systems, and CPs 0, 1, 2, and 3 on CRAY T94 systems) and the DD2 option (interface to the external error logger channel). The source of the sanity tree is the IO module that has the master sanity code generator. The DM option on the IO module that has the master sanity code generator routes error logger data words to the DD2 option (Figure 14).

I/O channel and SECDED errors received from each DC option are routed to the DM option using a rotating priority scheme. The priority order is for errors received from DC0, DC1, DC2, DC3, and then parity errors. After the DM option reports an error from one DC (or a parity error), errors from that DC have the lowest priority, which gives outstanding errors from other DCs a chance to be reported. The order of the priority does not change.

DC Option

The DC option has a 3-bit scanner that checks each type of IO module error to determine whether an error is waiting to be reported. (A 3-bit scanner is used because there are four types of errors detected on the IO module: write errors, read errors, write errors that are detected on the CP module and reported back to the DC, and channel errors.) On even counts, the scanner checks the error logger channel coming from the CP module. On odd counts, it checks one of the IO module error types. If there is an error, the scanner stops counting until the error is reported, and then the scanner continues counting with the next error type. No error type is excluded. As a minimum, every other reported error is from a CP module (assuming that a CP module has errors waiting to be reported).

Each error type has two holding ranks that hold errors waiting to be reported to the error logger channel. The opposite error type has priority; if two single-bit errors are being held and a double-bit error comes in, the double-bit error overwrites the last single-bit error. If two double-bit errors are being held and a single-bit error comes in, the single-bit error overwrites the last double-bit error. If a single-bit error and a double-bit error are being held, the newly arriving error is lost.

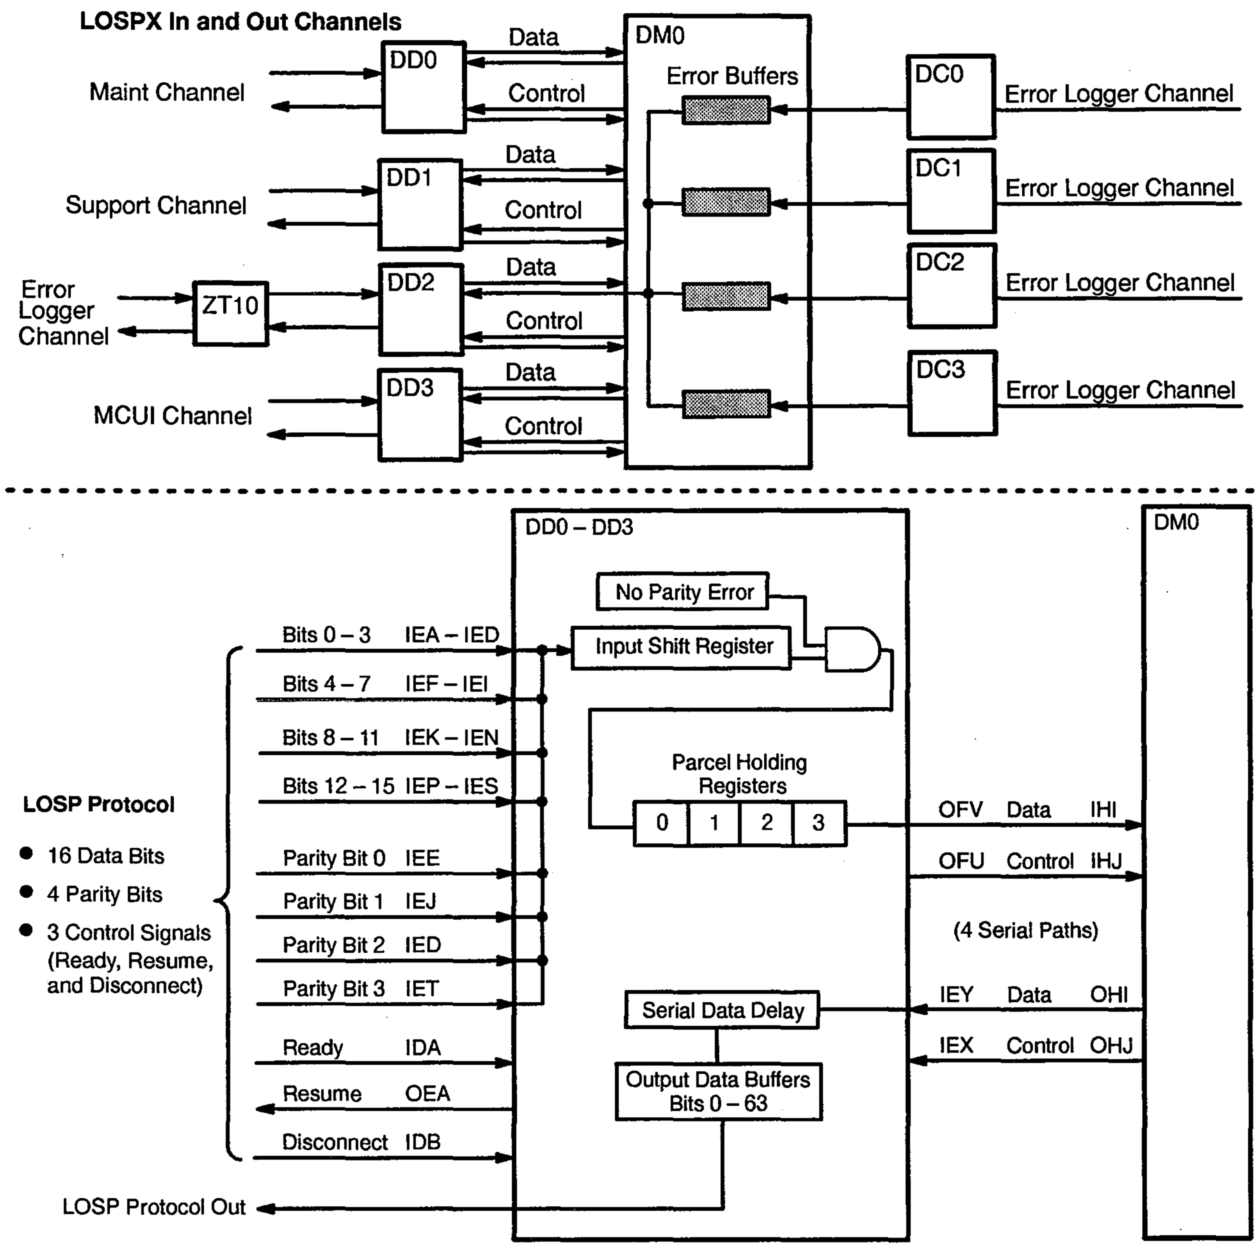

Figure 14. DD Option LOSPX Interface to IO Module

DM-to-DD Protocol

The channel protocol between the DM and the DD options has four lines: control in, data in, control out, and data out. Figure 14 shows the LOSPX interfaces to the DD options.

NOTE: LOSPX refers to either Error Logger, Maintenance, Support, or MCUI LOSP channels that connect to the Support System.

The control lines use the following three 5-bit serial codes to control the sending and receiving of data words to and from the DD options:

- Attention - 35
- Data Tag - 36
- Reply - 37

Control is sent in clock periods 1 through 5. If data is being sent, it is sent starting in clock period 6. There must be at least 11 clock periods between control packets. A data tag is sent before each data word. A reply must be returned for every data tag sent. A new data word is not sent until a reply is received. A reply can be sent any time after a data tag is received.

An attention code must be sent from the DD to the DM to initially clear the channel. An attention code takes the DM out of data mode. All data sent to the DM before the attention code is received will be transferred out the serial maintenance channel.

If the first data word received after an attention code has bits 54 through 63 set, then that word is a boundary scan module function. The DM ignores this word and any word sent after it until another attention code is received. The attention code returns the DM option to normal operation.

SECDED Errors on I/O Data

SECDED is performed on the I/O channels, and errors are detected as follows (refer to Figure 15 and Figure 17):

I/O write data error correction.

I/O write data is corrected twice. Initially, the DA/DC options on the IO module correct the write data. Then, the HA0 and HA1 options on the CP module correct the write data. The HA options report the errors back to the DC option on the IO module on a path shared with I/O commands. These errors do not use the error logger path to the shared module.

I/O read data error correction.

