Memory Module (CM03)
-184-

Memory Module (CM03)

(CRAY T90™ Series)

HTM-184-0

Cray Research Proprietary

Cray Research, Inc.

Record of Revision

REVISION	DESCRIPTION
	June 1996. Original printing.

Any shipment to a country outside of the United States requires a letter of assurance from Cray Research, Inc.

Requests for copies of Cray Research, Inc. publications should be directed to:

CRAY RESEARCH, INC.
Customer Service Logistics
1100 Lowater Road
P.O. Box 4000
Chippewa Falls, WI 54729-0078
USA

Comments about this publication should be directed to:

CRAY RESEARCH, INC.
Service Publications and Training
890 Industrial Blvd.
P.O. Box 4000
Chippewa Falls, WI 54729-0078
USA

MEMORY MODULE (CM03)

Figures

Figures (continued)

Figures (continued)

Tables

MEMORY MODULE (CM03)

CM03 Module Description

Crahy Research designed the CM03 memory module to replace the CM02 memory module in a CRAY T90 series system. The CM03 can be built with either **asynchronous** SRAM chips, as in the CM02, or with **synchronous** SRAM chips. When populated with asynchronous SRAMs, the performance of the CM03 is nearly identical to that of the CM02. However, when populated with synchronous SRAMs, the performance of the CM03 is increased considerably.

In synchronous mode, the CM03 module has 16 banks in each of two sections of memory. Each section of memory can be accessed by four CPUs, and the memory array in each section can provide up to 4 words per clock period. Thus, on average, each CPU can get 1 memory word per clock period.

Like the CM02 memory module, a write reference requires a 2-packet transfer, but a read reference requires only 1 packet.

Error detection and correction is done on the CPU module.

Capacity

The CM03 module has the same capacity as the CM02 module. When populated with **synchronous memory chips** the CM03 has 32 banks of 1 million locations per bank, for a total of 32 million words per module. When populated with **asynchronous memory chips** the CM03 has 16 banks of 2 million locations per bank, also for a total of 32 million words.

In synchronous mode, the memory chip is used as 1 million locations by 4 bits. In the asynchronous mode, the memory chip is used as 2 million locations by 2 bits.

Memory Organization

Central memory is organized into sections, subsections and banks. Remember, however, that the CRAY T94 system does not have subsections. All CRAY T90 series computers normally have 8 sections of memory; but memory may be degraded with the SCE program by logically eliminating failing components until time is scheduled to replace the failing module.

Section

A memory section is the range of components that can be the destination of a request from a CPU through a single path. The components that compose the section are spread out over all the memory modules that compose the module stack in the CRAY T916 and CRAY T932 systems. The components that compose the section in the CRAY T94 system are all located on one module. In fact, each memory module in the CRAY T94 has 2 full sections on it.

In the CRAY T94 system, the CPU uses the section bits of the address to steer the memory reference to a particular connector on the appropriate memory module. In the CRAY T916 and CRAY T932 systems, the section bits are used to steer the reference to the appropriate network module.

The section is made up of 8 subsections in all models except the CRAY T94 system, which has no subsections.

Subsection

There is a separate path from the network module to each subsection. The network module decodes the subsection bits and steers the memory reference to the appropriate memory module once any subsection conflicts that may exist have been resolved.

Each subsection contains 8 or 16 banks when the module is populated with synchronous memory chips, depending on the model of mainframe. When populated with asynchronous memory chips, each subsection contains 4 or 8 banks.

Bank

A bank contains either 2 million or 1 million 76-bit words (64 data and 12 check bits), depending on whether it is populated with asynchronous or synchronous memory chips. A bank always resides on one module.

Memory Configurations

The three types of mainframes (CRAY T932, CRAY T916, and CRAY T94) normally have 8 sections of memory.

- A CRAY T94 system includes only one memory module stack; the module stack contains either two or four memory modules. Refer to page 14 for a more detailed explanation.

- A CRAY T916 system includes two module stacks; each module stack contains either four or eight memory modules. Refer to page 20 for a more detailed explanation.

- A CRAY T932 system can include two or four module stacks; each module stack contains either two, four, or eight memory modules. Refer to page 26 for a more detailed explanation.

Table 1 shows the various memory configurations of the three types of systems. The numbers in this table do not reflect partitioned or degraded memory. The number of banks is based on the use of synchronous memory chips. If the module is populated with asynchronous memory chips, the number of banks is halved. Remember that banks populated with synchronous memory chips are 1 million words and that banks populated with asynchronous memory chips are 2 million words.

Table 1. Memory Configurations by System Type

Model	CRAY T94	CRAY T916	CRAY T932
Sections	8 or 4	8	8
Subsections/section	1	8 or 4	8 or 4 or 2
Banks/subsection	16	8	16
Total banks/system	128 or 64	512 or 256	1024 or 512 or 256

Memory Partitioning and Degradation

When you degrade memory, you are bypassing areas in memory where failures occur so that the system can continue to operate. By degrading memory, you are forcing selected section, subsection, and bank bits. Therefore, the memory address referenced by the CPU is different from the address received by memory. For specifics and examples, refer to “Memory Degradation” in the *SCE User Guide*, publication number HDM-069-A.

Software address mapping through the maintenance channel controls memory configuration. The System Configuration Environment (SCE) provides the interface for selecting sections, subsections, banks, and groups in a particular configuration to control memory degrades and/or logical memory partitions.

Memory Address Mapping

Memory addressing depends on the configuration of the CRAY T90 series system. The CP, network, and memory modules determine how memory addressing occurs. Table 2 describes the address map used in the CRAY T90 series CPU. This address map accommodates future CRAY T90 series systems and is capable of addressing up to 16 Gwords of memory. The “CRAY T94 Memory,” “CRAY T916 Memory,” and “CRAY T932 Memory” subsections include address maps for the specific CRAY T90 series systems.

Table 2. Address Map

Address Bits	Function
0 – 2	Section select
3 – 5	Subsection select
6 – 9	Bank select
10 – 33	Word select †

SCE uses bits 10 and 11 as group profile and select bits 0 and 1 when partitioning memory.

Module Components

The CM03 memory module contains the central memory that is common to all processors in a CRAY T90 series mainframe. The module consists of a printed circuit board (PCB) with memory stacks and logic options on board 1 and logic options on board 2. Actually, the PCB is composed of two PCBs laminated together. The term *board 1* refers to one side of the PCB, and *board 2* refers to the other side of the PCB. Memory stacks contain the memory chips that store the data. Refer to Figure 1 and Figure 2 for drawings of the boards.

Memory modules are arranged in stacks within the mainframe. This is referred to as a **memory module stack**. Do not confuse this with a **memory stack**, which is a stack of memory chips that is mounted on board 1 of the memory module.

Logic Options and Connectors

There are five types of logic options on the CM03 module. They are:

- 16 MF options (8 per section), which interface with the network or CPU modules
- 16 MC options (8 per section), which steer the address, data and control to one of four 4-bank groups
- 32 MB options (16 per section), which steer address and data to one of four banks in the group
- 1 MZ option for maintenance functions
- 1 TZ option for clock fanout

There are three types of logic connectors. They are:

- 8 ZA connectors, between the network or CPU modules
- 8 ZC connectors, between the memory stacks in one of the two sections on the module
- 8 ZD connectors, between the memory stacks in the other section on the module

Figure 1. CM03 Memory Module – Board 1

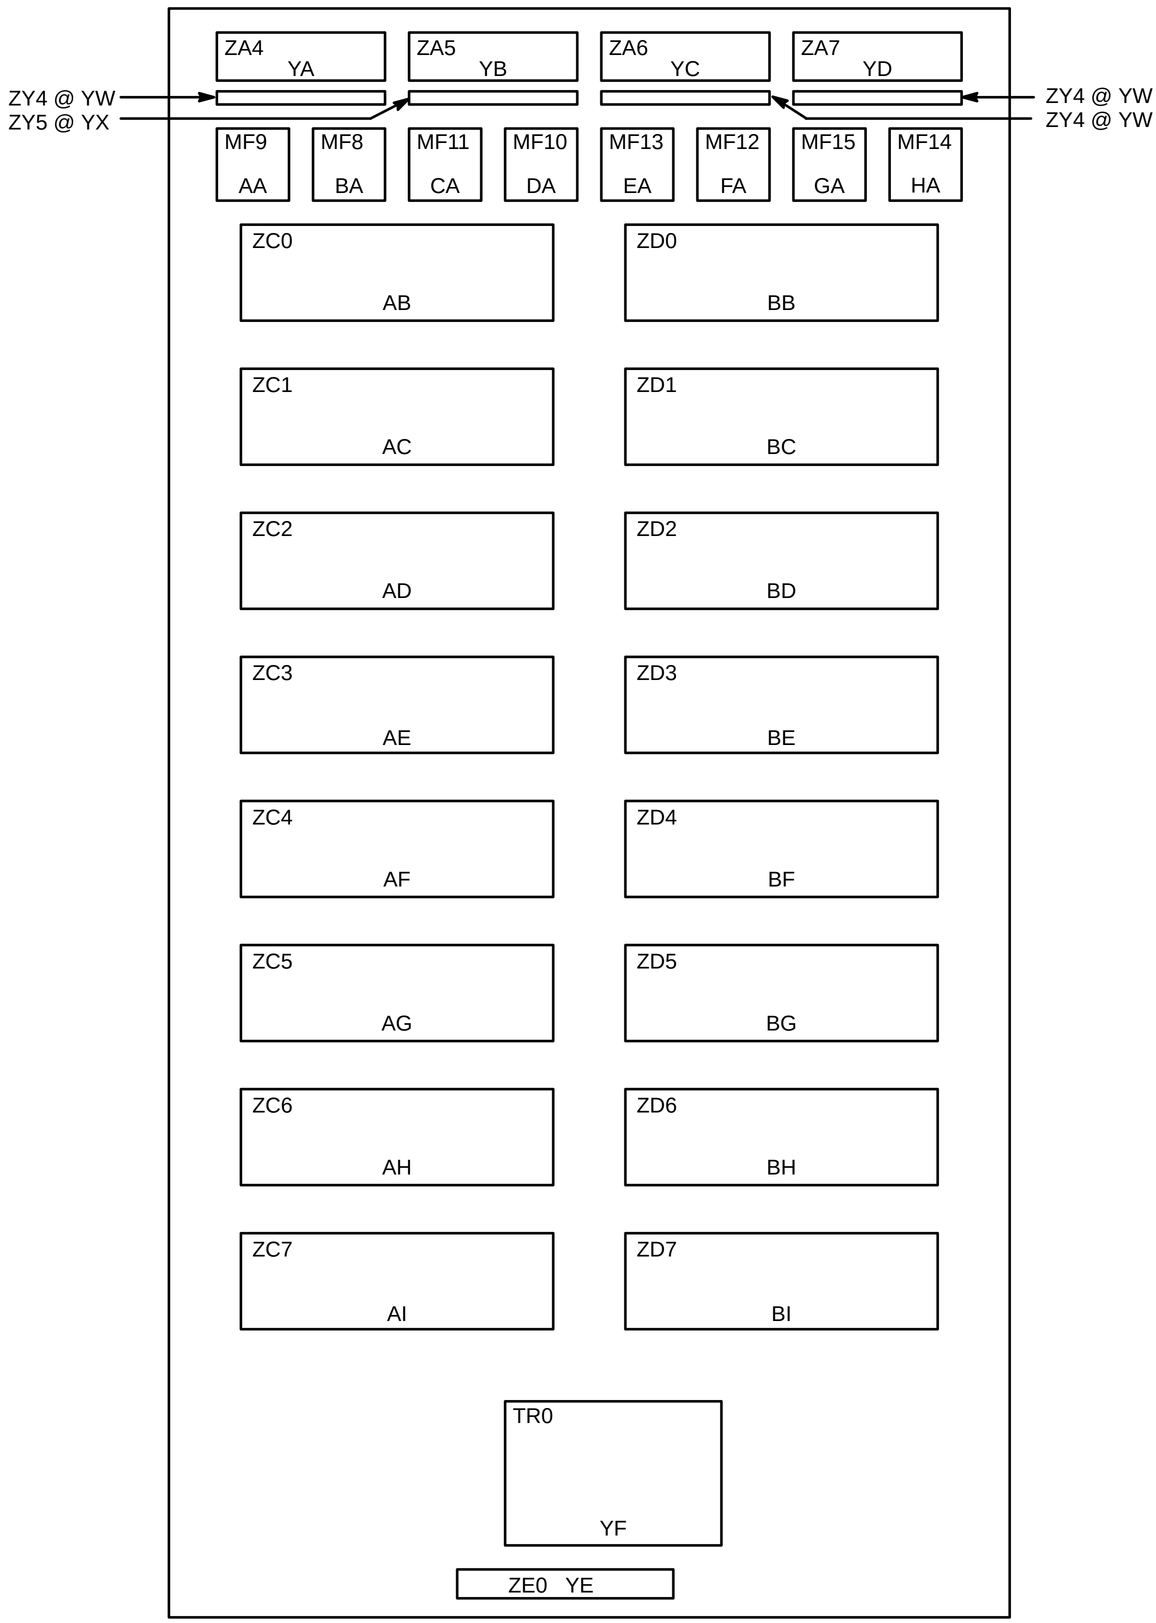

Figure 2. CM03 Memory Module – Board 2

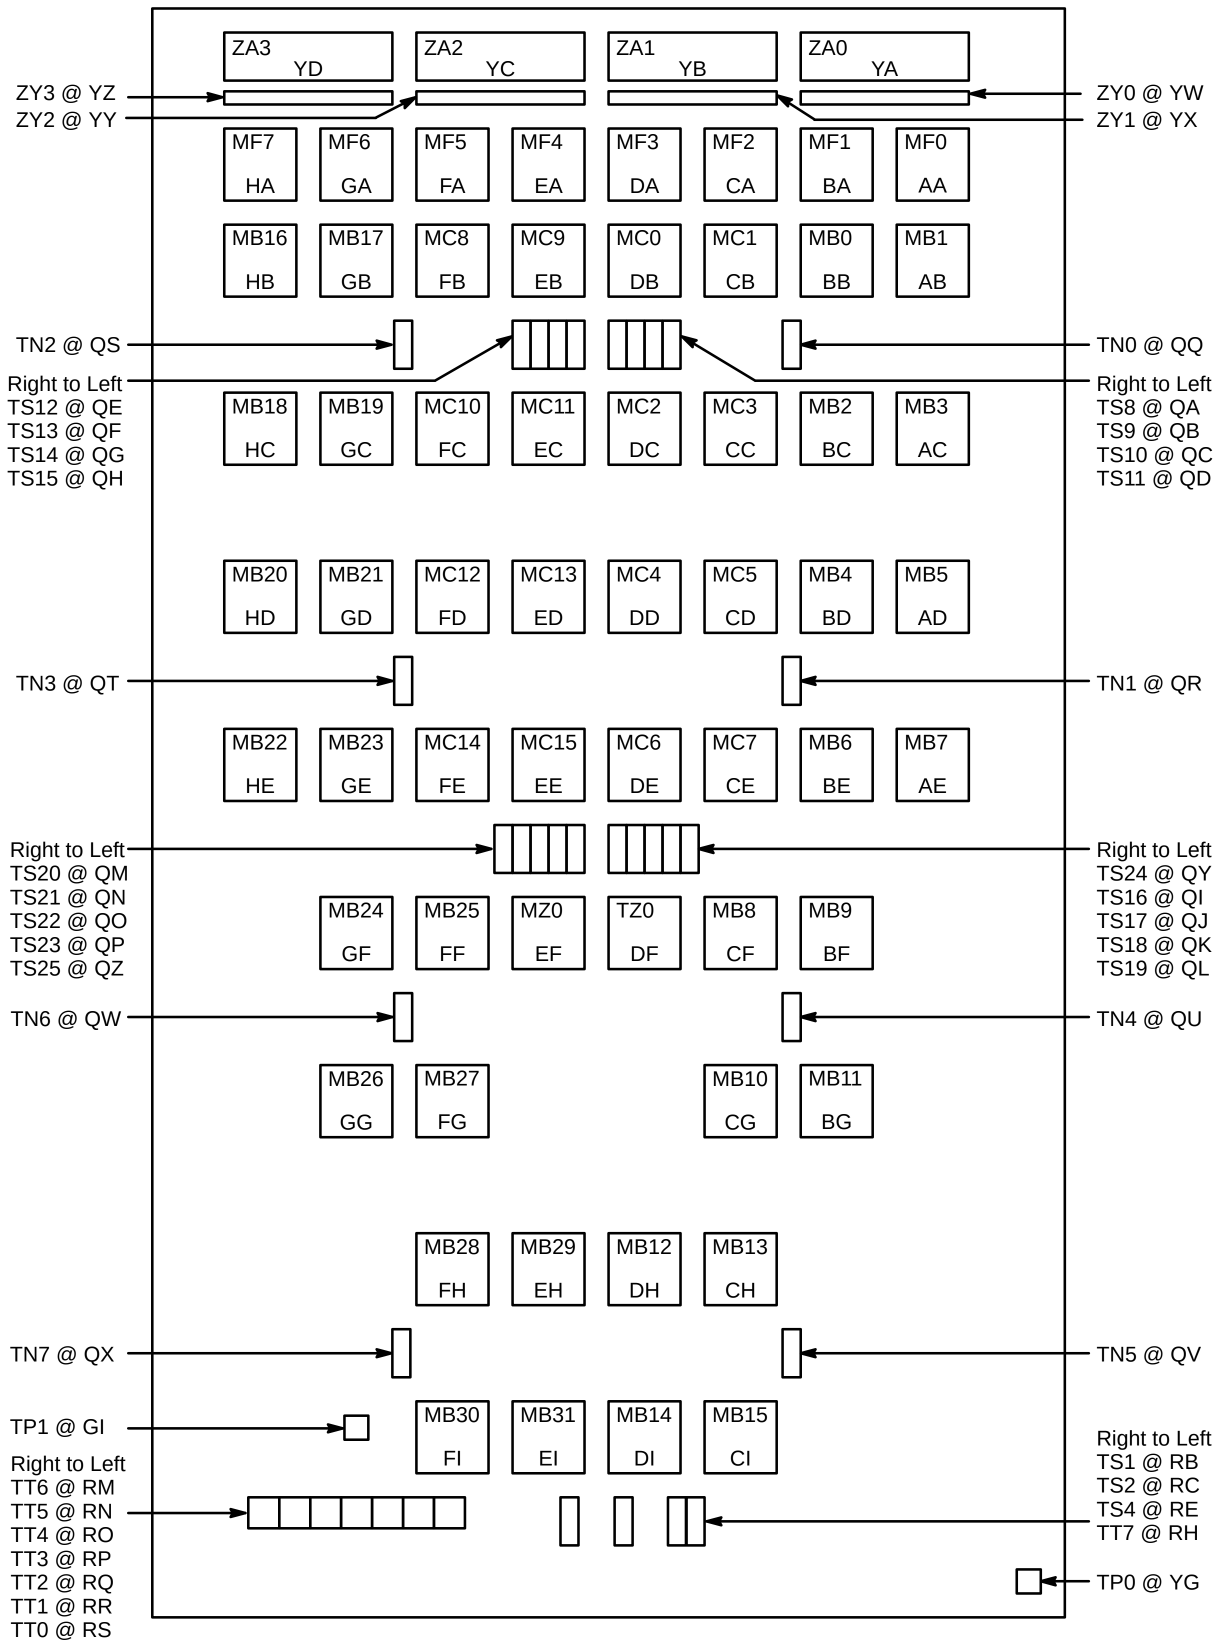

Memory Stacks

There are 16 memory stacks in each memory module. Each memory stack has 2 half-stacks. These half-stacks hold either the upper or lower 32 data bits, plus 6 check bits. Each half of the memory stack has 19 memory chips for data and check bits, and 1 spare chip for that half, for a total of 40 memory chips. These memory chips may be of either the synchronous type or the asynchronous type. Each type of memory chip has 4 million memory cells. The asynchronous memory chip is used as a 2 meg by 2 bit part, and the synchronous memory chip is used as a 1 meg by 4 bit part. That is, the memory stack populated with asynchronous chips has half as many banks as the memory stack populated with synchronous memory chips, but each asynchronous bank has twice the locations of a synchronous bank. Refer to Figure 3 for the bank layout for each type of stack.

The bits for a bank are always spread out over two memory stacks. Therefore, it takes 2 memory stacks to make up 2 or 4 complete banks, depending on the type of memory chip used. Refer to Figure 4 for a drawing of a CM03 memory stack.

Figure 3. Synchronous and Asynchronous Memory Stack Bit Layouts

Bank N+4	Bank N
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Bits	2 Bits
2 Check Bits	2 Check Bits
2 Check Bits	2 Check Bits
2 Check Bits	2 Check Bits
2 Spare Bits	2 Spare Bits

Memory Stack
Asynchronous Chips

Bank N+4	4 Bits	4 Bits	Bank N
Bank N+14	4 Bits	4 Bits	Bank N+10
Bank N+4	4 Bits	4 Bits	Bank N
Bank N+14	4 Bits	4 Bits	Bank N+10
Bank N+4	4 Bits	4 Bits	Bank N
Bank N+14	4 Bits	4 Bits	Bank N+10
Bank N+4	4 Bits	4 Bits	Bank N
Bank N+14	4 Bits	4 Bits	Bank N+10
Bank N+4	4 Bits	4 Bits	Bank N
Bank N+14	4 Bits	4 Bits	Bank N+10
Bank N+4	4 Bits	4 Bits	Bank N
Bank N+14	4 Bits	4 Bits	Bank N+10
Bank N+4	4 Bits	4 Bits	Bank N
Bank N+14	4 Bits	4 Bits	Bank N+10
Bank N+4	4 Bits	4 Bits	Bank N
Bank N+14	4 Bits	4 Bits	Bank N+10
Bank N+4	4 Check Bits	4 Check Bits	Bank N
Bank N+14	4 Check Bits	4 Check Bits	Bank N+10
Bank N+4	4 Check Bits	4 Check Bits	Bank N
	4 Spare Bits	4 Spare Bits	

Memory Stack
Synchronous Chips

Figure 4. CRAY T90 Series Memory Stack

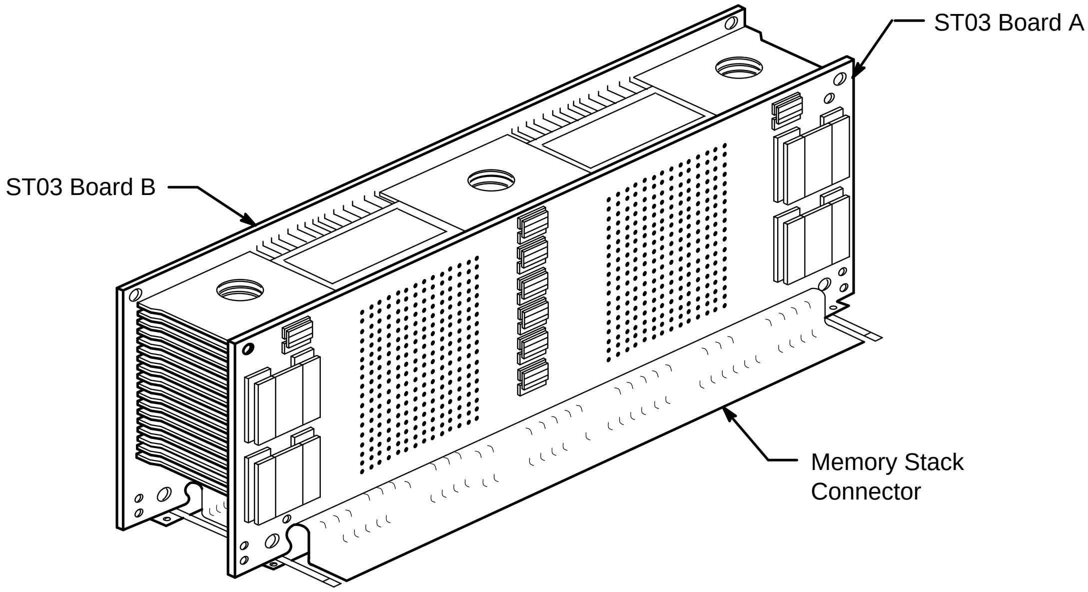

Spare Chips

There are two spare memory chips in each memory stack, one for each half of the stack. A bad memory chip can be flawed out, and the data in the remainder of the stack shifted down to the next chip, with the last chip shifting its bits to the spare chip for that half-stack. The flawing is accomplished with the SCE maintenance program. For a detailed explanation of spare chips and the flawing of memory chips refer to “Chip Flawing” on page 32 of this document.

