Network Module

(CRAY T90™ Series)

HTM-192-B

Cray Research Proprietary

Cray Research, Inc.

Record of Revision

REVISION	DESCRIPTION
	August 1995. Original printing.
A	October 1995. Corrections were made regarding the inverted CP modules.
B	May 1996. Correction was made to Table 1 indicating correct network module position in relation to memory sections.

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

NETWORK MODULE

General Description 3
Physical Description 6
Options 9
LA Option 9
LB Option 9
LC Option 9
LM Option 9
Sanity Code in the Network Module 10
Write Sequence 11
Read Sequence 15

Figures

Figure 1. CRAY T932 Chassis and Module Map 3
Figure 2. CPU and Memory Interface 4
Figure 3. Network Module Board 1 7
Figure 4. Network Module Board 2 8
Figure 5. General Flow through Network Module during a Write Sequence 11
Figure 6. Write Data Flow through Network Module 13
Figure 7. Address Flow through Network Module 14
Figure 8. General Flow through Network Module during a Read Sequence 15
Figure 9. Read Data Flow through Network Module 17

Tables

Table 1. Network Modules in CRAY T916 Systems 5
Table 2. Network Modules in CRAY T932 Systems 5

This page intentionally left blank.

General Description

The network module provides the interface between the CPUs and the memory modules. The CRAY T916 system contains eight network modules; the CRAY T932 system contains 16 network modules. The CRAY T94 systems do not contain network modules. The network modules are arranged in groups of four in chassis locations C, G, K, and O, as shown in the chassis map in Figure 1. The CRAY T916 chassis is similar, except quadrants 2 and 3 are empty.

Figure 1. CRAY T932 Chassis and Module Map

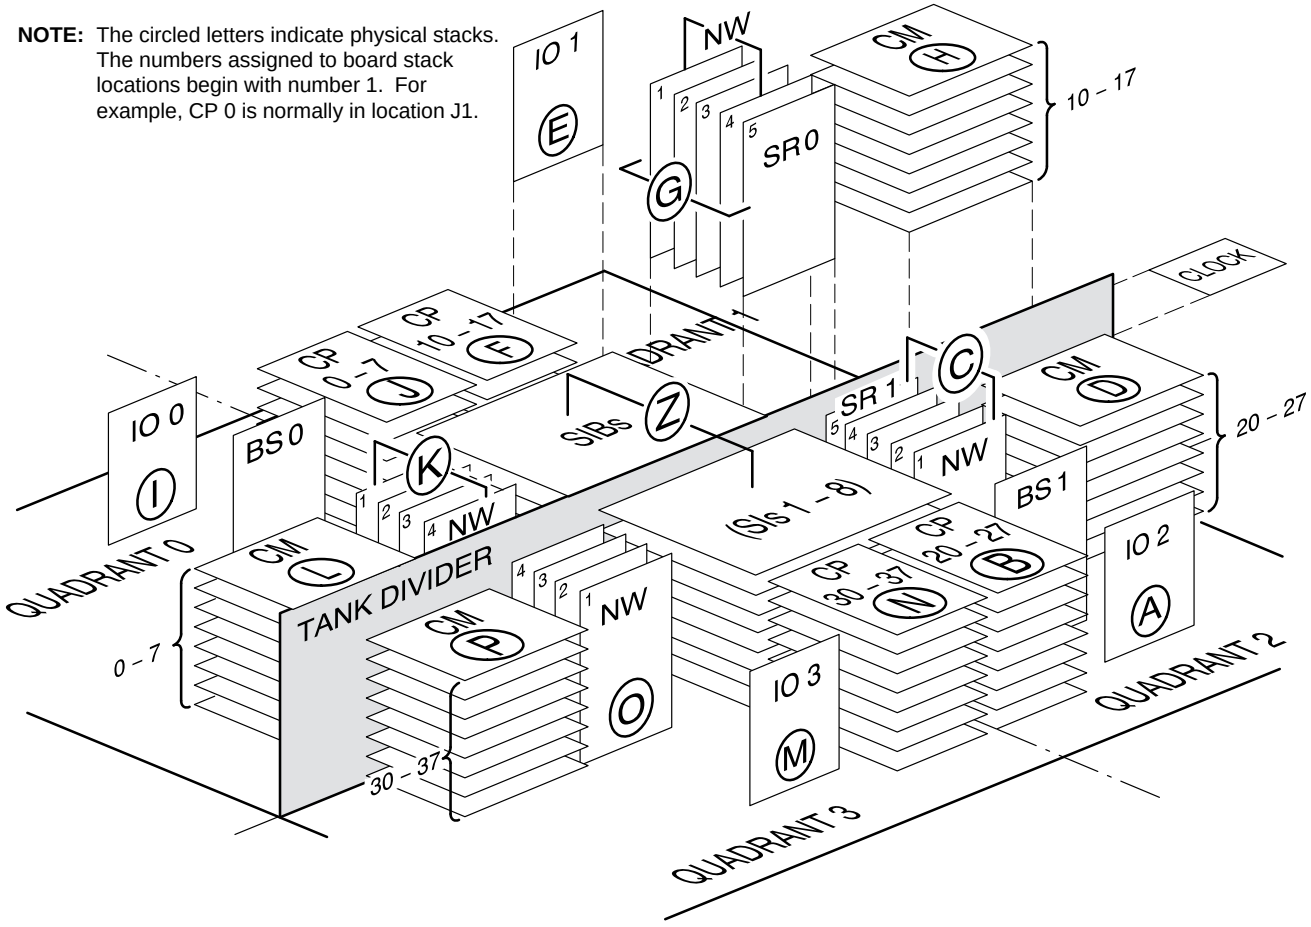

Each network module is responsible for steering data between eight CPUs and the eight subsections within two sections of memory, as shown in Figure 2. One side of the network module handles data for the even section; the other side of the module handles data for the odd section.