The HA2 and HA3 options on the CP module generate check bits on data being read from memory, but no error correction is performed in these options. The CJ options on the CP module perform SBCDBD. The DA and DC options on the IO module perform SECDED on data read from memory.

I/O Write Data Flow (Input)

I/O data from the channels is sent into the DR options. The DR options assemble channel data into 72-bit words (64 data bits and an 8-bit checkbyte). On LOSP references, parity is checked on the DR option. Data is sent through SECDED logic on the DA and DC options. For LOSP references, the DA and DC options generate a checkbyte before the data is forwarded to the CP module; no SECDED is performed on the DA/DCs for LOSP references. On HISP and VHISP references, data and check bits can be corrected. Check bits are forwarded along with data bits to memory.

After checking data, the DAs forward data bits 00 through 15, 32 through 47, and all 8 check bits. The DCs forward data bits 16 through 31 and 48 through 63. The HB and HC options send 72-bit words to the HA0 and HA1 options every clock period. The "HB/HC to HA0/HA1 Data Word Flow" subsection on page 29 describes how these transfers are made.

Data and check bits coming into the HA0/1 options are written into one of eight 8-word buffers. The 72-bit words then move through SECDED logic. The corrected 64 data bits are next sent through SBCDBD generation, resulting in 76 bits (64 data and 12 check bits). The data and check bits are sent to the CH options and then to the CI options. The CI options are section drivers that transmit control and data from the CP module to the network and memory modules.

NOTE: I/O data going to memory is buffered twice and checked twice: once on the DA/DC options and once on the HA0/1 options.

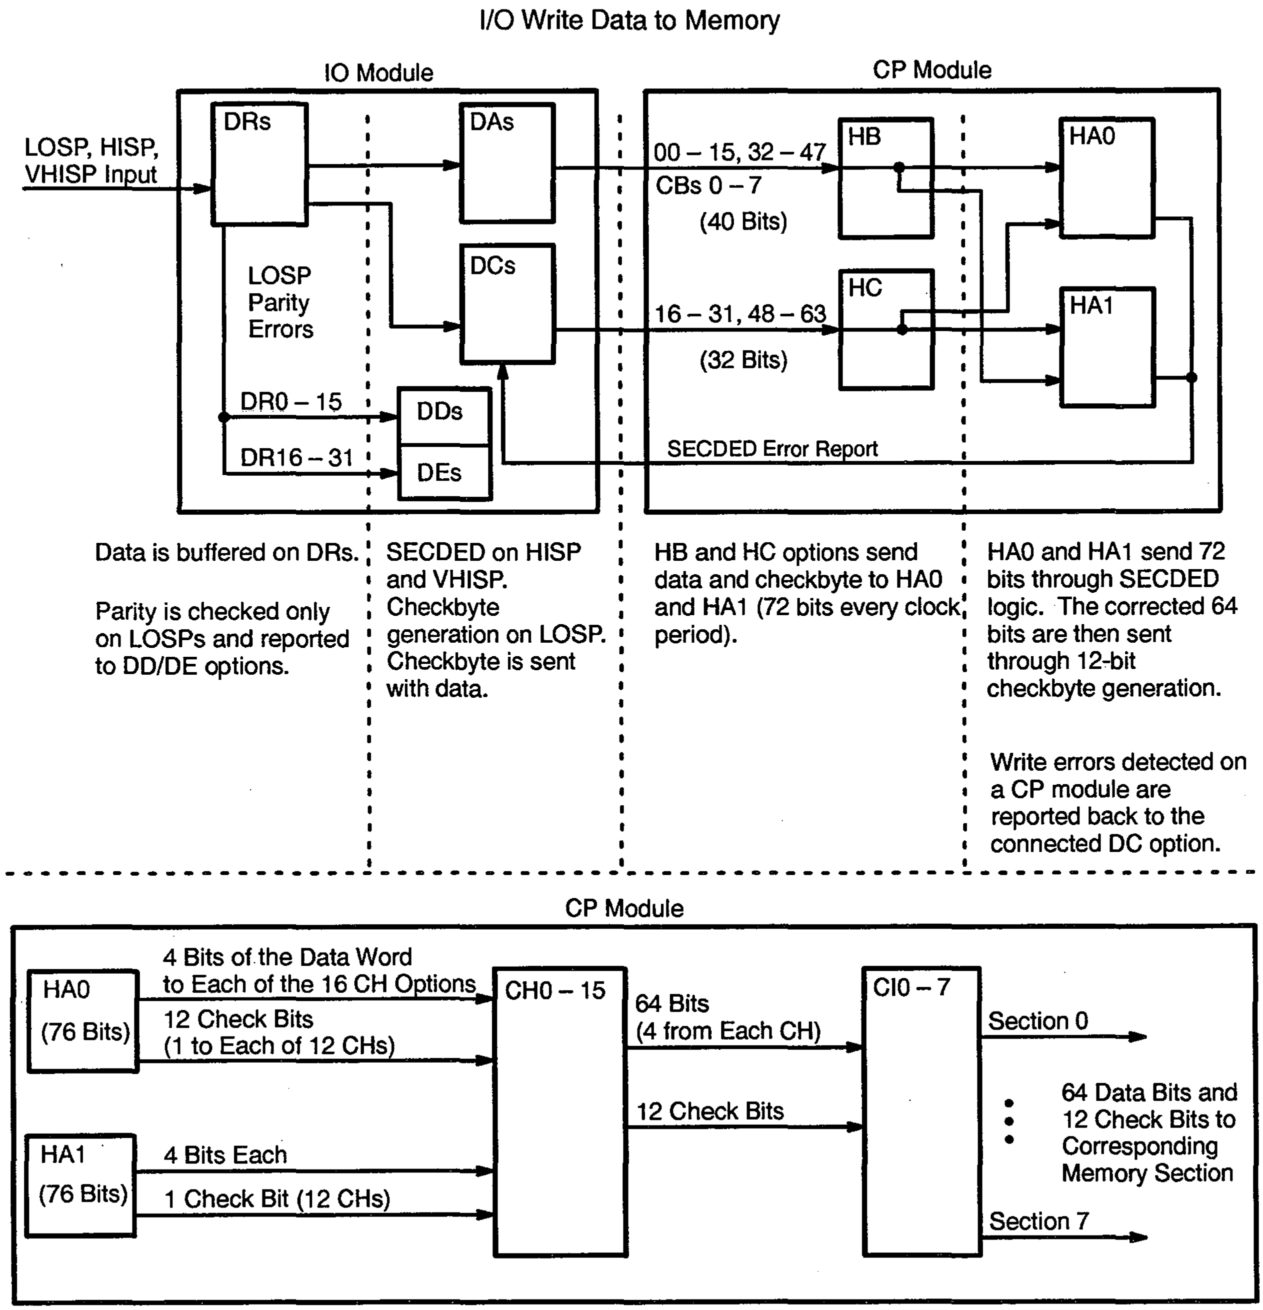

Figure 15. Error Correction on I/O Write Data to Memory2

HB/HC to HA0/HA1 Data Word Flow

The HB and HC options do not have enough spare pins to send out all data and check bits at one time to the HA0 and HA1 options. The HB/HC options send out bits to the HA options using the same sets of output pins on alternate clock periods.

Figure 16 illustrates data transfers from the HB and HC options. The HB/HC options receive a 72-bit word every clock period and send 72 bits to the HA0 and HA1 options every clock period. They alternate sending the upper and lower bits of every other data word. For example, during clock period 3 (refer to Figure 16), HB and HC send the upper bits of word 0 (32 through 63 and CBs 4 through 7) and the lower bits of word 1 (00 through 31 and CBs 0 through 3). The HB and HC act as halfword data paths that run parallel.