CRAY T94 Memory

CRAY T94 memory is organized by sections and banks. A fully configured system has 8 sections of memory. Each section contains 16 banks for a total of 128 banks. Figure 5 shows memory organization in a CRAY T94 system populated with synchronous memory chips. Remember that memory modules populated with asynchronous chips have half as many banks. That is, there are no banks 10 through 17.

Figure 5. CRAY T94 Memory Organization

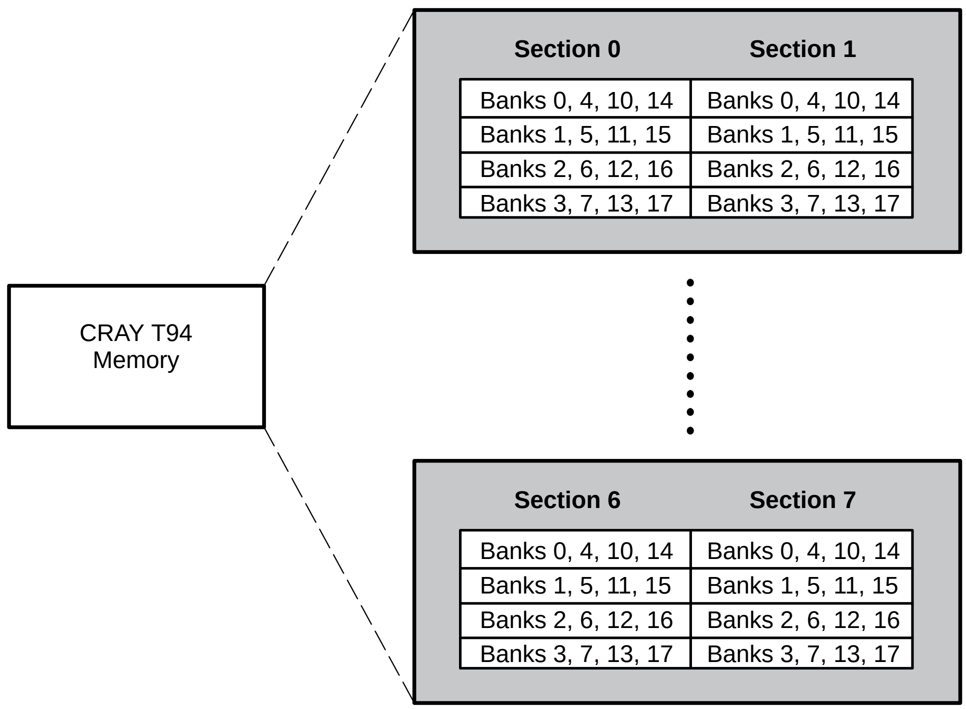

Figure 6 shows the addressing map for a fully configured CRAY T94 system populated with synchronous memory chips. It shows the module type that determines the section select, bank select, and word select. Bits 27 through 31 are presently not used; however, these bits may be used in future systems to address up to 4 Gwords of memory.

Figure 6. CRAY T94 Addressing for Synchronous Memory Chip
(No Partitioning or Degradation)

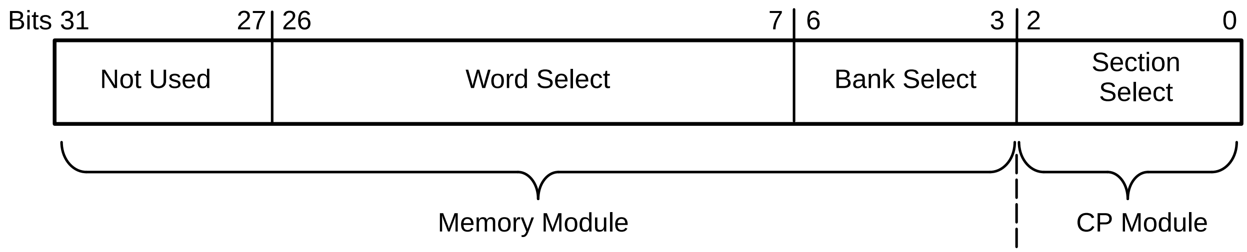

Figure 7 shows the addressing map for a fully configured CRAY T94 system populated with asynchronous memory chips. It shows the module type that determines the section select, bank select, and word select. Bits 27 through 31 are presently not used; however, these bits may be used in future systems to address up to 4 Gwords of memory.

Figure 7. CRAY T94 Addressing for Asynchronous Memory Chip
(No Partitioning or Degradation)

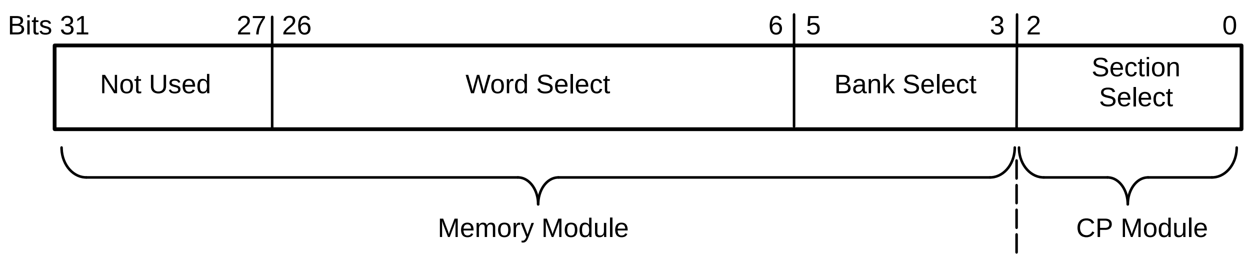

Table 3 lists the available CRAY T94 configurations. A CRAY T94 system has one memory module stack that contains either two or four memory modules. The memory modules connect directly to the CP modules. This table reflects bank counts for memory stacks populated with **synchronous** memory chips. Remember that memory stacks populated with asynchronous memory chips have half the number of banks but the same number of words.

Table 3. CRAY T94 System Configuration

Module Counts			Configuration			
Processor	Network	Memory	Sections	Subsects	Banks	MWords
1 to 4	0	4	8	1	128	128
1 to 4	0	4	8	1	64	64
1 to 4	0	2	4	1	64	64
1 to 4	0	2	4	1	32	32

Figure 8 shows the interconnections between the CP modules and memory modules.

Assuming the first reference is to section 0 and the memory references are sequential, the first two references go to sections 0 and 1 (bank 0) on the memory module at location C2, the next two references go to sections 2 and 3 on the memory module at location C4, and so on, until each module in the stack receives two references. References nine through sixteen follow the same pattern but address bank 1; references 17 through 24 address bank 2, and so on, until all banks (64 in a fully configured system) have been referenced.

Figure 8 also shows the section layout in a CRAY T94 system. Each memory module in a CRAY T94 system comprises 2 sections of memory.

Figure 8. CRAY T94 CPU-to-memory Interconnections

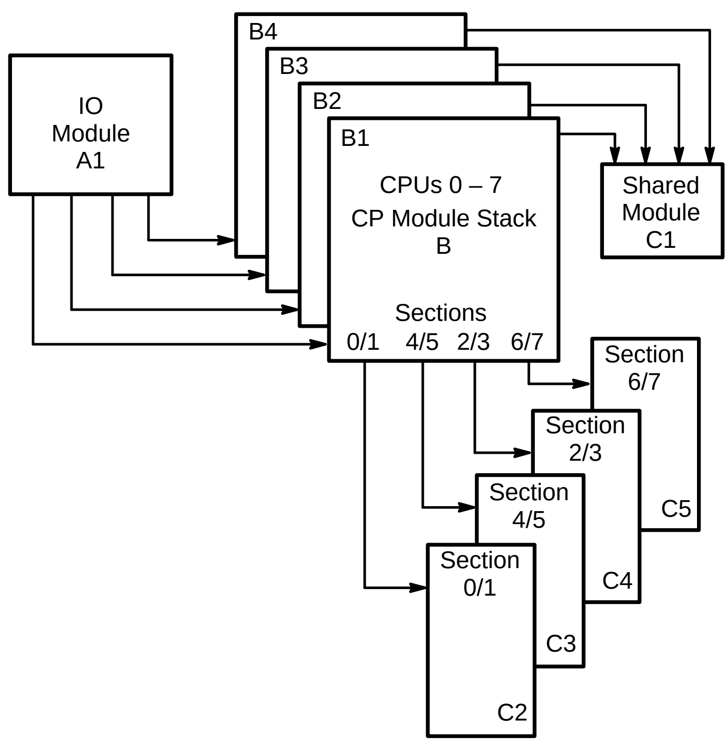

Figure 9 shows the stack layout on the memory module and the bit layout in the memory stack for the CRAY T94 system. The bit layout is for memory stacks populated with synchronous memory chips. Refer to Figure 10 for the bit layout of a memory stack populated with asynchronous chips.

Figure 9. CRAY T94 Memory Module Bank and Bit Layout for Synchronous Memory Chips

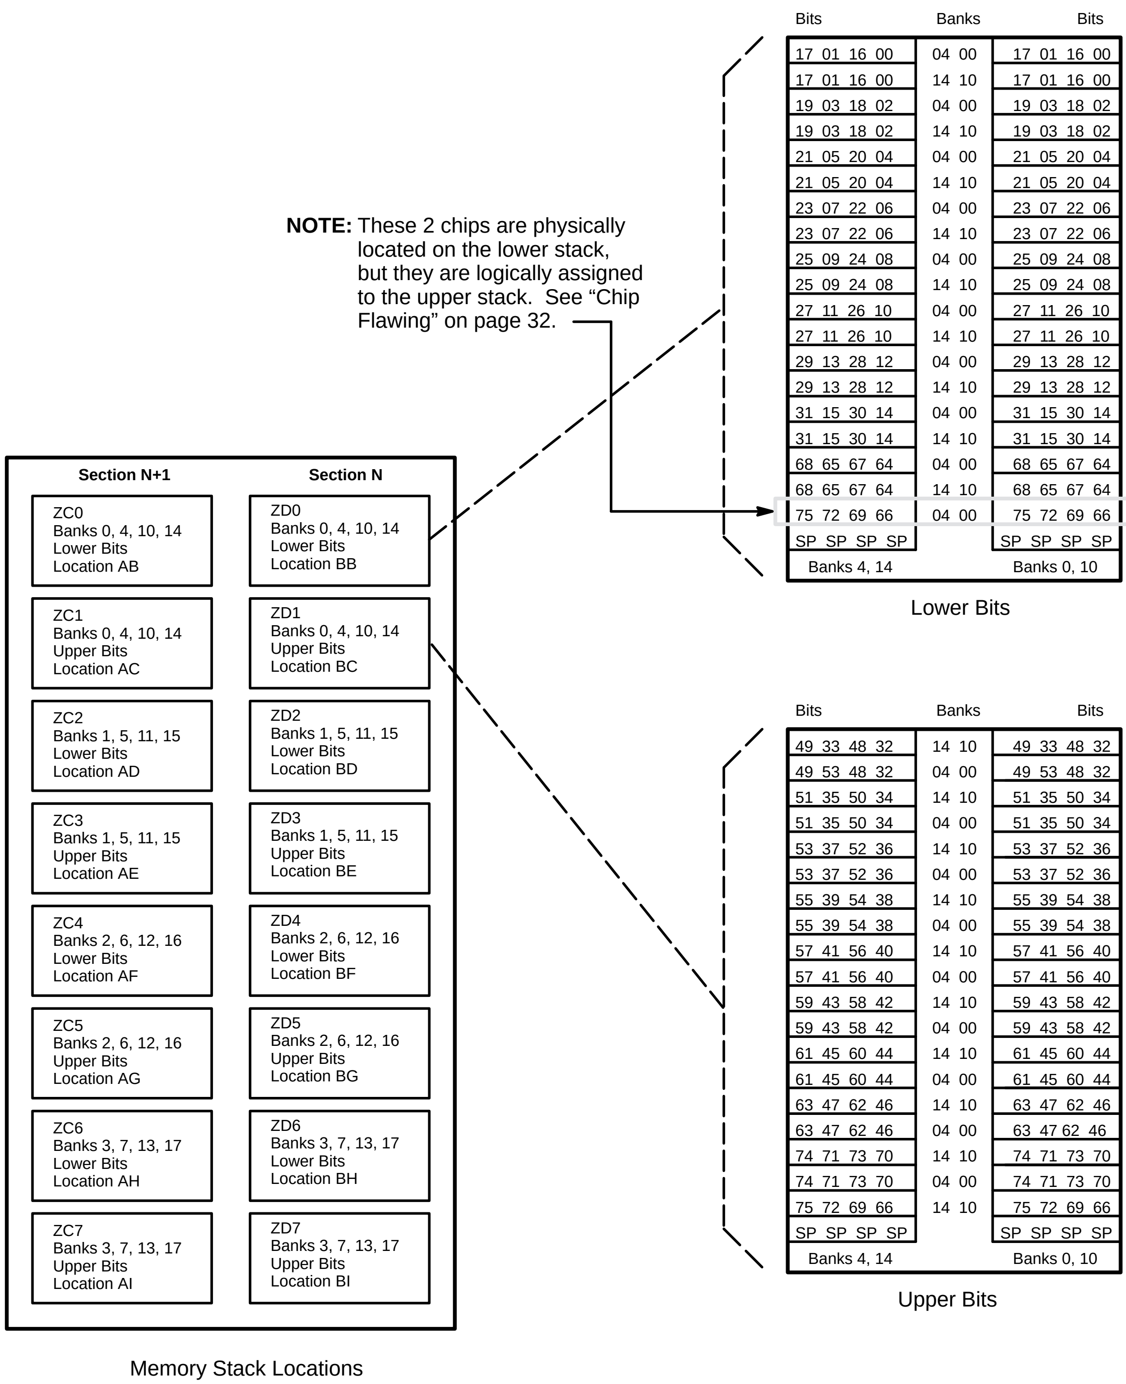

Figure 10. CRAY T94 Memory Module Bank and Bit Layout for Asynchronous Memory Chips

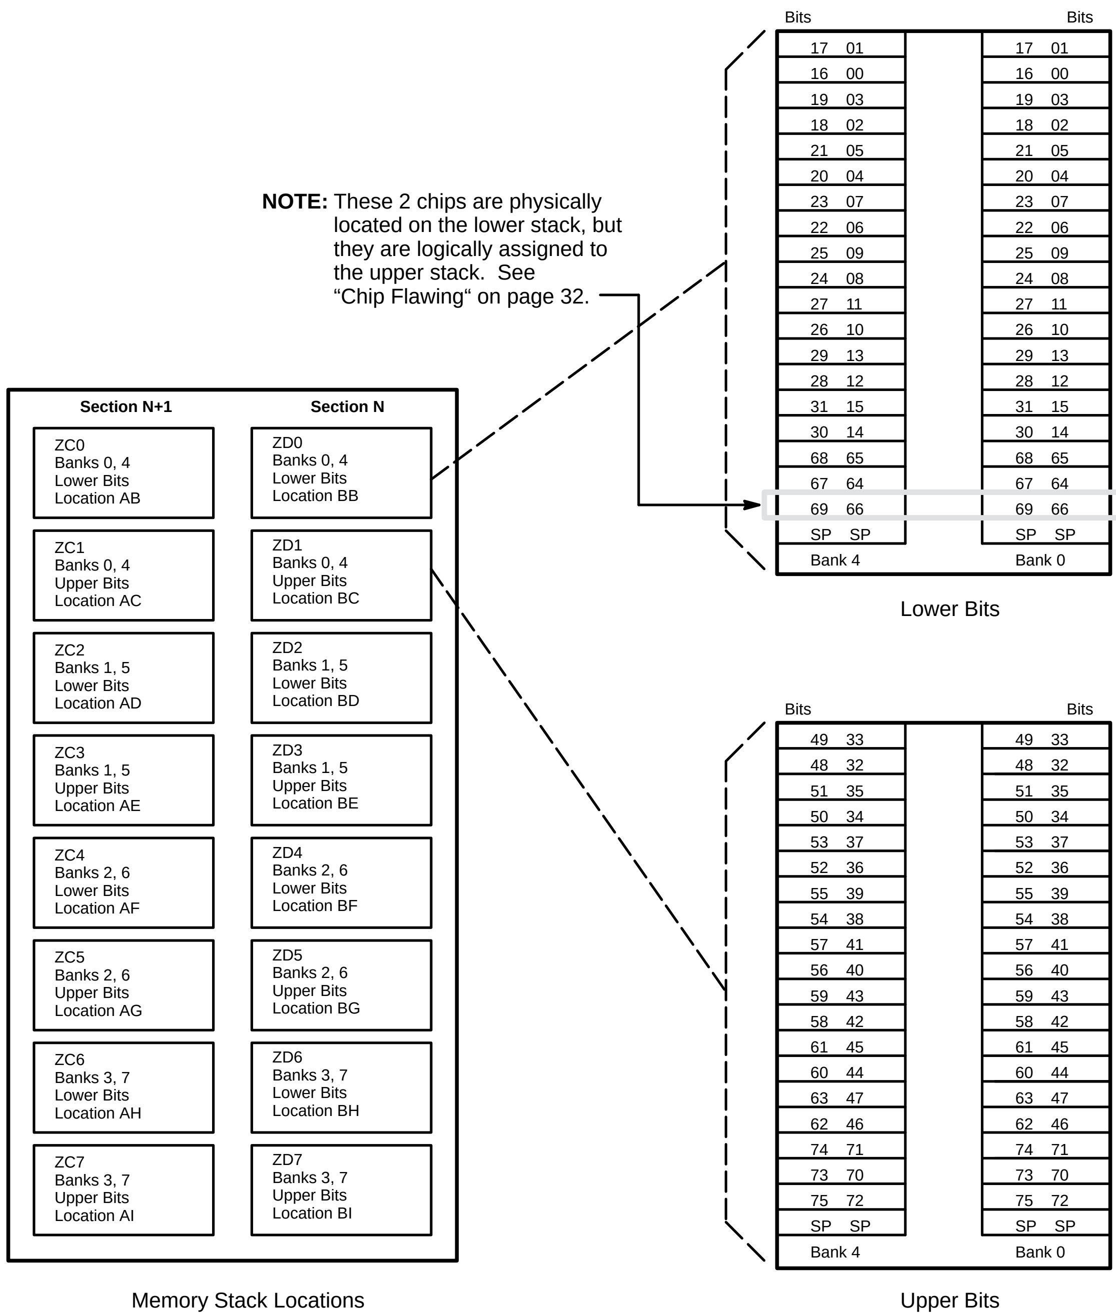

CRAY T916 Memory

CRAY T916 memory is organized by sections, subsections, and banks. A fully configured system has 8 sections of memory. Each section has 8 subsections and each subsection has 8 or 4 banks, depending on whether the memory stack is populated with synchronous or asynchronous memory chips, for a total of 512 or 256 banks, respectively. Figure 11 shows memory organization in a CRAY T916 system populated with synchronous memory chips. A system populated with asynchronous chips has only banks 0 through 3.

Figure 11. CRAY T916 Memory Organization

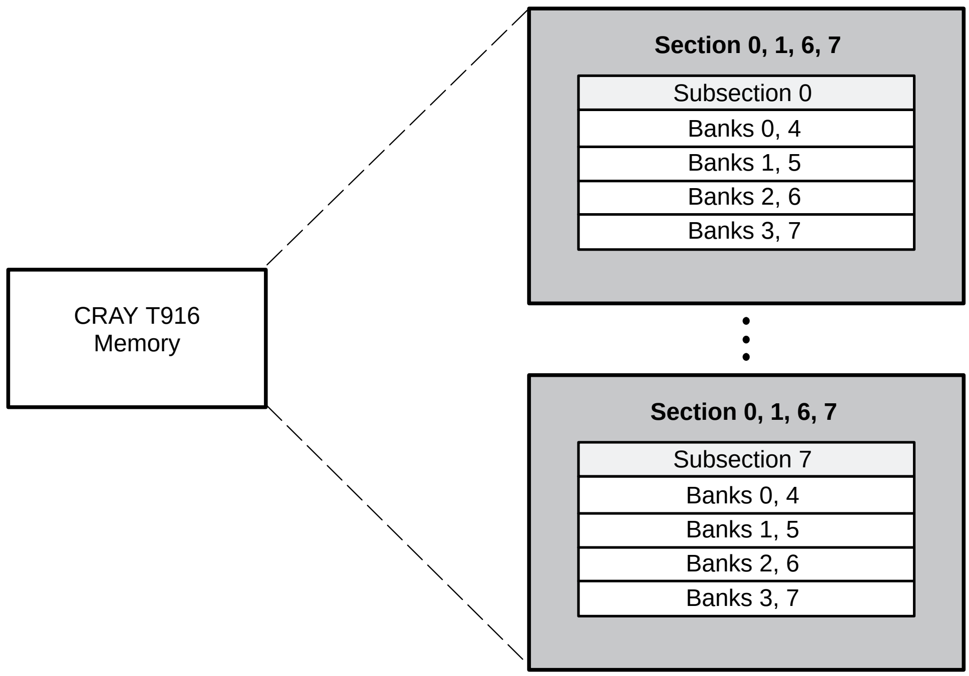

The CP and memory modules determine which section, bank, and word to address. Figure 12 shows the addressing map for a fully configured CRAY T916 system populated with synchronous memory chips. It also shows the module type that determines the section select, subsection select, bank select, and word select. Bits 29 through 33 are presently not used; however, these bits may be used in future systems to address up to 16 Gwords of memory.

Figure 12. CRAY T916 Addressing for Synchronous Memory Chip
(No Partitioning or Degradation)

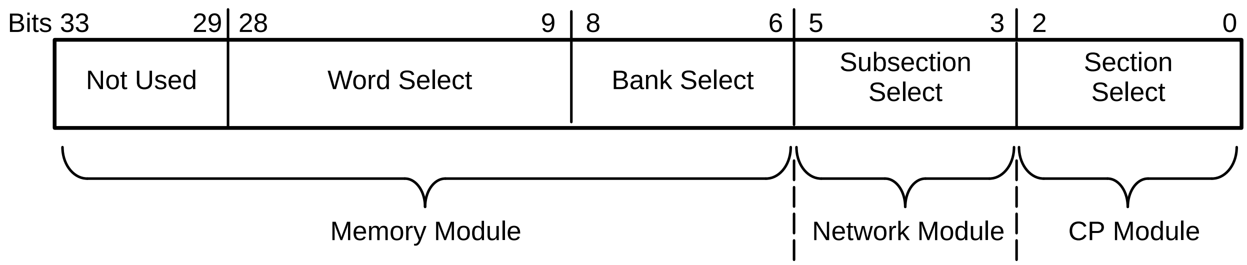

Figure 13 shows the addressing map for a fully configured CRAY T916 system populated with asynchronous memory chips. It also shows the module type that determines the section select, bank select, and word select. Bits 29 through 33 are presently not used; however, these bits may be used in future systems to address up to 16 Gwords of memory.

Figure 13. CRAY T916 Addressing for Asynchronous Memory Chip
(No Partitioning or Degradation)

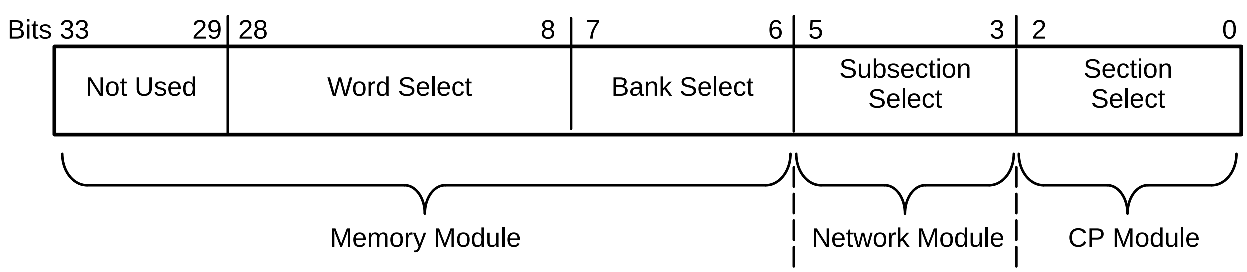

Table 4 lists the various configurations of CRAY T916 systems. A CRAY T916 mainframe has two memory module stacks that consist of four or eight memory modules each. The memory modules connect to the network modules, which connect to the CP modules. Connections between network and CP modules are actually made through the system interconnect board (SIB). This table reflects bank counts for memory stacks populated with **synchronous** memory chips. Remember that memory stacks populated with asynchronous memory chips have half the number of banks but the same number of words.