Figure 2. CPU and Memory Interface

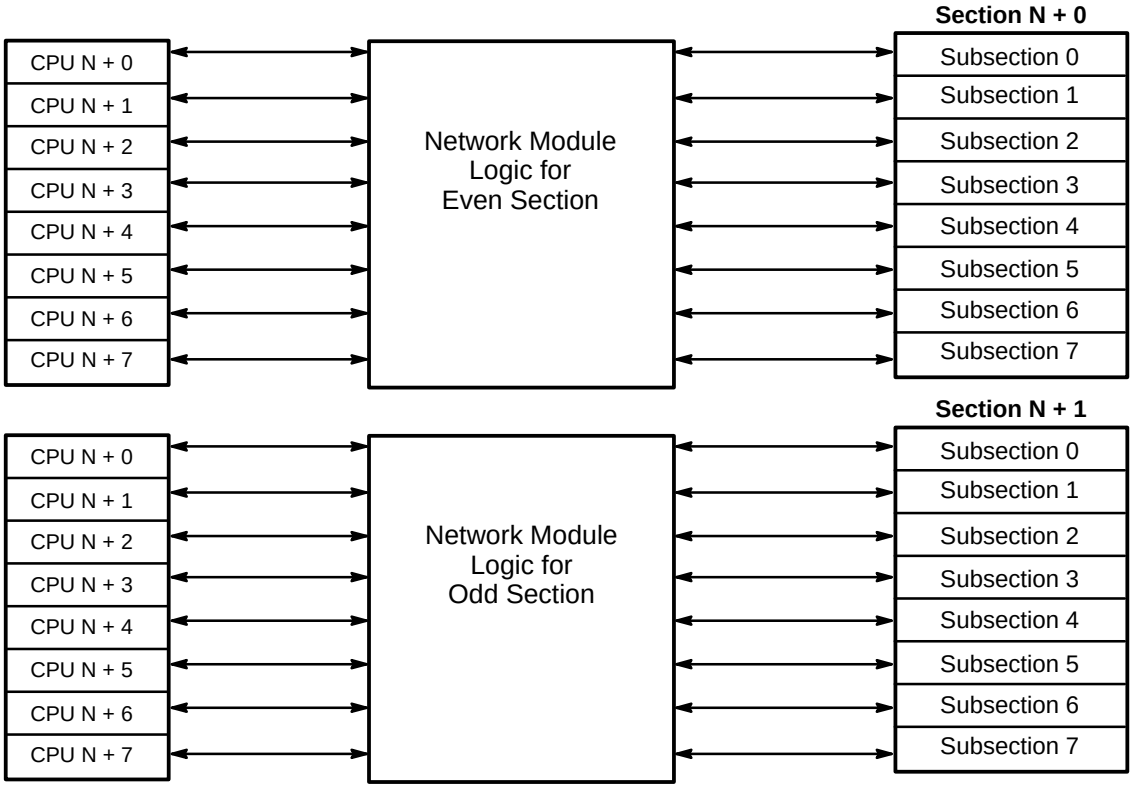

Table 1 and Table 2 define the CPUs and memory sections associated with each network module. The tables also identify the chassis location of each network module, as well as the locations of the CPUs associated with the network modules.

Table 1. Network Modules in CRAY T916 Systems

CPUs	CPU Module Chassis Location	Memory Sections	Network Module	Network Module Chassis Location
0 – 7	J1 – J8	0 & 1	NW03	K4
		2 & 3	NW06	G3
		4 & 5	NW04	G1
		6 & 7	NW01	K2
10 – 17	F1 – F8	0 & 1	NW02	K3
		2 & 3	NW07	G4
		4 & 5	NW05	G2
		6 & 7	NW00	K1

Table 2. Network Modules in CRAY T932 Systems

CPUs	CPU Module Chassis Location	Memory Sections	Network Module	Network Module Chassis Location
0 – 7	J1 – J8	0 & 1	NW16	O3
		2 & 3	NW13	C4
		4 & 5	NW04	G1
		6 & 7	NW01	K2
10 – 17	F1 – F8	6 & 7	NW00	K1
		4 & 5	NW05	G2
		2 & 3	NW12	C3
		0 & 1	NW17	O4
20 – 27	B1 – B8	0 & 1	NW14	O1
		2 & 3	NW11	C2
		4 & 5	NW06	G3
		6 & 7	NW03	K4
30 – 37	N1 – N8	6 & 7	NW02	K3
		4 & 5	NW07	G4
		2 & 3	NW10	C1
		0 & 1	NW15	O2

Physical Description

The network module is a single assembly that consists of one 8-layer printed circuit board (A board) laminated between two 22-layer printed circuit boards (boards 1 and 2). Logic chips are mounted on board 1 and board 2.

The network module also contains 16 connector pads for orthogonal interconnect module (OIM) connectors: 8 on each side of the module. One side of the network module connects to eight system interconnect boards (SIBs); the other side connects to eight memory modules. Figure 3 and Figure 4 illustrate the physical layout of the network module and identify which side of the module connects to the SIB and which side connects to the CM modules.

