

CRAY T3D™ SC Systems PE Configurations (S/N 6105 – 6199)

HMM-092-0

Cray Research Proprietary

PE Configurations (S/N 6105 – 6199)
-092-

Cray Research, Inc.

Record of Revision

REVISION	DESCRIPTION
----------	-------------

	May 1994. Original printing.
--	------------------------------

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

This document is the property of Cray Research, Inc. The use of this document is subject to specific license rights extended by Cray Research, Inc. to the owner or lessee of a Cray Research, Inc. computer system or other licensed party according to the terms and conditions of the license and for no other purpose.

Cray Research, Inc. Unpublished Proprietary Information — All Rights Reserved.


Autotasking, CF77, CRAY, Cray Ada, CRAY Y-MP, CRAY-1, HSX, SSD, UniChem, UNICOS, and X-MPEA are federally registered trademarks and CCI, CFT, CFT2, CFT77, COS, CRAY APP, CRAY C90, CRAY C90D, Cray C++ Compiling System, CRAY EL, Cray NQS, CRAY S-MP, CRAY T3D, CRAY X-MP, CRAY XMS, CRAY-2, Cray/REELlibrarian, CRInform, CRI/TurboKiva, CSIM, CVT, Delivering the power . . ., DGauss, Docview, EMDS, IOS, ND Series Network Disk Array, Network Queuing Environment, Network Queuing Tools, OLNET, RQS, SEGLDR, SMARTE, SUPERCLUSTER, SUPERLINK, Trusted UNICOS, and UNICOS MAX are trademarks of Cray Research, Inc.

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

CRAY RESEARCH, INC.
Logistics
6251 South Prairie View Road
Chippewa Falls, WI 54729

Comments about this publication should be directed to:

CRAY RESEARCH, INC.
Hardware Publications and Training
890 Industrial Blvd.
Chippewa Falls, WI 54729



CRAY T3D SC Systems PE Configurations (S/N 6105 - 6199)

Document Overview	2
Communication Links.....	2
Module Layout.....	5
Barrier Synchronization Circuits	5
CRAY T3D SC256 System	6
CRAY T3D SC256 Communication Links.....	7
CRAY T3D SC256 Module Layout.....	9
CRAY T3D SC256 Barrier Synchronization Circuits	11
CRAY T3D SC224 System	16
CRAY T3D SC224 Communication Links.....	17
CRAY T3D SC224 Module Layout.....	19
CRAY T3D SC224 Barrier Synchronization Circuits	21
CRAY T3D SC192 System	26
CRAY T3D SC192 Communication Links.....	27
CRAY T3D SC192 Module Layout.....	29
CRAY T3D SC192 Barrier Synchronization Circuits	31
CRAY T3D SC160 System	36
CRAY T3D SC160 Communication Links.....	37
CRAY T3D SC160 Module Layout.....	39
CRAY T3D SC160 Barrier Synchronization Circuits	41
CRAY T3D SC128 System	46
CRAY T3D SC128 Communication Links.....	47
CRAY T3D SC128 Module Layout.....	49
CRAY T3D SC128 Barrier Synchronization Circuits	51

1 Document Overview

This document describes the hardware characteristics that affect interconnect network routing, processing element node partitions, and barrier synchronization for each of the CRAY T3D single-cabinet (SC) systems.

This document describes the communication links, module layout, and barrier synchronization circuits for each of the CRAY T3D SC systems.

1.1 Communication Links

The Y-dimension communication links are shown first, followed by figures that show the X- and Z-dimension communication links. Figure 1 shows a sample diagram of the Y-dimension communication links.




Figure 1. Sample Y-dimension Communication Links

In each of the Y-dimension communication link figures, the nodes (represented by the oval shape) contain two numbers. The top number shows the logical node number. The logical node number is equivalent to the number read from a logical PE number (L_WHOAMI) register with bit 2^0 set to 0. For example, logical node 122_{16} contains the logical PEs 122_{16} and 123_{16} .

The bottom number shows the physical node number. The physical node number is equivalent to the number read from a physical PE number (P_WHOAMI) register with bit 2^0 set to 0. For example, physical node 122_{16} contains the physical PEs 122_{16} and 123_{16} .

Figure 2 shows a sample diagram of the X- and Z-dimension communication links. In each of the X- and Z-dimension communication link figures, the same numbering convention is followed as was described for the Y-dimension communication links.

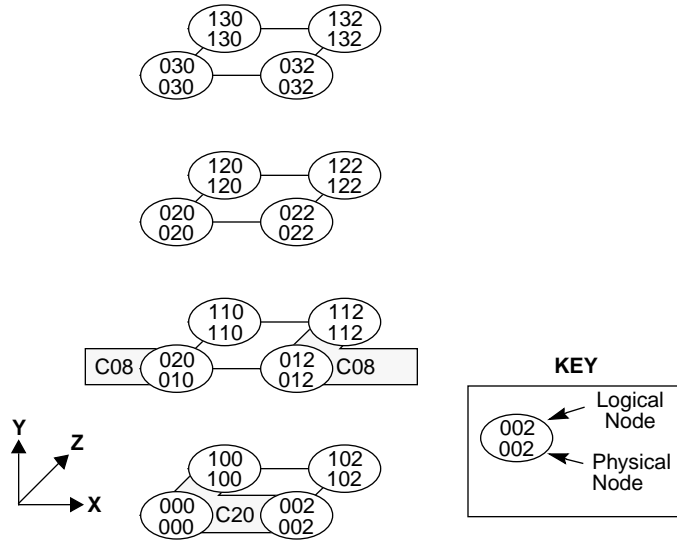


Figure 2. Sample X- and Z-dimension Communication Links

For clarity, the arrow heads on the communication links are removed and the communication links that complete the torus in the X- and Z-dimension are not shown in the figures. Figure 3 shows the torus communication links in the Y = 3 plane of nodes shown in Figure 2.




Figure 3. Sample Torus Communication Links

To keep the X- and Z-dimension communication link figures to a reasonable size, the exact positions of the I/O gateway nodes and spare processing element (PE) nodes are not shown. Instead, when a communication link between two PE nodes also connects to an I/O gateway or spare PE node, the communication link is highlighted and labeled with the physical node number of the I/O gateway input node or the physical node number of the spare PE node.

Figure 4 shows sample I/O gateway communication links. The most significant digit of the physical node number for an I/O gateway node is always C₁₆.




Figure 4. Sample I/O Gateway Communication Links

Figure 5 shows sample spare PE node communication links. The most significant digit of the physical node number for a spare PE node is always 9.

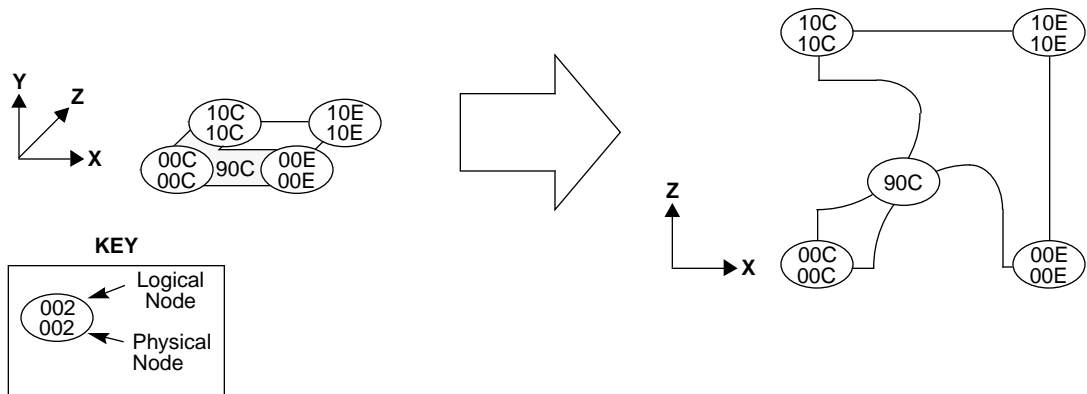


Figure 5. Spare PE Node Communication Links

Figure 6 shows sample communication links for PE nodes that connect to both I/O gateway nodes and a spare PE node.




Figure 6. Sample Spare PE Node and I/O Gateway Communication Links

1.2 Module Layout

The module layout figures show the slot number location of a module and show the location of the physical nodes on the module. The physical node number is equivalent to the number read from a P_WHOAMI register with bit 2^0 set to 0. For example, physical node 61C₁₆ contains the physical PEs 61C₁₆ and 61D₁₆.




Figure 7. Sample Module Layout

1.3 Barrier Synchronization Circuits

The barrier synchronization circuit figures show the bypass point connections for each of the CRAY T3D SC systems.

2 CRAY T3D SC256 System

The CRAY T3D SC256 system contains 256 PEs in 128 processing element nodes and is housed in one cabinet. The following subsections describe the communication links, module layout, and barrier synchronization bypass points for the CRAY T3D SC256 system.

NOTE: Throughout this document, physical PE numbers are represented as three-digit hexadecimal numbers that correspond to the physical PE numbers in the P_WHOAMI registers. For example, a PE may have a physical PE number of $10B_{16}$.

NOTE: Throughout this document, physical node numbers are represented as three-digit hexadecimal numbers. This number is equivalent to the number read from a P_WHOAMI register with bit 2^0 set to 0. For example, physical node $10A_{16}$ contains the physical PEs $10A_{16}$ and $10B_{16}$.

NOTE: Throughout this document, logical node numbers are represented as three-digit hexadecimal numbers. This number is equivalent to the number read from a L_WHOAMI register with bit 2^0 set to 0. For example, logical node $10A_{16}$ contains the logical PEs $10A_{16}$ and $10B_{16}$.

2.1 CRAY T3D SC256 Communication Links

Figure 8 shows the physical communication links between nodes in the Y dimension.

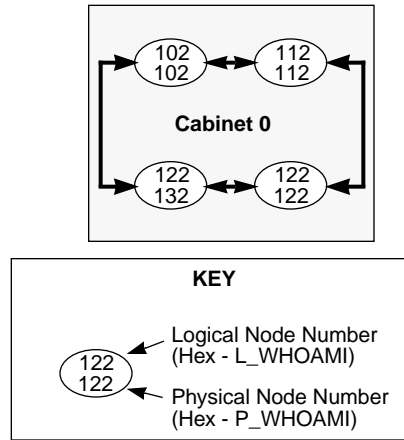


Figure 8. CRAY T3D SC256 Y-dimension Communication Links

Figure 9 shows the physical communication links between spare nodes in the Y dimension.




Figure 9. CRAY T3D SC256 Spare Node Y-dimension Communication Links

Figure 10 shows the physical communication links between the nodes in the X and Z dimensions. For clarity, the communication links that complete the torus in the X and Z dimensions are not shown.




Figure 10. CRAY T3D SC256 X- and Z-dimension Communication Links

2.2 CRAY T3D SC256 Module Layout

Figure 11 shows the module layout and physical node locations in the CRAY T3D SC256 system cabinet. In each figure, the physical node number is represented as a three-digit hexadecimal number that is equivalent to the number read from the P_WHOAMI register with bit 2^0 set to 0. For example, physical node 002_{16} contains the physical PEs 002_{16} and 003_{16} .




Figure 11. CRAY T3D SC256 Module Layout

2.3 CRAY T3D SC256 Barrier Synchronization Circuits

Figure 12 through Figure 15 show the barrier synchronization bypass points for all four physical barrier synchronization circuits in the CRAY T3D SC256 system. In these figures, the printed circuit board (PCB) bypass points and all other bypass points are shown.

The PCB bypass points are controlled by bit 2^{11} or bit 2^{12} of the network mode register (NODE_CSR). Every processing element node PCB contains four PCB bypass points (one for each physical barrier synchronization circuit).

Each PCB bypass point is given a name that corresponds to the physical node that controls the bypass point. For example, the PCB bypass point 402 in barrier synchronization circuit 0 is located on the PCB that contains physical PEs 402_{16} , 403_{16} , 412_{16} , and 413_{16} . This bypass is controlled by setting bit 2^{11} of the NODE_CSR in physical node 402 (which contains physical PEs 402_{16} and 403_{16}) to the appropriate value.

When bit 2^{11} or bit 2^{12} of the NODE_CSR is set to 0, the output of the AND gate in a PCB bypass point is redirected to the fanout block in the bypass point. When bit 2^{11} or bit 2^{12} of the NODE_CSR is set to 1, the output of the AND gate in a PCB bypass point is directed to a second-level bypass point.

The other bypass points shown in the following figures are controlled by bit 2^{13} of the NODE_CSR. For example, bypass point 0A in barrier synchronization circuit 0 of cabinet 0 is controlled by setting bit 2^{13} of the NODE_CSR in physical node 000_{16} (which contains physical PEs 000_{16} and 001_{16}) to the appropriate value.

When bit 2^{13} of the NODE_CSR is set to 0, the output of the AND gate in the bypass point is redirected to the fanout block in the bypass point. When bit 2^{13} of the NODE_CSR is set to 1, the output of the AND gate in the bypass point is directed to the next-level bypass point.