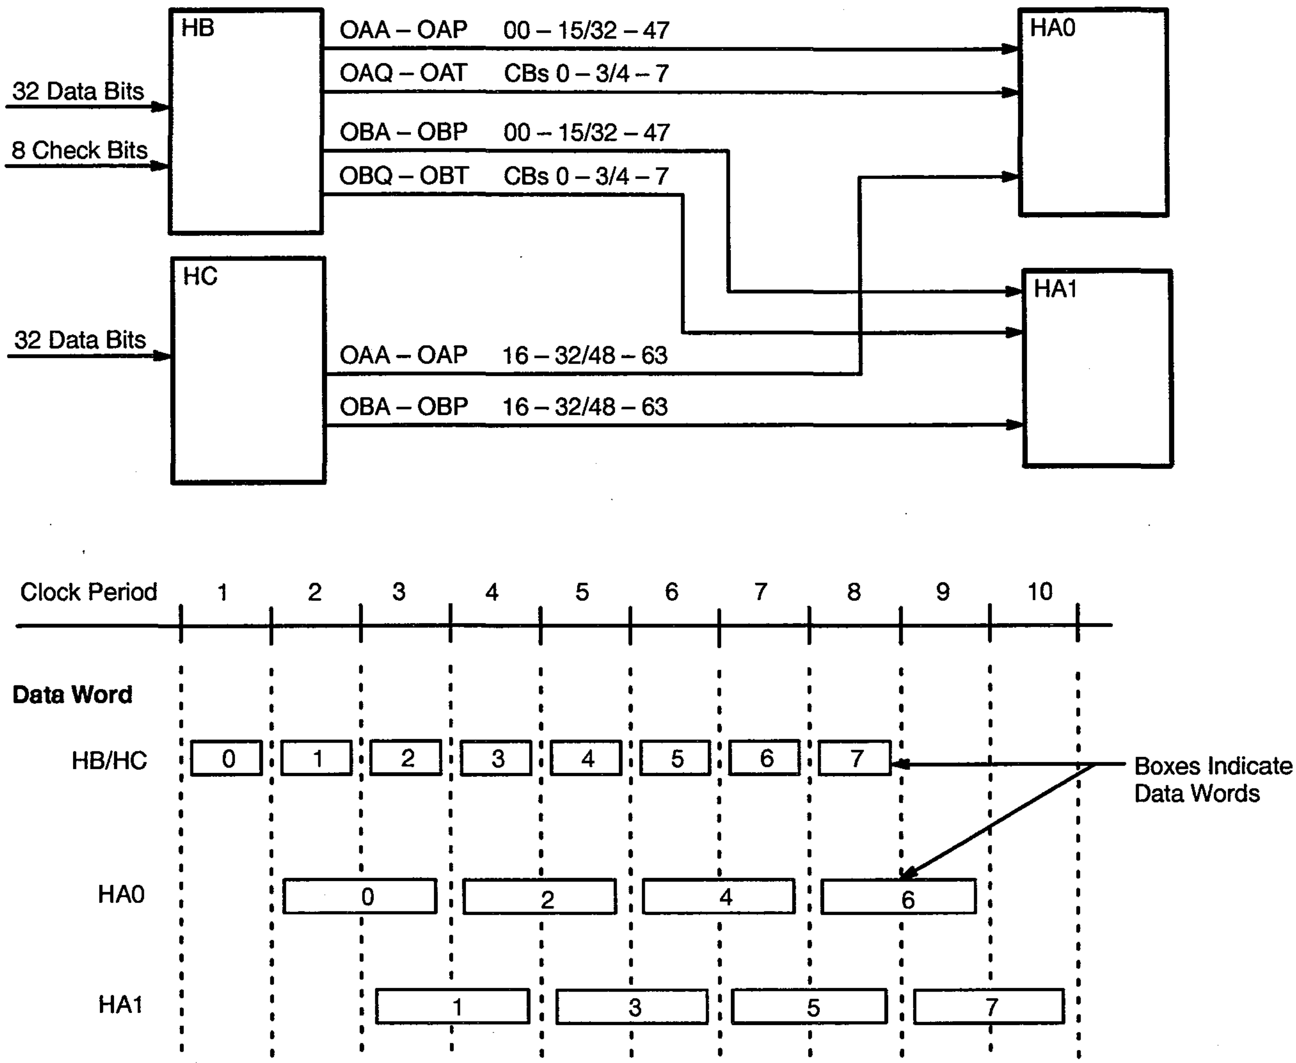

Figure 16. Data Transfers between HB/HC and HA0/1 Options

I/O Read Data Flow (Output)

Each CP module has 8 CJ options, one for each section of memory. The CJ options perform SBCDBD for data read from memory, as shown in Figure 17. The CJs use the 12 check bits to perform error correction on each adjacent 2 bits of the 64-bit data word. The CJs drop the 12 check bits and send the corrected 64 data bits to the CH options; 8 bits are sent to each of 8 different CH options. The CH options send all 64 bits to both the HA2 and HA3 options.

The HA2 and HA3 options each receive all 64 data bits in order to generate 8 check bits. Odd data and check bits are sent to a DA option, and even data and check bits are sent to a DC option, as shown in Figure 17.

I/O Read Data from Memory

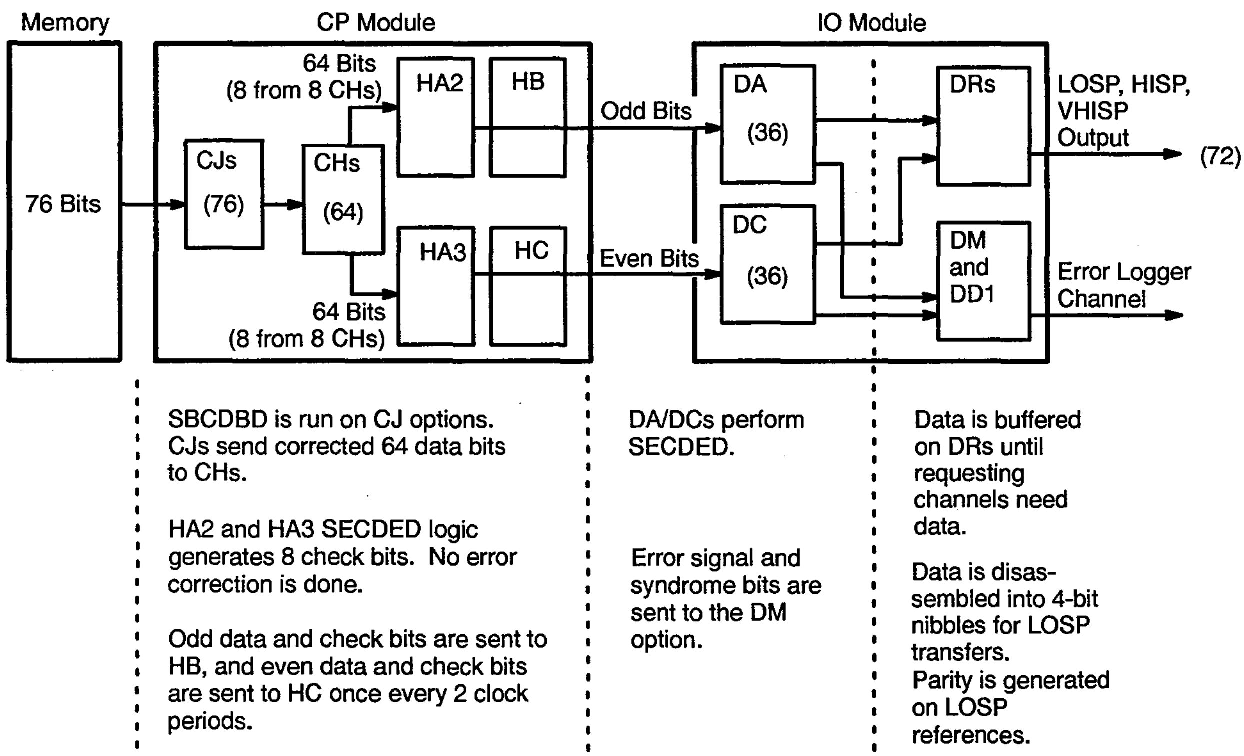

Figure 17. Error Correction on I/O Read Data from Memory

The DA/DCs pipe the data read from memory through SECDED logic that generates a checkbyte and performs error correction. If any uncorrectable memory errors are detected, an uncorrectable memory error signal and the associated syndrome bits are sent to the DM option.

Data leaves the DA and DC options and passes to the DR options. The DRs disassemble data into 4-bit nibbles. LOSP data parity is generated on each 4-bit nibble (1 parity bit per nibble). HISP and VHISP data is assembled in buffers that contain 16 data bits and 2 check bits. LOSP data transfers (20 bits) are sent out from 5 DR options. HISP 72-bit data transfers are sent out from 4 DR options ($16 + 2 = 18$ bits; $18 \times 4 = 72$ bits). VHISP channels transfer two 72-bit data words per clock period ($72 \times 2 = 144$ bits).