Figure 3. Network Module Board 1

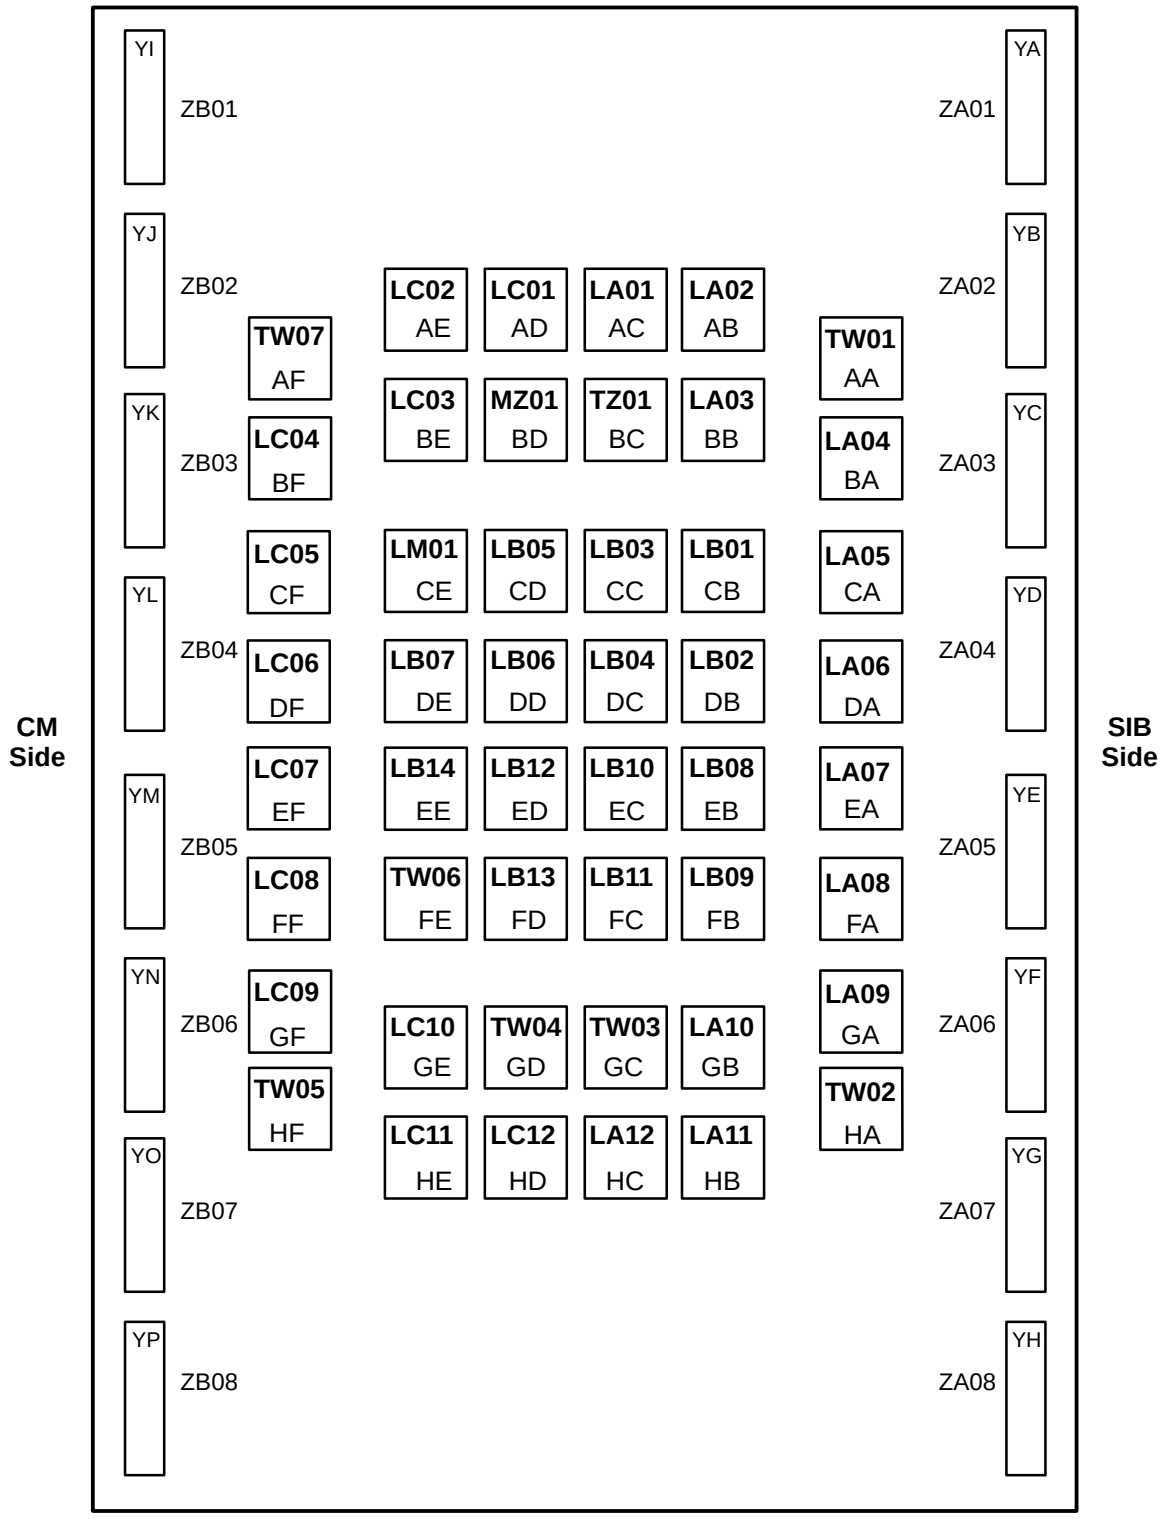

Figure 4. Network Module Board 2

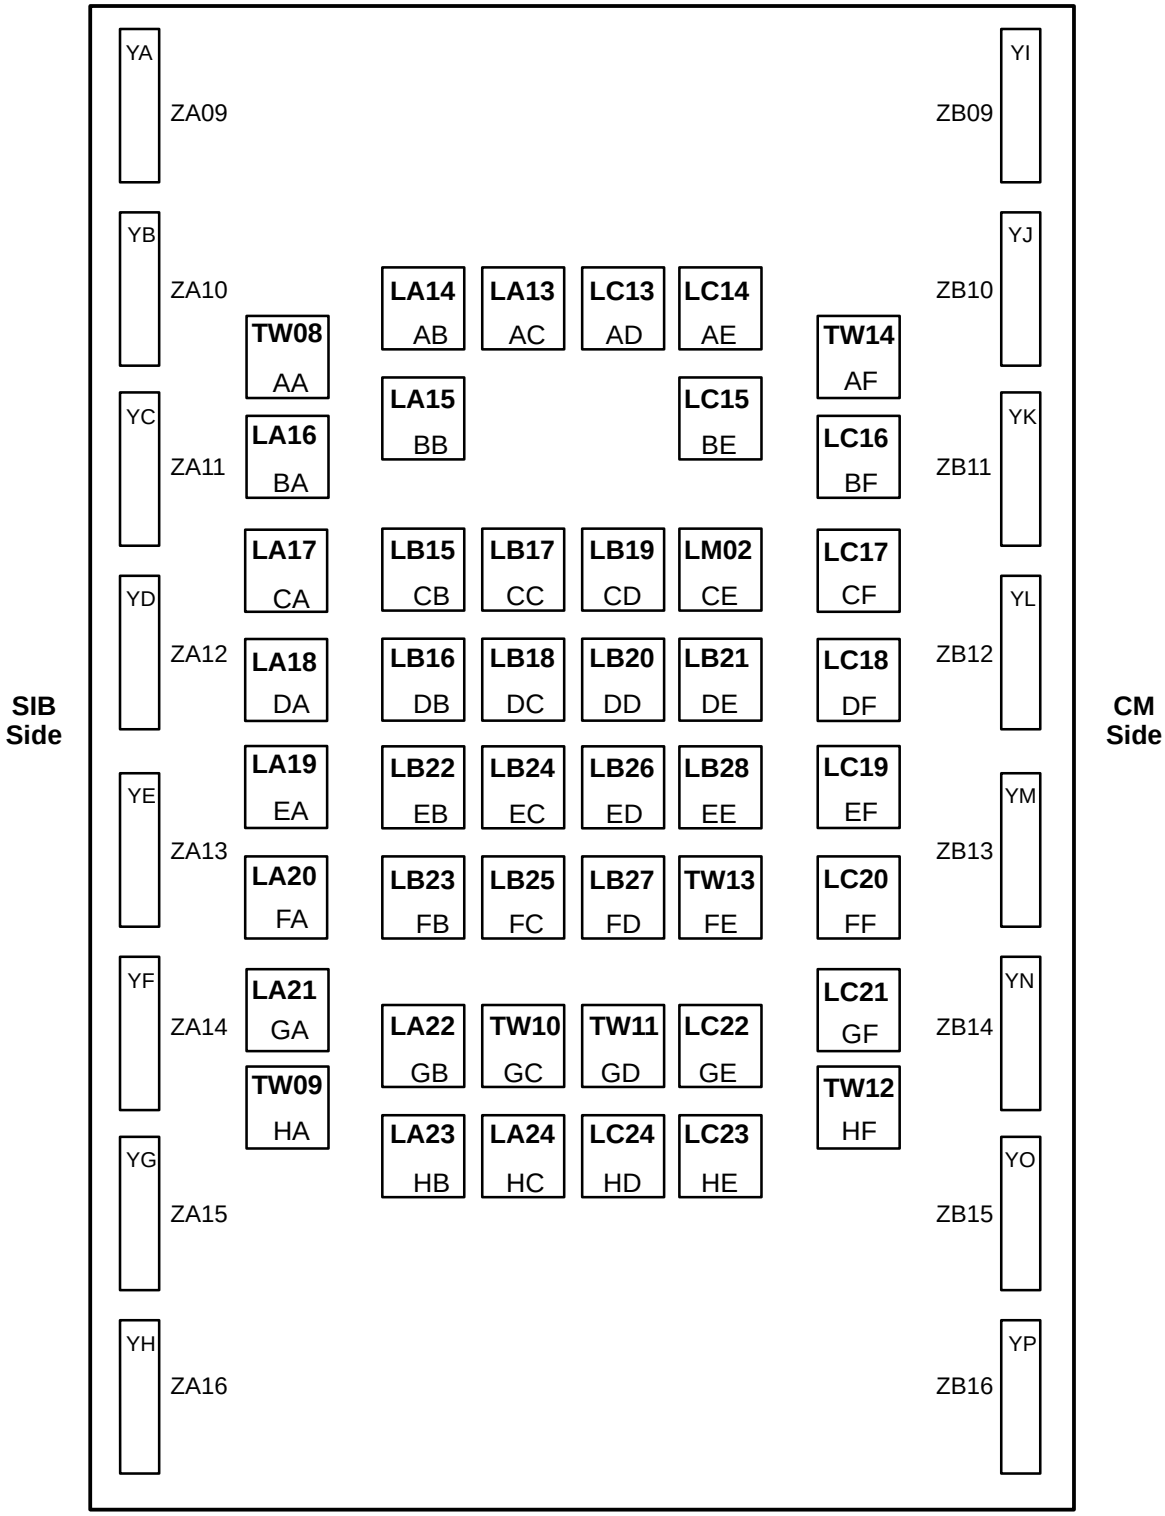

Options

There are three main option types that move the data through the network module: the LA, LB, and LC options.

LA Option

The network module contains 24 LA options. The LA option provides a 6-rank first-in-first-out (FIFO) buffer for write references. A memory write reference requires 1 1/2 LA options. The LA option provides a 2-by-2 multiplexer for read references. Read references require three LA options. There are four input FIFO buffers for each of the two inputs and one buffer for each of the two outputs.

LB Option

The network module contains 28 LB options. The LB option contains two pairs of 4-by-4 multiplexers for data, address, and control. One pair of multiplexers is used for CPU-to-memory references during a memory write operation; the other pair is used for memory-to-CPU references during a read operation. The read multiplexers operates independently of the write multiplexers, with the exception of the Master Clear signal, which clears both pairs of multiplexers. The LB option also handles subsection conflicts.

LC Option

The network module contains 24 LC options. The LC option contains a 4-rank FIFO with a 2-by-2 multiplexer for write operations. During read operations, the LC uses a 6-rank FIFO buffer to buffer the data coming from memory.

LM Option

The network module contains two LM options. The LM options are used for sanity code detection and for the logic monitor functions. If sanity code is not detected, the LM options force a Master Clear signal to the LA, LB, and LC options and the attached memory modules. When sanity code is detected, the Master Clear condition is cleared.