Table 4. CRAY T916 System Configuration

Module Counts			Configuration			
Processor	Network	Memory	Sections	Subsects	Banks	MWords
4 to 8	4	8	8	4	256	256
4 to 8	4	8	8	2	128	128
4 to 8	2	8	4	4	128	128
4 to 8	2	4	4	2	64	64
4 to 8	4	16	8	8	512	512
4 to 8	4	8	8	4	256	256
4 to 8	2	8	4	8	256	256
4 to 8	2	4	4	4	128	128
8 to 16	8	16	8	8	512	512
8 to 16	8	8	8	4	256	256
8 to 16	4	8	4	8	256	256
8 to 16	4	8	4	4	128	128

This indicates a degraded or partitioned system; the system is not usually sold with this configuration.

Figure 14 shows the logical interconnections between the CP modules and memory modules in a CRAY T916 system. Assuming the first reference is to section 0 and the memory references are sequential, reference one goes to section 0, subsection 0 at module location L1. Reference two goes to the same module but to section 1. References three through six go to sections 2, 3, 4, and 5 and subsection 0 on the memory module at location H1. References seven and eight go to sections 6 and 7, subsection 0 on the memory module at location L1. The next eight references follow the same sequence for subsection 1 but reference the modules at locations L3 and H3. This referencing pattern continues until all banks have been addressed.

Figure 14. CRAY T916 CPU-to-memory Interconnection

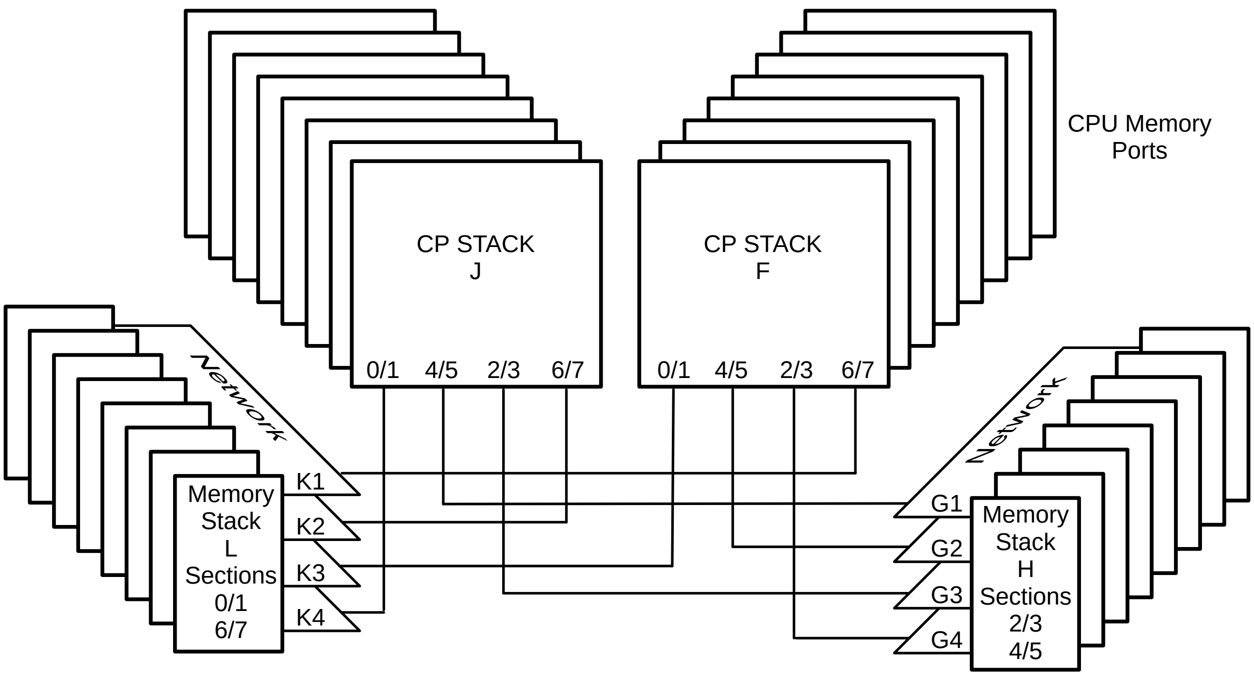

Figure 15 shows the memory module layout for a CRAY T916 system. Each memory module stack handles 4 sections of memory. Each memory module within the stack is 1 subsection for each of the 4 sections. Each subsection has 8 banks (synchronous memory chips) instead of 16 banks as in the CRAY T932 and CRAY T94 systems. Asynchronous memory is laid out the same way, but has only banks 0 through 3.

Figure 15 shows the module and bank layout for a CRAY T916 memory module that is populated with synchronous memory chips. Figure 16 shows the module and bank layout for a CRAY T916 memory module that is populated with asynchronous memory chips.

Figure 15. CRAY T916 Memory Module Layout – Synchronous Chips

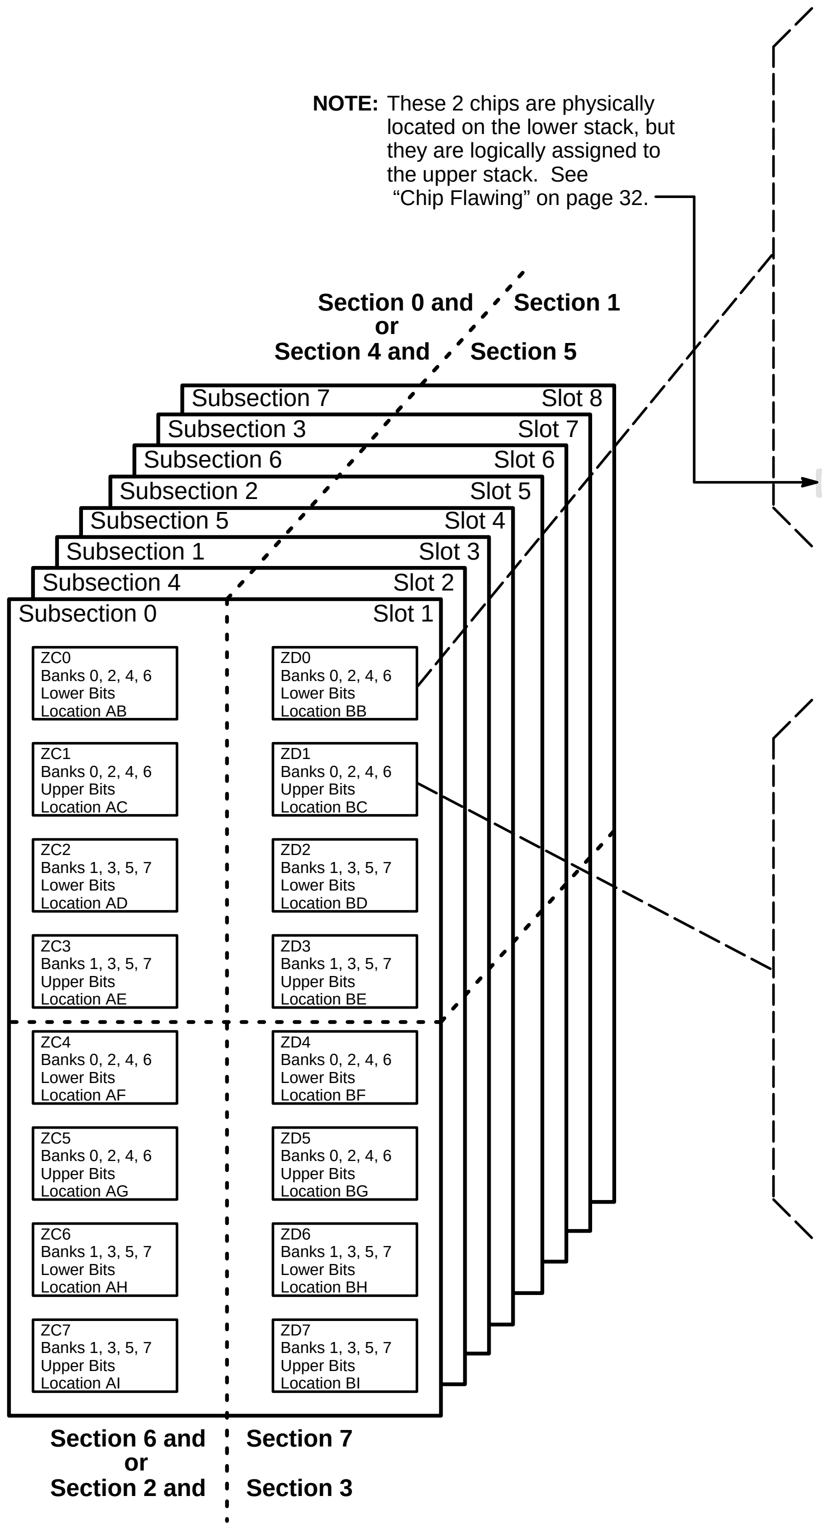

Bits	Banks	Bits
17 01 16 00	2 0	17 01 16 00
17 01 16 00	6 4	17 01 16 00
19 03 18 02	2 0	19 03 18 02
19 03 18 02	6 4	19 03 18 02
21 05 20 04	2 0	21 05 20 04
21 05 20 04	6 4	21 05 20 04
23 07 22 06	2 0	23 07 22 06
23 07 22 06	6 4	23 07 22 06
25 09 24 08	2 0	25 09 24 08
25 09 24 08	6 4	25 09 24 08
27 11 26 10	2 0	27 11 26 10
27 11 26 10	6 4	27 11 26 10
29 13 28 12	2 0	29 13 28 12
29 13 28 12	6 4	29 13 28 12
31 15 30 14	2 0	31 15 30 14
31 15 30 14	6 4	31 15 30 14
68 65 67 64	2 0	68 65 67 64
68 65 67 64	6 4	68 65 67 64
75 72 69 66	2 0	75 72 69 66
SP SP SP SP		SP SP SP SP
Banks 2, 6		Banks 0. 4

Lower Bits

Bits	Banks	Bits
49 33 48 32	6 4	49 33 48 32
49 53 48 32	2 0	49 33 48 32
51 35 50 34	6 4	51 35 50 34
51 35 50 34	2 0	51 35 50 34
53 37 52 36	6 4	53 37 52 36
53 37 52 36	2 0	53 37 52 36
55 39 54 38	6 4	55 39 54 38
55 39 54 38	2 0	55 39 54 38
57 41 56 40	6 4	57 41 56 40
57 41 56 40	2 0	57 41 56 40
59 43 58 42	6 4	59 43 58 42
59 43 58 42	2 0	59 43 58 42
61 45 60 44	6 4	61 45 60 44
61 45 60 44	2 0	61 45 60 44
63 47 62 46	6 4	63 47 62 46
63 47 62 46	2 0	63 47 62 46
74 71 73 70	6 4	74 71 73 70
74 71 73 70	2 0	74 71 73 70
75 72 69 66	6 4	75 72 69 66
SP SP SP SP		SP SP SP SP
Banks 2, 6		Banks 0, 4

Upper Bits

Figure 16. CRAY T916 Memory Module Layout – Asynchronous Chips

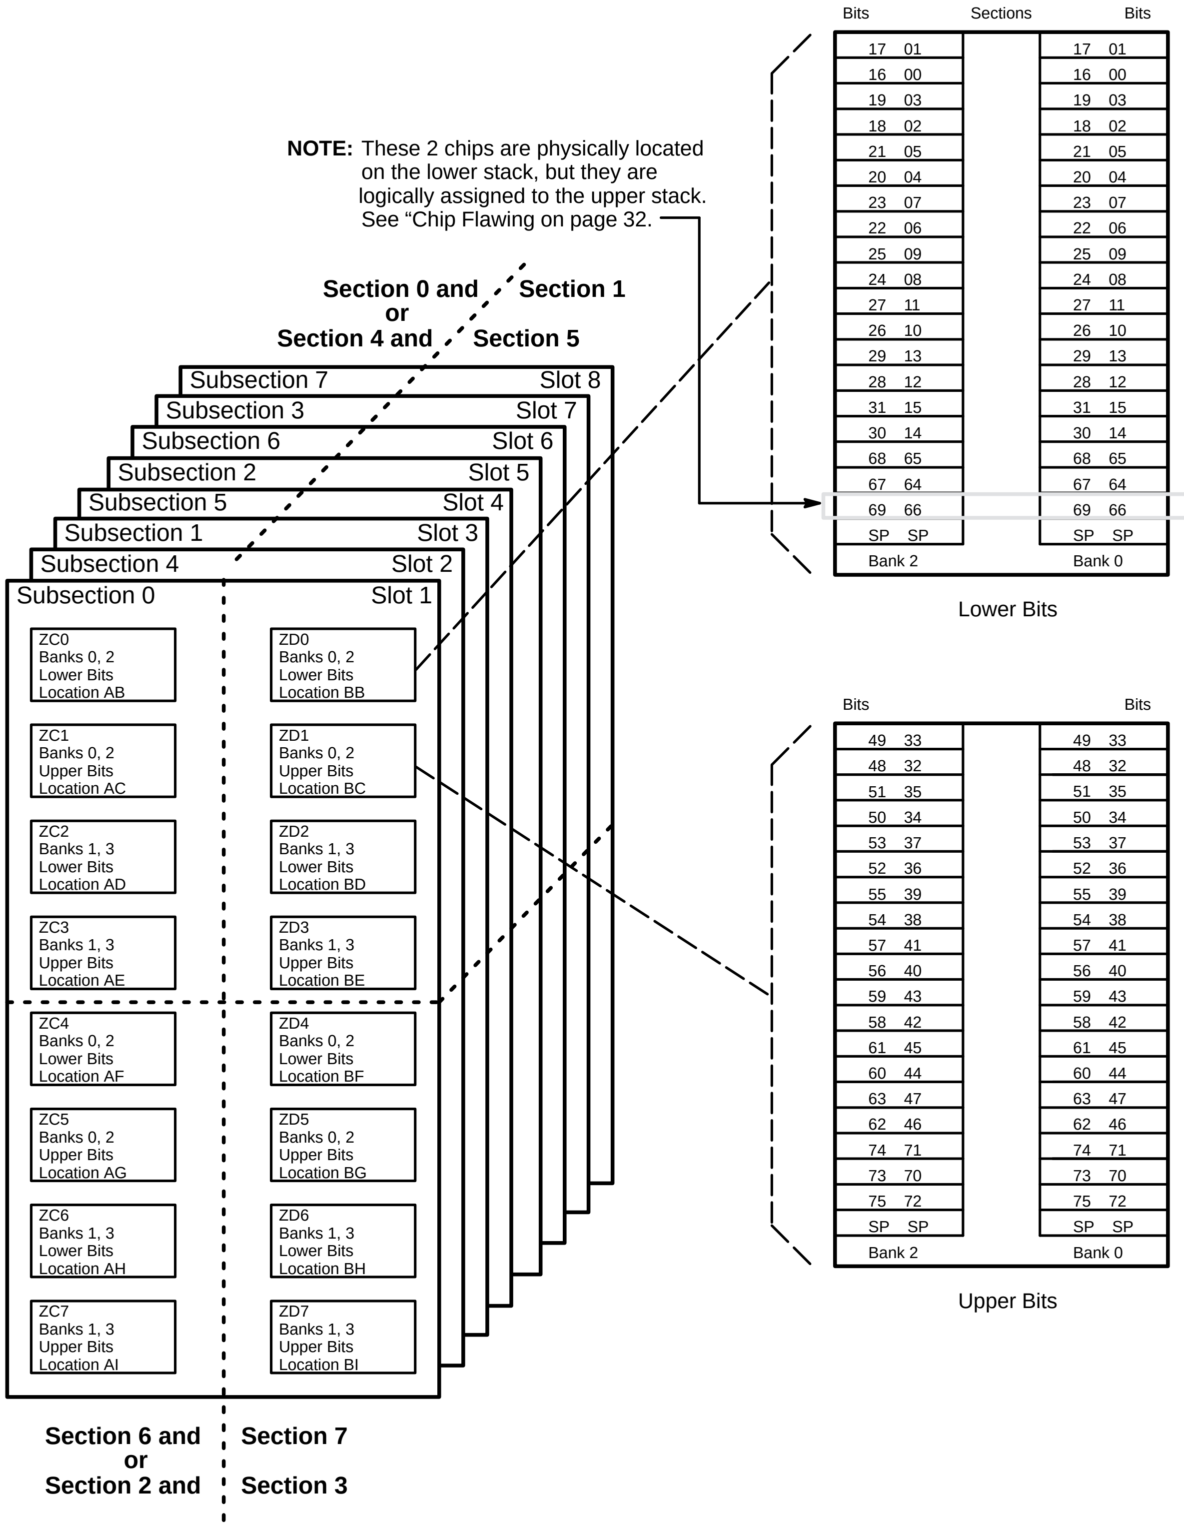

CRAY T932 Memory

CRAY T932 memory is organized by sections, subsections, and banks. A fully configured system has 8 sections of memory. Each section has 8 subsections and each subsection has 16 or 8 banks, depending on whether the memory stack is populated with synchronous or asynchronous memory chips, for a total of 1,024 or 512 banks, respectively. Figure 17 shows memory organization in a CRAY T932 system populated with synchronous memory chips.

Figure 17. CRAY T932 Memory Organization

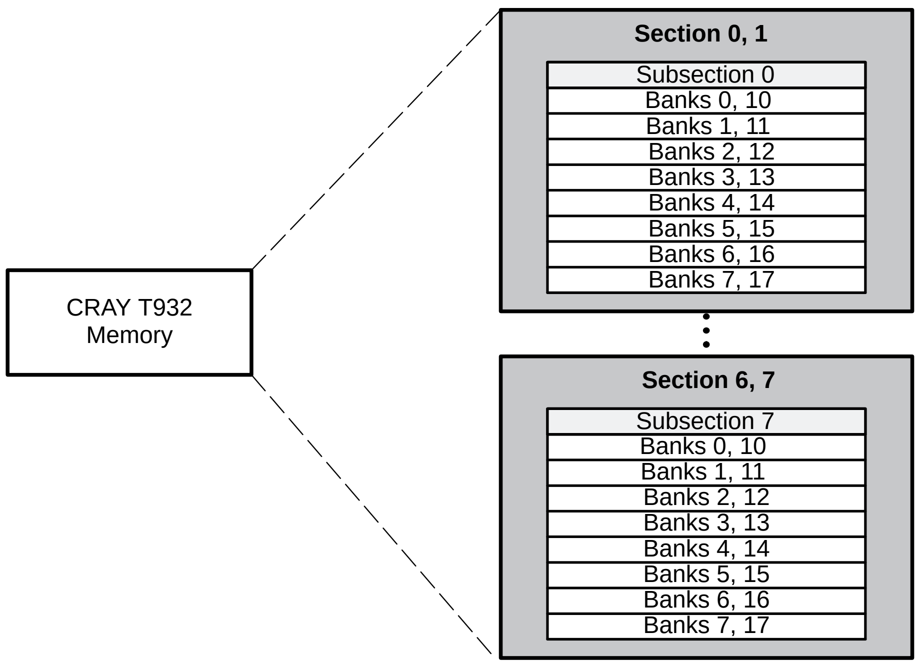

The CP, network, and memory modules determine which section, subsection, bank, and word in memory to address. Figure 18 shows the addressing map for a fully configured CRAY T932 memory populated with synchronous memory chips. It also shows the module type that determines the section select, subsection select, bank select, and word select. Bits 30 through 34 are presently not used; however, these bits may be used in future systems to address up to 32 Gwords of memory.

Figure 18. CRAY T932 Addressing for Synchronous Memory Chip
(No Partitioning or Degradation)

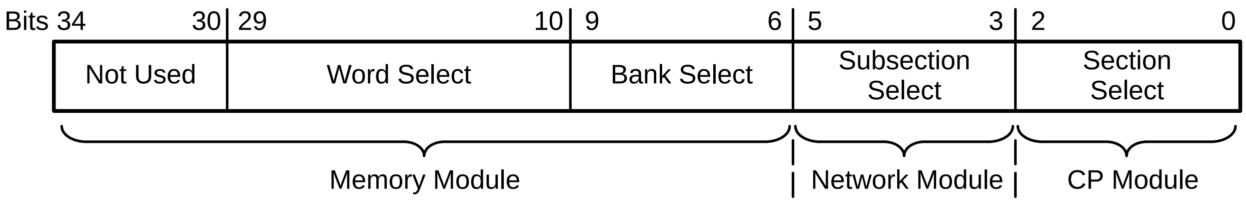

Figure 19 shows the addressing map for a fully configured CRAY T932 memory populated with asynchronous memory chips. It also shows the module type that determines the section select, subsection select, bank select, and word select. Bits 30 through 34 are presently not used; however, these bits may be used in future systems to address up to 4 Gwords of memory.

Figure 19. CRAY T932 Addressing for Asynchronous Memory Chip
(No Partitioning or Degradation)

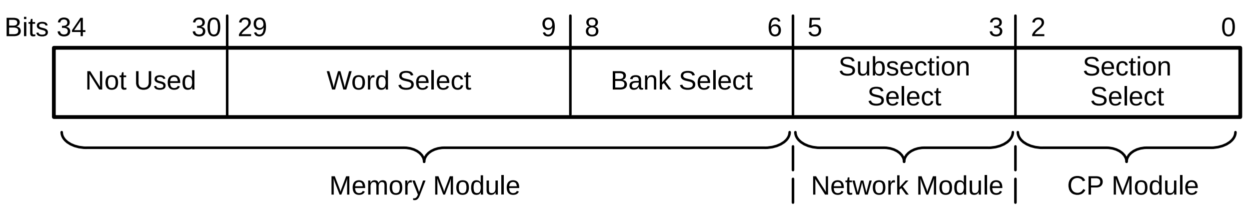

Table 5 lists the various configurations of CRAY T932 systems. A CRAY T932 system has two or four memory module stacks with each module stack containing either two, four, or eight memory modules. The memory modules connect to the network modules, which connect to the CP modules. The system interconnect board (SIB) connects these memory modules to the network modules and to the CP modules. This table reflects bank counts for memory stacks populated with **synchronous** memory chips. Remember that memory stacks populated with asynchronous memory chips have half the number of banks, but the same number of words.

Table 5. CRAY T932 System Configuration

Module Counts			Configuration			
Processor	Network	Memory	Sections	Subsects	Banks	MWords
8	4	32	8	8	1024	1024
8	4	16	8	4	512	512
8	4	8	8	2	256	256
8	8	8	8	2	256	256
8	8	8	8	2	128	128
8	4	4	4	2	128	128
8	4	4	4	2	64	64
8 to 16	8	16	8	4	512	512
8 to 16	8	16	8	4	256	256
8 to 16	8	8	8	2	256	256
8 to 16	4	8	4	4	256	256
16	16	16	8	4	512	512
16	16	16	8	4	256	256
16	16	8	8	2	256	256
16	8	8	4	4	256	256
16 to 32	16	32	8	8	1024	1024
16 to 32	16	32	8	8	512	512
16 to 32	16	16	8	4	512	512
16 to 32	8	16	4	8	512	512

This indicates a degraded or partitioned system; the system is not usually sold with this configuration.

Figure 20 shows the logical interconnections between the CP modules and memory modules in a CRAY T932 system. Assuming the first reference is to section 0 and the memory references are sequential, references one and two go to sections 0 and 1, subsection 0 at module location P1. References three and four go to sections 2 and 3, subsection 0 at module location D1. References five and six go to sections 4 and 5, subsection 0 at module location H1. References seven and eight go to sections 6 and 7, subsection 0 at module location L1. This sequence continues as a descending spiral through the module stacks and memory subsections until all subsections and banks have been addressed.