Data is buffered on the DR options until it is needed by the requesting channel, unless the data is being read from the support channel. For LOSPX support channel references, data and check bits are sent serially to the DM option and then to the DD1 option. On a LOSP reference, check bits are ignored and parity is generated on a DR option.

I/O Channel Errors

NOTE: The following I/O channel error-detection feature was designed to be used with future IOS/SSD systems. This feature was developed to enable the IO module to report these errors to the error logger to give users more information. Currently, these errors are reported only through software. The DM option has the control logic to perform this function. The DD/DE options have assigned output pins and signal paths to the DM. However, the DD/DE options do not contain the logic to perform this function; these options would need to be changed once new channel specifications are defined. This error-reporting feature was designed to minimize changes that would be needed for future systems.

The DD and DE options detect parity errors, block length errors, control sequence (ready, resume, and disconnect) errors, and other types of channel errors and report them to the error logger channel, as shown in Figure 18.

The DC options receive I/O channel errors from the corresponding DD or DE option and add them to the error logger channel. Channel errors are sent on the error logger channel to the DM option. The DM sends the error to the master IO module, which is the source that provided the sanity code (physical CP 2, 3, 4, 5, or the sanity code generator in the DM option). The DM option detects the master CPU and routes the errors to that CP module (physical CP 2, 3, 4, or 5). The error logger then passes through the HC option before it goes to the HG option. Errors are sent out externally on the error logger channel (rather than through a CP module) if the DM is on the IO module that owns the home port.

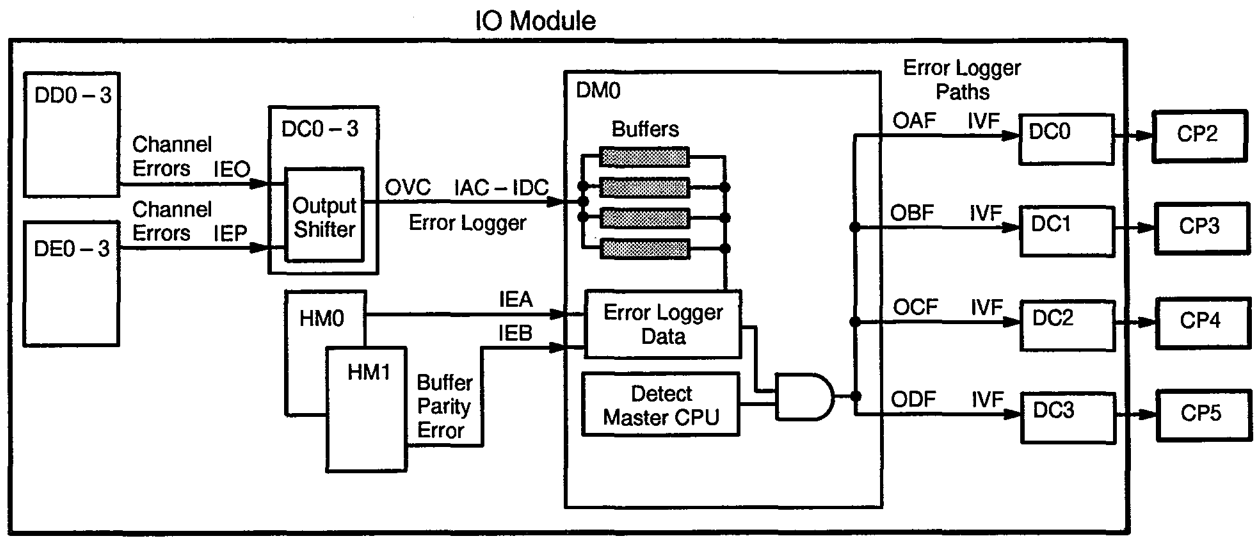

Figure 18. I/O Channel and Logic Monitor Parity Errors

I/O Logic Monitor Parity Errors

The DM option also adds parity errors sent from any of the eight buffers on each HM option to the error logger. The HM options make up the IO module logic monitor.

IO Module Error Words

Figure 19 and Figure 20 show the format of IO module error words.

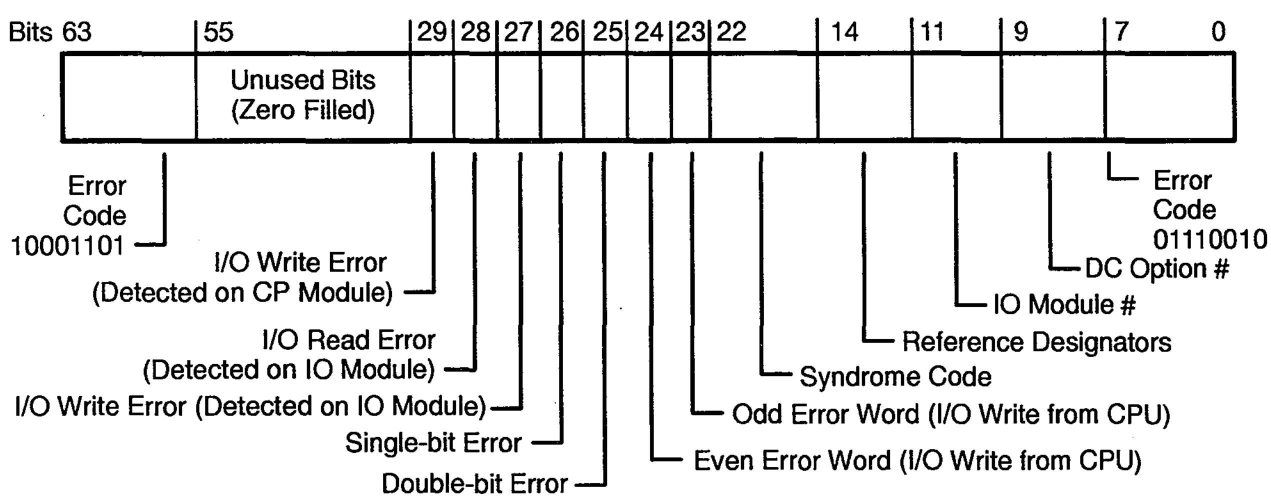

Figure 19. I/O Data SECDED Error Word Format

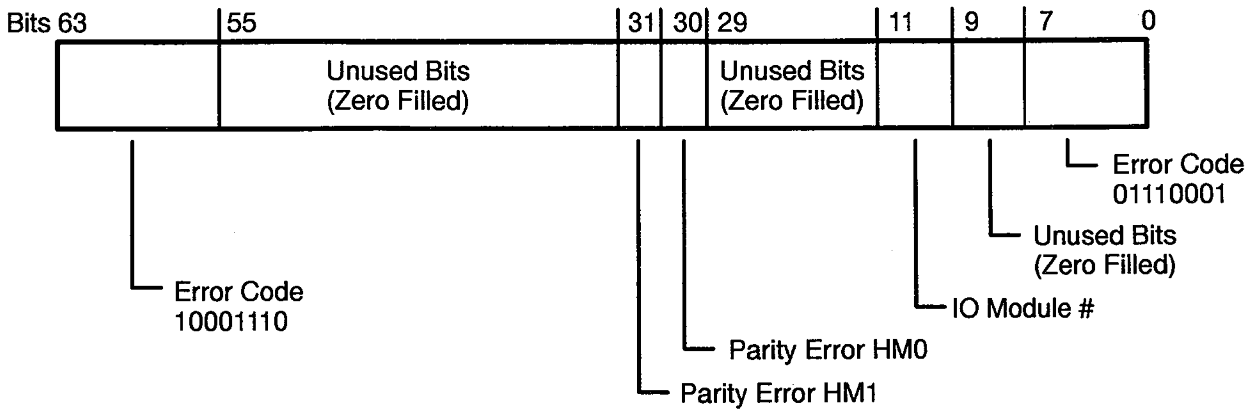

Figure 20. HM Parity Error Word Format

Shared Module Errors

The shared module receives errors from attached CP modules and from the remote shared module (in CRAY T932 systems) as shown in Figure 21. Errors from IO modules are sent to the shared module through one of the physical CP modules numbered 2, 3, 4, or 5 (CP 0 through 3 in CRAY T94 systems). The shared module may also receive parity errors from the HM options (shared module logic monitor) or from the SR options (shared B and T registers). Error in and out paths are shown in detail in Figure 25 through Figure 27.