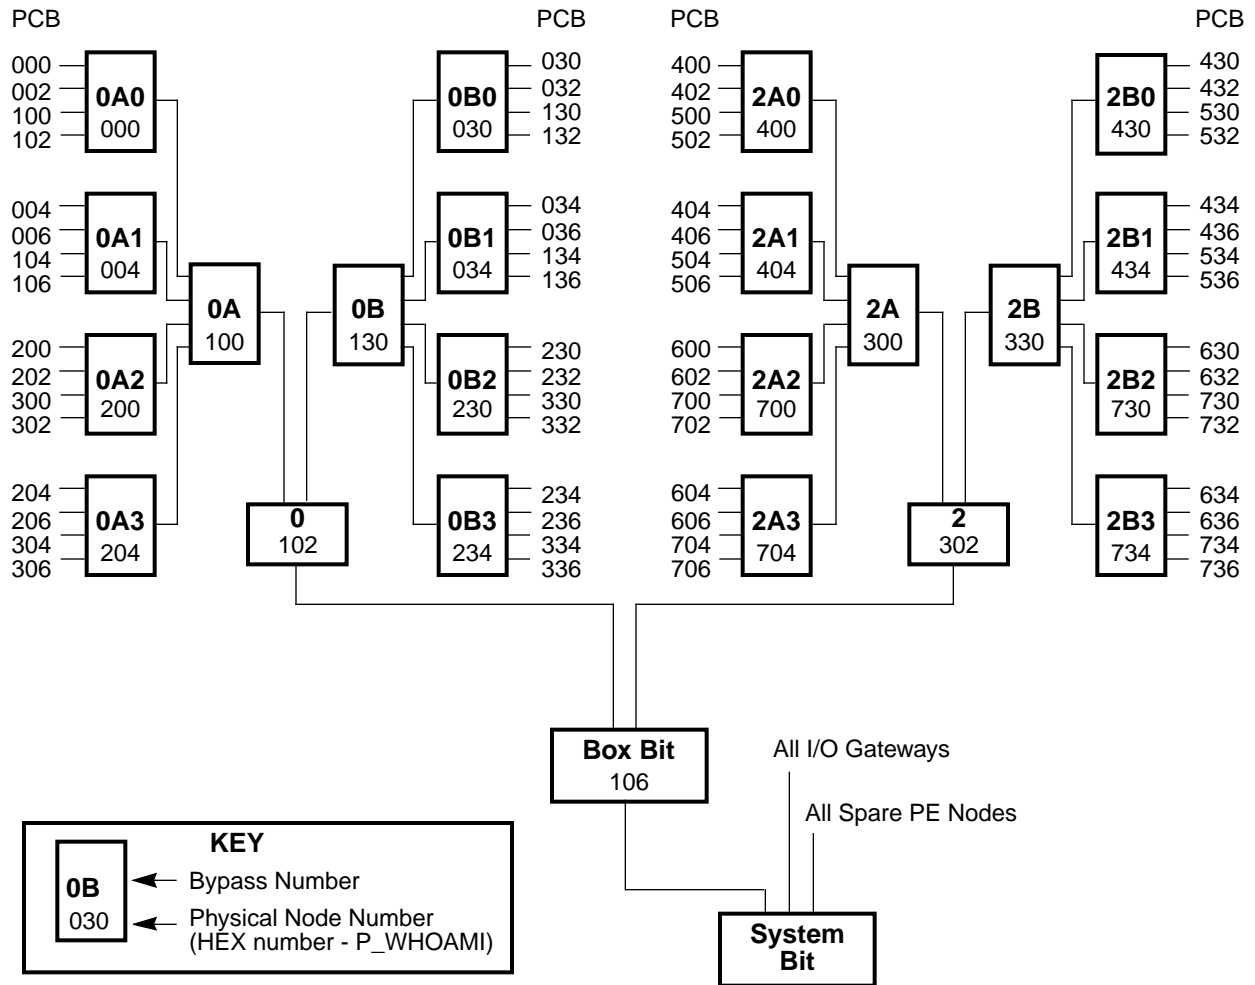
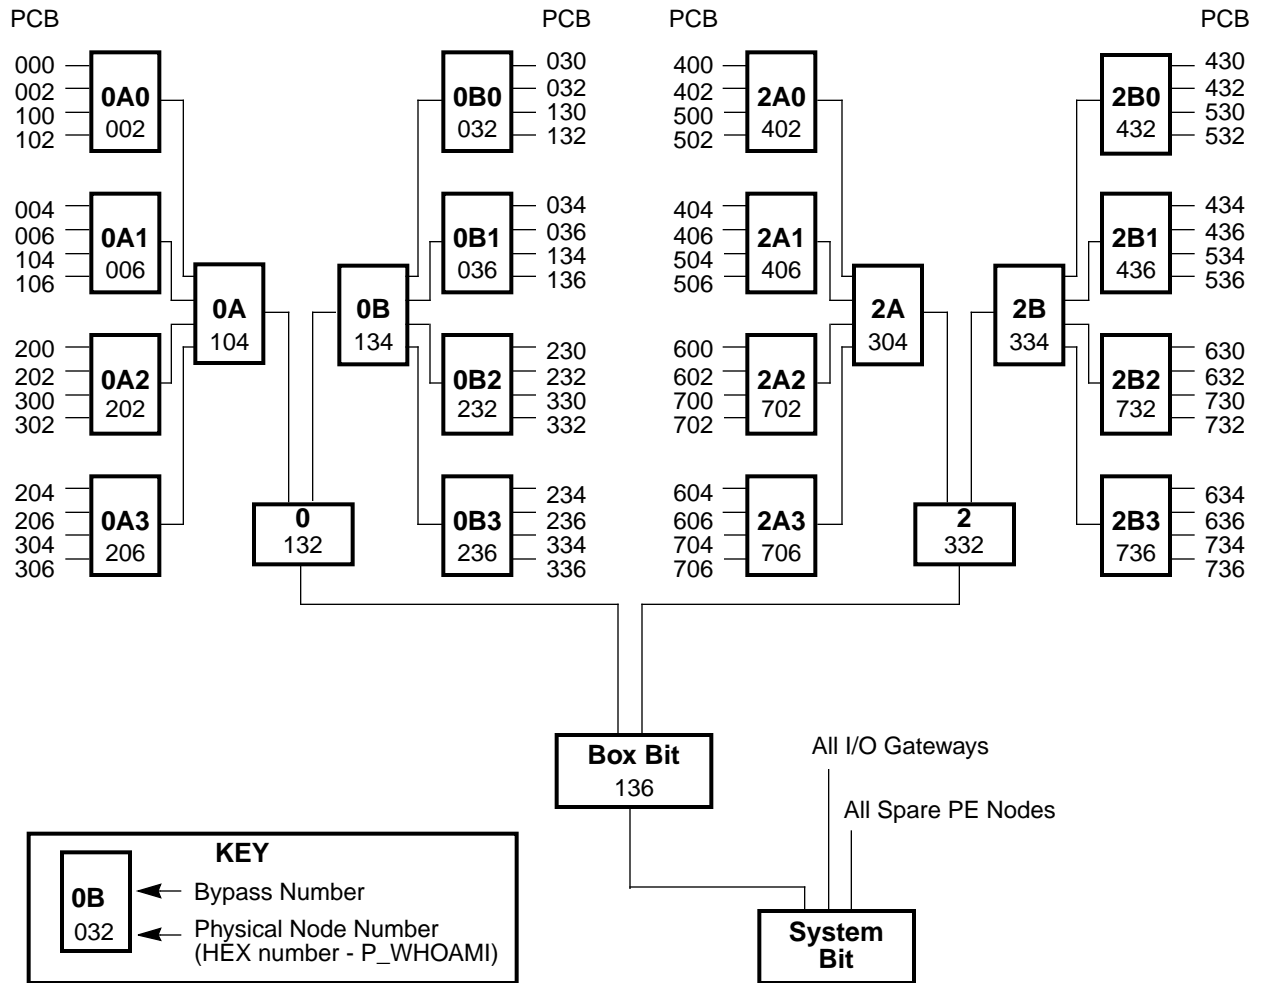


Figure 12. Barrier Synchronization Circuit 0 in CRAY T3D SC256 System



NOTE: The PCB bypass points shown in this figure are controlled by bit 2^{12} of the NODE_CSR. The other bypass points shown in this figure are controlled by bit 2^{13} of the NODE_CSR.

Figure 13. Barrier Synchronization Circuit 1 in CRAY T3D SC256 System

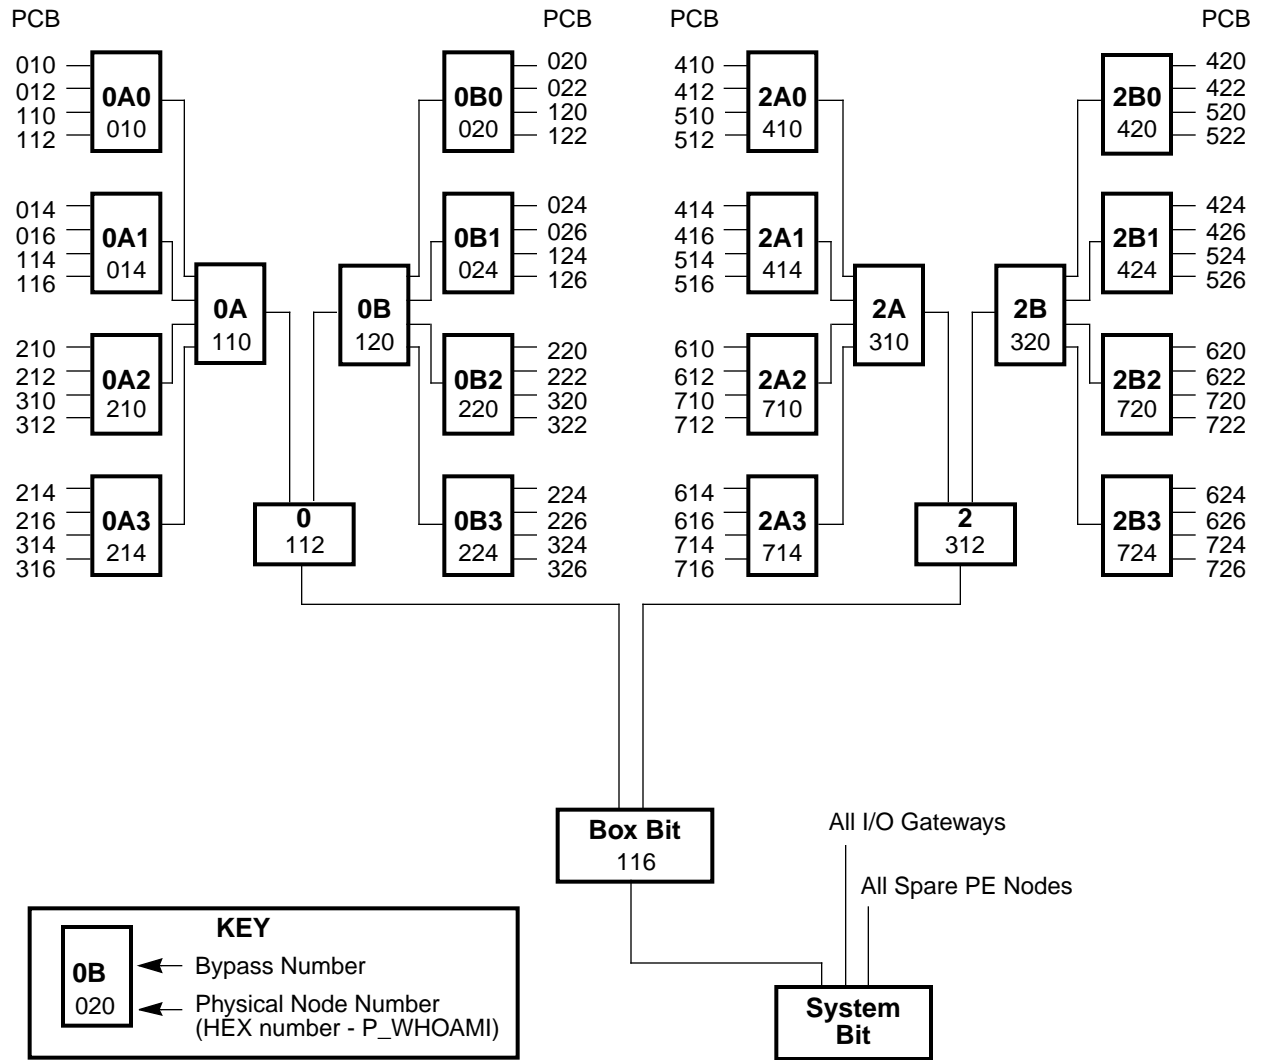


Figure 14. Barrier Synchronization Circuit 2 in CRAY T3D SC256 System

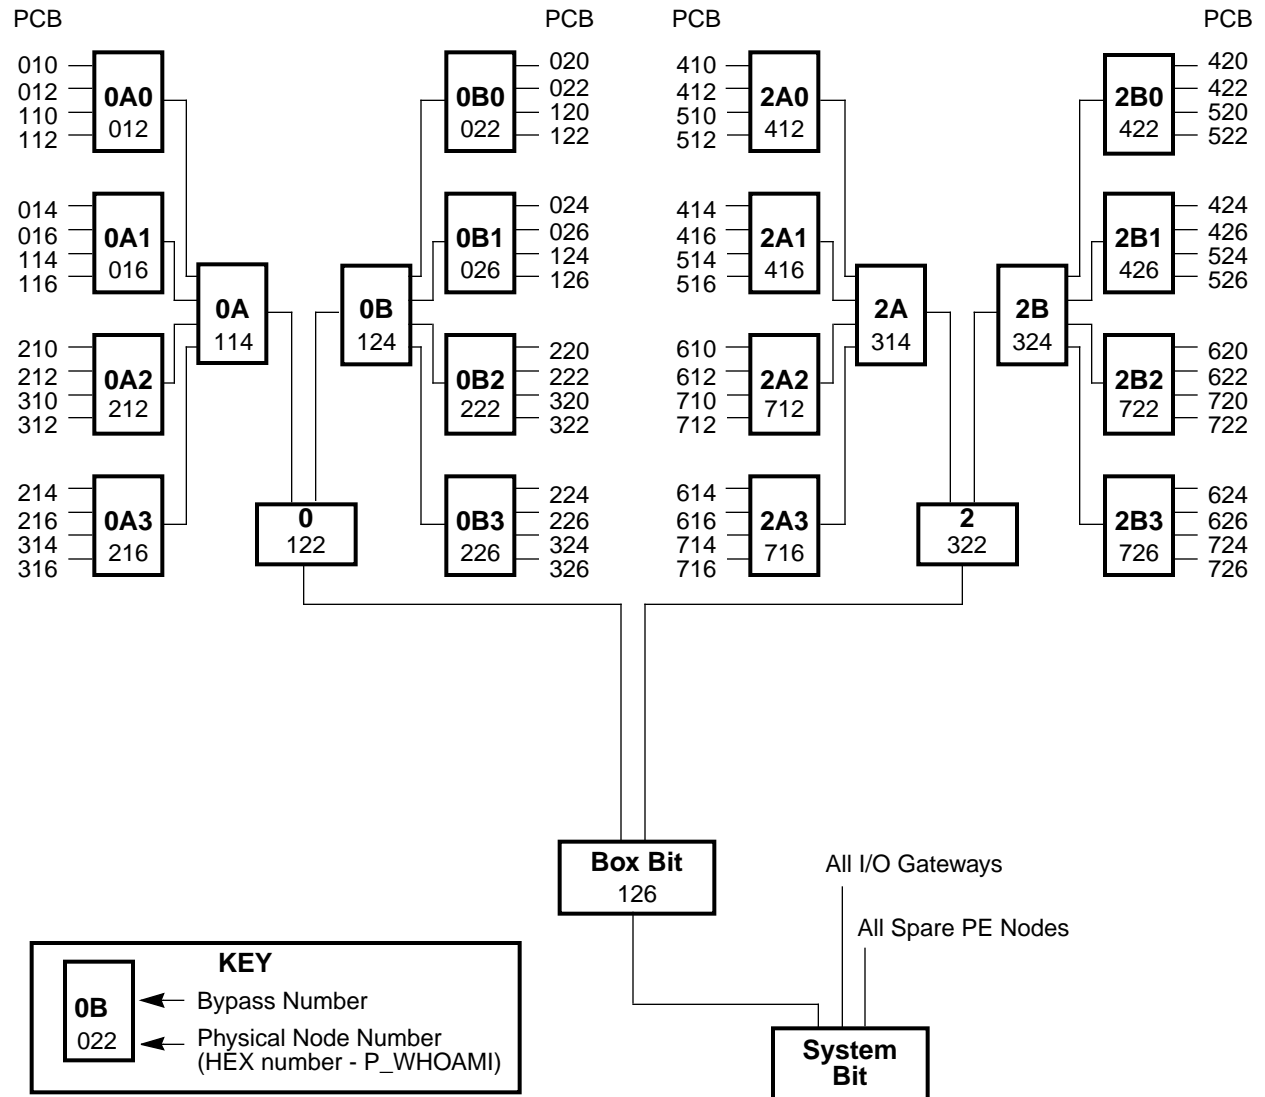


Figure 15. Barrier Synchronization Circuit 3 in CRAY T3D SC256 System

3 CRAY T3D SC224 System

The CRAY T3D SC224 system contains 224 PEs in 112 processing element nodes and is housed in one cabinet. The following subsections describe the communication links, module layout, and barrier synchronization bypass points for the CRAY T3D SC224 system.