Figure 20. CRAY T932 CPU-to-memory Interconnections

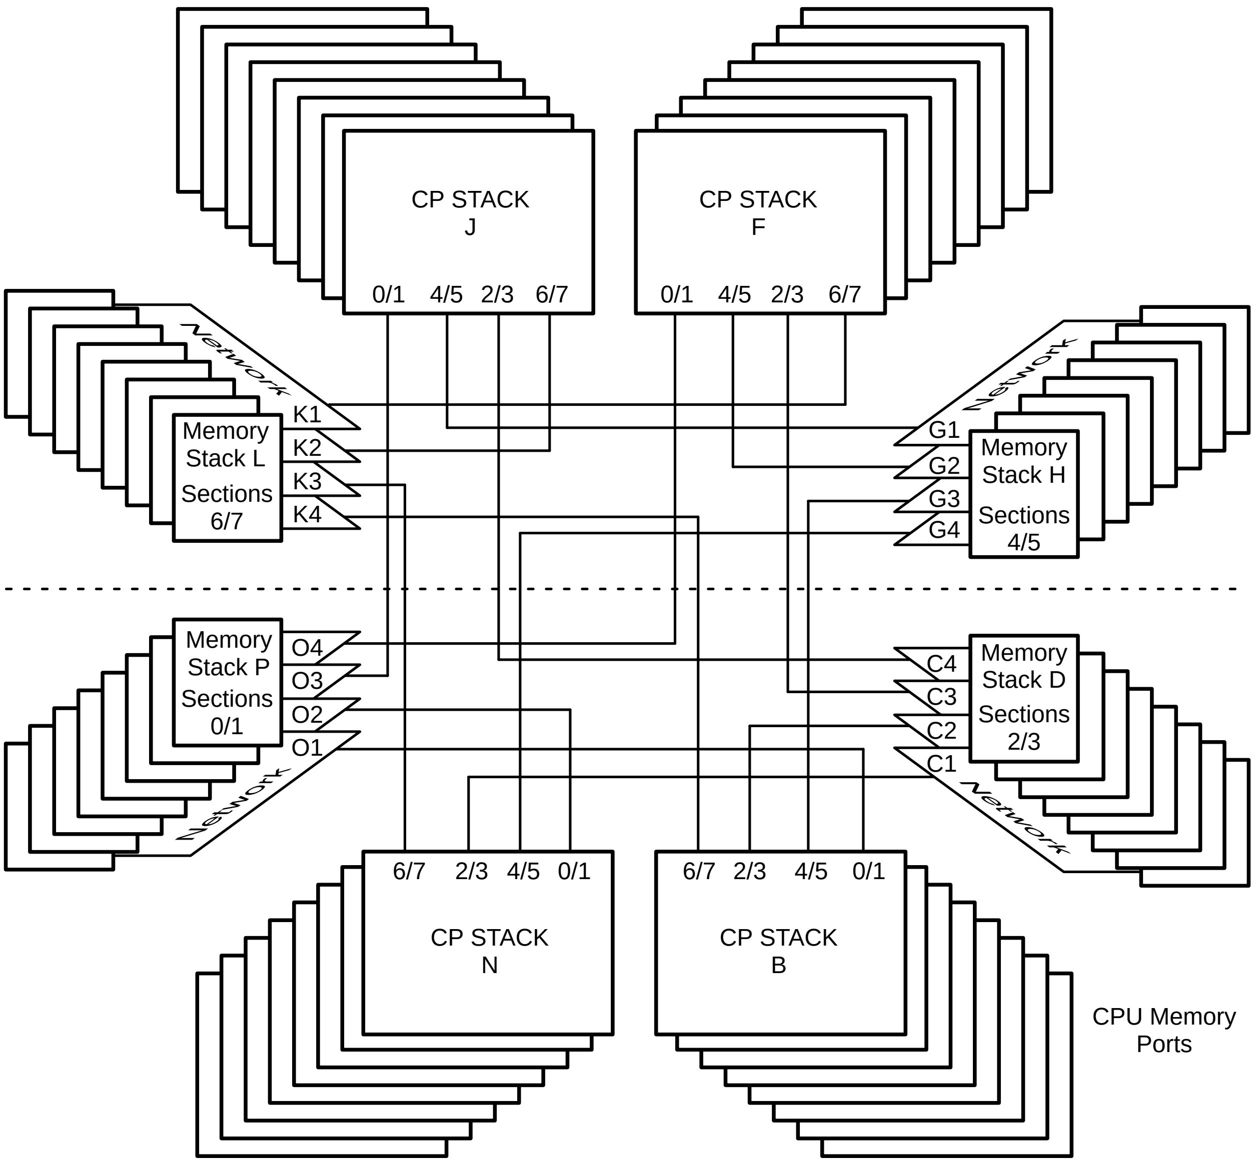

Figure 21 shows the module layout for a CRAY T932 system with the memory modules populated with synchronous memory chips. Each memory module within the memory module stack is 1 subsection for each of 2 sections.

Figure 21. CRAY T932 Memory Module Layout – Synchronous Chips

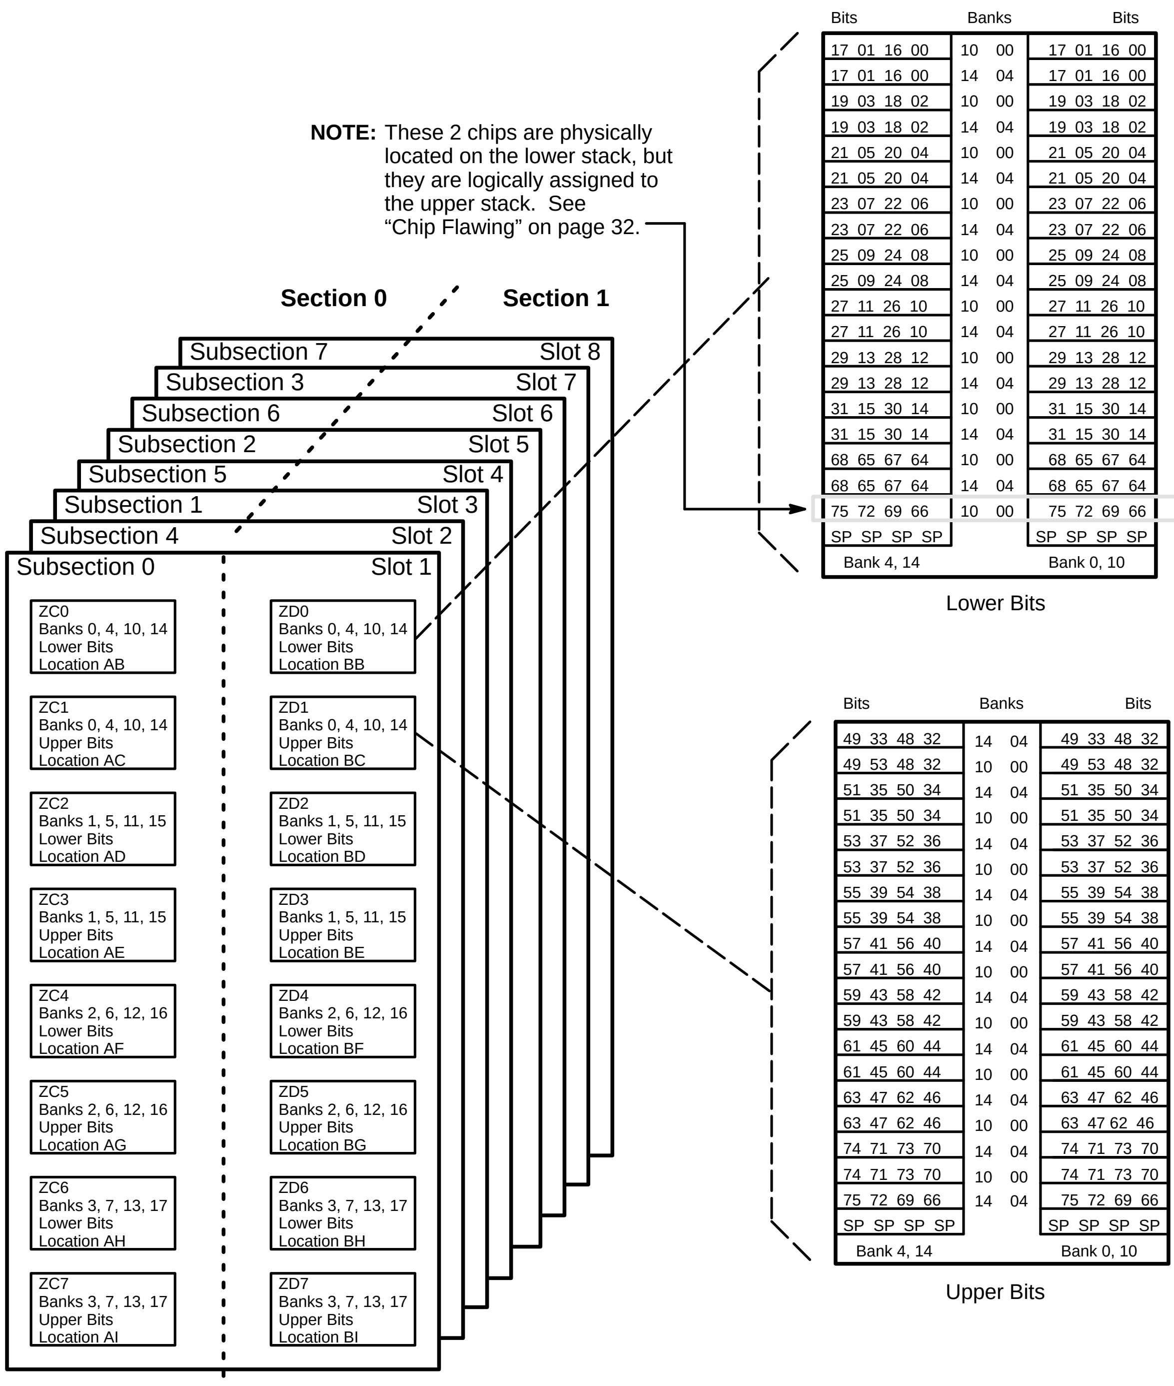

Figure 22 shows the module layout for a CRAY T932 system with the memory modules populated with asynchronous memory chips. Each memory module within the memory module stack is 1 subsection for each of 2 sections.

Figure 22. CRAY T932 Memory Module Layout – Asynchronous Chips

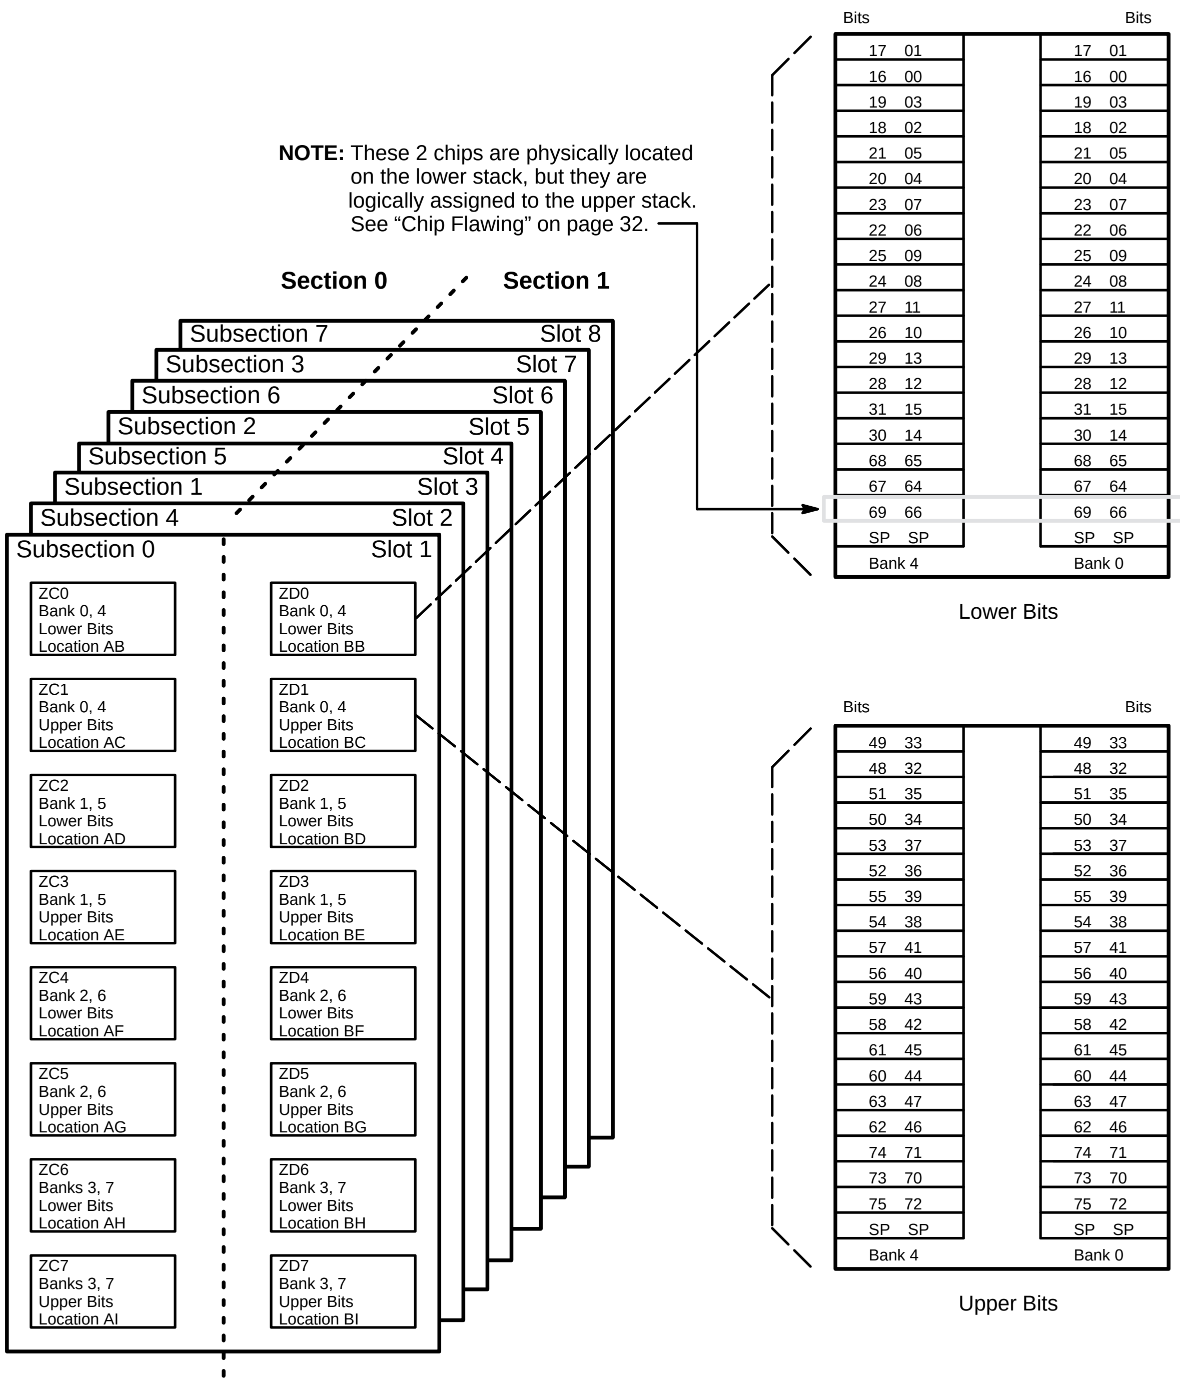

Chip Flawing

Each memory stack has 2 spare memory chips, one for each half of the stack, whether using synchronous or asynchronous memory chips. This feature enables you to flaw out a bad memory chip from that half of the memory stack and shift the data up the spare chip. Use the maintenance program SCE to perform chip flawing. Refer to “Spare Chip Memory Management” in the *SCE User Guide*, publication number HDM-069-A, for information on how to enter or remove a flaw.

The spare bit configuration for a synchronous memory stack is different from an asynchronous memory stack; therefore, this document describes each memory stack separately.

Error Correction

Error detection/correction is done on the CPU module. The CPU module must be able to accommodate a number of configurations: the CM02 module, which handles 2 bits per memory chip; and the CM03 module, which handles either 2 or 4 bits per memory chip and 1 or 2 banks per half-stack.

The CPU module is capable of correcting only the 2 **data bits** that are 16 bits apart on a memory chip. For example, bits 0, 16, 1, and 17 are on 1 memory chip on a synchronous memory chip. Bits 0 and 16 are correctable **or** bits 1 and 17 are correctable, but not bits 0 and 1, or 0 and 17, and so on. The 2 correctable **check bits** on 1 chip are only 3 bits apart, bits 70 and 73 for example.

Synchronous Memory Stack

Each memory stack handles half the bits of a word for 4 banks, and each half of the stack handles half the bits for banks n and n+10. Refer to Figure 23 for a drawing of how the bits are **physically** laid out.

Each memory half-stack contains 20 memory chips; 19 chips for data and 1 spare chip. The nineteenth memory chip on each lower half-stack (chip 23_8) is logically assigned to the corresponding half-stacks on the upper memory stack. Refer to Figure 24. This means that the spare chip on each lower half-stack accommodates one flaw for the 18 chips in that half-stack and that the spare chip on the corresponding upper half-stack accommodates one flaw for 20 chips assigned to that half-stack.

The spare chip can accommodate one flaw. If, for example, the memory chip that handles bits 0, 16, 1, and 17 for bank 0 is flawed out, then there are no remaining flaws for this half of the stack: not for bank 0 or for bank 10.

Figure 23. Physical Bit Layout of Synchronous Memory Stacks

Bits	Banks	Bits	Chip Number	Bits	Banks	Bits
49 33 48 32	14 10	49 33 48 32	00	17 01 16 00	04 00	17 01 16 00
49 53 48 32	04 00	49 53 48 32	01	17 01 16 00	14 10	17 01 16 00
51 35 50 34	14 10	51 35 50 34	02	19 03 18 02	04 00	19 03 18 02
51 35 50 34	04 00	51 35 50 34	03	19 03 18 02	14 10	19 03 18 02
53 37 52 36	14 10	53 37 52 36	04	21 05 20 04	04 00	21 05 20 04
53 37 52 36	04 00	53 37 52 36	05	21 05 20 04	14 10	21 05 20 04
55 39 54 38	14 10	55 39 54 38	06	23 07 22 06	04 00	23 07 22 06
55 39 54 38	04 00	55 39 54 38	07	23 07 22 06	14 10	23 07 22 06
57 41 56 40	14 10	57 41 56 40	10	25 09 24 08	04 00	25 09 24 08
57 41 56 40	04 00	57 41 56 40	11	25 09 24 08	14 10	25 09 24 08
59 43 58 42	14 10	59 43 58 42	12	27 11 26 10	04 00	27 11 26 10
59 43 58 42	04 00	59 43 58 42	13	27 11 26 10	14 10	27 11 26 10
61 45 60 44	14 10	61 45 60 44	14	29 13 28 12	04 00	29 13 28 12
61 45 60 44	04 00	61 45 60 44	15	29 13 28 12	14 10	29 13 28 12
63 47 62 46	14 10	63 47 62 46	16	31 15 30 14	04 00	31 15 30 14
63 47 62 46	04 00	63 47 62 46	17	31 15 30 14	14 10	31 15 30 14
74 71 73 70	14 10	74 71 73 70	20	68 65 67 64	04 00	68 65 67 64
74 71 73 70	04 00	74 71 73 70	21	68 65 67 64	14 10	68 65 67 64
75 72 69 66	14 10	75 72 69 66	22 23	75 72 69 66	04 00	75 72 69 66
SP SP SP SP		SP SP SP SP		SP SP SP SP		SP SP SP SP
Bank 4, 14		Bank 0, 10		Bank 4, 14		Bank 0, 10
Upper Bits				Lower Bits		

There is no chip 22 on the lower stack. Chip 23, which physically resides on the lower stack is logically assigned to the upper stack. Refer to Figure 24 for a **logical** picture of the memory stacks. The lower-bit stack spares 18 memory chips, and the upper-bit stack spares 20 memory chips.

Figure 24. Logical Bit Layout of Synchronous Memory Stacks

63 47 62 46	04 00	63 47 62 46	17	31 15 30 14	14 10	31 15 30 14
74 71 73 70	14 10	74 71 73 70	20	68 65 67 64	04 00	68 65 67 64
74 71 73 70	04 00	74 71 73 70	21	68 65 67 64	14 10	68 65 67 64
75 72 69 66	14 10	75 72 69 66	22	SP SP SP SP		SP SP SP SP
75 72 69 66	04 00	75 72 69 66	23			
SP SP SP SP		SP SP SP SP				
Bank 4, 14		Bank 0, 10		Bank 4, 14		Bank 0, 10
Upper Bits				Lower Bits		

Synchronous Bit-shift Patterns

Once a chip is flawed out, its bits are shifted on the MB options to the next chip for that bank. Refer to Figure 28. Notice on the **lower half-stack** that chip 0's bits are shifted to chip 2, chip 2 is shifted to chip 4, and so on, until chip 20 is reached. There is no chip 22 on this half-stack, so chip 20's bits are shifted to the spare chip on this half-stack. If chip 1 is flawed out, its bits are shifted to chip 3, chip 3's bits are shifted to chip 5, and so on, until chip 21 is reached. Because chip 23 is assigned to the upper memory stack, chip 21's bits are shifted to the spare chip on this half-stack.

If chip 0 on the **upper half-stack** is flawed out, it follows the same pattern as for the lower half-stack until it gets to chip 20. There is a chip 22 on this half-stack, so chip 20 shifts to chip 22, and chip 22 shifts to the spare chip on this half-stack. If chip 1 is flawed out, it, like the lower half-stack, shifts to chip 3, and so on, until chip 21 is reached. Chip 21 shifts its bits to chip 23. Remember that chip 23, while logically assigned to this half-stack, resides on the lower half-stack. Therefore, the bits are shifted over to the lower half-stack to chip 23, and chip 23 is shifted to the spare chip on the upper half-stack.

Figure 28 shows both the logical and the physical bit shifts.