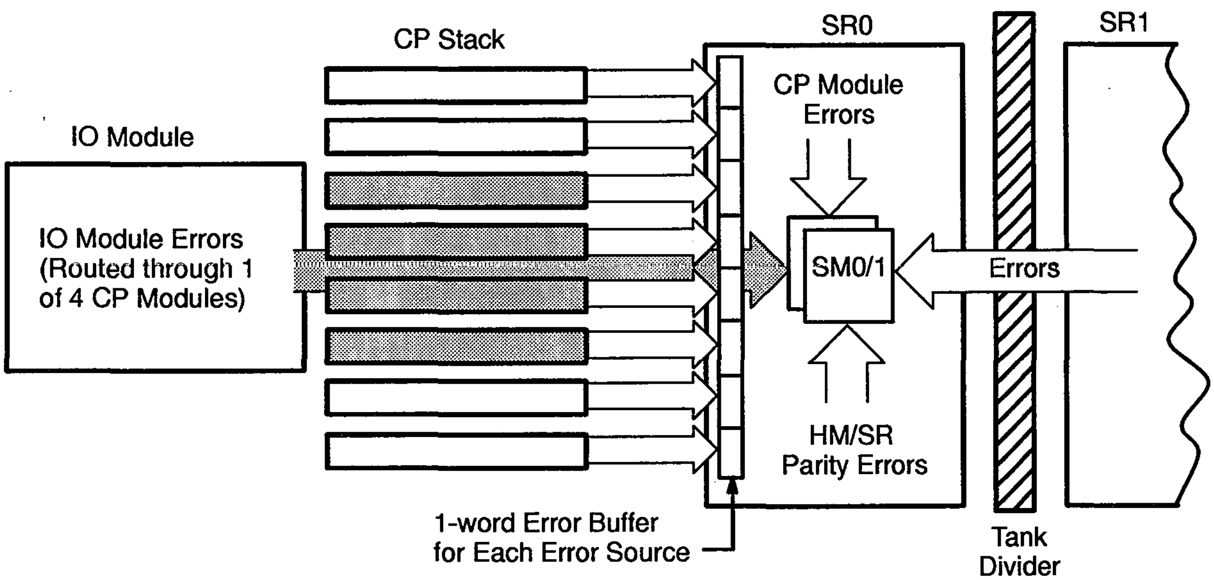

Figure 21. Shared Module Error Input Sources

Error Logger Routing from the Shared Module

The various input sources of error logger messages to the shared module arbitrate for the Error Logger path back to the home port on the IO module where the sanity tree was initiated. If the home port is on an IO module on the other side of the CRAY T932 chassis, errors are sent to the other shared module (from SR1 to SR0 in Figure 21).

Figure 22 shows the various error sources entering the SM options. Each SM option must arbitrate between these sources for the Error Logger path back to the home port on the IO module.

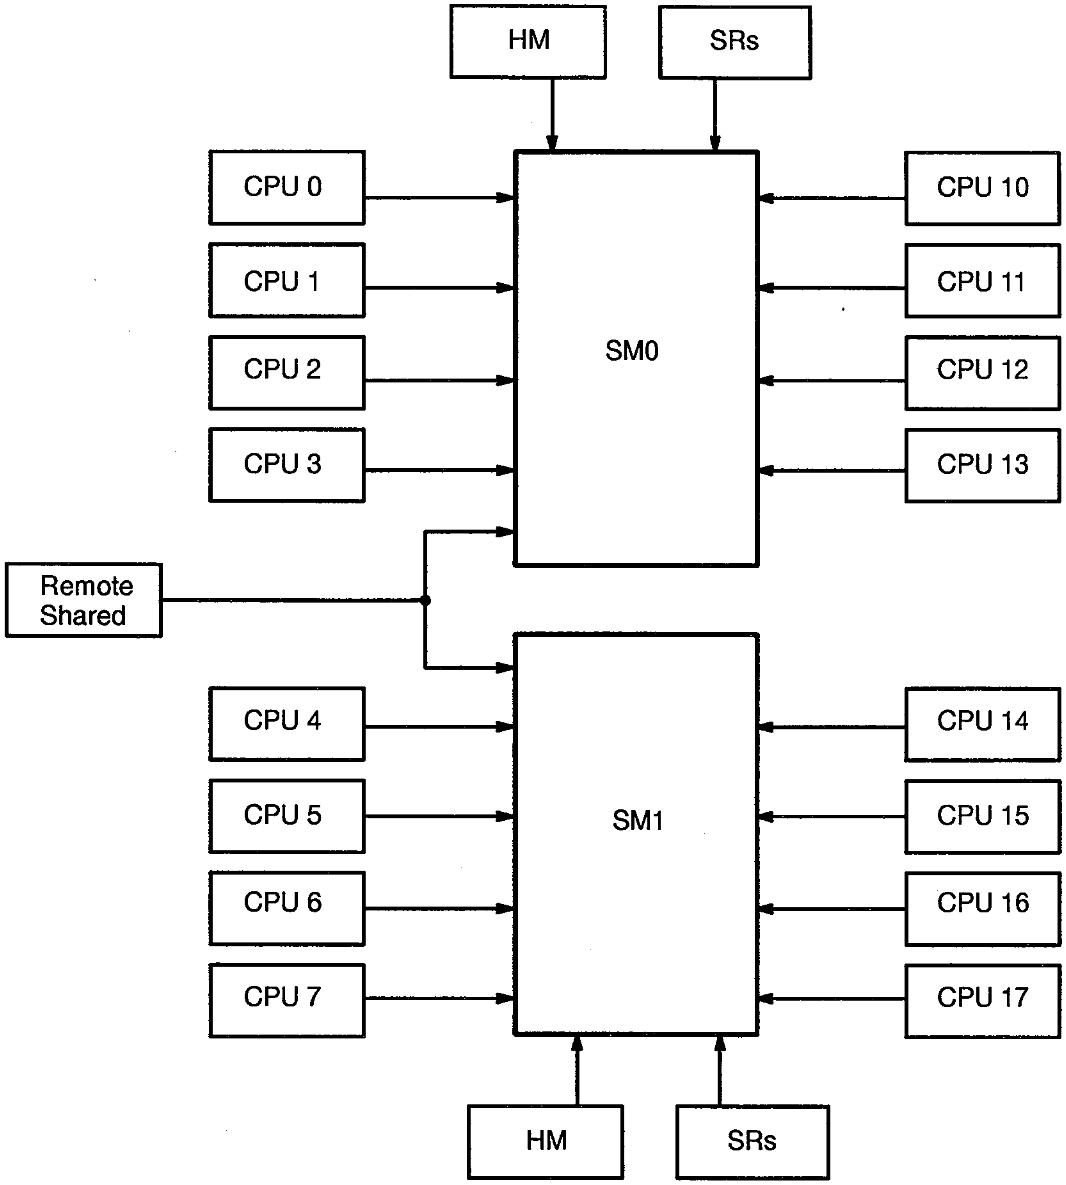

Figure 22. Error Input Sources to the SM Options

Each SM option arbitrates among nine sources to send out error logger messages as shown in Figure 22. There are eight sources for local CPs (four possible from each of two stacks) and one source for the logic monitor (HM) and shared register (SR) options. The port for each error logger input has a 1-word queue for waiting error messages.

The SM option sends a new error message to the home port only after all of the three following situations occur: (1) An Error Logger Acknowledge start bit has been received, (2) 80 clock periods have elapsed since the last error message was initiated, and (3) there is a new message to send.

Arbitration Process

When an error message arrives at the SM option, it tries once to arbitrate for output; if it is successful, it sends the error word to the home port, bypassing the queue. If arbitration is not successful, the entire 64-bit error word is loaded into a shift register queue, and the request process is repeated until arbitration is granted. The 64 bits of error data and 5-bit 37 code are sent through a CP module on the way to the home port located on an IO module.

Arbitration is weighted to ensure that each error source has an equal chance of reporting errors. The remote shared module is given priority 50% of the time to give the other half of a CRAY T932 system equal time to report errors. The other SM option is given priority 25% of the time, and the remaining 25% of the time is divided equally among the eight CP modules, the logic monitor (HM option), and the shared registers (SR options) that are connected to the SM option.