NOTE: Throughout this document, physical PE numbers are represented as three-digit hexadecimal numbers that correspond to the physical PE numbers in the P_WHOAMI registers. For example, a PE may have a physical PE number of $10B_{16}$.

NOTE: Throughout this document, physical node numbers are represented as three-digit hexadecimal numbers. This number is equivalent to the number read from a P_WHOAMI register with bit 2^0 set to 0. For example, physical node $10A_{16}$ contains the physical PEs $10A_{16}$ and $10B_{16}$.

NOTE: Throughout this document, logical node numbers are represented as three-digit hexadecimal numbers. This number is equivalent to the number read from a L_WHOAMI register with bit 2^0 set to 0. For example, logical node $10A_{16}$ contains the logical PEs $10A_{16}$ and $10B_{16}$.

3.1 CRAY T3D SC224 Communication Links

Figure 16 shows the physical communication links between nodes in the Y dimension.

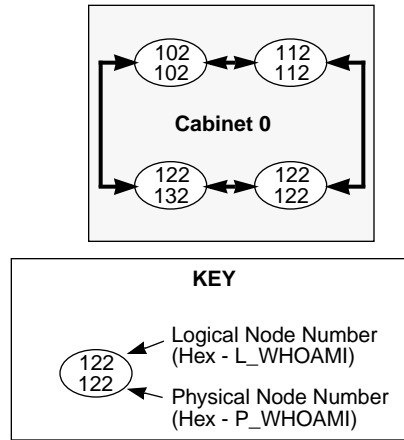


Figure 16. CRAY T3D SC224 Y-dimension Communication Links

Figure 17 shows the physical communication links between spare nodes in the Y dimension.

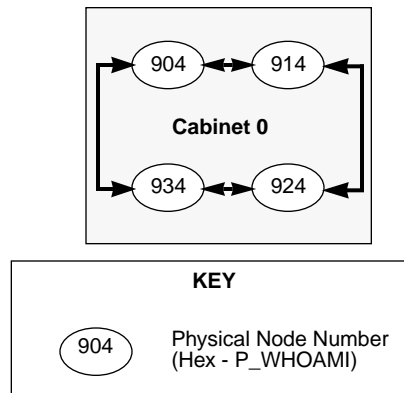


Figure 17. CRAY T3D SC224 Spare Node Y-dimension Communication Links

Figure 18 shows the physical communication links between the nodes in the X and Z dimensions. For clarity, the communication links that complete the torus in the X and Z dimensions are not shown.




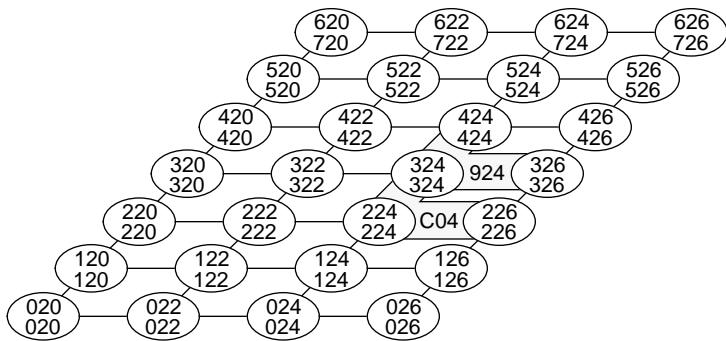
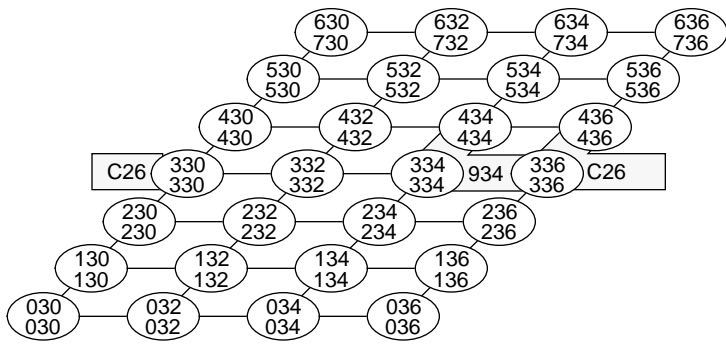



Figure 18. CRAY T3D SC224 X- and Z-dimension Communication Links

3.2 CRAY T3D SC224 Module Layout

Figure 19 shows the module layout and physical node locations in the CRAY T3D SC224 system cabinet. In each figure, the physical node number is represented as a three-digit hexadecimal number that is equivalent to the number read from the P_WHOAMI register with bit 2^0 set to 0. For example, physical node 002_{16} contains the physical PEs 002_{16} and 003_{16} .



KEY
 230 This is the HEX node number. (P_WHOAMI)

Figure 19. CRAY T3D SC224 Module Layout

3.3 CRAY T3D SC224 Barrier Synchronization Circuits

Figure 20 through Figure 23 show the barrier synchronization bypass points for all four physical barrier synchronization circuits in the CRAY T3D SC224 system. In these figures, the printed circuit board (PCB) bypass points and all other bypass points are shown.

The PCB bypass points are controlled by bit 2^{11} or bit 2^{12} of the network mode register (NODE_CSR). Every processing element node PCB contains four PCB bypass points (one for each physical barrier synchronization circuit).

Each PCB bypass point is given a name that corresponds to the physical node that controls the bypass point. For example, the PCB bypass point 402 in barrier synchronization circuit 0 is located on the PCB that contains physical PEs 402_{16} , 403_{16} , 412_{16} , and 413_{16} . This bypass is controlled by setting bit 2^{11} of the NODE_CSR in physical node 402 (which contains physical PEs 402_{16} and 403_{16}) to the appropriate value.

When bit 2^{11} or bit 2^{12} of the NODE_CSR is set to 0, the output of the AND gate in a PCB bypass point is redirected to the fanout block in the bypass point. When bit 2^{11} or bit 2^{12} of the NODE_CSR is set to 1, the output of the AND gate in a PCB bypass point is directed to a second-level bypass point.

The other bypass points shown in the following figures are controlled by bit 2^{13} of the NODE_CSR. For example, bypass point 0A in barrier synchronization circuit 0 of cabinet 0 is controlled by setting bit 2^{13} of the NODE_CSR in physical node 000_{16} (which contains physical PEs 000_{16} and 001_{16}) to the appropriate value.

When bit 2^{13} of the NODE_CSR is set to 0, the output of the AND gate in the bypass point is redirected to the fanout block in the bypass point. When bit 2^{13} of the NODE_CSR is set to 1, the output of the AND gate in the bypass point is directed to the next-level bypass point.

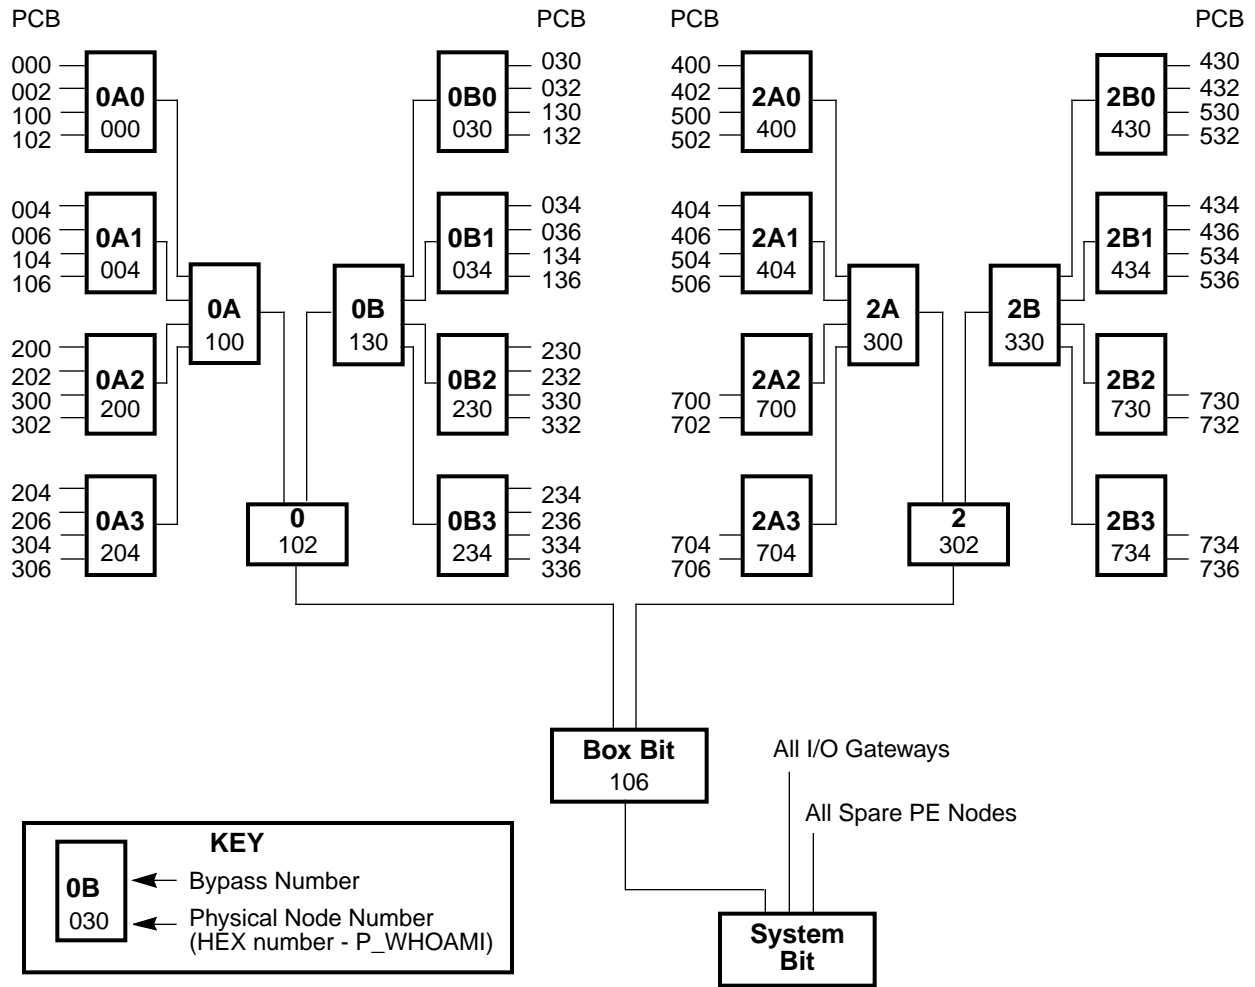
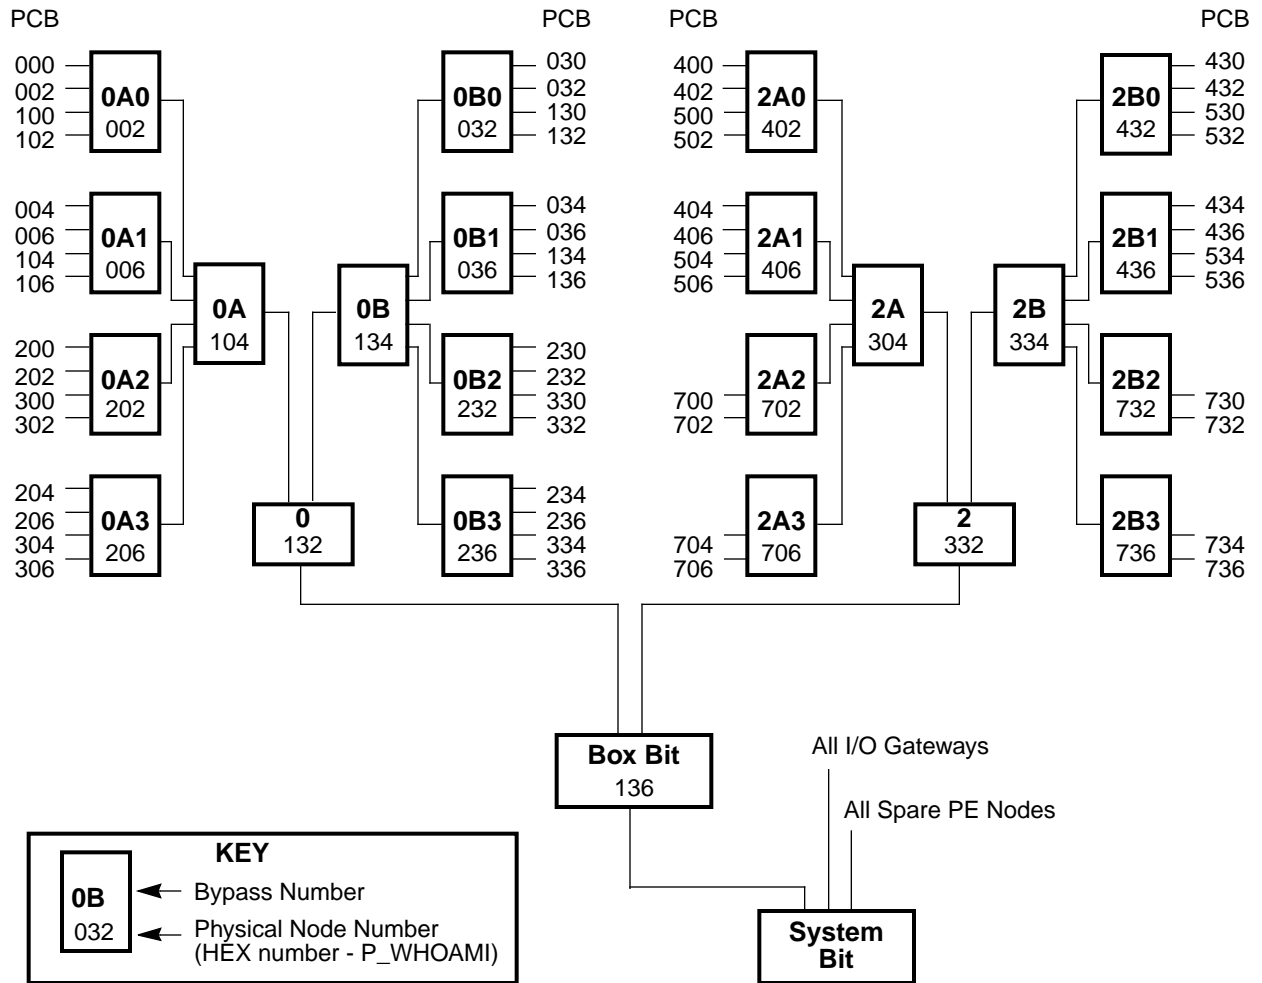
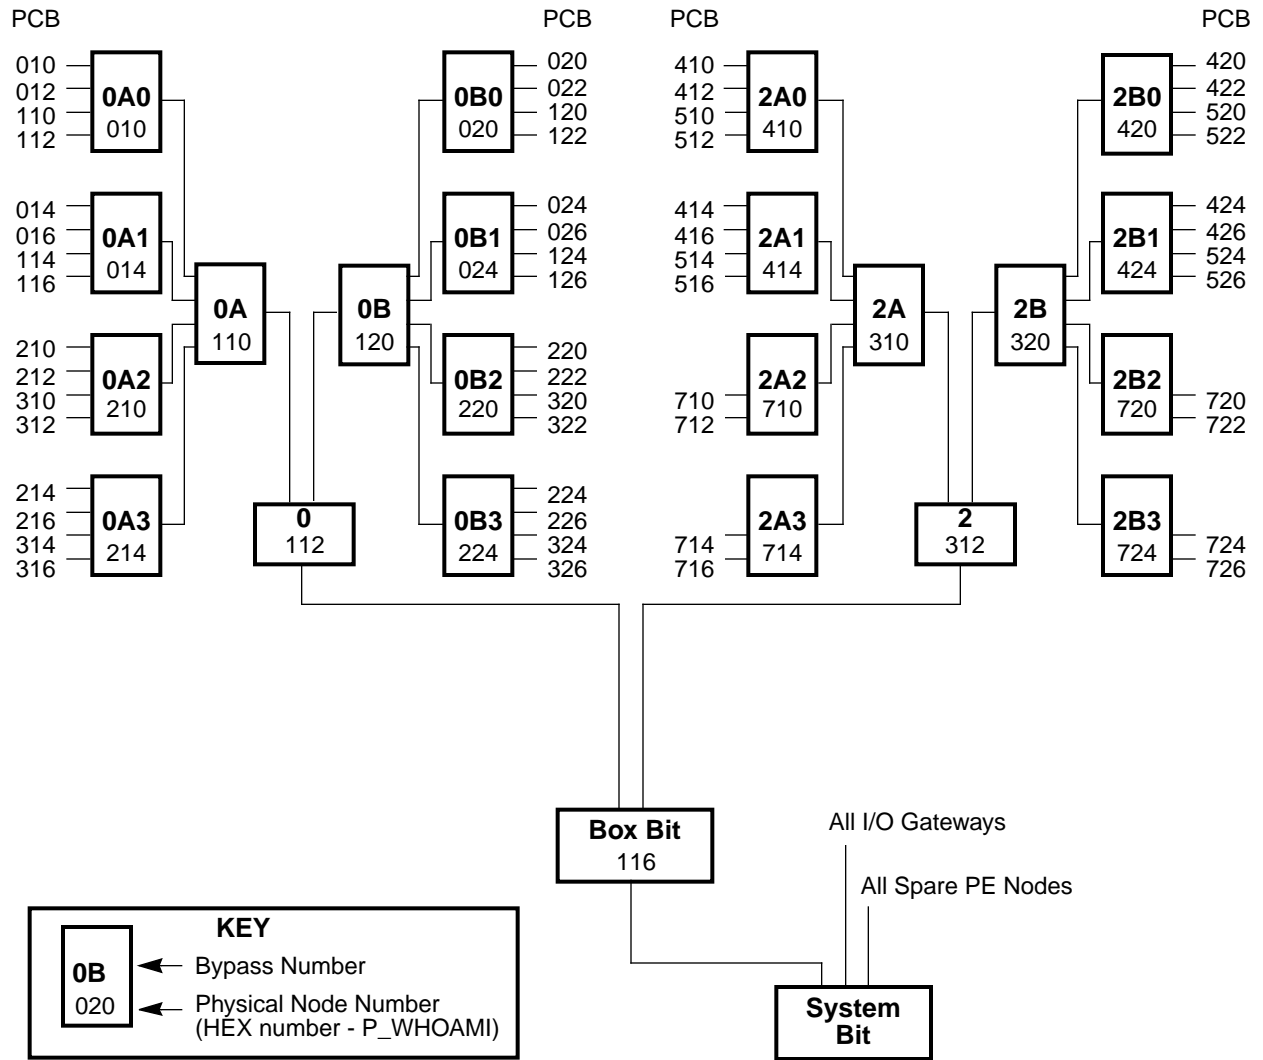