Figure 25. Logical and Physical Bit Shifts - Asynchronous Memory Stacks

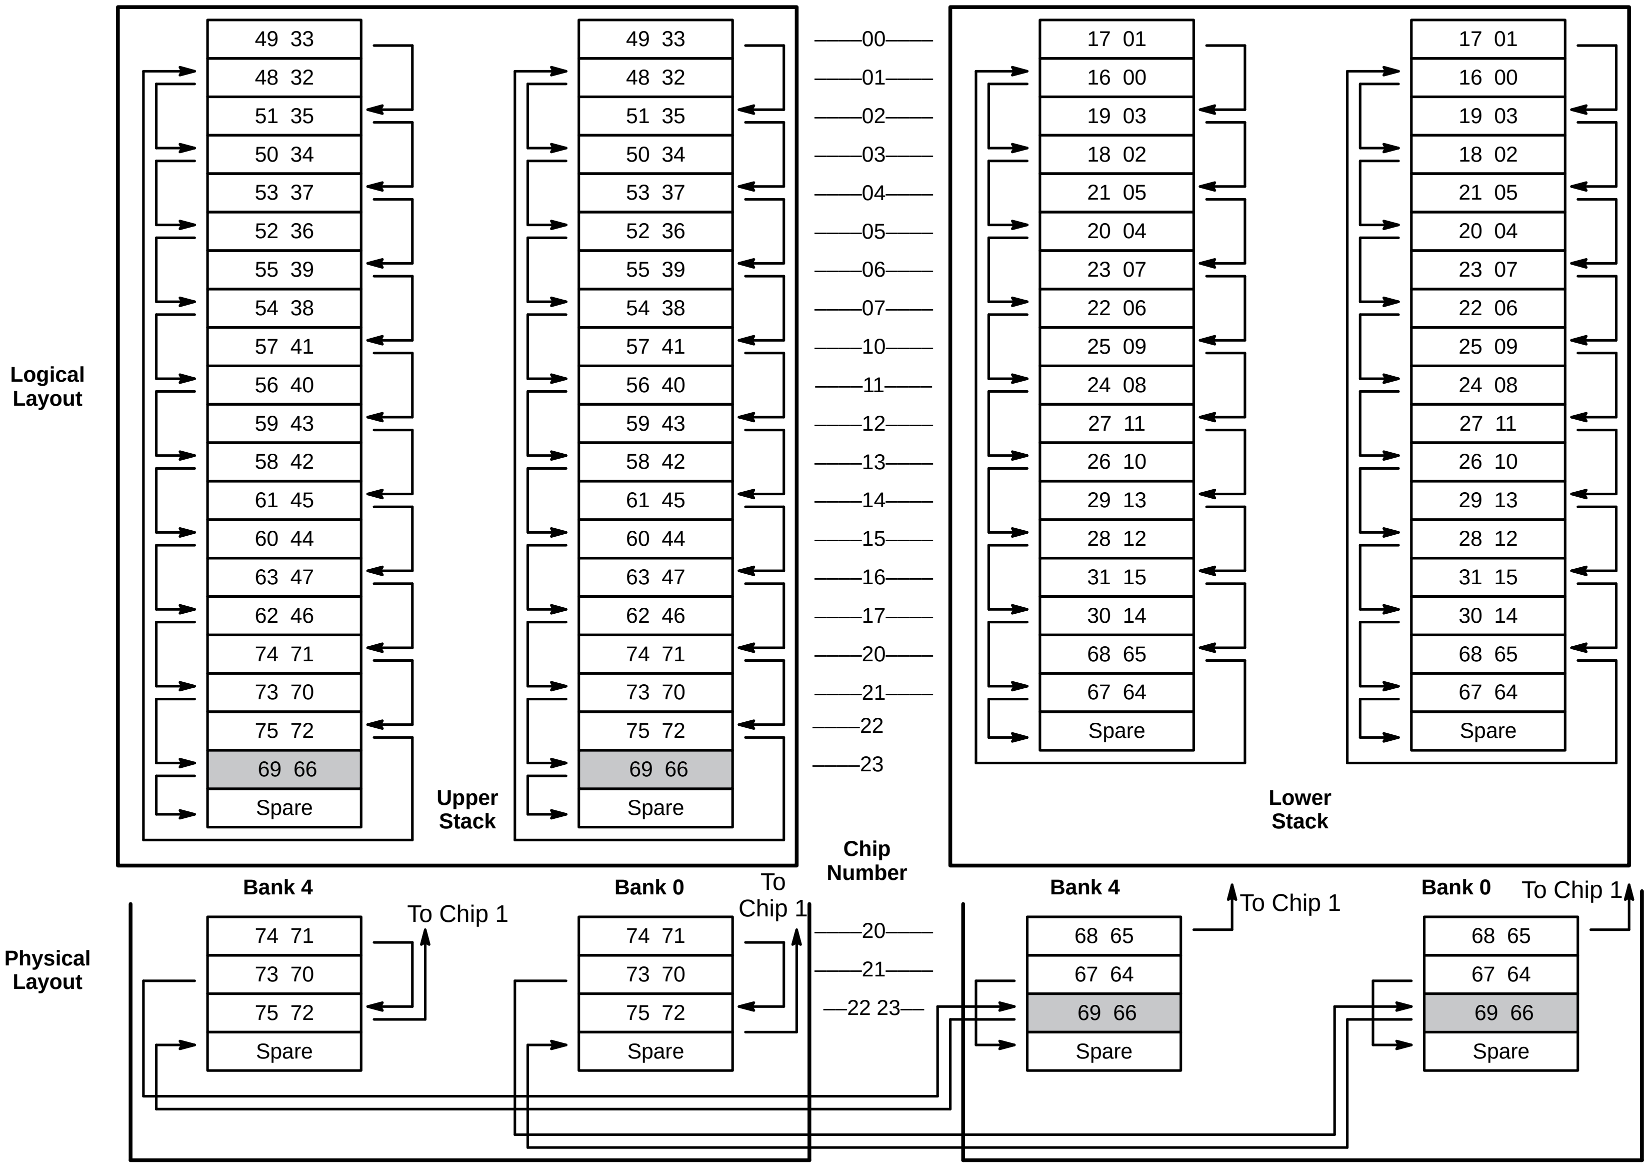

Asynchronous Memory Stack

Each memory stack handles half the bits of a word for 2 banks, and each half of the stack handles half the bits for 1 bank. Refer to Figure 26 for a drawing of how the bits are **physically** laid out.

Each memory half-stack contains 20 memory chips; 19 chips for data and 1 spare chip. The nineteenth memory chips on the lower half stack (chip 23_8) are logically assigned to the corresponding half-stacks on the upper memory stack. Refer to Figure 26. This means that the spare chip on each lower half-stack accommodates one flaw for the 18 chips assigned to that half-stack and that the spare chip on the corresponding upper half-stack accommodates one flaw for 20 chips assigned to that half-stack.

The spare chip can accommodate one flaw. If, for example, the memory chip that handles bits 1 and 17 for bank 0 is flawed out, there are no remaining flaws for this half of the stack.

Figure 26. Physical Bit Layout of Asynchronous Memory Stacks

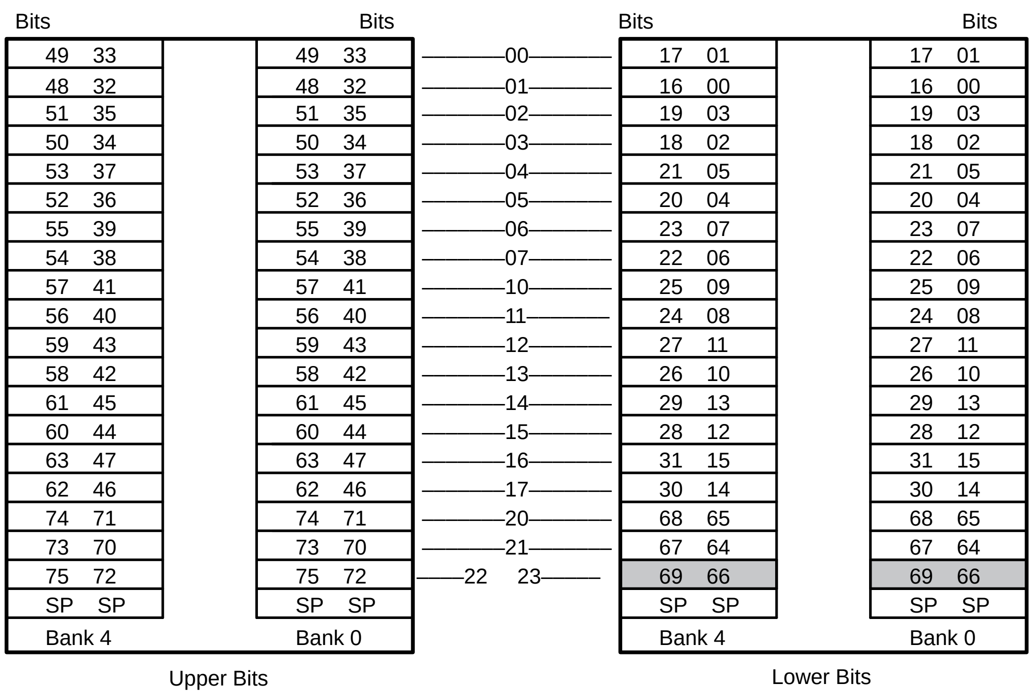

There is no chip 22 on the lower stack. Chip 23, which physically resides on the lower stack is logically assigned to the upper stack. Refer to Figure 27 for a **logical** picture of the memory stacks. The lower-bit stack spares 18 memory chips, and the upper-bit stack spares 20 memory chips.

Figure 27. Logical Bit Layout of Asynchronous Memory Stacks

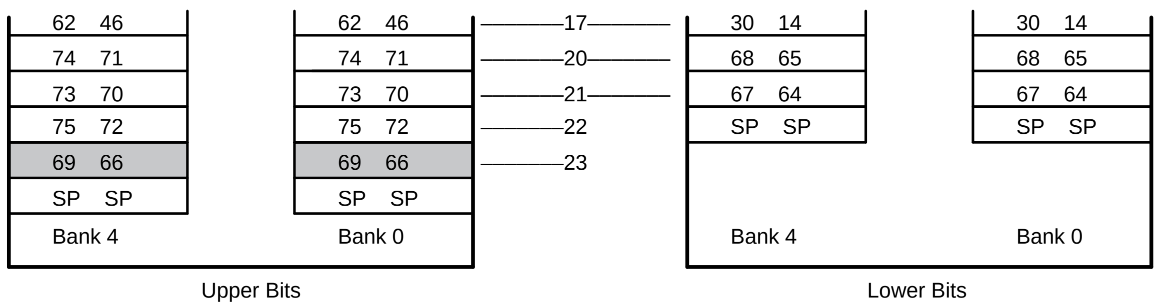

Asynchronous Bit-shift Patterns

Once a chip is flawed out, its bits are shifted on the MB options to the chip that would have received the bits during a synchronous shift, except when the bits are shifted into the spare. Refer to Figure 25. Notice on the **lower half-stack** that chip 0's bits are shifted to chip 2, chip 2 is shifted to chip 4, and so on, until chip 20 is reached. There is no chip 22 on this half-stack, so chip 20's bits are shifted to chip 1, which is shifted into chip 3, and so on, until chip 21 is reached. Chip 21 is then shifted to the spare chip on this half-stack. If chip 1 is flawed out, its bits are shifted to chip 3, chip 3's bits are shifted to chip 5, and so on, until chip 21 is reached. Because chip 23 is assigned to the upper memory stack, chip 21's bits are shifted to the spare chip on this half-stack.

If chip 0 on the **upper half-stack** is flawed out, it follows the same pattern as for the lower half-stack until it gets to chip 20. There is a chip 22 on this half-stack, so chip 20 shifts to chip 22, and chip 22 shifts to chip 1, which is shifted to chip 3, and so on, until chip 23 is reached. Chip 23 is then shifted into the spare chip on this half-stack. If chip 1 is flawed out, it, like the lower half-stack, shifts to chip 3, and so on, until chip 21 is reached. Chip 21 shifts its bits to chip 23. Remember that chip 23, while logically assigned to this half-stack, resides on the lower half-stack. Therefore, the bits are shifted over to the lower half-stack to chip 23, and chip 23 is shifted to the spare chip on the upper half-stack.

Figure 25 shows both the logical and the physical bit shifts.

Figure 28. Logical and Physical Bit Shifts - Synchronous Memory Stacks

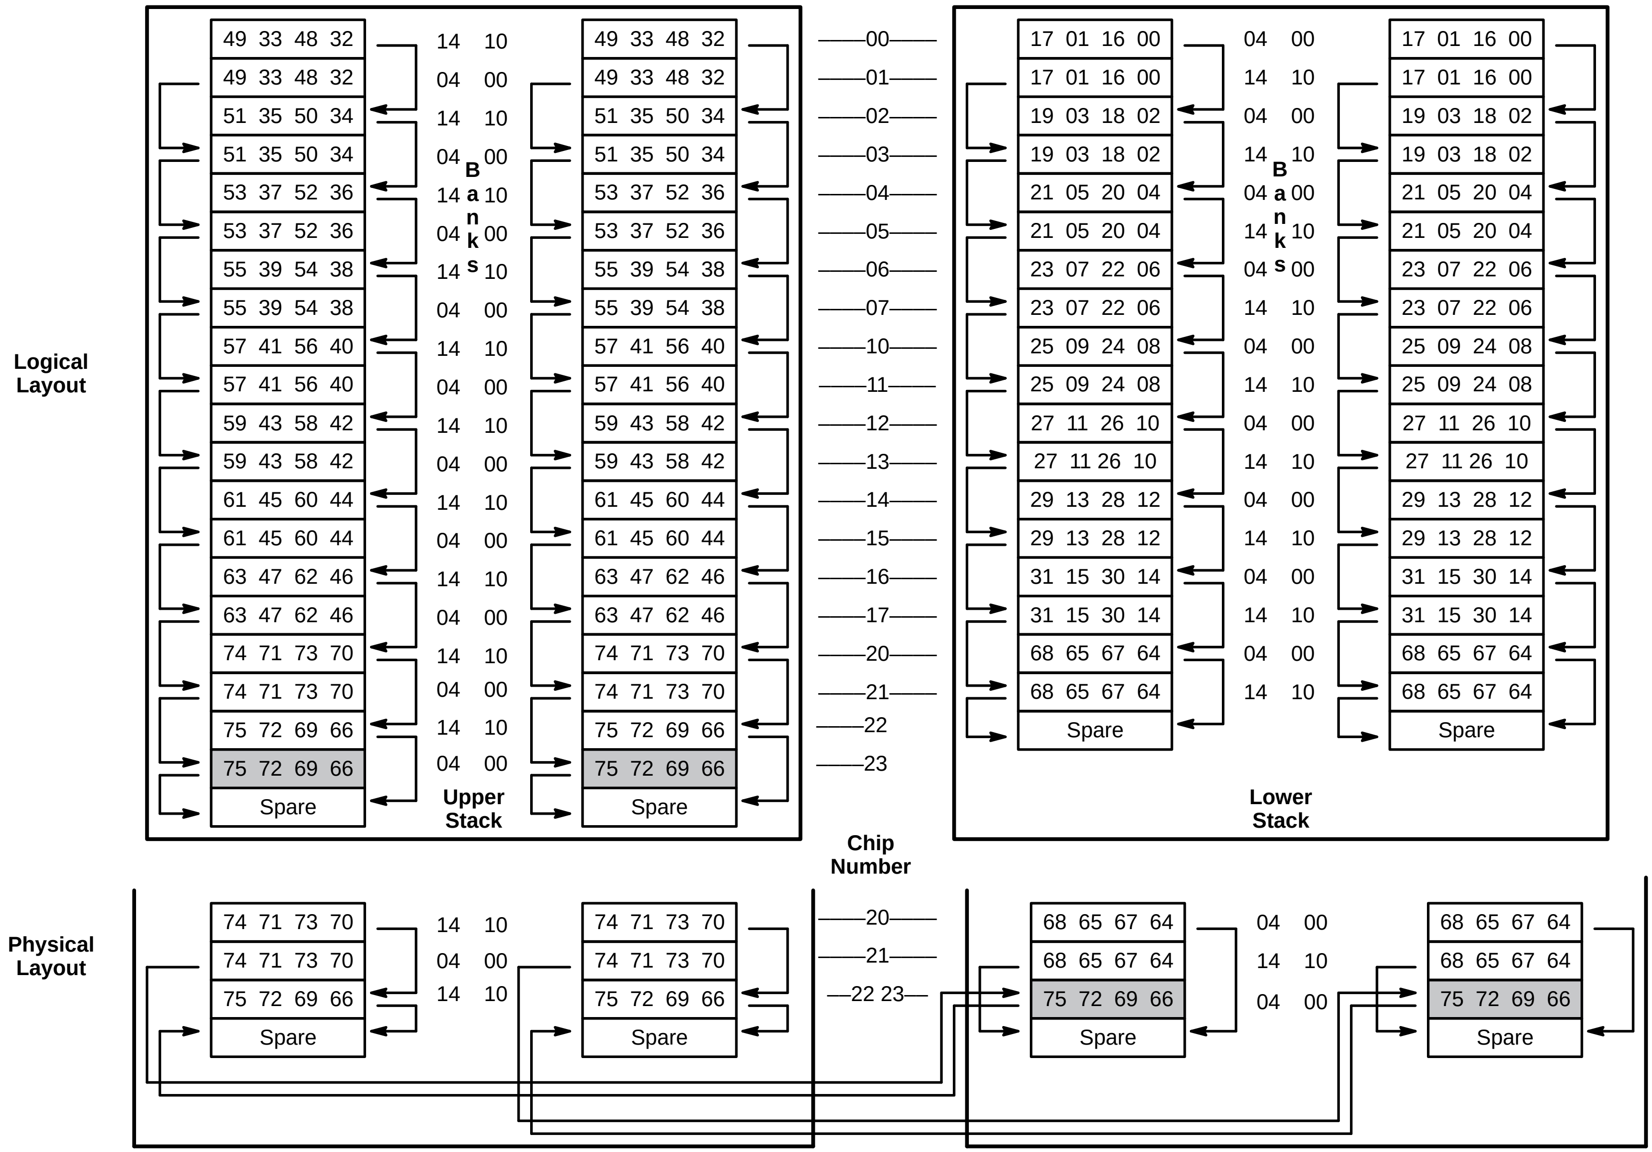

Memory Overview

The communication protocol, or *handshaking*, that occurs between a CP module and memory module and between the options on the memory module uses Valid and Resume signals. Each transmitting option contains a counter. Each receiving option contains buffers. The transmitting option can send as many valid signals to the receiving option as there are buffers in that option. With each Valid signal sent, the counter in the transmitting option increments by one. When the reference advances to the next stage, the receiving option sends a Resume signal to the transmitting option. The counter in the transmitting option then decrements by one. Figure 29 illustrates the handshaking that occurs between options.

Figure 29. Handshaking between Transmitting and Receiving Options

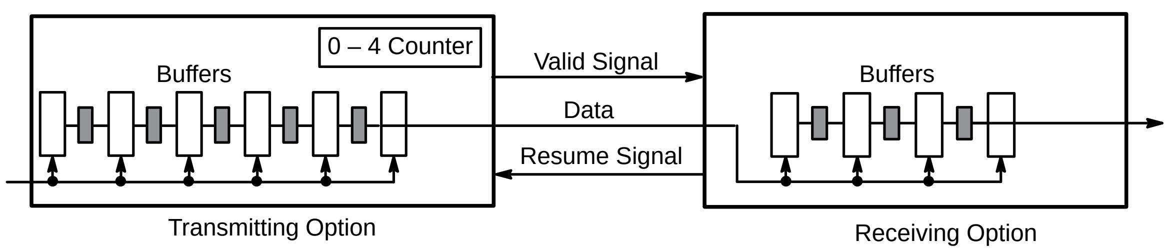

CP-to-memory Module Communications

Figure 30 shows communications between a CP and memory module during a memory reference. When a CPU references memory, it sends valid signals to memory followed by control signals and then address and data. Control information includes **steering bits**, a **Write Reference** signal, and a **Reference Code**. The steering bits direct data to the appropriate bank, the Write Reference signal notifies memory whether the reference is a read or a write reference, and the Reference Code notifies memory of the type of reference coming from the CPU. During a read reference, the CPU also sends a **Destination Code** that gets buried into the write data field. This Destination Code is sent back to the CPU with the data to ensure that the reference arrives at the appropriate location.

During a write operation, the CPU transfers data to memory in two data packets. Each data packet is 38 bits wide; 2 packets equal 1 data word. The data packets are sent in 2 clock periods (CPs). The first data packet is sent in CP 1, and the second data packet is sent in CP 2. During a read operation, the full 76-bit word is transferred to the CPU in 1 packet.

Figure 30. CPU-to-memory Communications

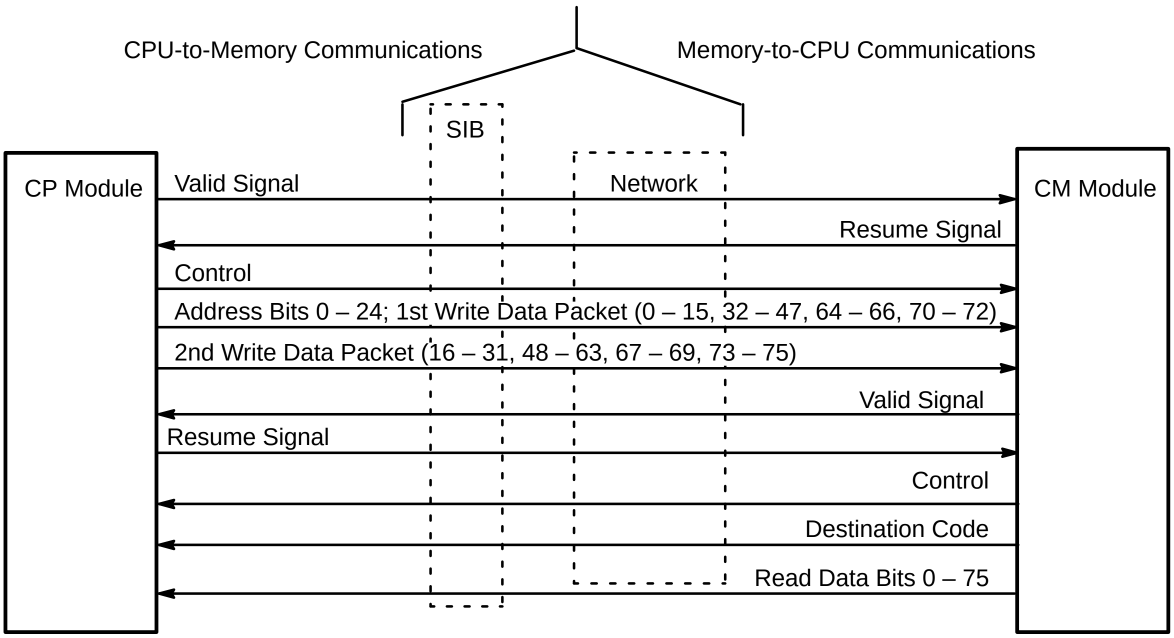

Memory Module Operations

Control, address, and data pass through one of eight ZA connectors to arrive on the memory module. Each ZA connector is associated with one CPU access path. During a memory write operation, the ZA connectors distribute control, address, and data to the MF options. The MF options buffer the control, address, and data and send it to the MC options, which multiplex it and select its bank destination. The MB options steer the control, address, and data to the bank specified by the MC option.

On a read from memory, the ZA connectors distribute the control and address to the MF options. The MF options direct the address to the MC options that select the bank to address. The MB options steer the address to the bank specified by the MC option. Figure 31 is a block diagram of the memory module that shows the path through the options on the memory module. The following subsections explain the function of each option.

Figure 31. Option Path on the Memory Module

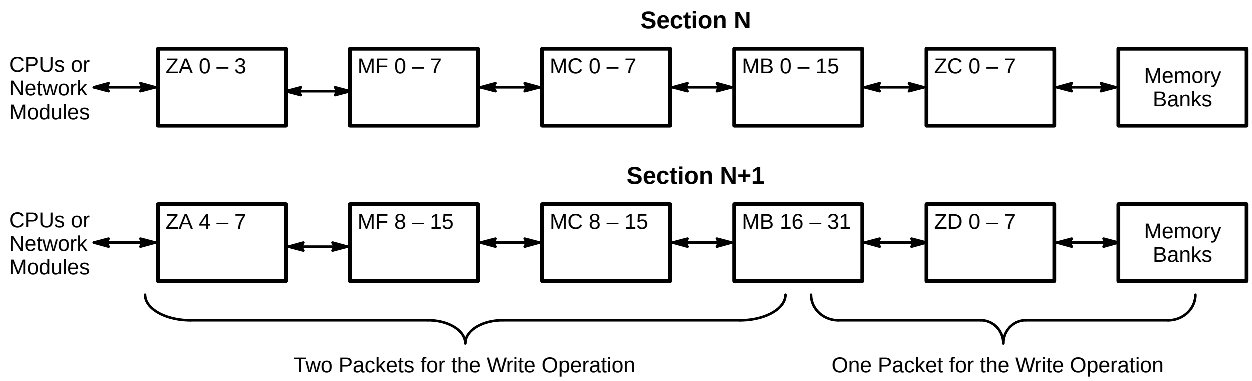

MF Option

The MF options are located adjacent to the ZA connectors (refer to Figure 1 and Figure 2); 2 MF options are associated with one ZA connector and are grouped by processor path, 2 MF options per processor path. Each option handles 32 bits of the data word, 6 of the check bits, and one-half of the address bits, steering bits, and Reference Code. The even-numbered option handles the lower bits (0 – 31 and half of the check bits), and the odd-numbered option handles the upper bits (32 – 63 and half of the check bits). Each MF option relays the data to 4 MC options during a memory write operation and outputs the data to a ZA connector during a memory read operation.

Figure 32 is a block diagram of the MF option. The MF option provides a six-buffer relay between the CPU and memory and between memory and the CPU. Memory references arrive on the memory module from the processor and are latched in the MF options in a first-in-first-out order. The reference is advanced through the six buffers by a Valid signal from the CPU (CI option) and a Resume signal from the MF option. Memory references can be made as long as a buffer is available.