Error data from the SR options is sent as a 5-bit packet to the SM option (a start bit followed by 4 data bits). No further data is expected from that SR option until the SM sends a Cluster Error Logger Acknowledge signal to that SR option.

If the SM receives a start bit and all 4 data bits are 0's, it sends an error message with those 0 bits. The first bit received must be a 1. At bit position 72, the start bit mechanism is re-armed, so messages can be received once every 72 clock periods. Errors can be sent only once every 80 clock periods from the SM. The SM queues errors as necessary.

Error Logger Acknowledge Messages

As soon as an error logger message is sent from the SM option (either to the home port or to the other SM option on its way to the home port), an Error Logger Acknowledge signal is generated and sent to the CP module (or remote shared module) that sent the error message. (However, no Error Logger Acknowledge signal is generated if an error occurs within the SM option.) When the home port receives the error word, it sends an acknowledgement (5-bit 37 code followed by 64 zeroes) on the Error Logger Acknowledge path.

Because Error Logger Acknowledge messages occur in synchronization with error messages being sent to the home port, they can occur, at most, once every 80 clock periods. The requesting CP module (or remote shared module) can generate a new error message as soon as it begins to receive an Error Logger Acknowledge message. The CP or remote shared module can send a new error message into the 1-word queue while simultaneously shifting out the last error message.

The SM option waits for an Error Logger Acknowledge message after it sends an error message to the home port. Because Error Logger Acknowledge messages are not passed by the SM option, it is not possible for them to have any content.

Shared Module Error Words

Figure 23 shows the error word format for shared module errors.

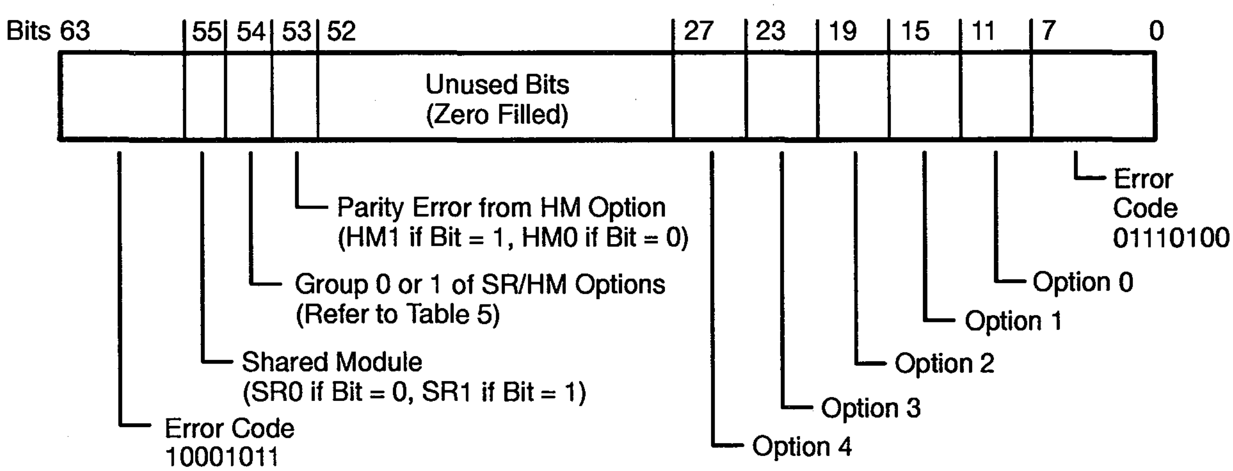

Figure 23. Shared Module Error Word Format

Table 5 identifies the failing SR or HM option when bit 54 is a 0 or 1. Figure 24 identifies the section of bits in which the error occurred.

Table 5. Failing Option Identification

Field	Bits	Group 0 (Bit 54 = 0)	Group 1 (Bit 54 = 1)
Option 0	8 – 11	SR00	SR05
Option 1	12 – 15	SR01	SR06
Option 2	16 – 19	SR02	SR07
Option 3	20 – 23	SR03	SR08
Option 4	24 – 27	SR04	—
HM Error	53	HM00	HM01

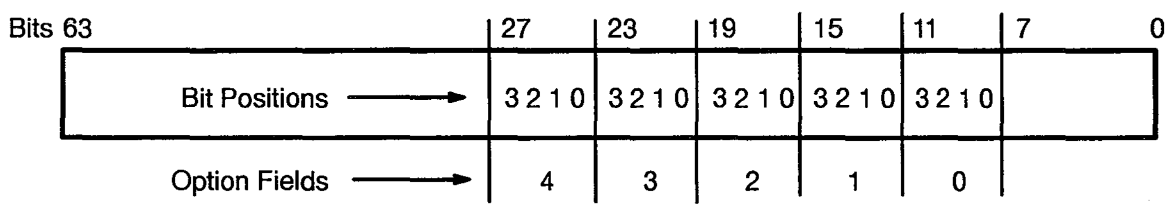

NOTES: The bit position within each option field indicates a failure in the specified section of data bits.

For example, if bit 13 is set (option 1 field, bit position 1), it indicates a failure in data bits 32 – 47.