The TW options are placed in unused option locations to enable the master clock signal to be routed to the blank locations.

Sanity Code in the Network Module

The network modules receive sanity code from the CP modules. Configuration codes in the CPU enable sanity code to be sent to the memory modules (via the network modules). The LA options in the network module receive the sanity code from the CI options in the CPU. After the LA option has verified the sanity code, it automatically returns the sanity code to the CPU and forwards it to the LM option.

When the LM option receives the sanity code from an LA option, it removes the Master Clear condition from the network module and from the connected memory modules. It then forwards the sanity code to the LC options. The LC options send the sanity code to the memory modules. When a memory module receives the sanity code, it automatically returns the sanity code to the LC option from which it received the sanity code.

If a CPU drops its sanity connection to the network module, the LM option searches for a source of sanity code from another CPU. This is done without any interruption in the performance of the network module, as long as another source of sanity code is found. If no other sanity code is detected, the LM option sends a Master Clear signal to the options on the network module and to the connected memory modules.

Write Sequence

During a write reference, the network module receives data from one of eight CPUs. The write data comes from the CPU (via the SIB) in two 38-bit packets. The CPU performs the section steering so that the data going to the odd section arrives on one board of the network module, and the data going to the even section arrives on the other board. Figure 5 illustrates the general flow of data and control through the network module.