Figure 32. MF Option Block Diagram

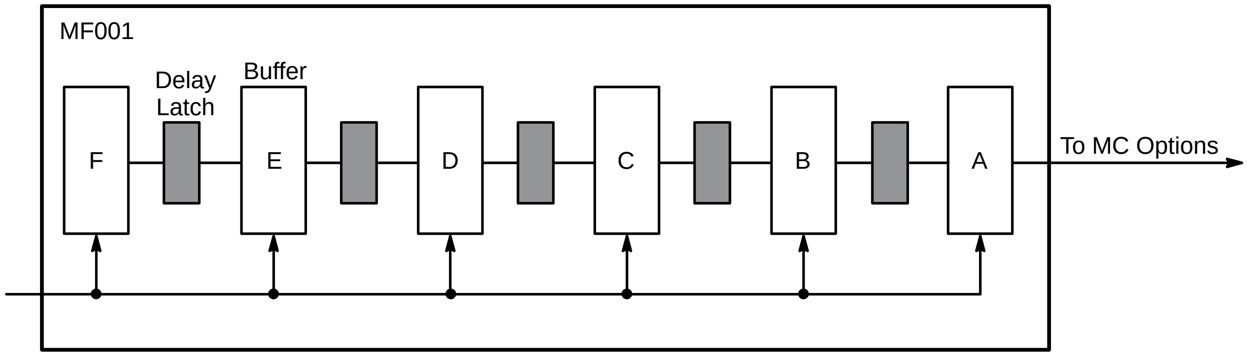

MC Option

Each memory module contains 16 MC options. The MC options handle bank and processor access. The priority is handled on a first-come, first-serve basis. When a reference gains access to a bank or processor, that bank or processor gets last priority the next time it is accessed. The MC options have four input holding buffers that enable them to stack four memory references if the output request buffers are busy.

The MC options receive 4 bank bits. Two of the bits (bank bits 0 and 1) determine which output buffer to select. The MC option sends the other 2 bits to the MB options, which use it to determine the bank to address. Refer to Figure 33.

During a write reference, each MC option receives control, address, and data from 4 MF options, performs a 4 x 4 multiplex of the control, address, and data, and then sends it to 4 MB options. The MC options receive data in two packets; the first packet contains the lower data bits, and the second packet contains the upper data bits.

During a read reference, the MC option does a 4 x 4 multiplex of the **data** and **Destination Code** coming from the MB options. **Bank bits**, from the MB option, are sent back to the CPU along with the data. The **return path** bits are used by the MC option to determine the CPU port to which the reference data must be steered.

Figure 33. MC Option Block Diagram

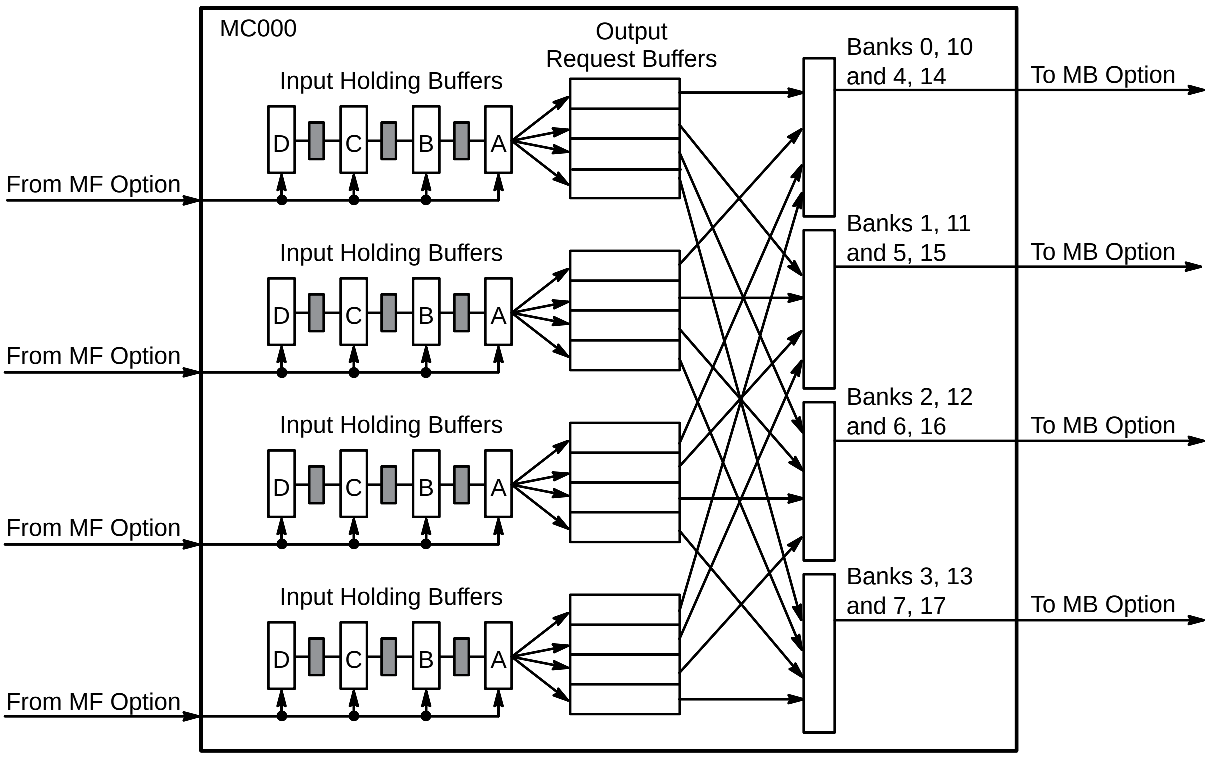

MB Option

The MB option provides all the control signals necessary to access 4 banks of memory in synchronous mode. These 4 banks are divided into two interleaved bank pairs. The interleaved banks in each pair are 180 degrees out of phase with respect to each other in terms of timing. The MB option also provides a spare chip selection feature that enables the maintenance port to configure around a bad memory chip.

When the asynchronous mode is selected on the MB option, it runs with two-bit asynchronous parts, eliminating the need to support the CM02 memory module.

The MB option receives 13 bits of address, 12 bits of packeted write data (24 bits total), 4 bits of reference code, and 5 steering and control bits from the MC multiplexer.

Each MB option controls one-fourth of the data for a 4-bank group, which requires four MBs to work together on one bank group (4 banks synchronous, 2 banks asynchronous).

The write reference information coming into the MB option first goes into a 4-rank FIFO. If the bank requested by the reference in rank A of the FIFO is available, the reference information gets stored in a bank holding register, one for each of the 4 banks, until the appropriate bank is available. This allows the next reference in the FIFO access to its bank or bank holding register without having to wait for the bank cycle slot of the first reference, unless the second reference is requesting the same bank as the first reference. Refer to Figure 34.

The MB option also receives 20 bits of readout data from each of the 4 banks, plus 4 readout spare bits, 2 bits from each bank pair. There is a 9-rank FIFO in which the readout data and the corresponding reference information is stored before it is sent out to the MC output multiplexer. The corresponding reference information includes steering, control, and destination code, which are all delayed on the MB option to line up with the appropriate readout data.

Figure 34. MB Option Block Diagram

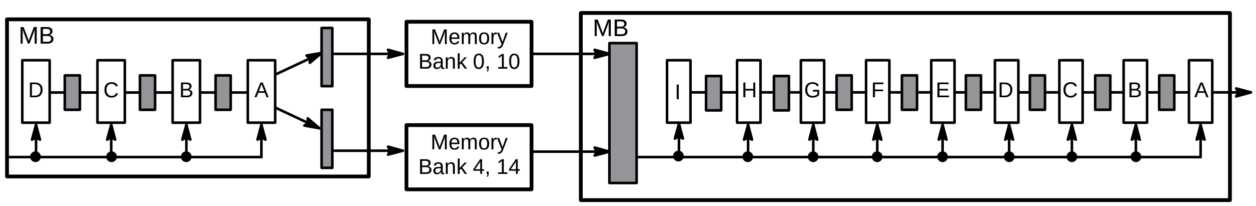

Address Distribution

CRAY T90 series systems can use up to 35 address bits to determine which section, subsection, bank, and word in memory to address; however, not all address bits are used at this time (refer to “CRAY T94 Memory,” “CRAY T916 Memory,” and “CRAY T932 Memory” on pages 14, 20, and 26 respectively, for information on addressing).

The CP module examines the lower 3 address bits (0 through 2) to determine which memory section receives the data and address. The remaining address bits enter the network module through a connector. The network module examines bits 3 through 5 to determine which subsection receives the address and data (for CRAY T916 and CRAY T932 systems only). The remaining address bits enter the memory module that determines which bank and word receives the address and data.

As address bits are distributed through the CP, network, and memory modules, the modules strip off some of the address bits after they determine the section, subsection, bank, or word to address. Then, some of the bits are replaced with Return Path bits that are used to steer the reference back to the CP module that sent the reference. Figure 35 shows the distribution of address bits through the CP, network, and memory modules.

Figure 35. Address Distribution

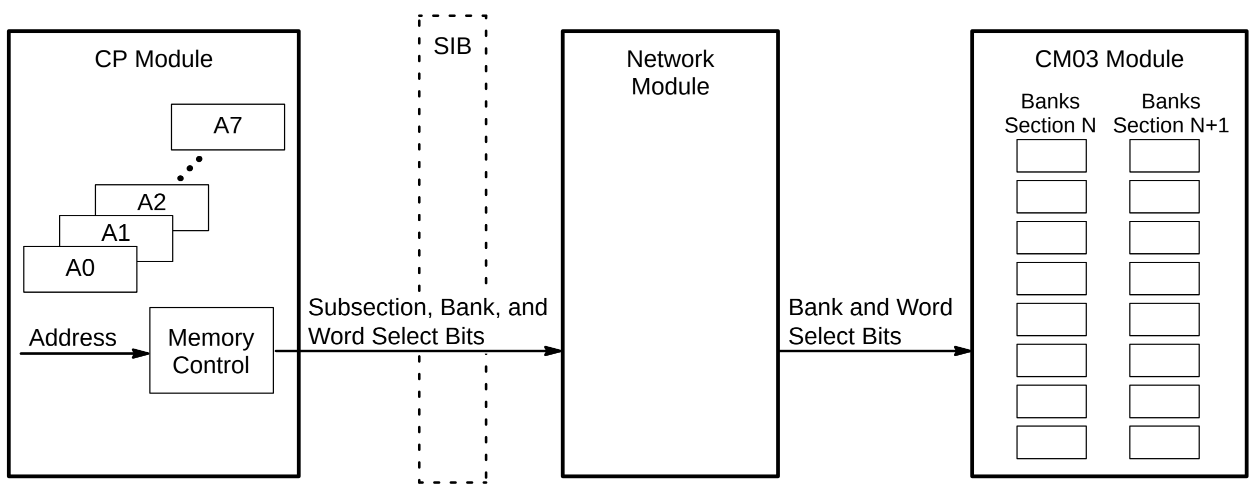

When address bits arrive on the memory module, the even-numbered MF options receive address bits 0 through 11, and the odd-numbered MF options receive address bits 12 through 24. Figure 39 is a block diagram that shows the flow of address bits through the memory module. Table 9,

Table 10, and Table 11 show the address bit assignment for the MF, MC, and MB options, respectively. As the address bits flow through each option, the options strip off the bits that are no longer needed to determine which word in memory to address. The MB option requires 20 address bits to determine the word select.

Memory Write Operation: A Block Diagram Description

When a CPU writes data into memory, it sends a **Valid** signal followed by control signals and then address bits and data. A data word is 76 bits: 64 data bits and 12 error-correction bits. The write reference arrives on the memory module through one of eight ZA connectors that are associated with a CPU access path. The ZA connectors pass the reference to the MF options, which pass the data to the MC and then to the MB options before they arrive on the memory stack.

Write Control Signals

A write reference uses the following control signals: **Bank Bits (0, 1, 2,** and **3)**, **Subsection Bits (0, 1, 2,** and **3),** a **Write Reference** signal, and **Reference Type (0, 1, 2, and 3)** bits. The MC and MB options use the bank bits to steer the address and data to the appropriate bank in memory. The **Subsection Bits** only pass through the memory module options; the network module used these bits to determine which memory section to address.

The **Write Reference** signal informs each memory module option that the reference is a write reference. This signal is needed because data is written to memory in packet form, with each option receiving 2 data packets. The first data packet contains the lower data bits, and the second data packet contains the upper data bits. If the **Write Reference** signal is equal to 1, the reference is a write reference and the memory module options can expect to receive 2 data packets. If the **Write Reference** signal is equal to 0, then the reference is a read reference and the options can expect to receive 1 data packet.

The CPU generates the **Reference Type** bits and sends them to the memory module with every reference. Table 6 lists the MB option input terms associated with the **Reference Type** bits and decodes these bits by CPU function. Refer to Figure 38.

Figure 39 shows the control signals associated with a memory write operation; however, the figure illustrates only the options associated with a write operation for banks 0, 4, 10, and 14 being sent from processor 0. Sending a complete data word to memory requires 2 MF options, 8 MC options, and 4 MB options.

Table 6. MB Option Reference Type Bit Decode

Function	ICD	ICC	ICB	ICA
Abort	0	0	0	0
CPU or I/O read	0	0	0	1
CPU write	0	0	1	0
Not valid function	0	0	1	1
Set spare chip config.	0	1	0	0
Not valid function	0	1	0	1
Read spare chip config.	0	1	1	0
Not valid function	0	1	1	1
I/O write	1	X	X	X

NOTE: X = either a 1 or a 0.

If the reference is an I/O write reference, the MB option sends 3 of the **Reference Type** bits (ICA, ICB, and ICC) with 4 forced bits back to the CPU as the **Destination Code**. If the reference is a CPU write reference, then the MB option forces 7 **Destination Code** bits and passes them back to the CPU. The CPU uses the **Destination Code** bits to determine that the reference actually completed. Table 7 lists the Destination Code bits forced by the MB option. Figure 40 shows how the MB distributes the **Reference Type** bits and how the **Destination Code** bits are sent back to the CP module via the MC and MF options.

Table 7. MB Option Destination Code Forced Bits

Function	MB0 OGD	MB1 OGC	MB0 OGC	MB1 OGB	MB0 OGB	MB1 OGA	MB0 OGA
Read/abort	IAD	IAC	IAC	IAB	IAB	IAA	IAA
I/O write	0	0	0	ICC	ICB	ICA	1
Processor write	0	0	0	0	0	1	0
Set bad bit (Reconfigure)	0	0	0	0	1	1	0

The CPU counts the number of references that it sent out and then checks this count against the number of references that actually completed. This information is sent to the J series options on the CP module, which determines if memory is quiet or if the CPU should generate a hold issue condition.

Write reference control signals pass through the MF options and enter the MC options. Table 9 lists the control and data bits that leave each MF option. Two MF options are needed to direct all the control signals from one processor to memory.

The MC options use the **Bank Bits** to steer the address and data to the appropriate bank in memory. The MC options use **Bank Bits 0** and **1** to steer the reference to one of the four 4-bank groups. **Bank Bit 2** and **3** signals are passed along to the MB options where they are used to determine which bank within the bank group to address. Once the MC options determine which bank group to address, they drop **Bank Bits 0** and **1** and create **Return Path Bits 0 and 1**. The **Return Path Bit** is a 2-bit code that is forced to either a 0 or a 1 to designate which processor connector the reference came from. The MB options pass the **Return Path Bits** back to the MC options where they are used to steer the reference response back to the initiating CPU.

The other write reference control signals pass through the MC options and enter the MB options. Table 10 lists the control, address, and data bits that leave each MC option. Eight MC options handle the control, address, and data for four processors. Options MC000 through MC007 handle control, address, and data for processor connectors 0 through 3; options MC008 through MC015 handle control, address, and data for processor connectors 4 through 7.

The 16 MB options handle control, address, and data for four processors. Options MB000 through MB015 handle control, address, and data for processor accesses 0 through 3; options MB016 through MB031 handle control, address, and data for processor accesses 4 through 7. Table 11 lists the control, address, and data received by each MB option.

The MB options generate two control signals that enter the memory stack: **Write Enable**, and **Clock**. The **Chip Select** signal is forced. The **Write Enable** signal must be present for a write reference to complete.

Write Data Path

Figure 36 shows the write data path for processors 0 through 7. Refer to this illustration for detailed descriptions of the data flow between options on a memory module.

Data enters the MF options in packet form; the lower bits (0 through 31 and 64 through 69) arrive on the even-numbered MF options in 2 clock periods, and the upper bits (32 through 63 and 70 through 75) arrive on the odd-numbered MF options. Each MF option sends data to 4 MC options. The MC options distribute the data to the MB options, which direct it to the appropriate memory bank.

Write Completion

The terms listed below are sent back to the CPU that originated the write reference in order to inform the CPU that the write reference has completed.

Return Path bits 0 and 1 steer the reference response back to the proper CPU access port (ZA0 through ZA7). The bits are forced on the MC option to the value of the port number after the reference leaves the MF options for the MC options. The return path bits travel with the reference to the MB options and then back to the MC options to steer the response to the proper MF options.

The **Valid** signal is sent back to the CPU to inform it that the reference has completed. It is originated by the MB option and sent back with the other response bits.

Destination Code bits are formed on the MB option and returned to the CPU with the other response bits. To learn how these bits are formed, refer to page 48.

The **Subsection** and **Bank** bits are also returned to the CPU with the other response bits.

Memory Read Operation: A Block Diagram Description

When a CPU reads data from memory, it sends a **Valid** signal, control signals, and an address to the memory module. The **Valid** and control signals, address and destination code arrive on the memory module through one of eight ZA connectors that are associated with a CPU access path. The read destination code occupies the first packet of what would be the write data. These signals flow through the MF, MC, and MB options and steer the read reference to the designated location in the memory stack to retrieve the data.

Data leaves the memory stack and enters the MB options. The MB options send the data and control signals back to the CPU through the MC and MF options.

Read Control Signals

Figure 40 shows the control signals associated with a memory read operation; however, the figure illustrates only the options associated with a read reference for banks 0,10 and 4,14 for processor 0.

A read reference uses the same control bits as a write reference to get to the memory stack: **Subsection** bits 0 through 3, **Bank** bits 0 through 3, **Reference Type** bits 0 through 3, and a **Write Reference** signal (which should be zero on a read reference).

On the way back to the CPU a read reference carries with it the following control signals: **Subsection** bits 0 through 3, **Bank** bits 0 through 6, and **Destination Code** bits 0 through 13.

On the way to the memory module, the subsection bits were stripped off by the network module and replaced with return information. They flow through the options on the memory module and are used on the way back by the network module to steer the reference back to the originating CPU. On the way to memory, bank bits 0 through 3 are used by the MC and MB options to steer the reference to the correct bank. They are then returned with 3 extra bits forced **Bank bits 4 through 6** so that the CPU can determine which bank the data came from.

The CPU generates the **Reference Type** bits**.** The MB options pull the **Destination Code** out of the write data field on a read reference. Table 6 lists the terms associated with the **Reference Type** bits and decodes the bits by CPU function. When the MB option detects a read reference, it checks the write data field and sends the **Destination Code** back to the CPU with the return data. Table 7 lists the **Destination Code** bits that leave the MB options. The CPU uses the **Destination Code** to determine what to do with the data. Refer to the *CPU Module (CP02)* document, publication number HTM-003-0, for more information on how the CPU decodes the Destination Code bits.

The MB options also receive a **Return Path** bit. The **Return Path** is a 2-bit code that steers the reference back to the CPU that sent the reference. Read reference control signals pass through the MB options and enter the MC options. Table 11 lists the control and data bits that leave each MB option.

The MC options use the **Return Path** bits to determine which processor the reference came from; the MC options then drop these bits and add two forced **Bank Bits.** The CPU uses the **Bank Bits** to determine which memory bank the data came from.

During a read reference, bank bits 0 through 6 are reported on options MC0 through 3, and MC5, and on MC8 through 11 and MC13. These bank bits are used by the CPU for error reporting. At this time, bank bits 4 through 6 (MC0, MC5, MC8, and MC13) are unused and forced to zeros. Bank bits 0 and 1 were dropped on the way out to memory, but are re-created by forcing inputs on the MC option, depending on the 4-bank group from which the reference is returning. Table 8 shows how these bank bits are forced on the inputs to MC0 and MC9.

Table 8. Bank Bits for Error Correction

MC001, MC009	MC001, MC009	Bank Number
Bank Bit 1	Bank Bit 0	
Forced 0 (IEN)	Forced 0 (IEM)	Bank 0, 10, 4, 14
Forced 0 (IFN)	Forced 1 (IFM)	Bank 1, 11, 5, 15
Forced 1 (IGN)	Forced 0 (IGM)	Bank 2, 12, 6, 16
Forced 1 (IHN)	Forced 1 (IHM)	Bank 3, 13, 7, 17

Control signals leave the MC options and enter the MF options. Table 10 lists the data and control signals that leave the MC options. Table 9 lists the data and control signals that enter the MF options. The MF options send the data and control signals back to the CPU.

Read Data Path

Figure 37 shows the read data path for processors 0 through 7.

Data leaves the memory stack as a 76-bit data word and enters the MB options. The spare bit is configured in the memory stack when a bad memory chip exists. If this bit is used, the MB options shift the bits to bypass the bad chip during a write operation and shift them back during a read operation. Refer to “Chip Flawing” on page 32 for more information.

Data leaves the MB options and flows through the MC and MF options before it leaves the memory module through the ZA connectors.