Figure 20. Barrier Synchronization Circuit 0 in CRAY T3D SC224 System



NOTE: The PCB bypass points shown in this figure are controlled by bit 2^{12} of the NODE_CSR. The other bypass points shown in this figure are controlled by bit 2^{13} of the NODE_CSR.

Figure 21. Barrier Synchronization Circuit 1 in CRAY T3D SC224 System



NOTE: The PCB bypass points shown in this figure are controlled by bit 2¹¹ of the NODE_CSR. The other bypass points shown in this figure are controlled by bit 2¹³ of the NODE_CSR.

Figure 22. Barrier Synchronization Circuit 2 in CRAY T3D SC224 System

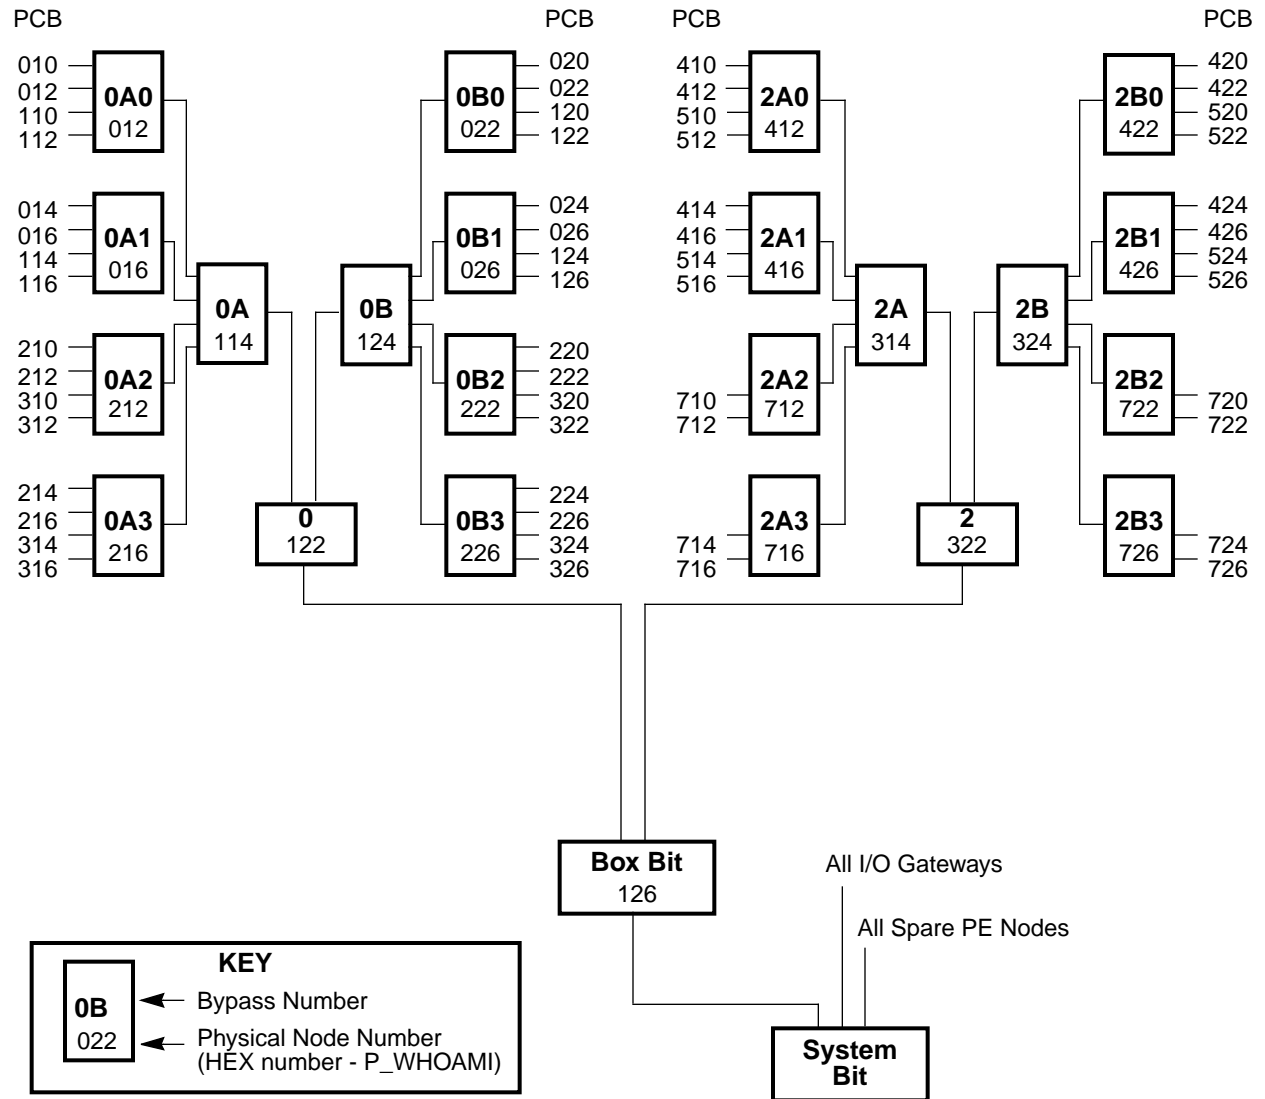


Figure 23. Barrier Synchronization Circuit 3 in CRAY T3D SC224 System

4 CRAY T3D SC192 System

The CRAY T3D SC192 system contains 192 PEs in 96 processing element nodes and is housed in one cabinet. The following subsections describe the communication links, module layout, and barrier synchronization bypass points for the CRAY T3D SC192 system.

NOTE: Throughout this document, physical PE numbers are represented as three-digit hexadecimal numbers that correspond to the physical PE numbers in the P_WHOAMI registers. For example, a PE may have a physical PE number of $10B_{16}$.

NOTE: Throughout this document, physical node numbers are represented as three-digit hexadecimal numbers. This number is equivalent to the number read from a P_WHOAMI register with bit 2^0 set to 0. For example, physical node $10A_{16}$ contains the physical PEs $10A_{16}$ and $10B_{16}$.

NOTE: Throughout this document, logical node numbers are represented as three-digit hexadecimal numbers. This number is equivalent to the number read from a L_WHOAMI register with bit 2^0 set to 0. For example, logical node $10A_{16}$ contains the logical PEs $10A_{16}$ and $10B_{16}$.

4.1 CRAY T3D SC192 Communication Links

Figure 24 shows the physical communication links between nodes in the Y dimension.




Figure 24. CRAY T3D SC192 Y-dimension Communication Links

Figure 25 shows the physical communication links between spare nodes in the Y dimension.




Figure 25. CRAY T3D SC192 Spare Node Y-dimension Communication Links

Figure 26 shows the physical communication links between the nodes in the X and Z dimensions. For clarity, the communication links that complete the torus in the X and Z dimensions are not shown.


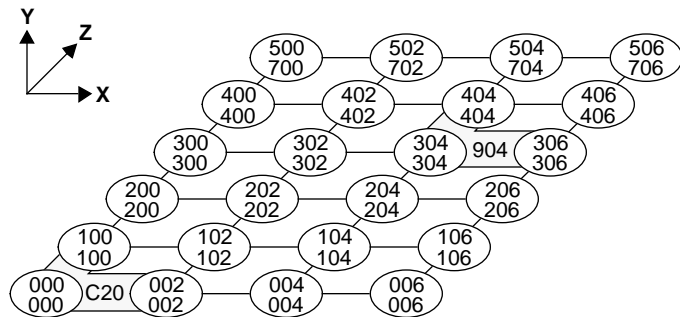

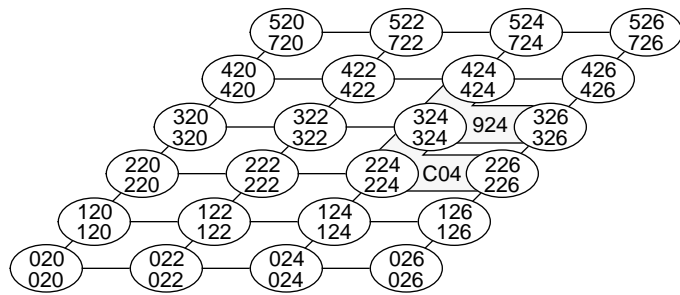



Figure 26. CRAY T3D SC192 X- and Z-dimension Communication Links

4.2 CRAY T3D SC192 Module Layout

Figure 27 shows the module layout and physical node locations in the CRAY T3D SC192 system cabinet. In each figure, the physical node number is represented as a three-digit hexadecimal number that is equivalent to the number read from the P_WHOAMI register with bit 2^0 set to 0. For example, physical node 002_{16} contains the physical PEs 002_{16} and 003_{16} .




Figure 27. CRAY T3D SC192 Module Layout

4.3 CRAY T3D SC192 Barrier Synchronization Circuits

Figure 28 through Figure 31 show the barrier synchronization bypass points for all four physical barrier synchronization circuits in the CRAY T3D SC192 system. In these figures, the printed circuit board (PCB) bypass points and all other bypass points are shown.