Figure 5. General Flow through Network Module during a Write Sequence

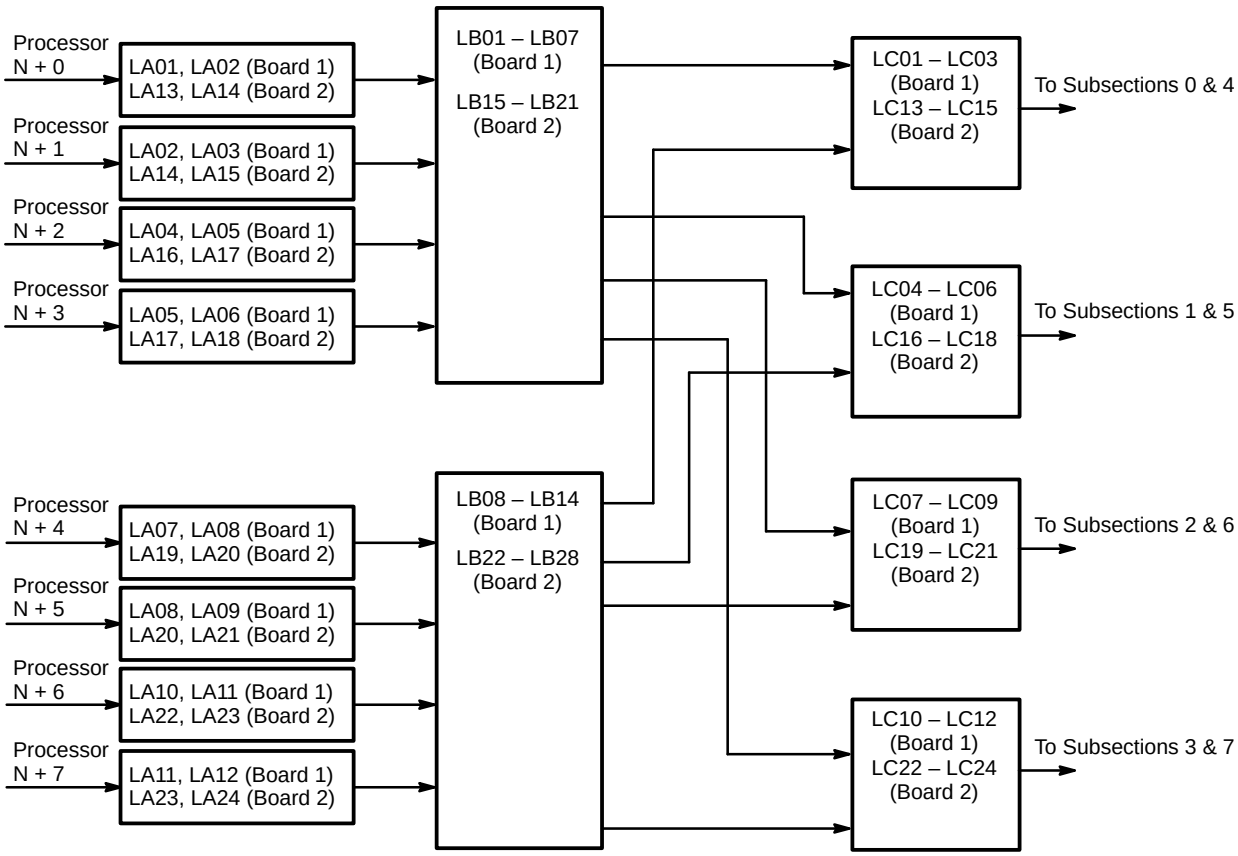

The following paragraphs describe a write sequence from CPU 0 to memory section 0, subsection 4. Figure 6 illustrates the write data flow through the options on the network module; Figure 7 illustrates the address flow through the network module. Refer to these figures when reading the following paragraphs.

As the data enters the network module, it is latched into the first available rank within two LA options, as shown in Figure 6. The LA options also receive address and control information from the CI option in the CPU to steer the data to the correct location in memory.

The data leaves the LA options and enters the LB options. The LB options handle subsection conflict resolution. If more than one CPU in a stack requests the same subsection in the same clock period, the LB options give priority to the lower-numbered CPU's request and delay each subsequent request for 1 clock period. LB01 through LB07, and LB15 through LB21 handle conflicts among processors 0 through 3; options LB08 through LB14, and LB22 through LB28 handle conflicts among processors 4 through 7, as shown in Figure 5.

The LB options use the subsection bits to steer the data to the correct group of LC options. Each group of three LC options is associated with two subsections, as shown in Figure 5. The LC options handle subsection conflicts between upper (4 through 7) and lower (0 through 3) CPUs. The LC options steer data to one of two subsections. Data leaving one board goes to the subsections in the odd memory section; data leaving the other board goes to the subsections in the even memory section.