Bit Position	Failing Data Bits
0	48 – 63
1	32 – 47
2	16 – 31
3	00 – 15

Figure 24. Shared Module Error Failing Bit Identification

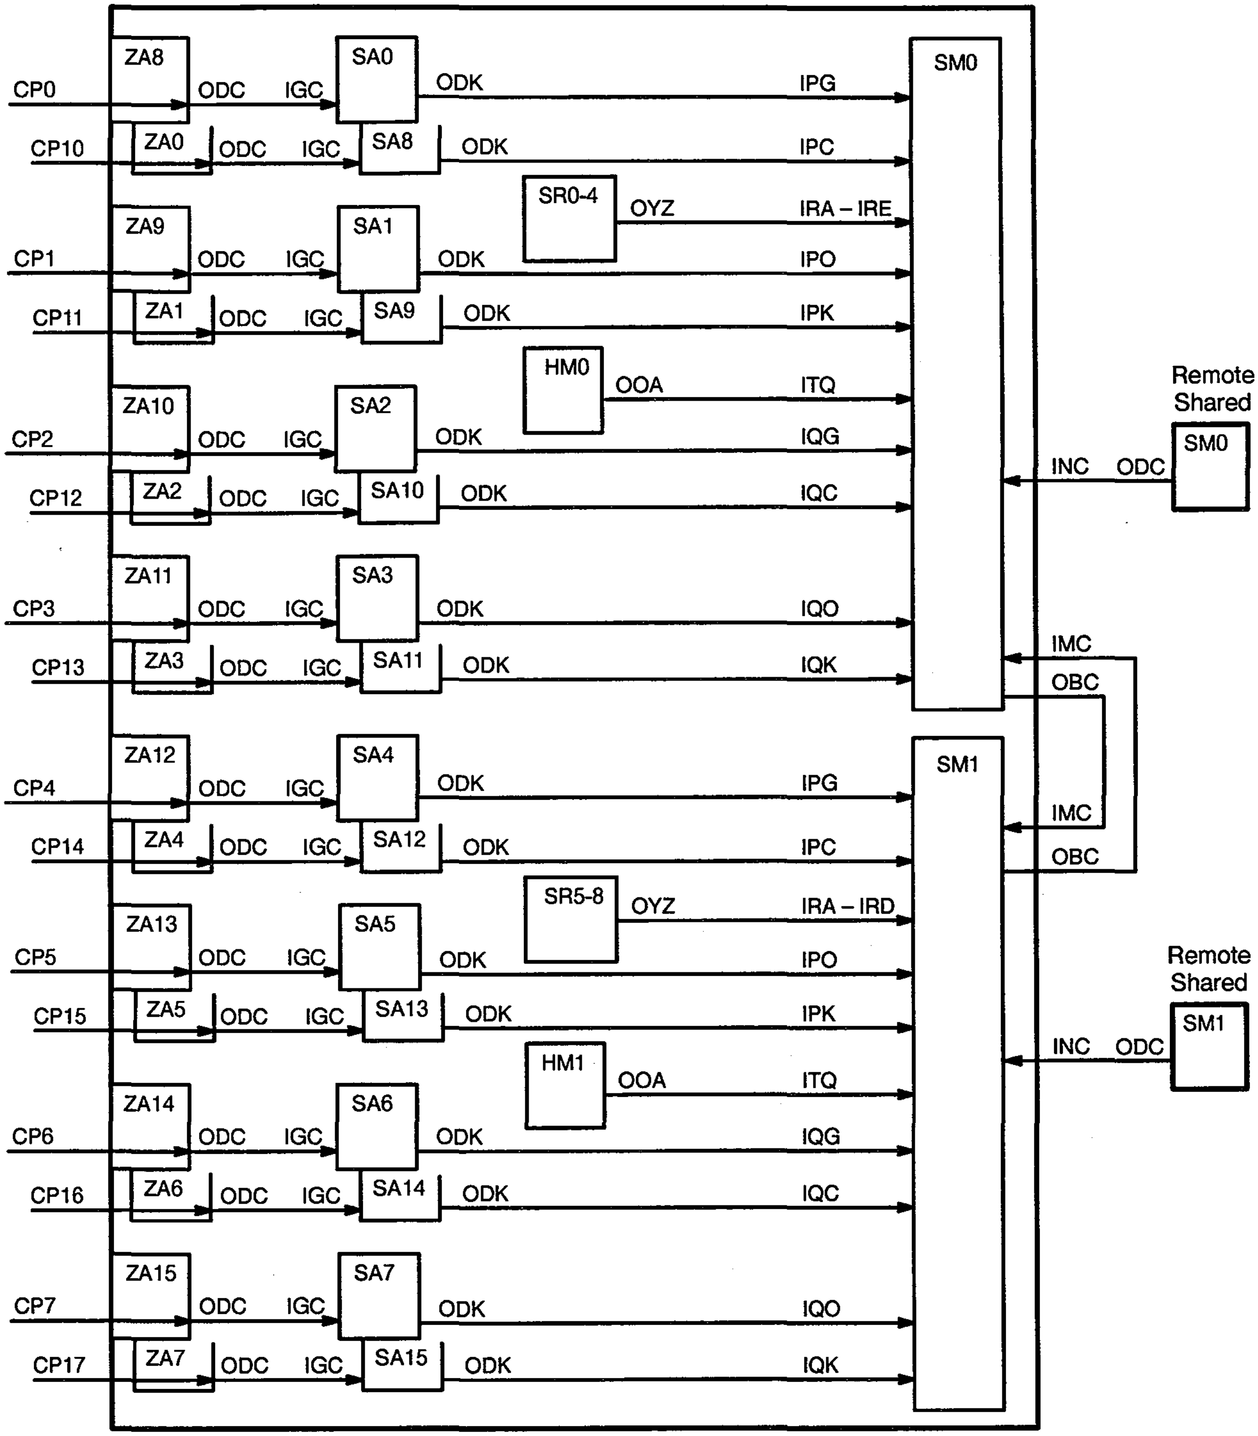

NOTES: The open boxes indicate options on board 2 of the shared module.