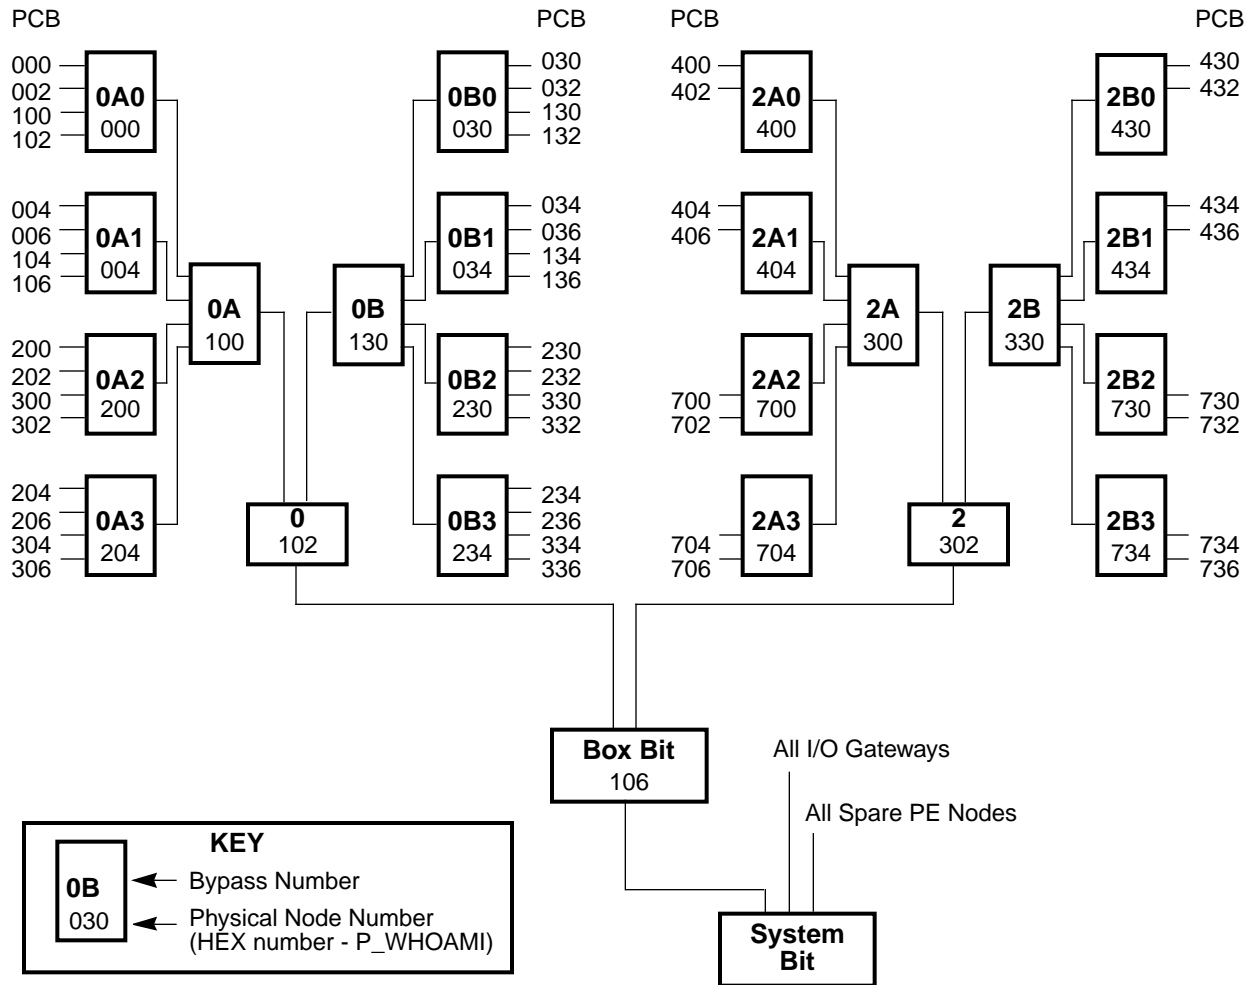
The PCB bypass points are controlled by bit 2^{11} or bit 2^{12} of the network mode register (NODE_CSR). Every processing element node PCB contains four PCB bypass points (one for each physical barrier synchronization circuit).

Each PCB bypass point is given a name that corresponds to the physical node that controls the bypass point. For example, the PCB bypass point 402 in barrier synchronization circuit 0 is located on the PCB that contains physical PEs 402_{16} , 403_{16} , 412_{16} , and 413_{16} . This bypass is controlled by setting bit 2^{11} of the NODE_CSR in physical node 402 (which contains physical PEs 402_{16} and 403_{16}) to the appropriate value.

When bit 2^{11} or bit 2^{12} of the NODE_CSR is set to 0, the output of the AND gate in a PCB bypass point is redirected to the fanout block in the bypass point. When bit 2^{11} or bit 2^{12} of the NODE_CSR is set to 1, the output of the AND gate in a PCB bypass point is directed to a second-level bypass point.

The other bypass points shown in the following figures are controlled by bit 2^{13} of the NODE_CSR. For example, bypass point 0A in barrier synchronization circuit 0 of cabinet 0 is controlled by setting bit 2^{13} of the NODE_CSR in physical node 000_{16} (which contains physical PEs 000_{16} and 001_{16}) to the appropriate value.

When bit 2^{13} of the NODE_CSR is set to 0, the output of the AND gate in the bypass point is redirected to the fanout block in the bypass point. When bit 2^{13} of the NODE_CSR is set to 1, the output of the AND gate in the bypass point is directed to the next-level bypass point.



NOTE: The PCB bypass points shown in this figure are controlled by bit 2¹¹ of the NODE_CSR. The other bypass points shown in this figure are controlled by bit 2¹³ of the NODE_CSR.