Figure 6. Write Data Flow through Network Module

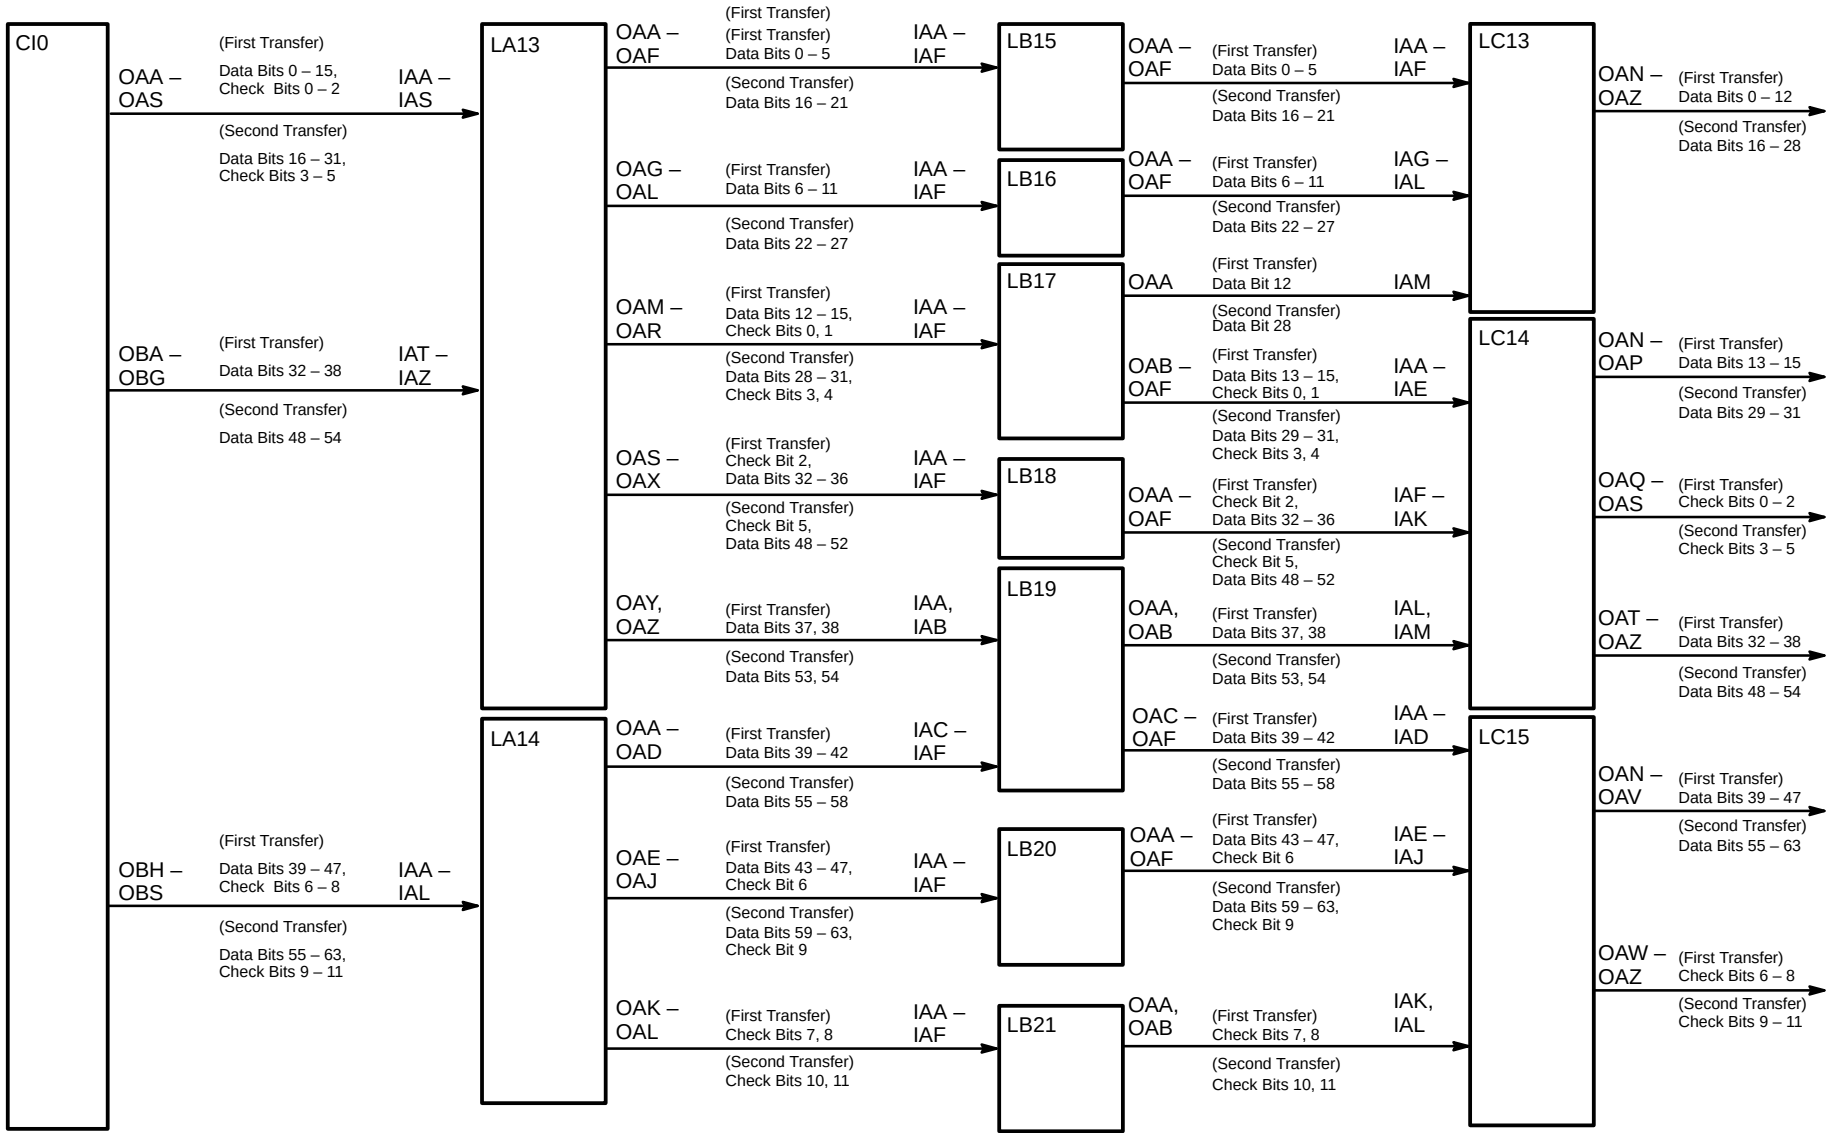

Figure 7. Address Flow through Network Module

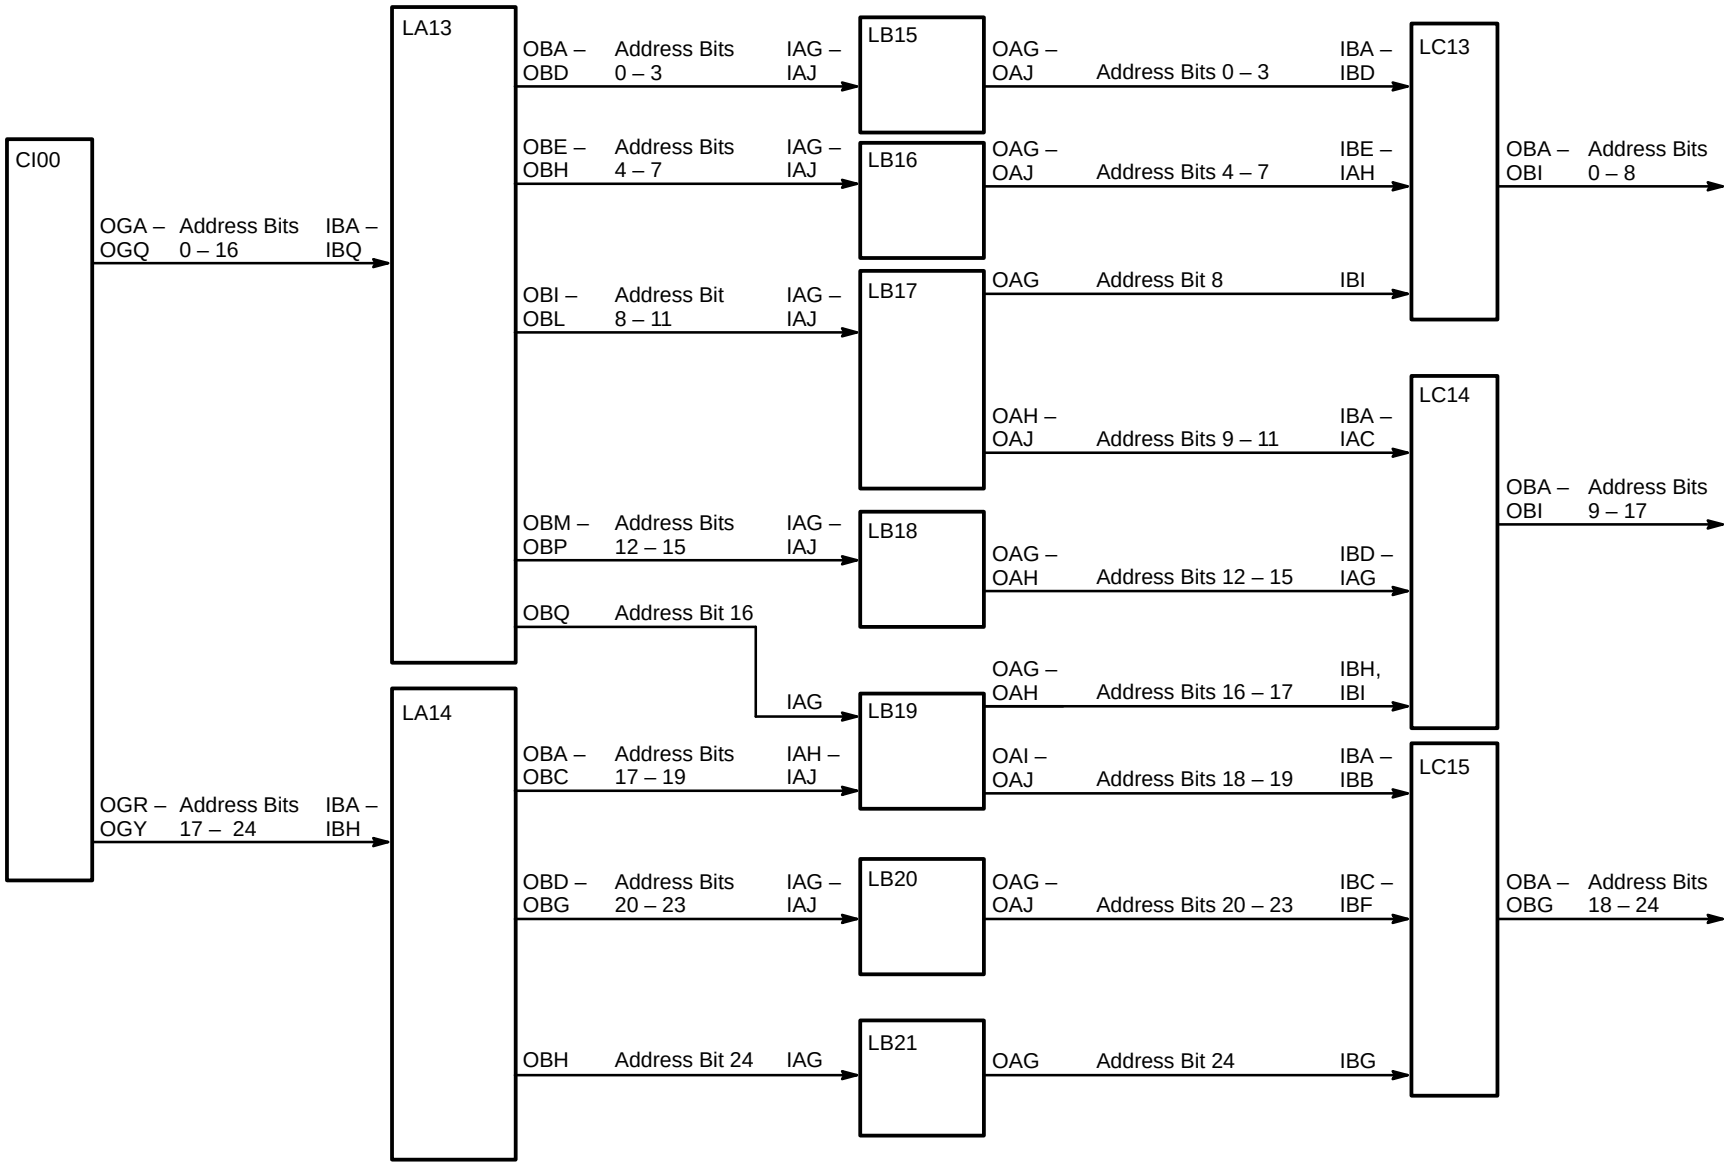

Figure 8. Read Data Flow through Network Module

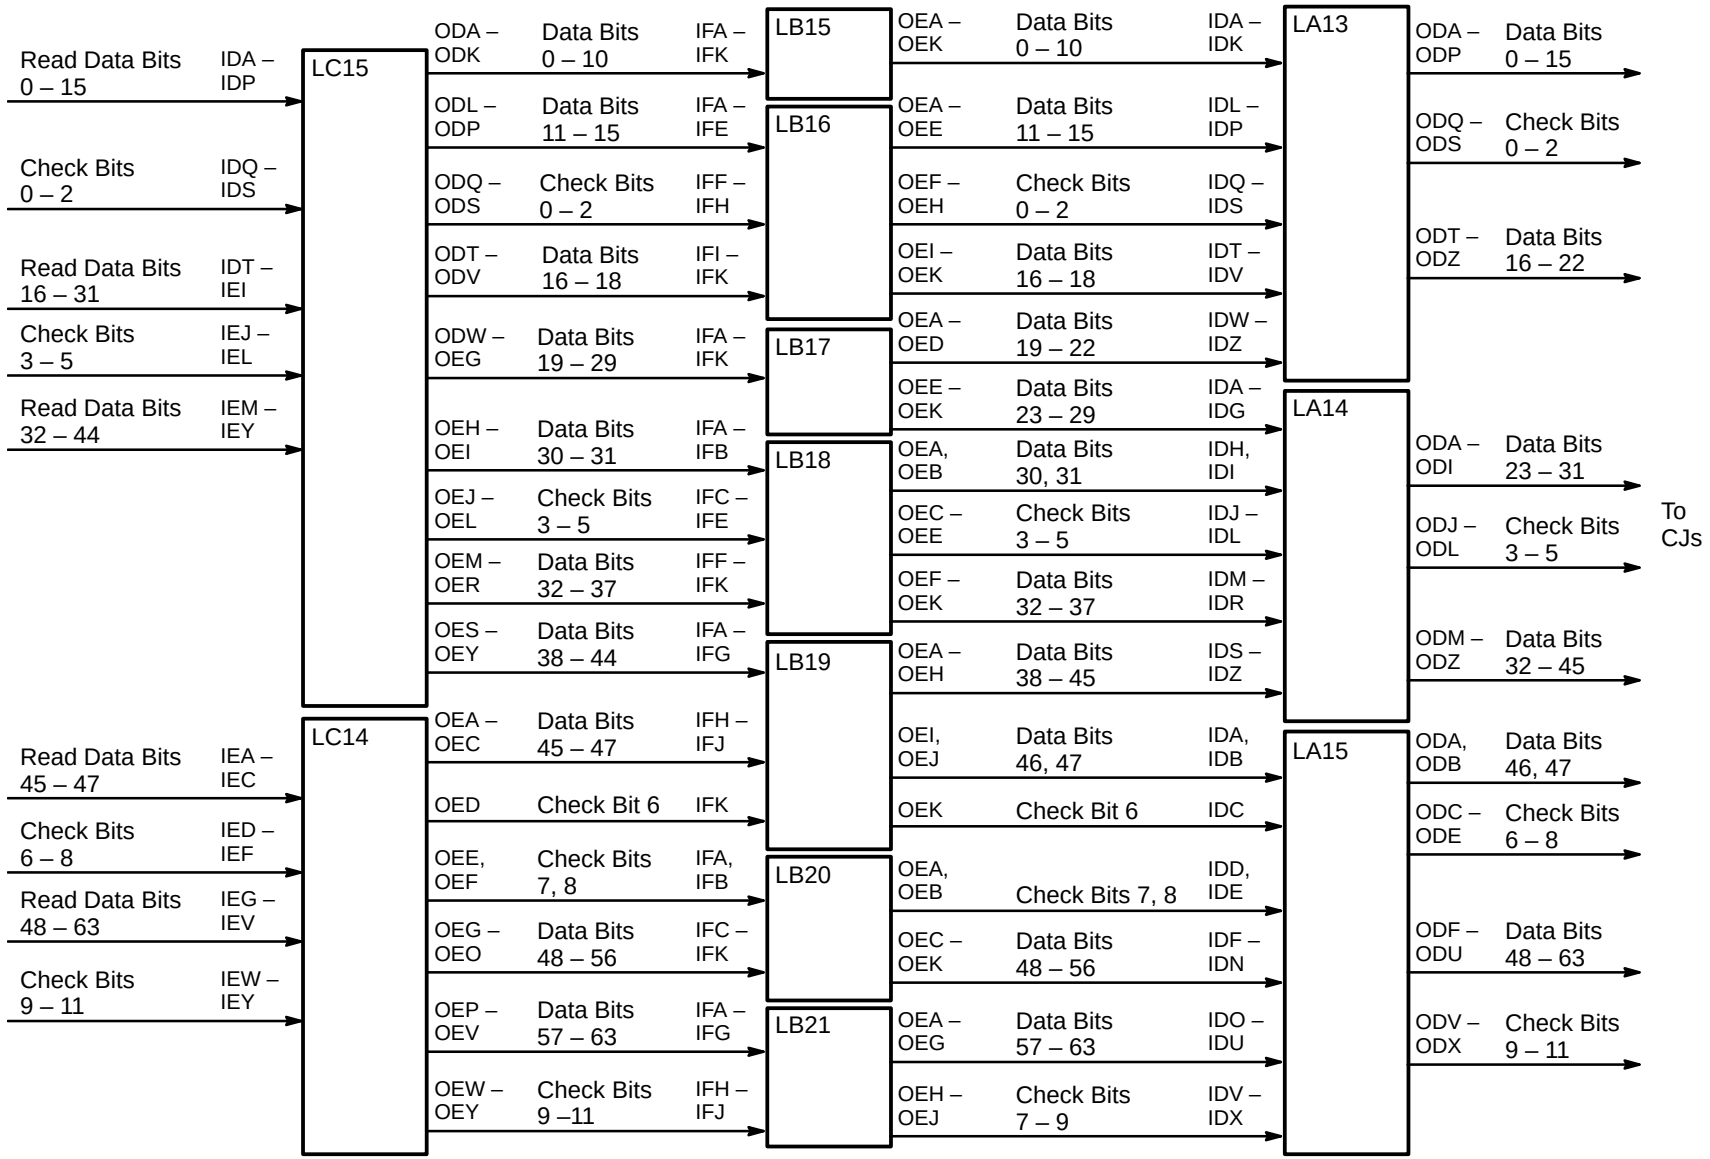

Read Sequence

During a read reference, the network module receives data from one subsection. The data from the odd section of memory enters on one board, and the data from the even section enters on the other board of the network module. Figure 9 illustrates the general flow of data and control through the network module during a read sequence.