Table 9. MF Option Bit Assignments

Processor 0	
MF 000	MF 001
Processor 1	
MF 002	MF 003
Processor 2	
MF 004	MF 005
Processor 3	
MF 006	MF 007
Processor 4	
MF 008	MF 009
Processor 5	
MF 010	MF 011
Processor 6	
MF 012	MF 013
Processor 7	
MF 014	MF 015
Reference to Memory Signals	
Write Data 0 – 31, 64 – 69/Read Dest. 0 – 6	Write Data 32 – 63, 70 – 75/Read Dest 7 – 13
Address bits 0 – 11	Address bits 12 – 23
Subsection bits 0 – 01	Subsection bits 02 – 03
Bank bits 00, 02	Bank bits 01, 03
Reference type 0 – 01	Reference type 02 – 03
Write reference bit for lower bits	Write reference bit for upper bits
Valid in for lower bits	Valid in for upper bits
Resume from lower bits	Resume from upper bits
Reference from Memory Signals	
Read data 0 – 31, 64 – 69	Read data 32 – 63, 70 – 75
Address bits 0 – 11	Address bits 12 – 23
Subsection bits 0 – 01	Subsection bits 02 – 03
Bank bits 0 – 2	Bank bits 3 – 6
Destination bits 0 – 6	Destination bits 7 – 13
Valid out from lower bits	Valid out from lower bits
Resume in from lower bits	Resume in from upper bits

Table 10. MC Option Bit Assignments

	Lower Bits				Upper Bits			
Proc. 0 – 3	MC00	MC01	MC02	MC03	MC04	MC05	MC06	MC07
Proc. 4 – 7	MC8	MC9	MC10	MC11	MC12	MC13	MC14	MC15
Write Data	Even Bits	Odd Bits	Even Bits	Odd Bits	Even Bits	Odd Bits	Even Bits	Odd Bits
1st Data Packet	00 – 08	01 – 09	10 – 14 64	11 – 15 65	32 – 40	33 – 41	42 – 46 66, 70	43 – 47 71, 72
2nd Data Packet	16 – 24	17 – 25	26 – 30 67	27 – 31 68	48 – 56	49 – 57	58 – 62 69, 73	59 – 63 74, 75
Addr Bits †	18 – 23	06 – 11	12 – 17	00 – 05	18 – 23	06 – 11	12 – 17	00 – 05
Ref Type	2 – 3	2 – 3	0 – 1	0 – 1	2 – 3	2 – 3	0 – 1	0 – 1
SubS Bit			1	0			3	2
Bank Bit	2, 3	2, 3			2, 3	2, 3		
Return Bit ‡			1	0			1	0
Write Ref.	X	X	X	X	X	X	X	X
Read Data	0 – 9	10 – 15 64 – 65	16 – 25	26 – 31 67, 68	32 – 41	42 – 47 66, 70 – 72	48 – 57	58 – 63 69 73 – 75
Destination	0, 2	4,6	1,3	5	7, 9	11, 13	8, 10	12
Return Bit	0, 1	0, 1	0, 1	0, 1	0, 1	0, 1	0, 1	0, 1
SubS Bit			1	0			3	2
Bank Bit			3	2				
Force Bank Bit	4, 5	0, 1				6		

† Chip address

‡ Forced

Table 11. MB Option Bit Assignments

	Processors 0 – 3				Processors 4 – 7			
Banks 0, 4, 10, 14	MB00	MB01	MB02	MB03	MB16	MB17	MB18	MB19
Banks 1, 5, 11, 15	MB04	MB05	MB06	MB07	MB20	MB21	MB22	MB23
Banks 2, 6, 12, 16	MB08	MB09	MB10	MB11	MB24	MB25	MB26	MB27
Banks 3, 7, 13, 17	MB12	MB13	MB14	MB15	MB28	MB29	MB30	MB31
Write Data	Even 00 – 30	Odd 01 – 31	Even 32 – 62	Odd 33 – 63	Even 00 – 30	Odd 01 – 31	Even 32 – 62	Odd 33 – 63
	64, 66, 67, 69	65, 68, 72, 75	70, 73	71, 74	64, 66, 67, 69	65, 68, 72, 75	70, 73	71, 74
Extra Copies for logical chip sparing	70, 73	71, 74	66, 69	72, 75	70, 73	71, 74	66, 69	72, 75
Read Data	0 – 15 64 – 65	16 – 31 67 – 68	32 – 47 66, 70 – 72	48 – 63 69, 73 – 75	0 – 15 64 – 65	16 – 31 67 – 68	32 – 47 66, 70 – 72	48 – 63 69, 73 – 75
Destination	0, 2, 4, 6	1, 3, 5	7, 9, 11, 13	8, 10, 12	0, 2, 4, 6	1, 3, 5	7, 9, 11, 13	8, 10, 12
Return	1	0	1	0	1	0	1	0
SSec Bit	1	0	3	2	1	0	3	2
Bank Bit	3	2			3	2		

Figure 36. CM03 Memory Module Write Data Path

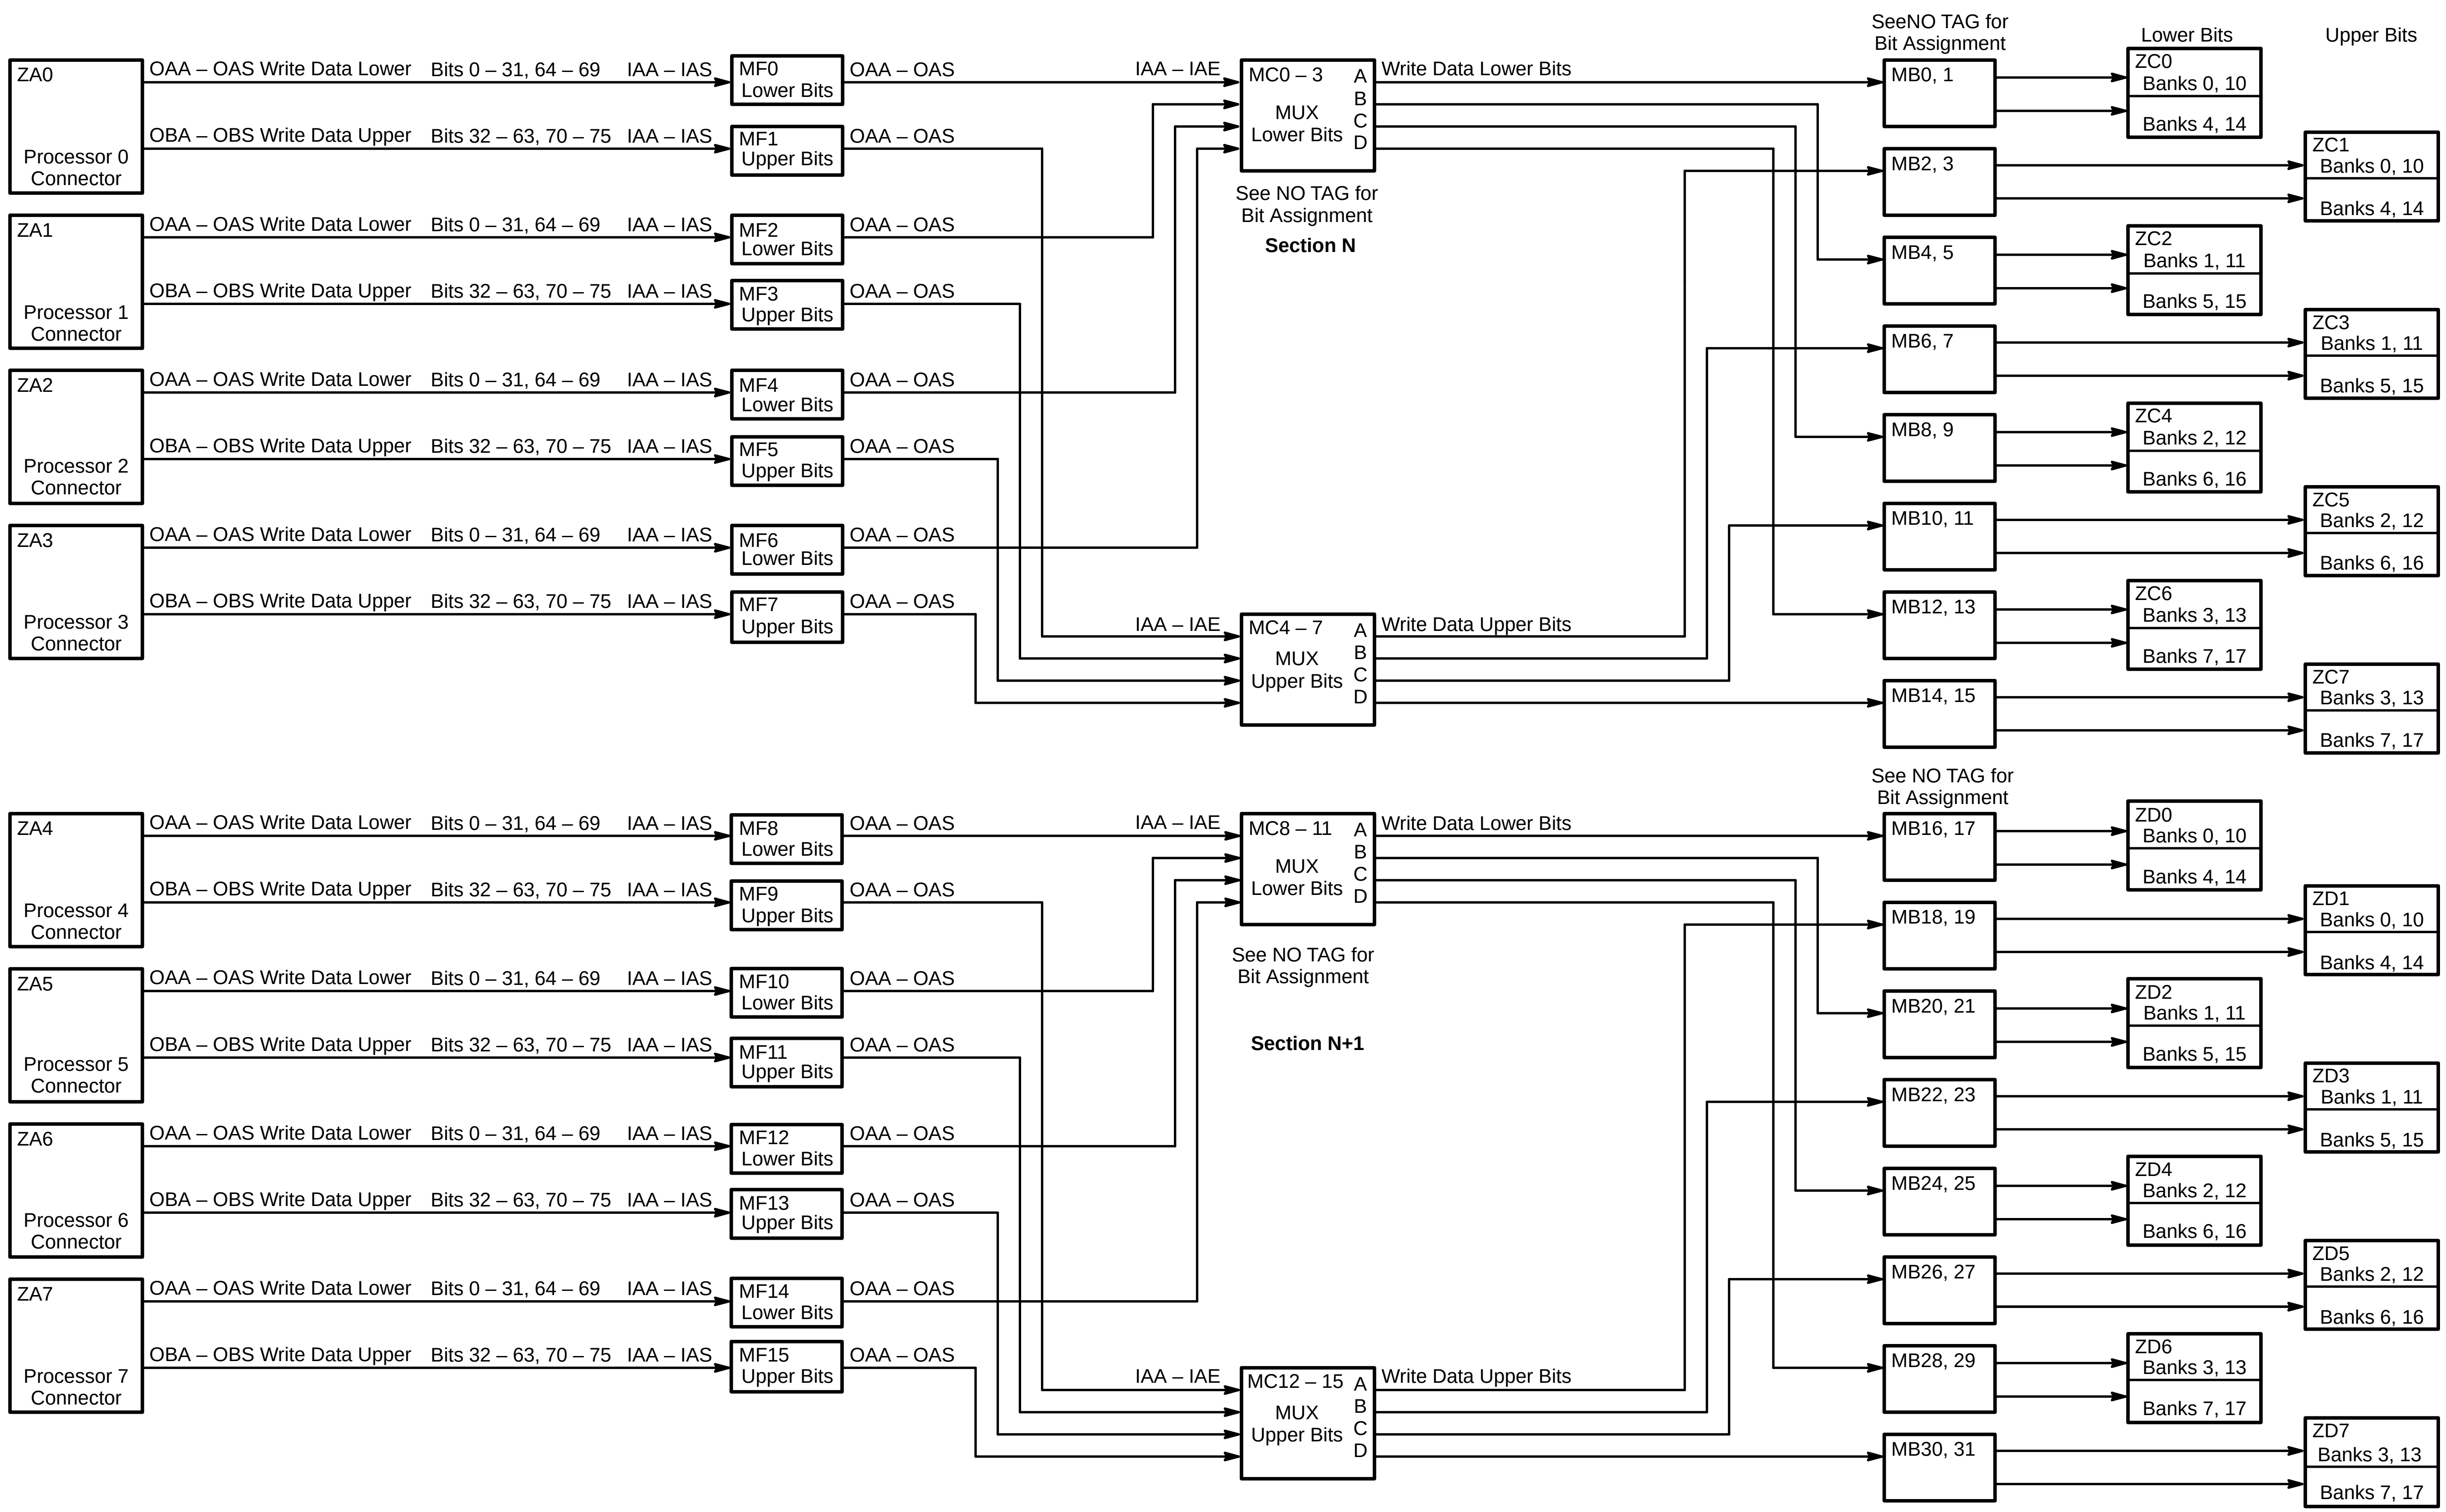

Figure 37. CM03 Memory Module Read Data Path

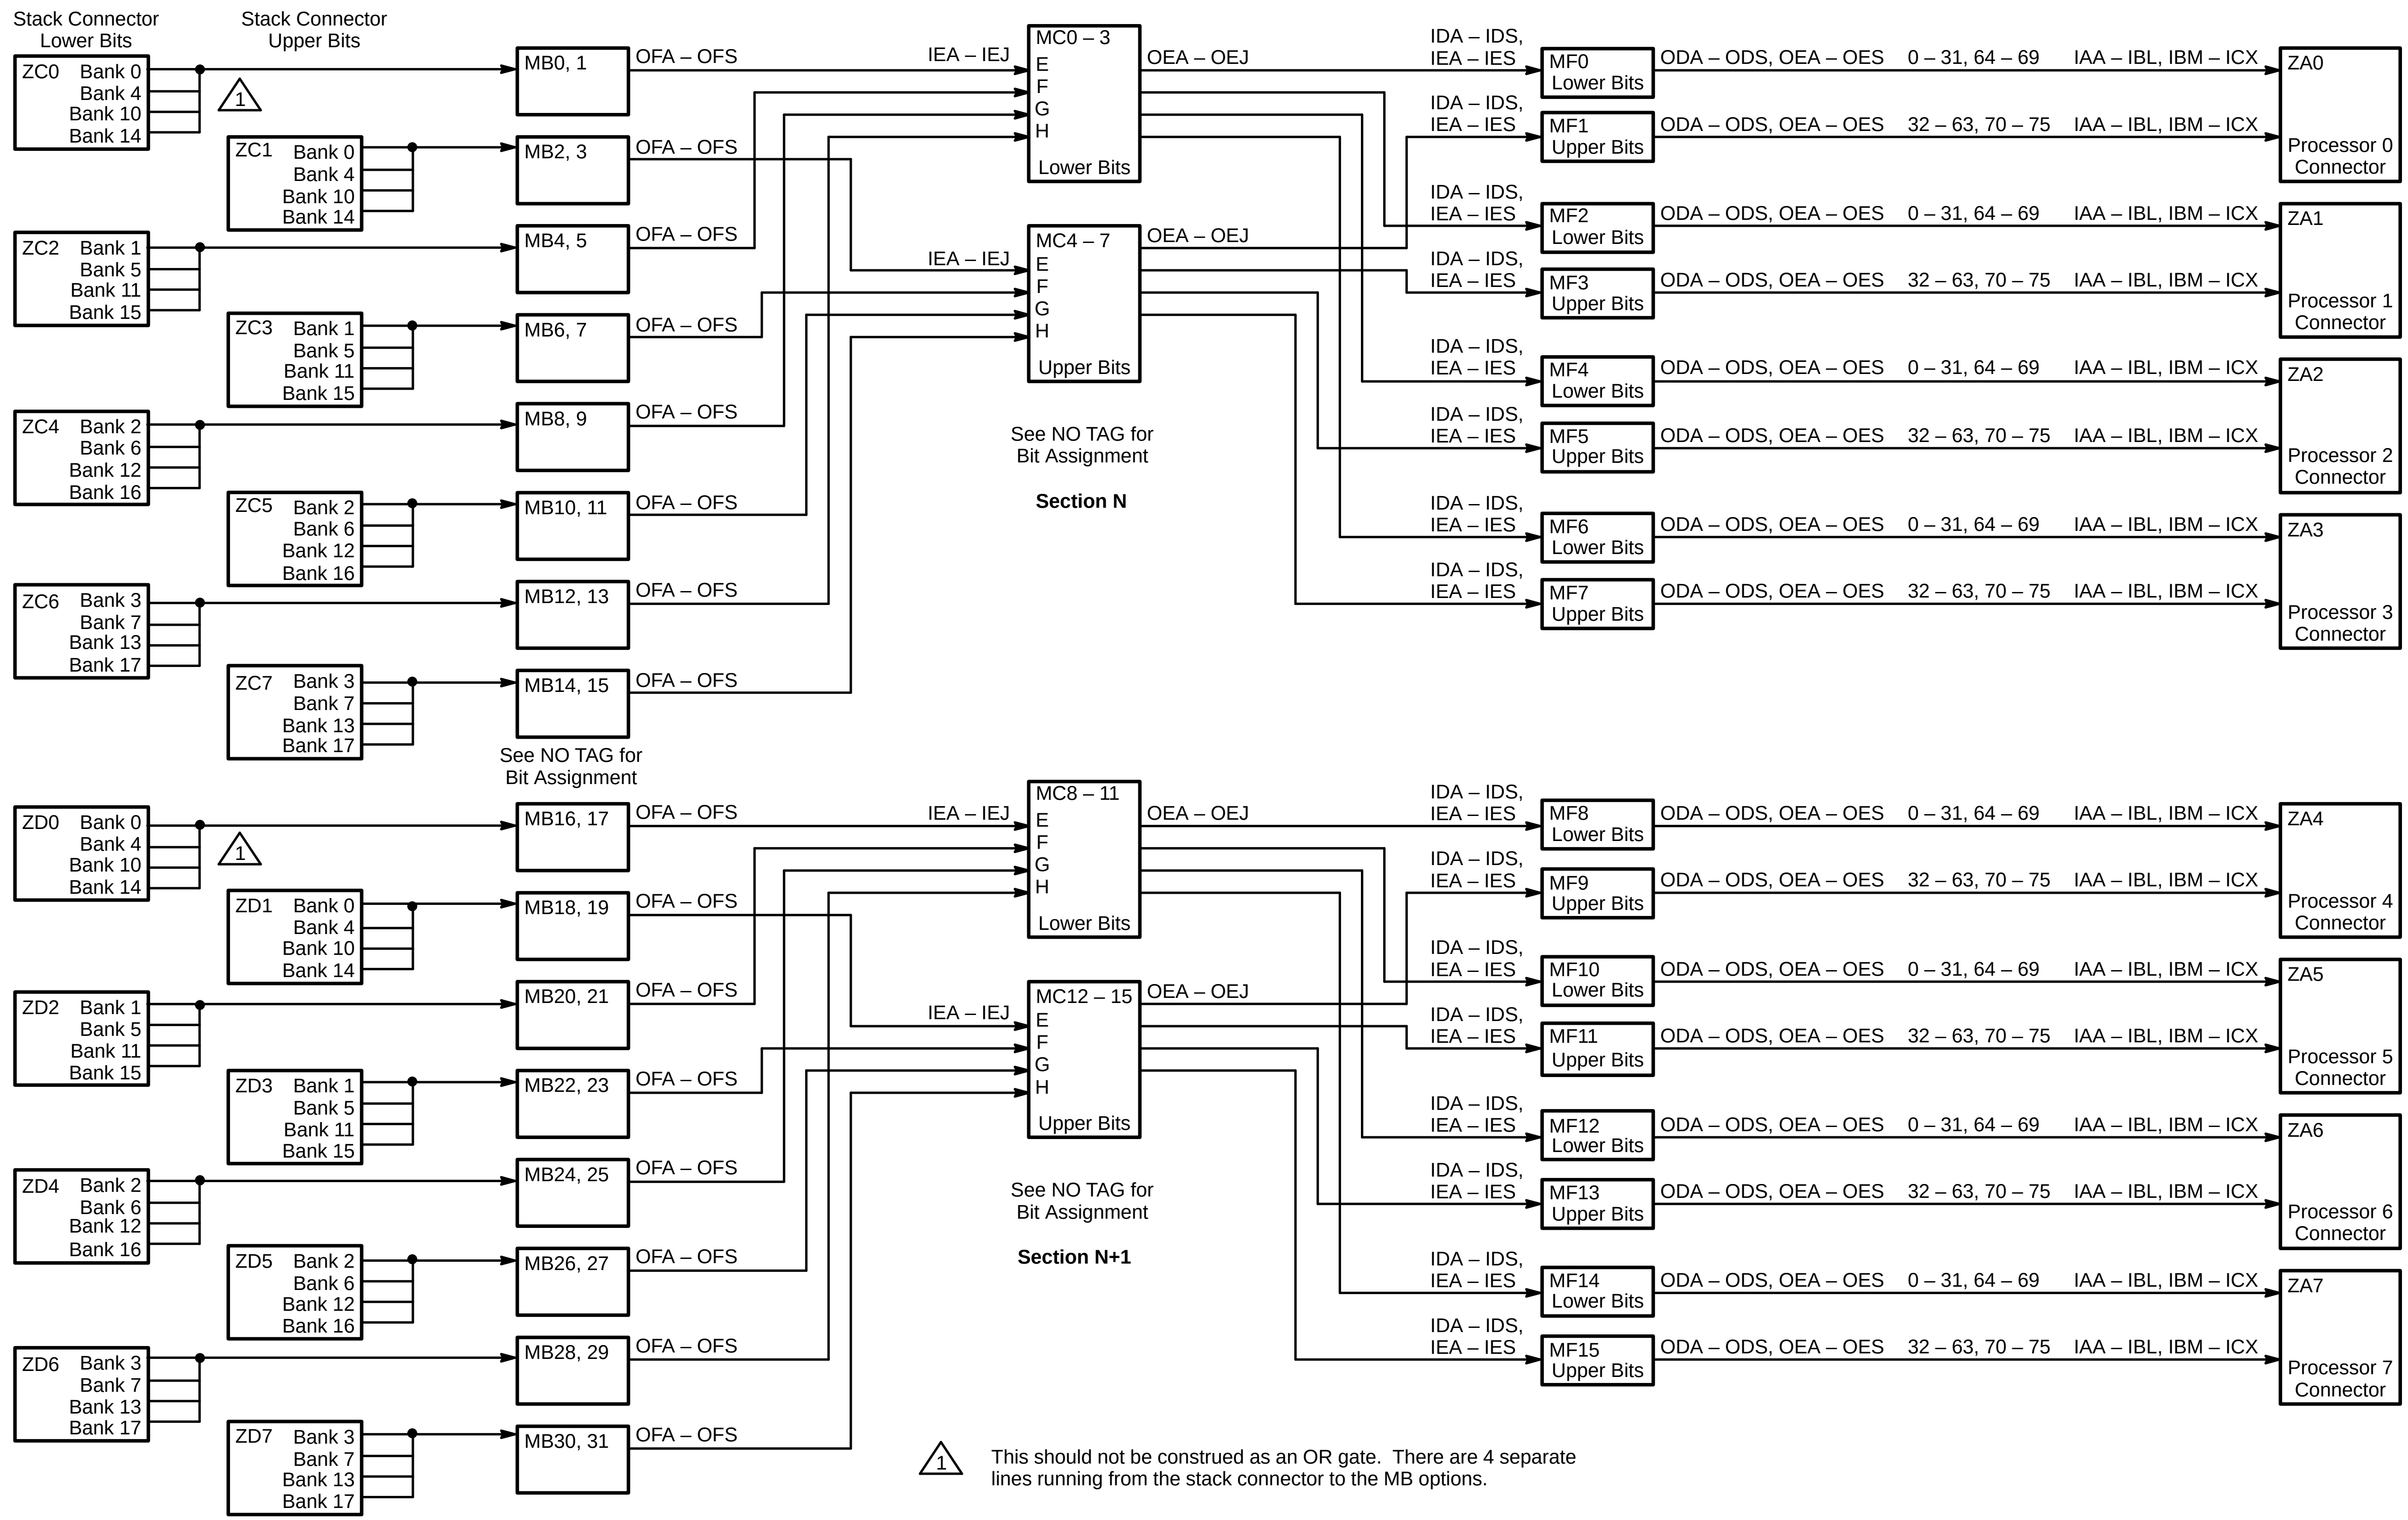

Figure 38. Reference Type Bit Generation on CI option (CP02 Module)