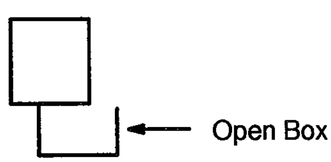

Figure 25. SM Option Error Logger Input Sources

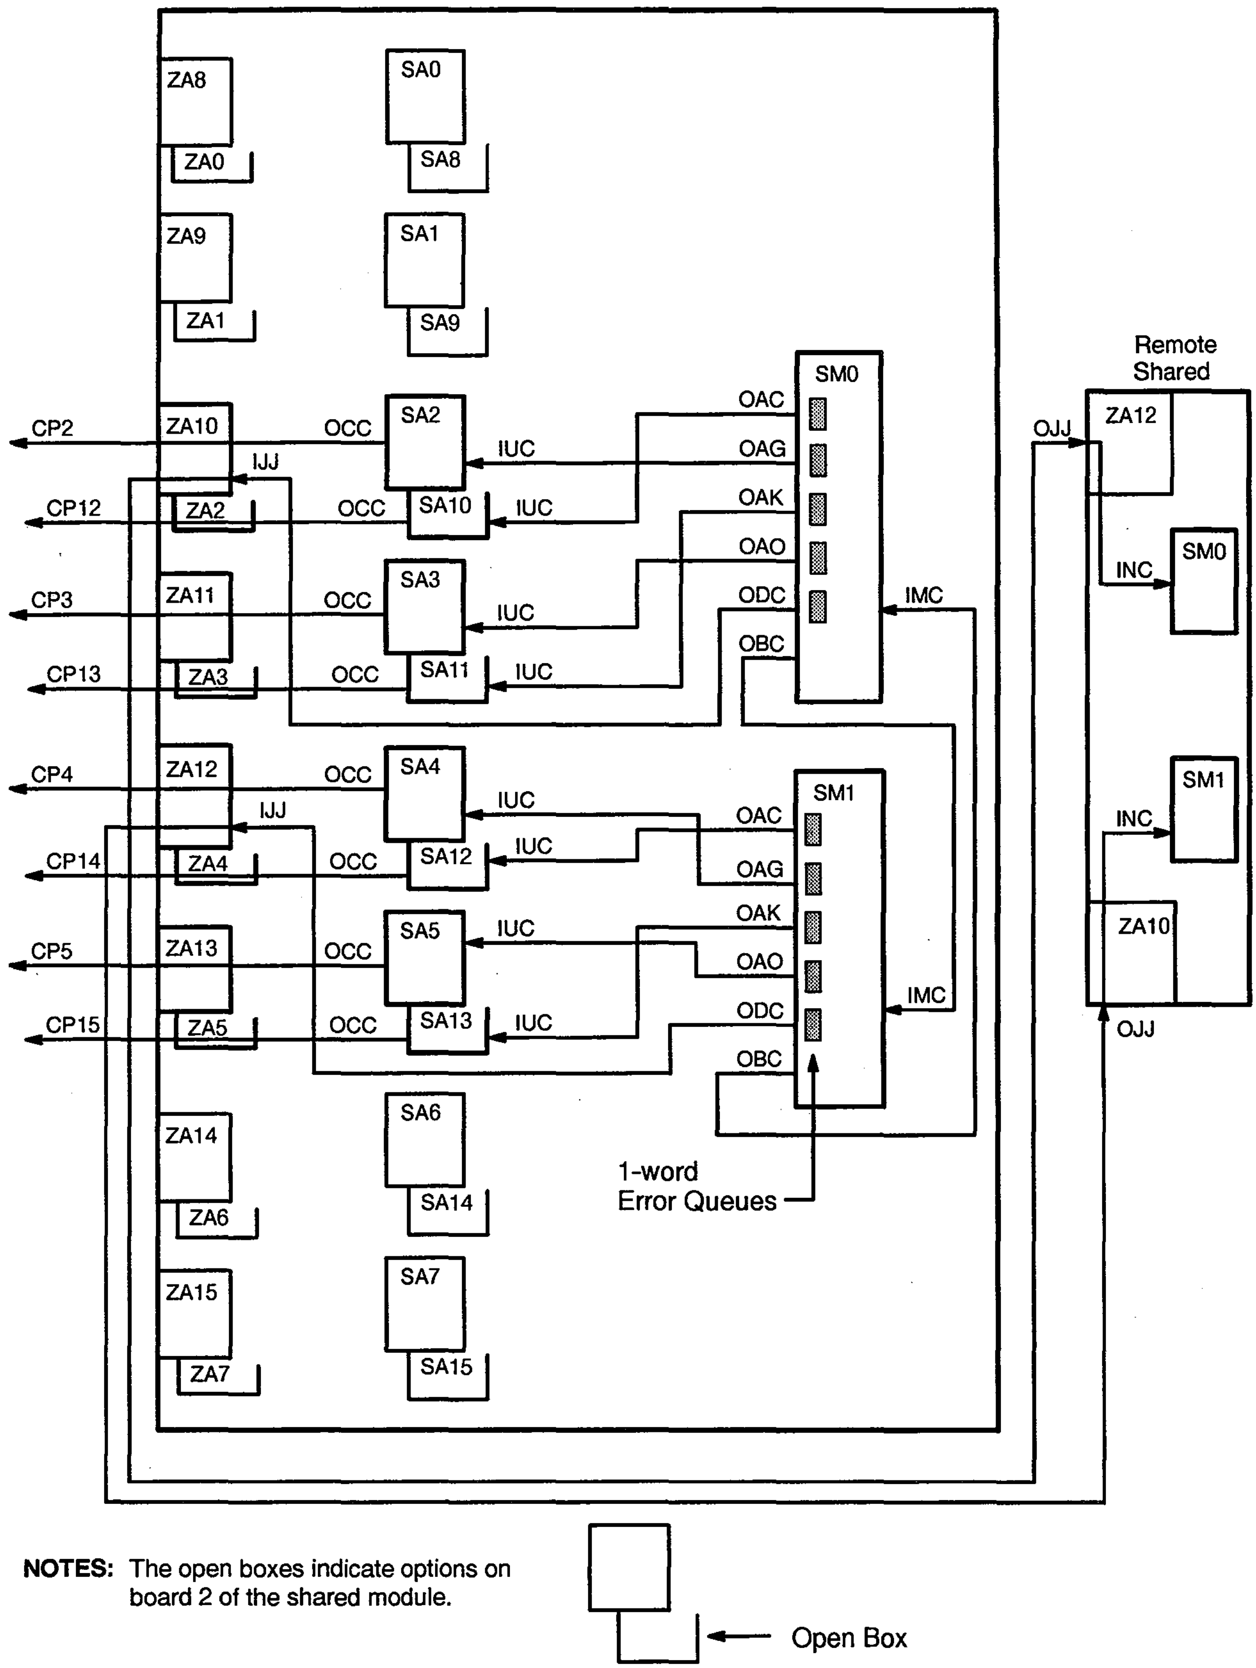

NOTES: The open boxes indicate options on board 2 of the shared module.

Figure 26. SM Option Error Logger Output Paths

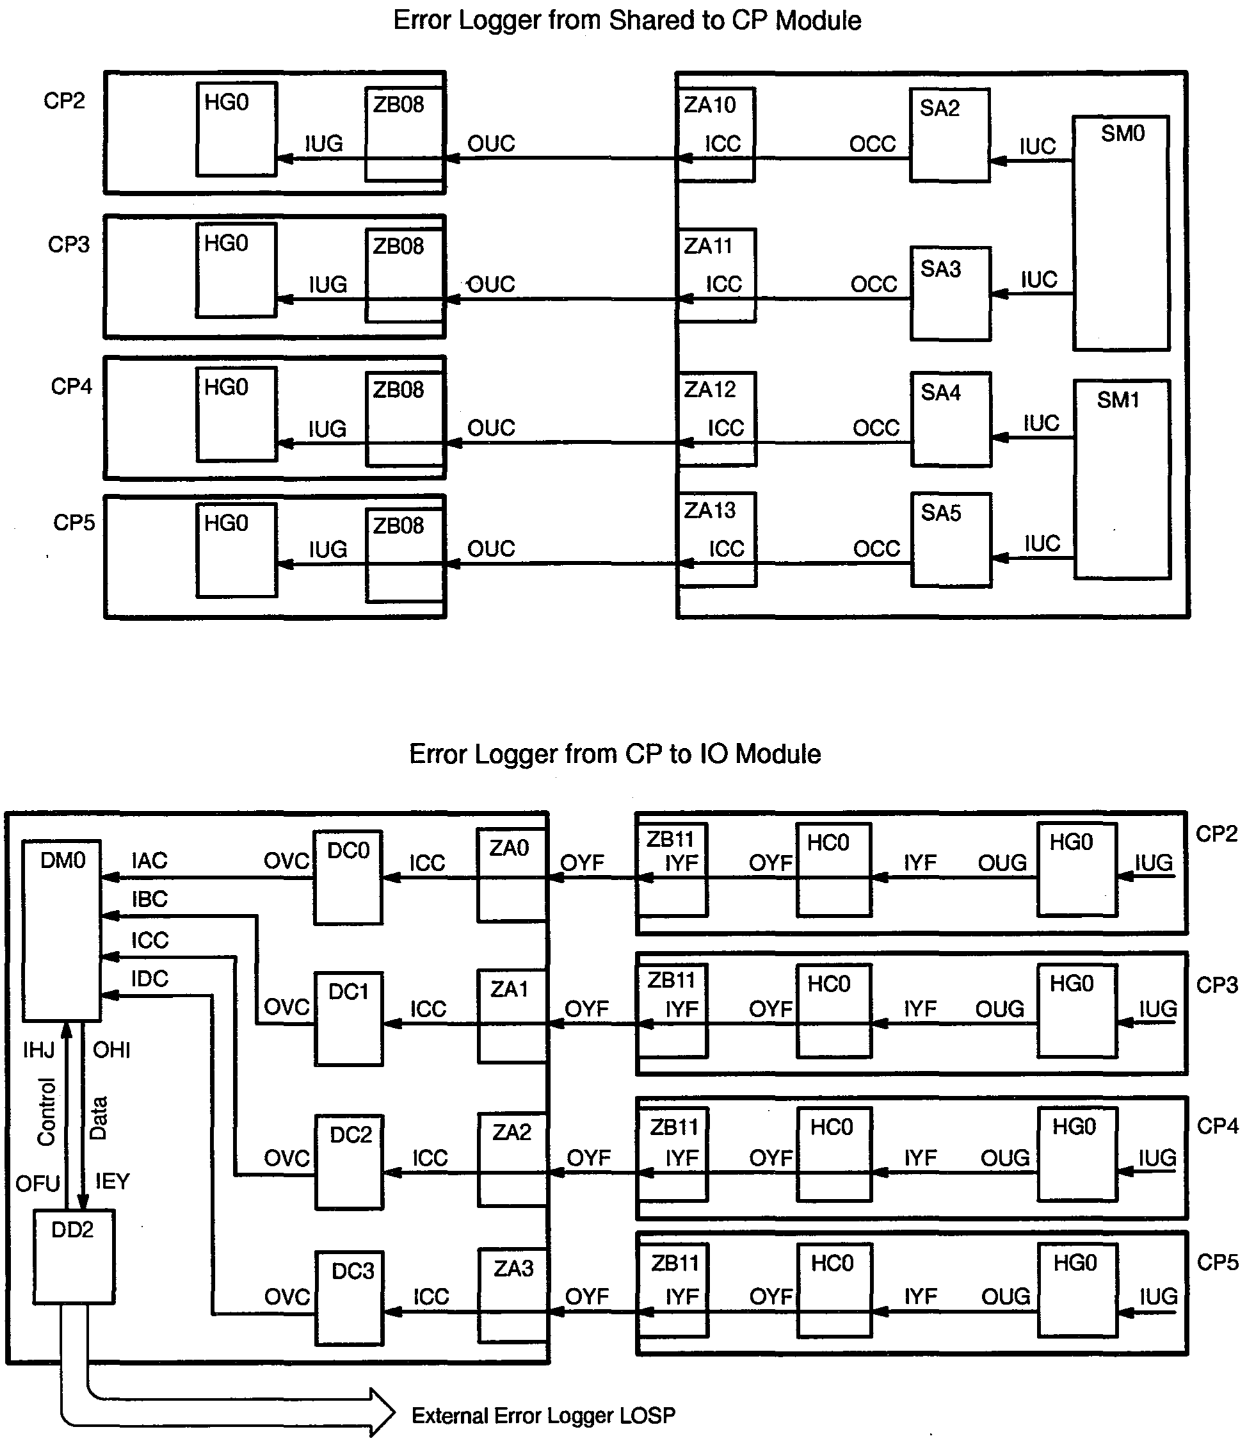

NOTES: These figures show error logger output paths through only physical CP2 – CP5. Only the ZA and ZB connectors will differ for other CP paths; the I and O terms are the same.

Figure 27. Error Logger Paths Out from the Shared Module

Reader Comment Form

Title: Error Logger Channel
(CRAY T90™ Series)

Number: HTM-007-B

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 and explain your answer)

______Too much information ______Too little information ______Correct amount

____
____
____

You may write additional comments in the space below and mail them to the address on the back of this form, or fax them to us at 715-726-4353. You may also E-mail your comments to us at *hpt@cray.com.* When possible, please give specific page and paragraph references. We will respond to your comments in writing within 48 hours.

____
____
____
____
____

NAME ______________________________
JOB TITLE______________________________
E-MAIL ADDRESS______________________________
SITE/LOCATION______________________________
TELEPHONE______________________________
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