Figure 9. General Flow through Network Module during a Read Sequence

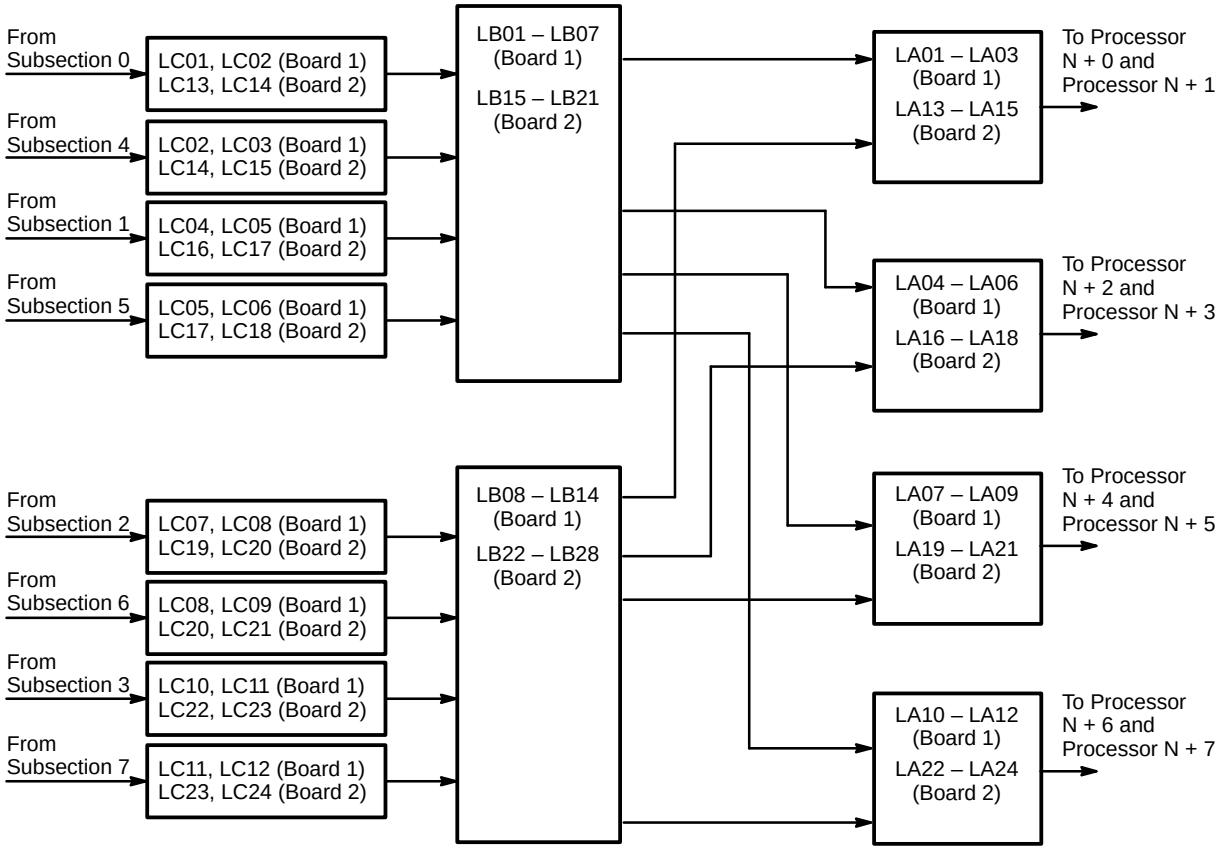

The following paragraphs describe the options involved during a memory read sequence from section 0, subsection 4 to CPU 0. Figure 8 illustrates the read data flow through the options on the network module; Figure 7 illustrates the address flow through the network module. Refer to these figures when reading the following paragraphs.

Read data leaves the memory module in 76-bit single-word transfers and is latched into the first available rank on the LC options. The LC options also receive a 5-bit destination code that steers the data to the requesting CPU. From the LC options, the data and destination code bits enter the LB options. The lower LB options (LB00 through LB07 and LB15

through LB21) handle conflicts among subsections 0, 4, 1, and 5. The upper LB options (LB08 through LB14 and LB22 through LB28) handle conflicts among subsections 2, 6, 3, and 7.

The LB options steer the data and destination code bits to a group of three LA options. The LA options handle conflicts between upper and lower LB options. Each group of three LA options communicates with two CPUs, as shown in Figure 9. The LA options send the data to the proper CPU via the SIB.