Figure 28. Barrier Synchronization Circuit 0 in CRAY T3D SC192 System

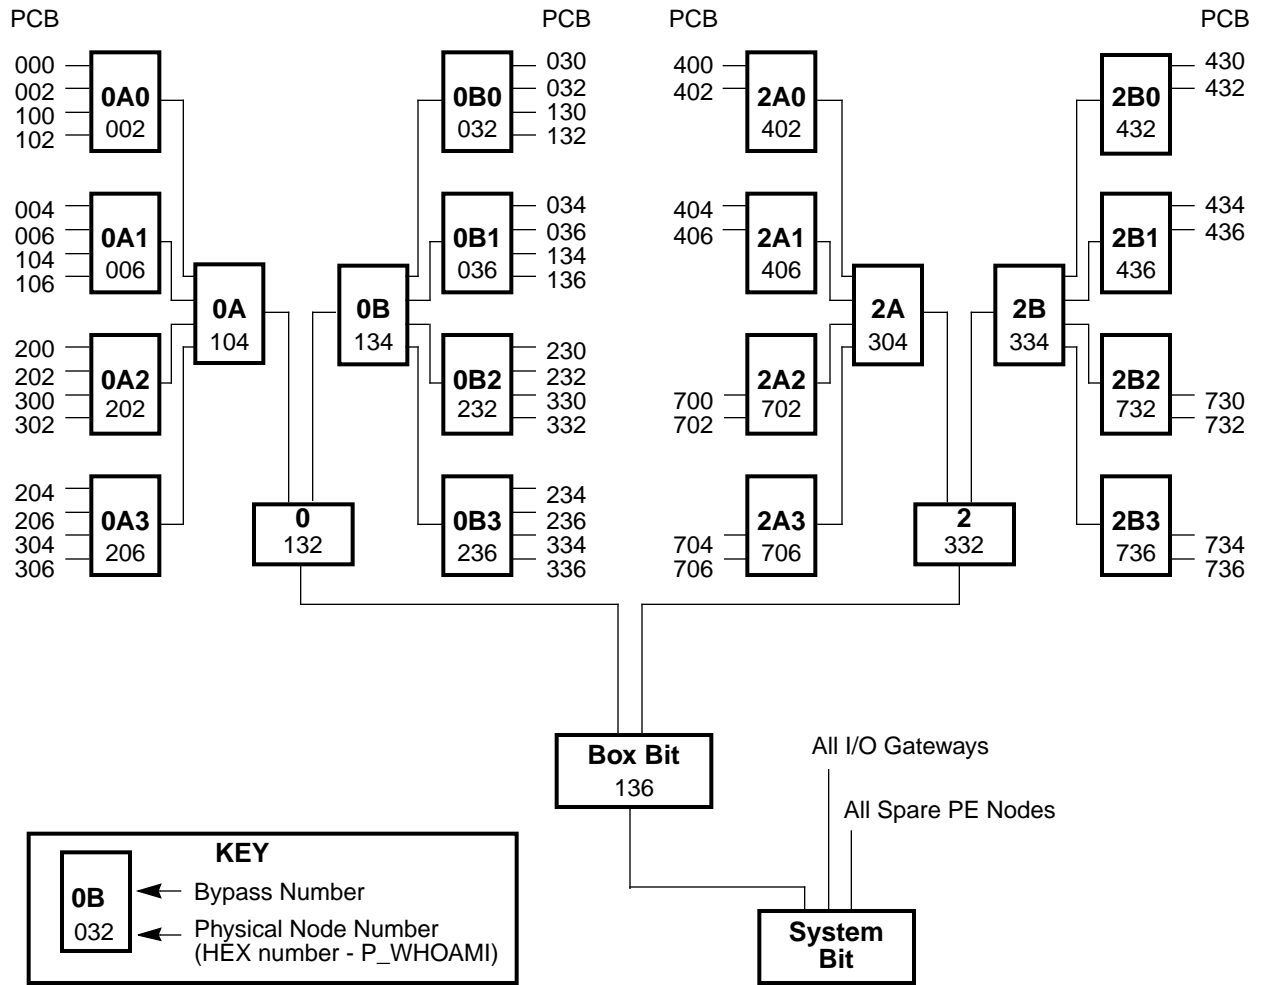


Figure 29. Barrier Synchronization Circuit 1 in CRAY T3D SC192 System

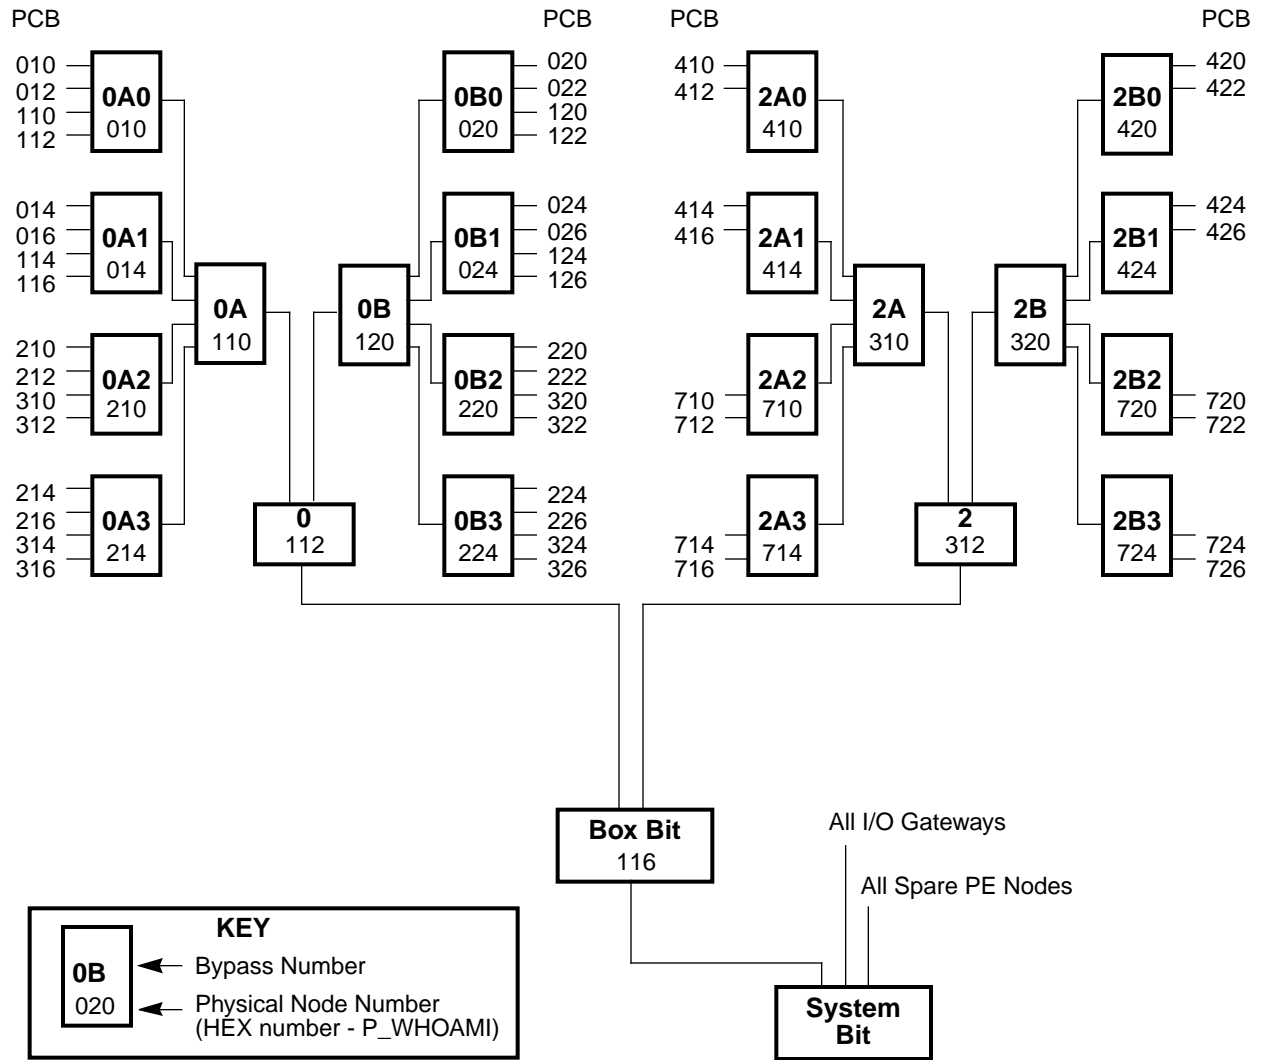


Figure 30. Barrier Synchronization Circuit 2 in CRAY T3D SC192 System

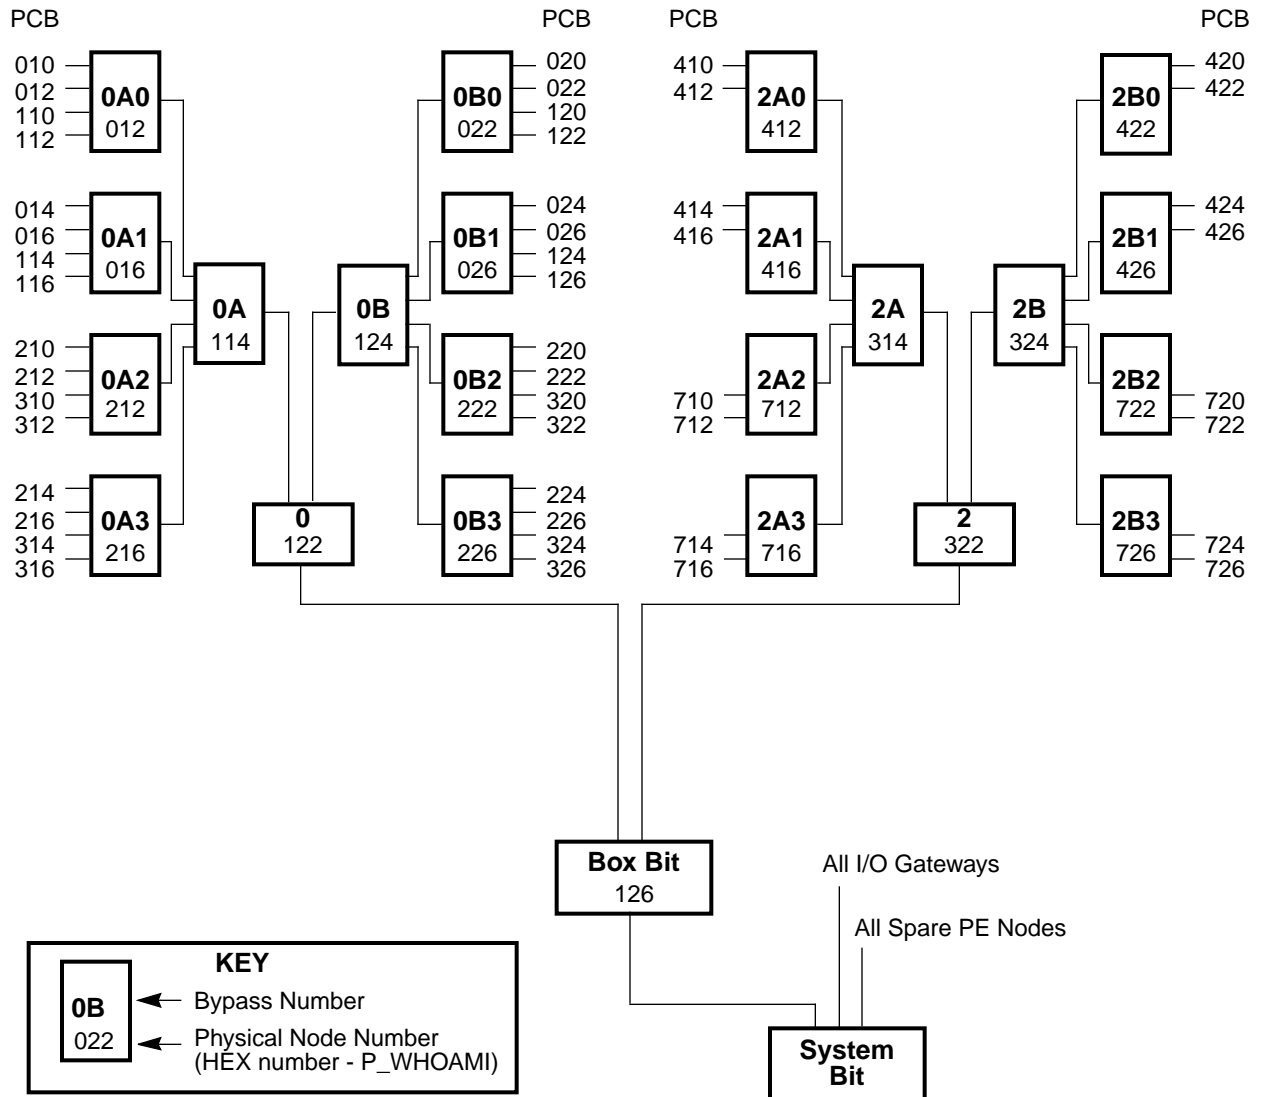


Figure 31. Barrier Synchronization Circuit 3 in CRAY T3D SC192 System

5 CRAY T3D SC160 System

The CRAY T3D SC160 system contains 160 PEs in 80 processing element nodes and is housed in one cabinet. The following subsections describe the communication links, module layout, and barrier synchronization bypass points for the CRAY T3D SC160 system.

NOTE: Throughout this document, physical PE numbers are represented as three-digit hexadecimal numbers that correspond to the physical PE numbers in the P_WHOAMI registers. For example, a PE may have a physical PE number of $10B_{16}$.

NOTE: Throughout this document, physical node numbers are represented as three-digit hexadecimal numbers. This number is equivalent to the number read from a P_WHOAMI register with bit 2^0 set to 0. For example, physical node $10A_{16}$ contains the physical PEs $10A_{16}$ and $10B_{16}$.

NOTE: Throughout this document, logical node numbers are represented as three-digit hexadecimal numbers. This number is equivalent to the number read from a L_WHOAMI register with bit 2^0 set to 0. For example, logical node $10A_{16}$ contains the logical PEs $10A_{16}$ and $10B_{16}$.

5.1 CRAY T3D SC160 Communication Links

Figure 32 shows the physical communication links between nodes in the Y dimension.




Figure 32. CRAY T3D SC160 Y-dimension Communication Links

Figure 32 shows the physical communication links between spare nodes in the Y dimension.




Figure 33. CRAY T3D SC160 Spare Node Y-dimension Communication Links

Figure 34 shows the physical communication links between the nodes in the X and Z dimensions. For clarity, the communication links that complete the torus in the X and Z dimensions are not shown.








Figure 34. CRAY T3D SC160 X- and Z-dimension Communication Links

5.2 CRAY T3D SC160 Module Layout

Figure 35 shows the module layout and physical node locations in the CRAY T3D SC160 system cabinet. In each figure, the physical node number is represented as a three-digit hexadecimal number that is equivalent to the number read from the P_WHOAMI register with bit 2^0 set to 0. For example, physical node 002_{16} contains the physical PEs 002_{16} and 003_{16} .




Figure 35. CRAY T3D SC160 Module Layout

5.3 CRAY T3D SC160 Barrier Synchronization Circuits

Figure 36 through Figure 39 show the barrier synchronization bypass points for all four physical barrier synchronization circuits in the CRAY T3D SC160 system. In these figures, the printed circuit board (PCB) bypass points and all other bypass points are shown.

The PCB bypass points are controlled by bit 2^{11} or bit 2^{12} of the network mode register (NODE_CSR). Every processing element node PCB contains four PCB bypass points (one for each physical barrier synchronization circuit).

Each PCB bypass point is given a name that corresponds to the physical node that controls the bypass point. For example, the PCB bypass point 402 in barrier synchronization circuit 0 is located on the PCB that contains physical PEs 402_{16} , 403_{16} , 412_{16} , and 413_{16} . This bypass is controlled by setting bit 2^{11} of the NODE_CSR in physical node 402 (which contains physical PEs 402_{16} and 403_{16}) to the appropriate value.

When bit 2^{11} or bit 2^{12} of the NODE_CSR is set to 0, the output of the AND gate in a PCB bypass point is redirected to the fanout block in the bypass point. When bit 2^{11} or bit 2^{12} of the NODE_CSR is set to 1, the output of the AND gate in a PCB bypass point is directed to a second-level bypass point.

The other bypass points shown in the following figures are controlled by bit 2^{13} of the NODE_CSR. For example, bypass point 0A in barrier synchronization circuit 0 of cabinet 0 is controlled by setting bit 2^{13} of the NODE_CSR in physical node 000_{16} (which contains physical PEs 000_{16} and 001_{16}) to the appropriate value.

When bit 2^{13} of the NODE_CSR is set to 0, the output of the AND gate in the bypass point is redirected to the fanout block in the bypass point. When bit 2^{13} of the NODE_CSR is set to 1, the output of the AND gate in the bypass point is directed to the next-level bypass point.