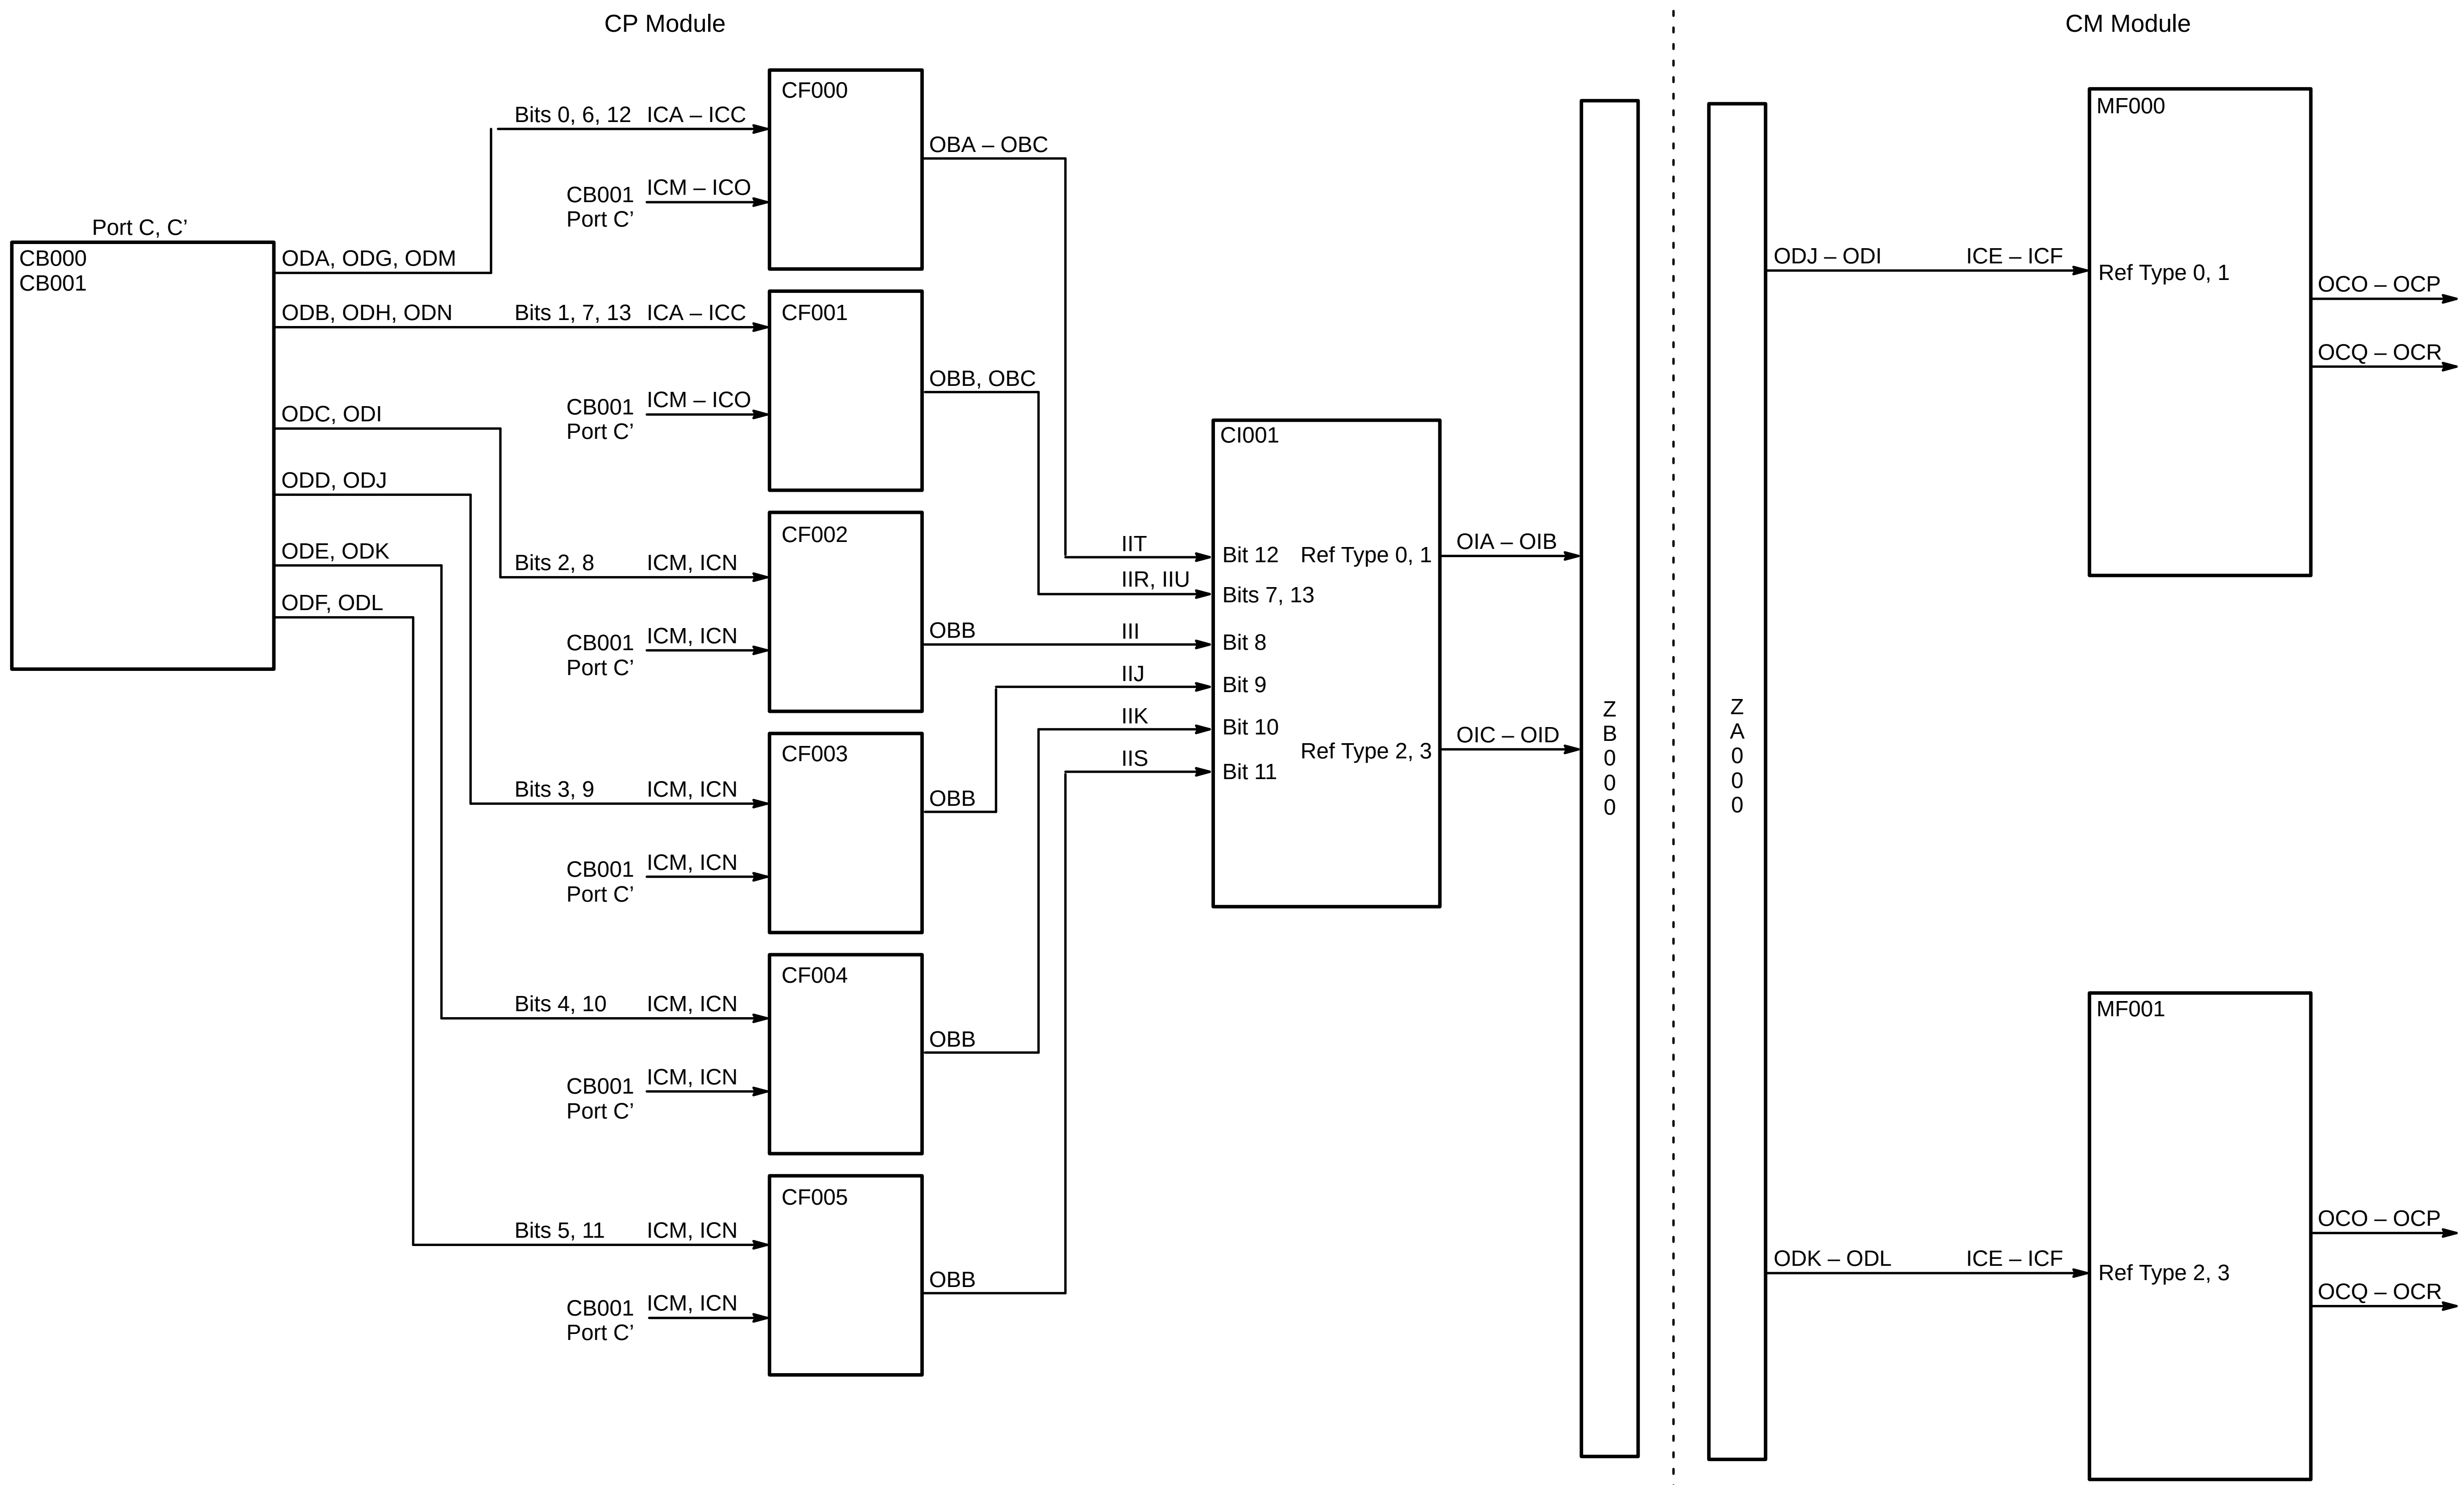

Figure 39. CM03 Memory Module Processor-to-memory Address and Control Signals for Processor 0 to Section N, Banks 0, 4, 10, 14 Only

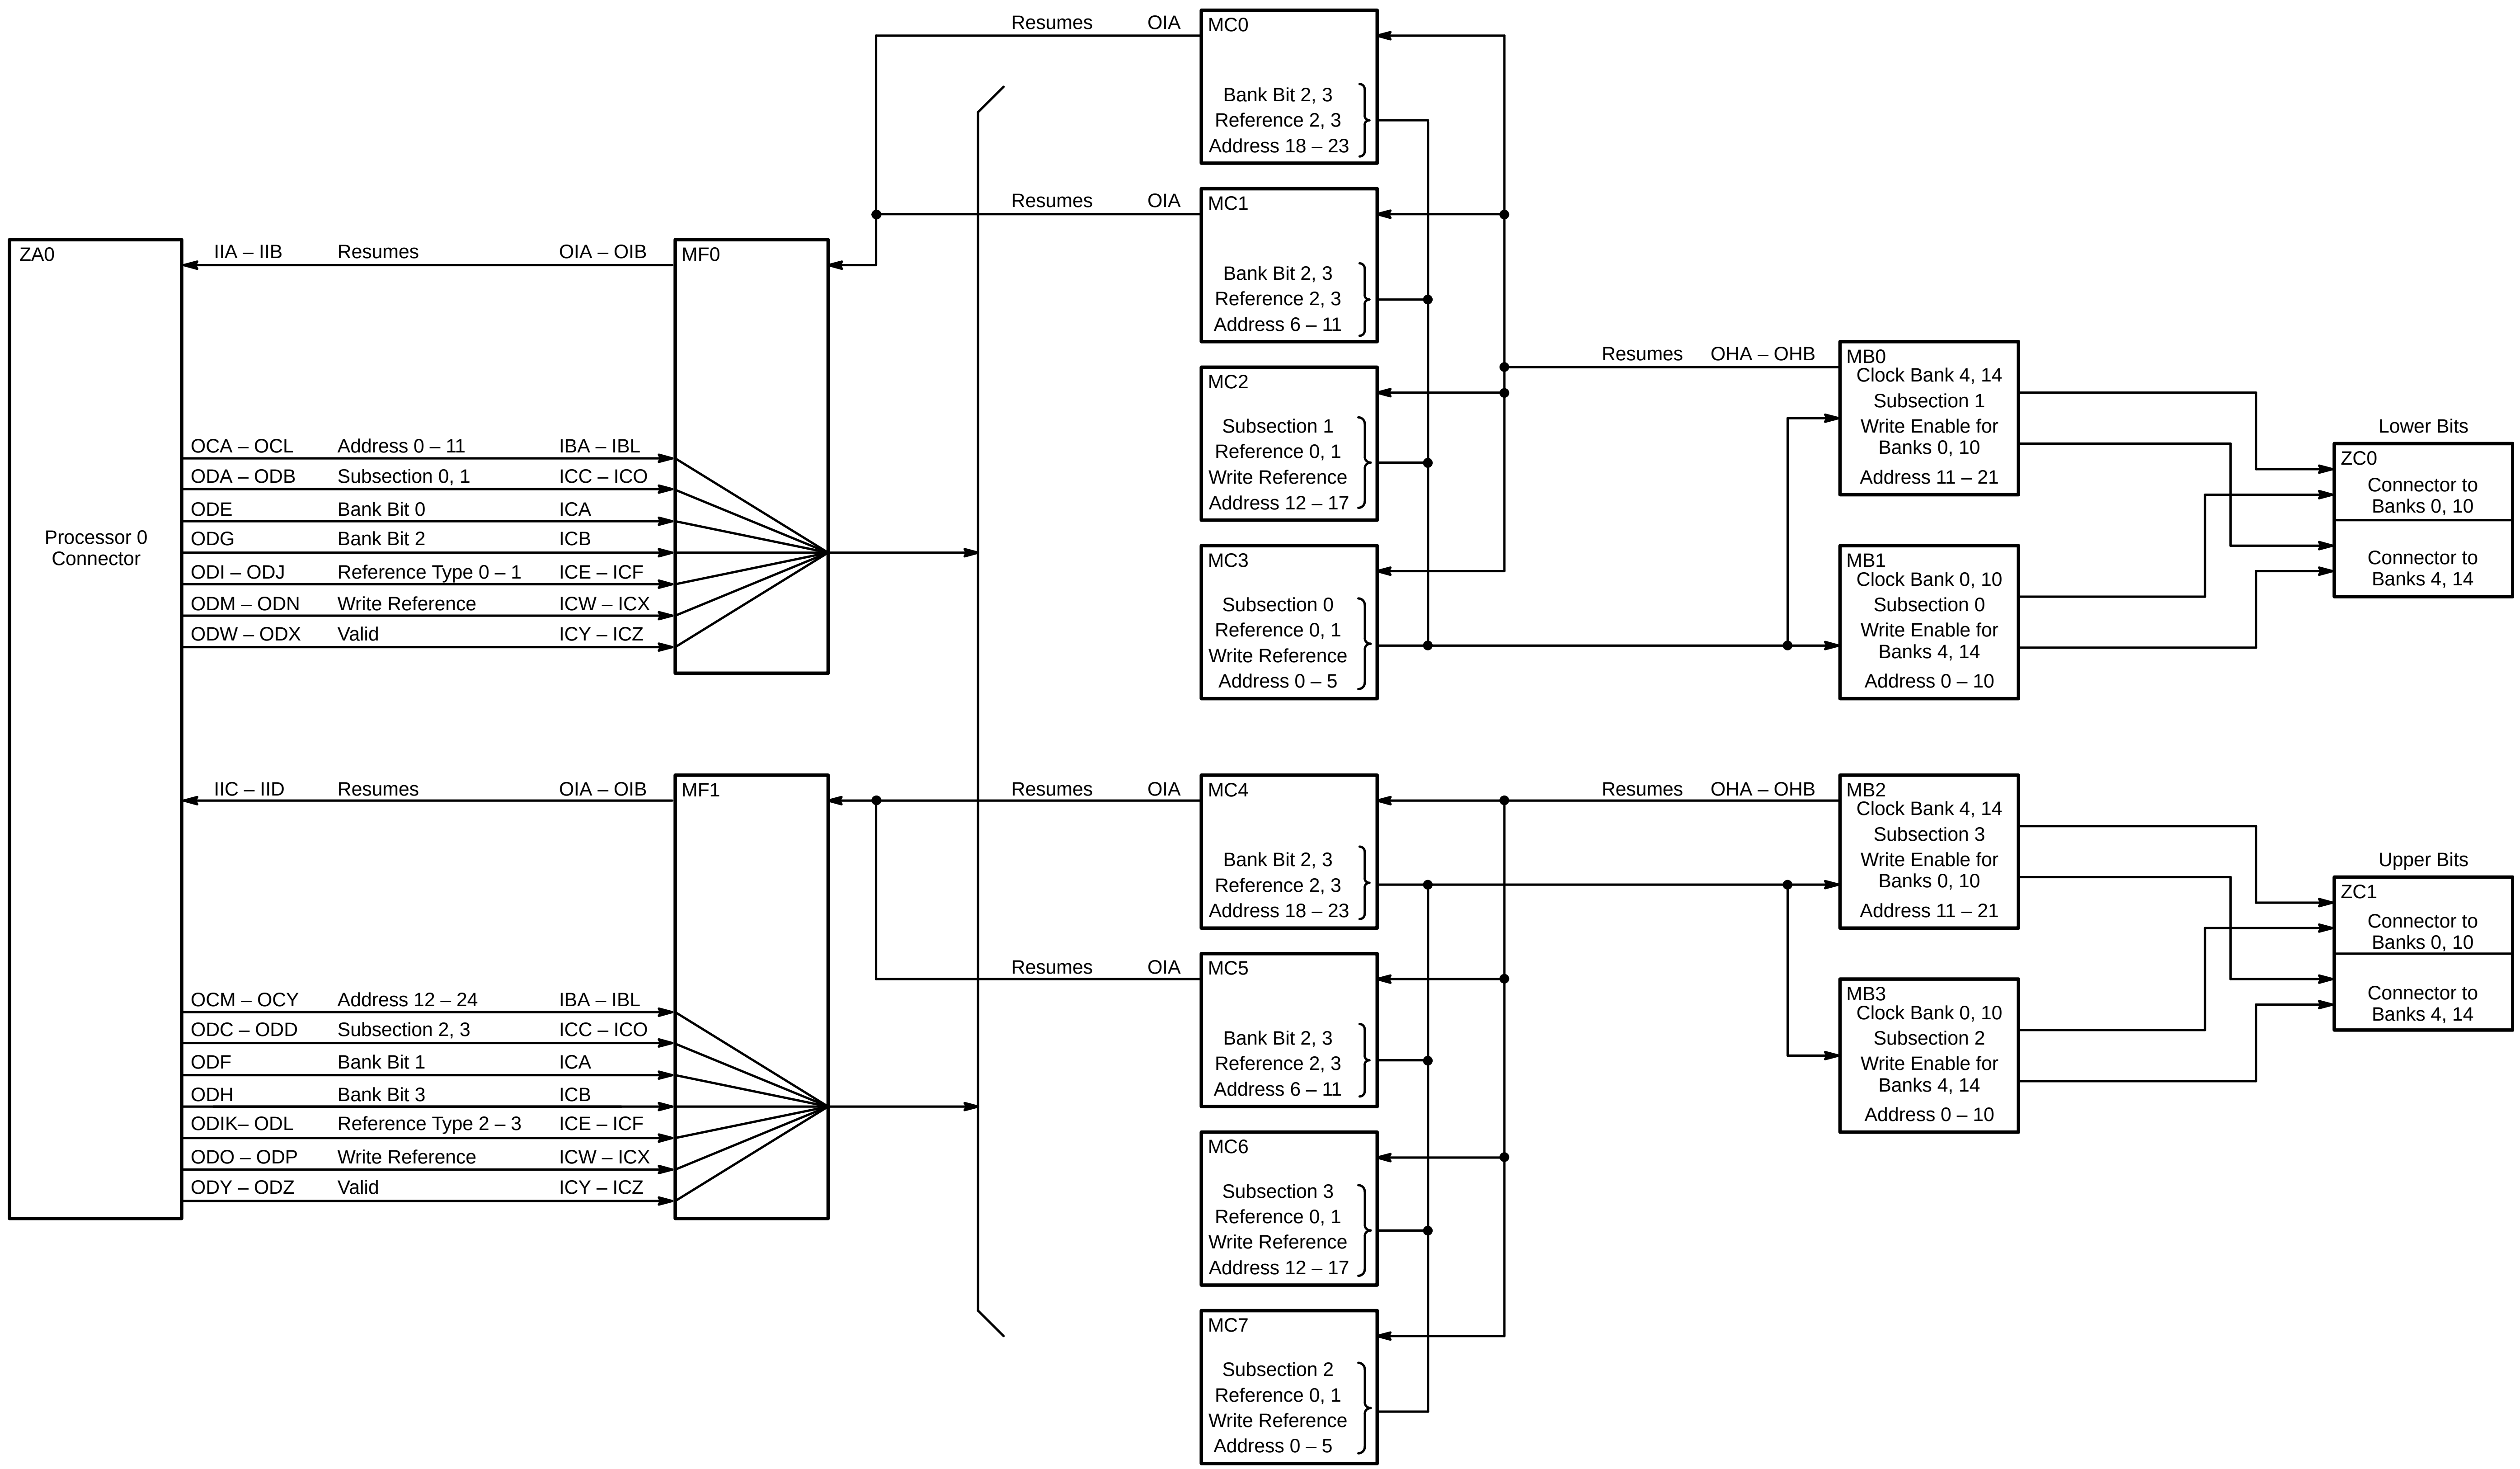

Figure 40. Memory-to-processor Control Signals (Banks 0, 10, 4, and 14 to Processor 0 Only)