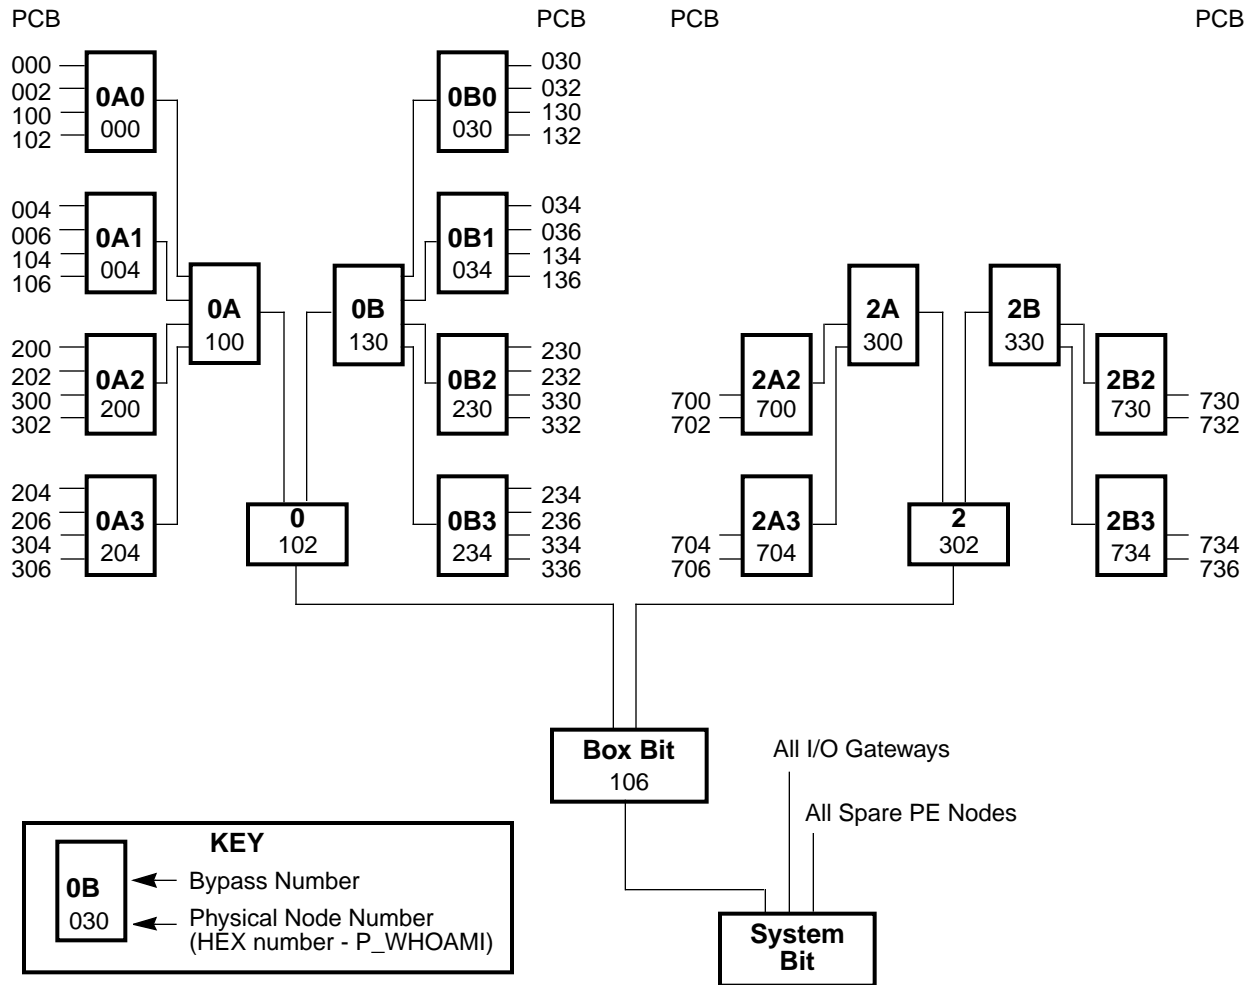


Figure 36. Barrier Synchronization Circuit 0 in CRAY T3D SC160 System

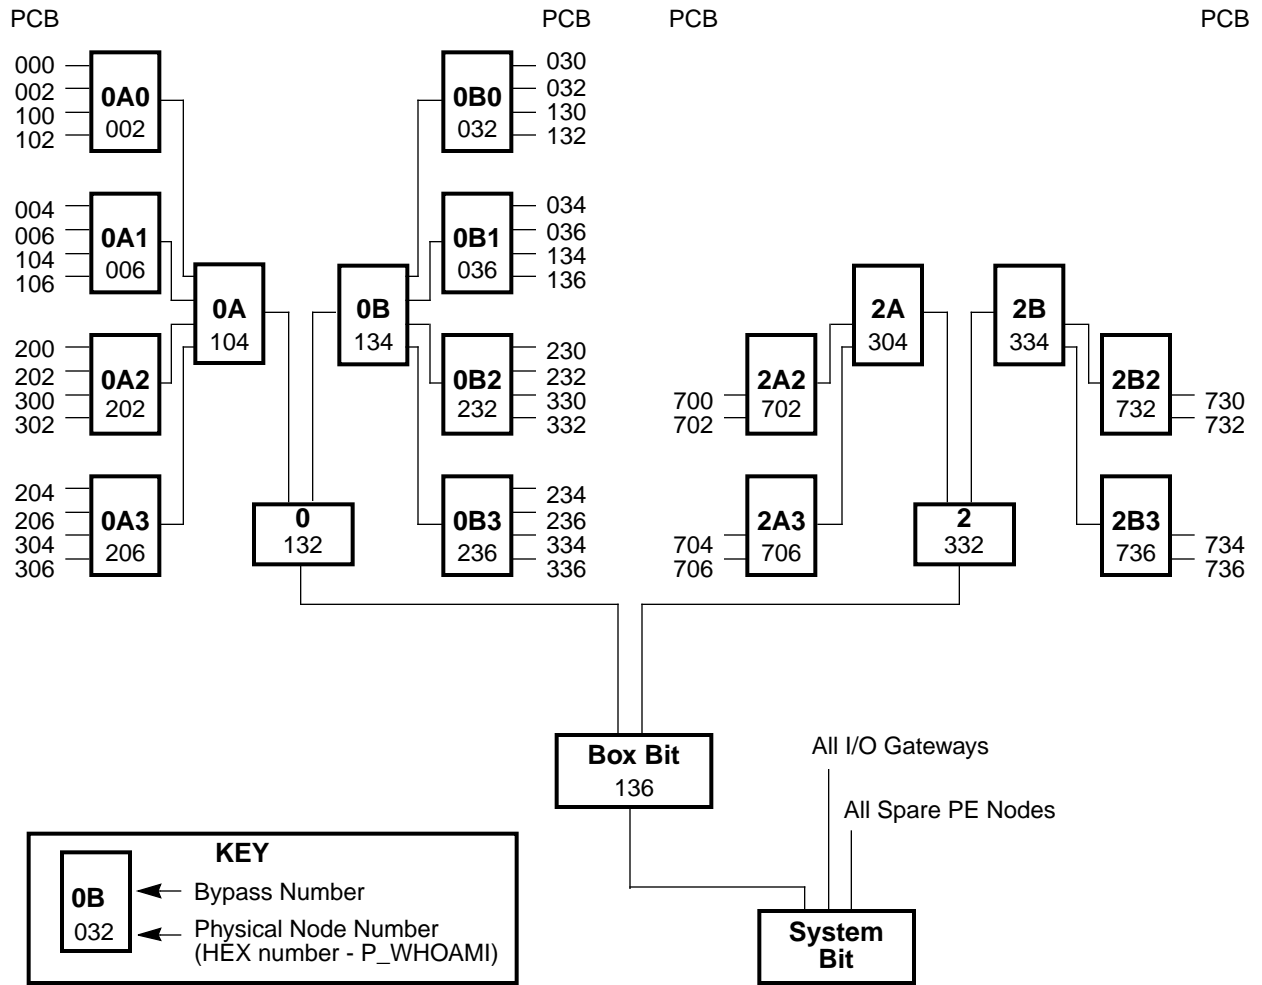


Figure 37. Barrier Synchronization Circuit 1 in CRAY T3D SC160 System

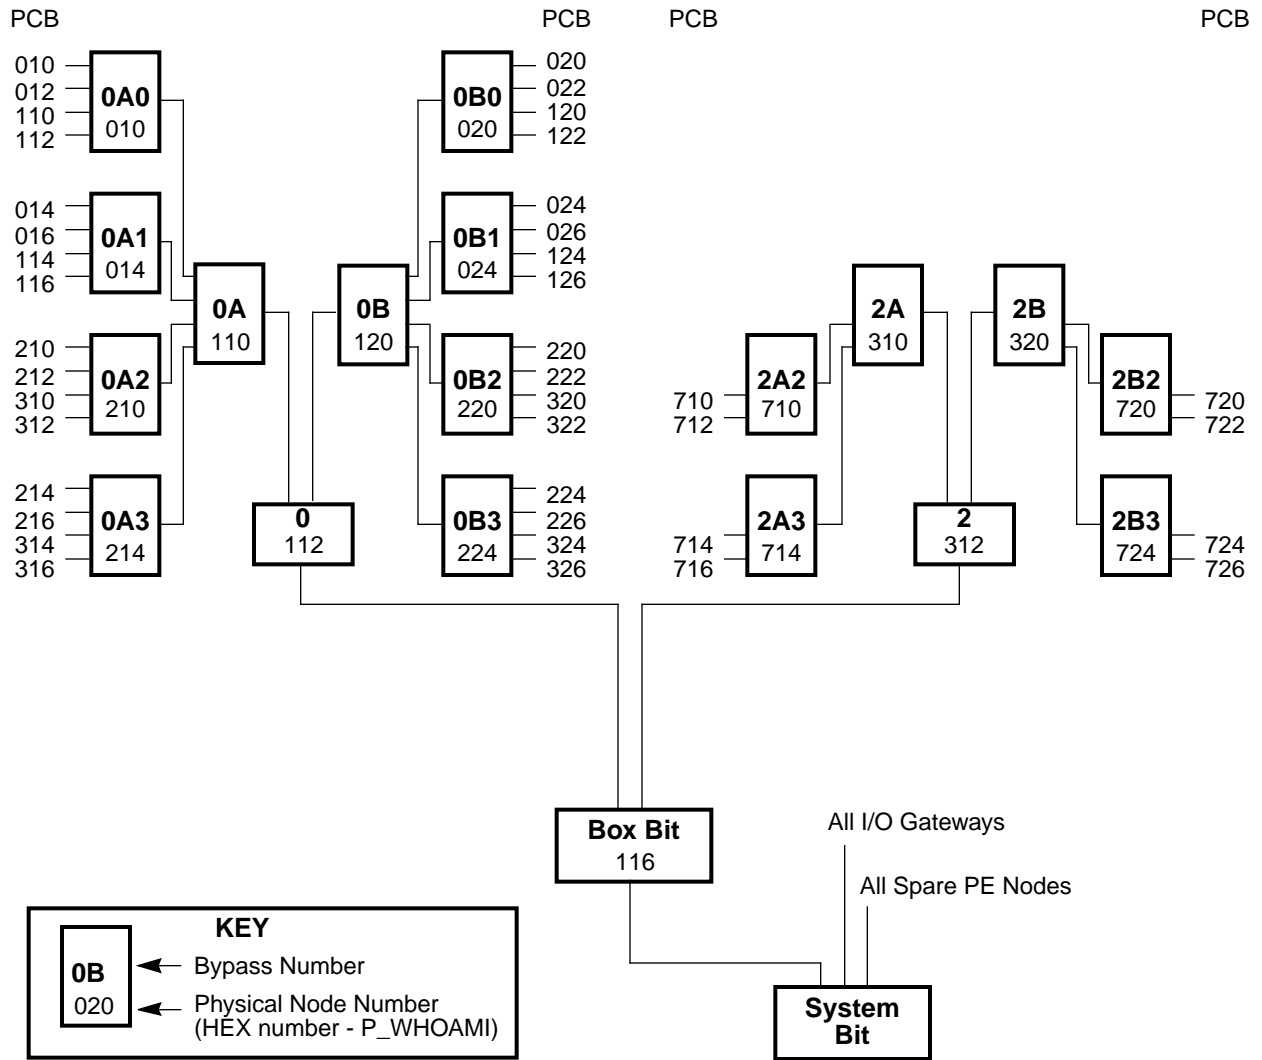
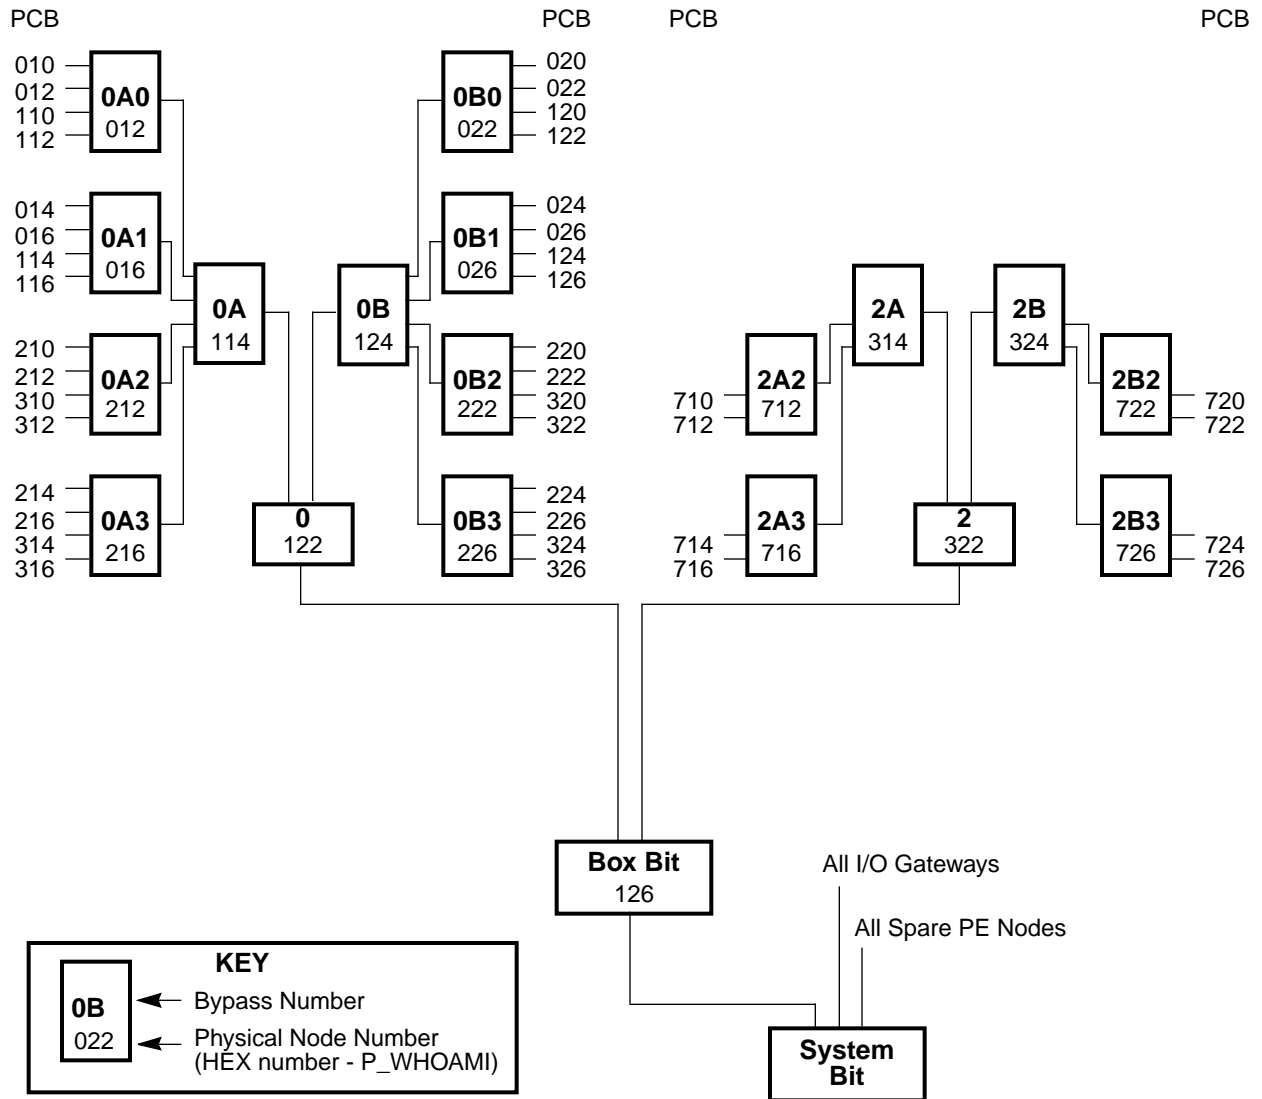


Figure 38. Barrier Synchronization Circuit 2 in CRAY T3D SC160 System



NOTE: The PCB bypass points shown in this figure are controlled by bit 2^{12} of the NODE_CSR. The other bypass points shown in this figure are controlled by bit 2^{13} of the NODE_CSR.

Figure 39. Barrier Synchronization Circuit 3 in CRAY T3D SC160 System

6 CRAY T3D SC128 System

The CRAY T3D SC128 system contains 128 PEs in 64 processing element nodes and is housed in one cabinet. The following subsections describe the communication links, module layout, and barrier synchronization bypass points for the CRAY T3D SC128 system.

NOTE: Throughout this document, physical PE numbers are represented as three-digit hexadecimal numbers that correspond to the physical PE numbers in the P_WHOAMI registers. For example, a PE may have a physical PE number of $10B_{16}$.

NOTE: Throughout this document, physical node numbers are represented as three-digit hexadecimal numbers. This number is equivalent to the number read from a P_WHOAMI register with bit 2^0 set to 0. For example, physical node $10A_{16}$ contains the physical PEs $10A_{16}$ and $10B_{16}$.

NOTE: Throughout this document, logical node numbers are represented as three-digit hexadecimal numbers. This number is equivalent to the number read from a L_WHOAMI register with bit 2^0 set to 0. For example, logical node $10A_{16}$ contains the logical PEs $10A_{16}$ and $10B_{16}$.

6.1 CRAY T3D SC128 Communication Links

Figure 40 shows the physical communication links between nodes in the Y dimension.

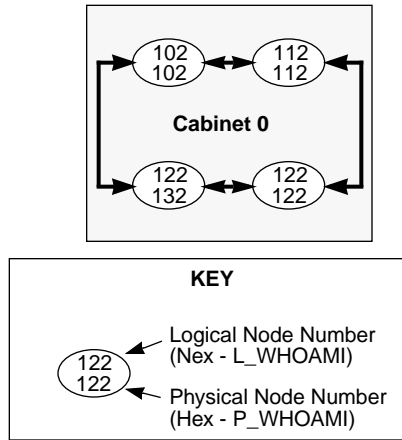


Figure 40. CRAY T3D SC128 Y-dimension Communication Links

Figure 41 shows the physical communication links between spare nodes in the Y dimension.




Figure 41. CRAY T3D SC128 Spare Node Y-dimension Communication Links

Figure 42 shows the physical communication links between the nodes in the X and Z dimensions. For clarity, the communication links that complete the torus in the X and Z dimensions are not shown.









Figure 42. CRAY T3D SC128 X- and Z-dimension Communication Links

6.2 CRAY T3D SC128 Module Layout

Figure 43 shows the module layout and physical node locations in the CRAY T3D SC128 system cabinet. In each figure, the physical node number is represented as a three-digit hexadecimal number that is equivalent to the number read from the P_WHOAMI register with bit 2^0 set to 0. For example, physical node 002_{16} contains the physical PEs 002_{16} and 003_{16} .




Figure 43. CRAY T3D SC128 Module Layout

6.3 CRAY T3D SC128 Barrier Synchronization Circuits

Figure 44 through Figure 47 show the barrier synchronization bypass points for all four physical barrier synchronization circuits in the CRAY T3D SC128 system. In these figures, the printed circuit board (PCB) bypass points and all other bypass points are shown.

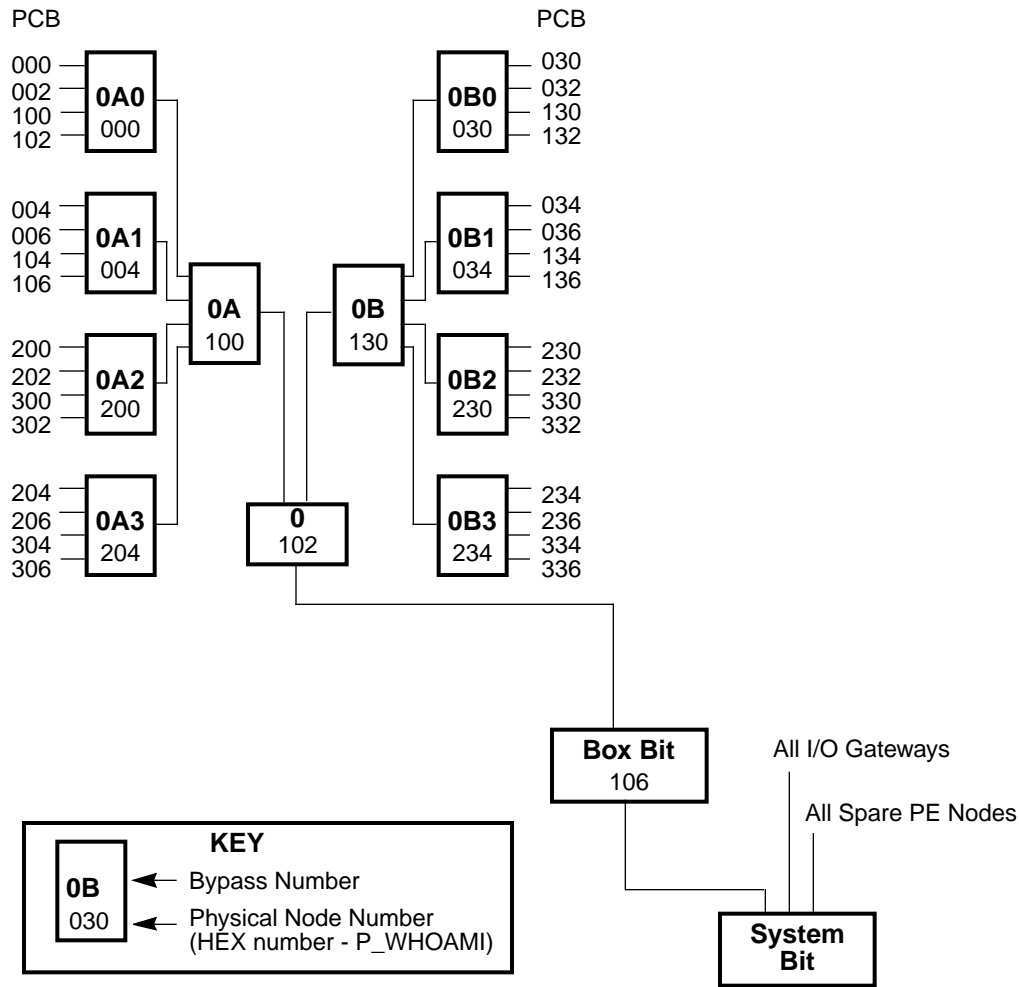
The PCB bypass points are controlled by bit 2^{11} or bit 2^{12} of the network mode register (NODE_CSR). Every processing element node PCB contains four PCB bypass points (one for each physical barrier synchronization circuit).

Each PCB bypass point is given a name that corresponds to the physical node that controls the bypass point. For example, the PCB bypass point 402 in barrier synchronization circuit 0 is located on the PCB that contains physical PEs 402_{16} , 403_{16} , 412_{16} , and 413_{16} . This bypass is controlled by setting bit 2^{11} of the NODE_CSR in physical node 402 (which contains physical PEs 402_{16} and 403_{16}) to the appropriate value.

When bit 2^{11} or bit 2^{12} of the NODE_CSR is set to 0, the output of the AND gate in a PCB bypass point is redirected to the fanout block in the bypass point. When bit 2^{11} or bit 2^{12} of the NODE_CSR is set to 1, the output of the AND gate in a PCB bypass point is directed to a second-level bypass point.

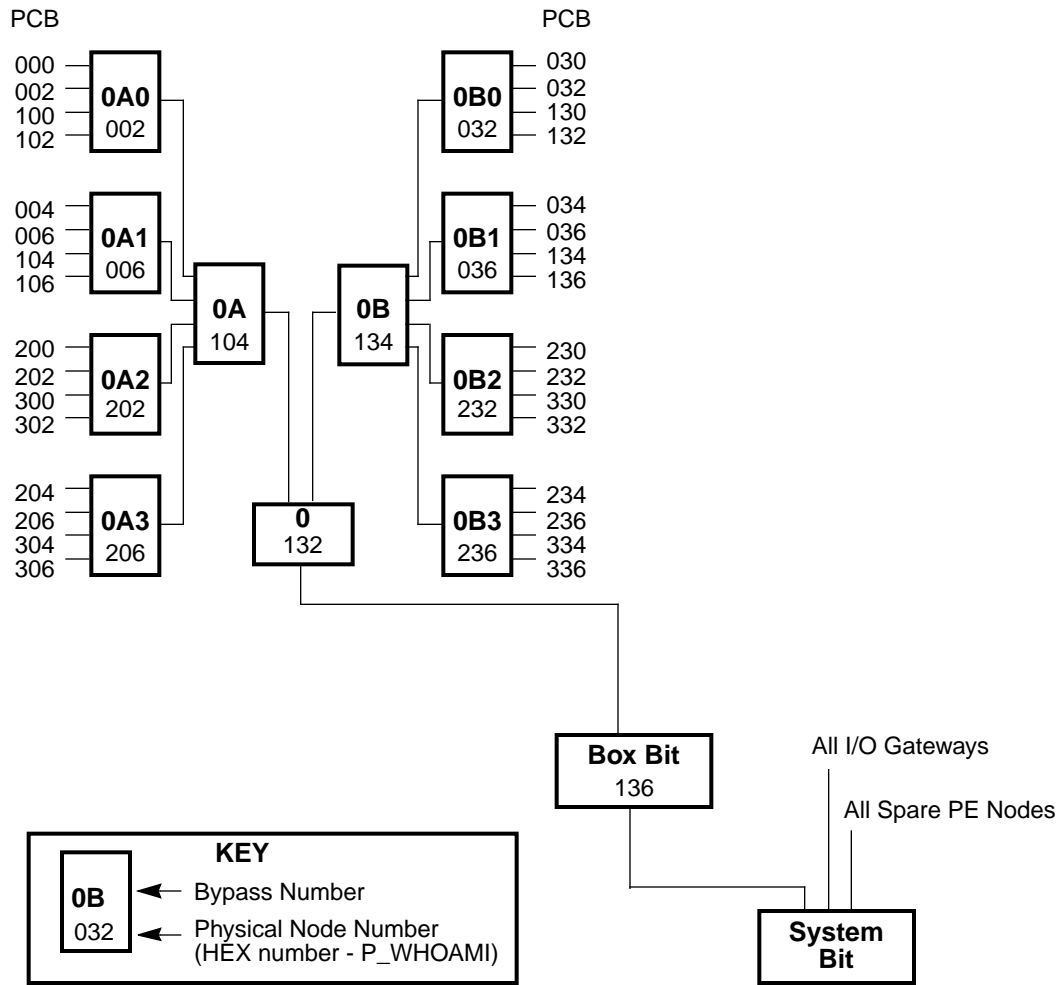
The other bypass points shown in the following figures are controlled by bit 2^{13} of the NODE_CSR. For example, bypass point 0A in barrier synchronization circuit 0 of cabinet 0 is controlled by setting bit 2^{13} of the NODE_CSR in physical node 000_{16} (which contains physical PEs 000_{16} and 001_{16}) to the appropriate value.

When bit 2^{13} of the NODE_CSR is set to 0, the output of the AND gate in the bypass point is redirected to the fanout block in the bypass point. When bit 2^{13} of the NODE_CSR is set to 1, the output of the AND gate in the bypass point is directed to the next-level bypass point.



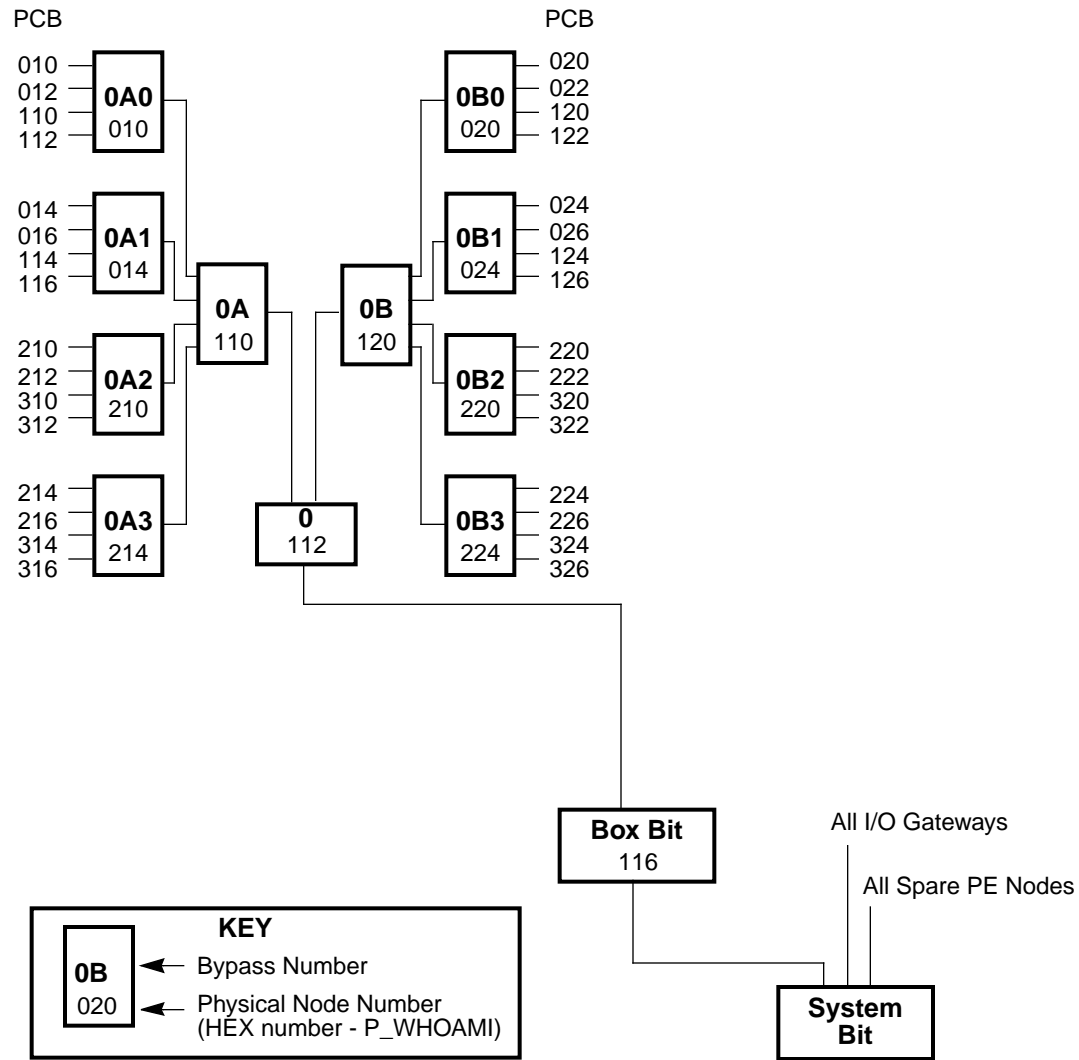
NOTE: The PCB bypass points shown in this figure are controlled by bit 2^{11} of the NODE_CSR. The other bypass points shown in this figure are controlled by bit 2^{13} of the NODE_CSR.

Figure 44. Barrier Synchronization Circuit 0 in CRAY T3D SC128 System



NOTE: The PCB bypass points shown in this figure are controlled by bit 2^{12} of the NODE_CSR. The other bypass points shown in this figure are controlled by bit 2^{13} of the NODE_CSR.

Figure 45. Barrier Synchronization Circuit 1 in CRAY T3D SC128 System



NOTE: The PCB bypass points shown in this figure are controlled by bit 2^{11} of the NODE_CSR. The other bypass points shown in this figure are controlled by bit 2^{13} of the NODE_CSR.

Figure 46. Barrier Synchronization Circuit 2 in CRAY T3D SC128 System

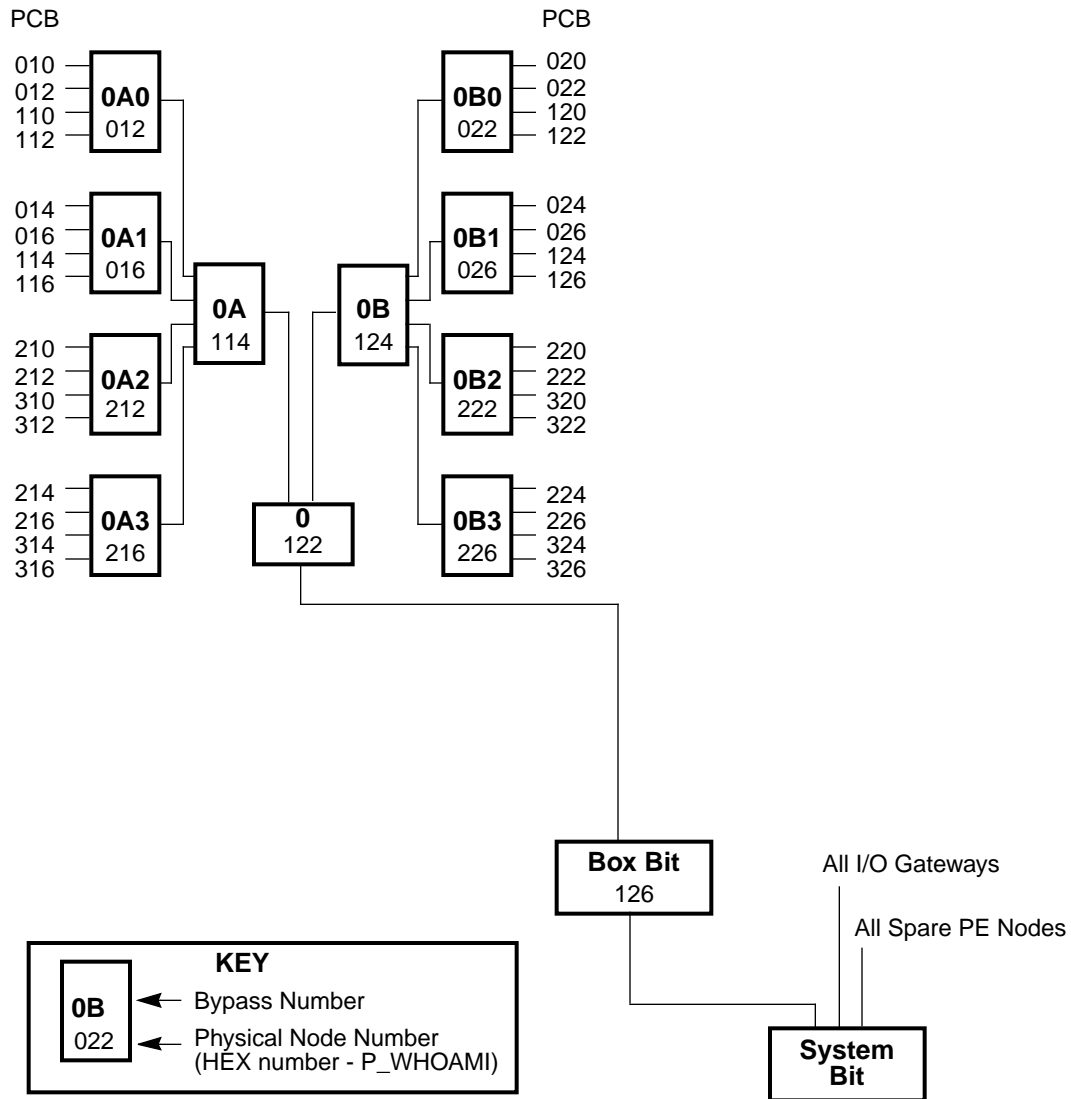


Figure 47. Barrier Synchronization Circuit 3 in CRAY T3D SC128 System