



Installation Guide for Dolphin PCI-SCI Adapters

Part Number: DI950-10353
Version: 1.5
Date: December 7th 2006

Dolphin Interconnect Solutions ASA
Olaf Helsets vei 6
P.O.Box 150, Oppsal
N-0619 Oslo, Norway
Email: pci-support@dolphinics.com
Phone: +47 23 16 70 00 Fax: +47 23 16 71 80

Table of Contents

CHAPTER 1	Introduction	7
1.1	General Overview	7
1.2	Product Components	7
1.3	PCI-to-SCI Adapter Features	7
CHAPTER 2	Installation of a PCI-to-SCI Adapter	8
2.1	Getting started	8
2.2	Installation steps	8
2.2.1	Installing the PCI-to-SCI Adapter	8
2.2.2	Inserting the Cable Connectors	9
2.2.3	Interconnecting Two Nodes	9
2.2.4	Interconnecting Two Nodes with Channel Bonding	10
2.2.5	Interconnecting Multiple Nodes in a Ringlet	11
2.2.6	Interconnecting Multiple Nodes with the D535	11
2.2.7	Interconnecting Multiple Nodes in an SCI Torus topology	12
2.2.8	Other supported topologies	16
CHAPTER 3	Troubleshooting	18
3.1	LED Indications	18
3.2	Power On/Off	18

List of Figures

Figure 1	A D330 Adapter	9
Figure 2	Interconnecting a 2 Node SCI Cluster with Channel Bonding	10
Figure 3	Dual Channel D350	10
Figure 4	4 Node SCI Ring Topology.	11
Figure 5	4 Node SCI Ring Topology with Interleaved Connections.	11
Figure 6	8 Node SCI Switch Topology	12
Figure 7	16 Node SCI Torus Topology (4x4)	13
Figure 8	2D (4x4) Torus Topology (4x4) showing nodeId and x-y location.	14
Figure 9	2D PCI-to-SCI Adapter	15
Figure 10	9 Node 2D SCI Torus Topology (3x3), Cabling Details.	15
Figure 11	NodeIds for a mixed 1D/2D cluster	16
Figure 12	Connecting a 4 Node 2D (2x2) SCI Cluster with Dual Fabric	17

1 Introduction

1.1 General Overview

The Dolphin PCI-SCI Adapter Cards support 32 and 64 bits regular PCI buses, running 33 or 66MHz clock rates [2] or PCI Express (PCIe) [3]. These Adapter cards are part of a high performance solution for clustering personal computers, workstations and servers using the Scalable Coherent Interface (SCI) standard [1]. Applications include clusters of systems for compute processing and high-availability database operations.

Dolphin has developed a whole range of different Adapter cards for different SCI configurations and different PCI environments, such as regular PCI, PMC, cPCI and PCI Express, and these Adapter Cards will be referred to as PCI-to-SCI Adapters throughout this installation guide. All the PCI-to-SCI Adapters can be interconnected in SCI systems.

1.2 Product Components

Dolphin is continuously developing and producing components for clustering. Further information on these products can be found on **www.dolphinics.com**.

1.3 PCI-to-SCI Adapter Features

- Regular PCI Adapters are compliant with PCI Specification 2.2, 32/64-bit, 33/66MHz PCI bus operation [2].
- PCIe Adapters use Intel[®] 41210 PCIe-to-PCI Bridge which is PCI Express Revision 1.0a compliant [3].
- Regular PCI Adapters support both 3.3V and 5V PCI slots (Universal card) [2].
- All Adapters are compliant with IEEE/ANSI 1596-1992 Scalable Coherent Interface specification.[1]
- Low latency and high bandwidth data transfer with remote memory access.
- SCI Interface with unidirectional links, each capable of 667 MByte/sec.

2 Installation of a PCI-to-SCI Adapter

2.1 Getting started

Open the shipping package to verify that you have received the items ordered. Inspect each item for damage.

CAUTION

Static electricity can damage integrated circuits on the card. Do not remove the Adapter from the anti-static bag until you have discharged any static electricity from your person by touching a grounded object. Maintain grounding by wearing the wrist strap provided, when handling the Adapter.

If everything is in order, continue the installation steps described in 2.2. If not, contact your Dolphin sales representative.

2.2 Installation steps

2.2.1 Installing the PCI-to-SCI Adapter

First some general PCIe considerations. The D350 Adapter is a x8 PCIe Card and requires a x8 or x16 slot, while the other D35x Adapters are x4 Cards and can also use a x4 slot.

Note that the PCIe standard allows up-plugging, but not down-plugging (physically prevented). Up-plugging means e.g. to plug a x4 Adapter into a x8 or x16 connector. However, according to the PCIe standard it is not guaranteed that a chipset is using the width of the Adapter, inserted in such a slot. The specification says that it should come up as the full width or as x1, but most chipsets support the width of the Adapter. If low bandwidth is experienced in such cases, it could be that the slot only supports PCIe x1.

Install the PCI-to-SCI Adapter in your computer by following these steps. General information about installing hardware in a computer is normally found in the installation manual for the equipment.

- 1) Turn OFF the power to the computer and peripheral devices.

CAUTION

Make sure the power is turned OFF before you continue. Wear a properly grounded wrist strap during the complete procedure of installing the Adapter.

- 2) Remove the cover from the computer.
- 3) Remove metal bracket covering the opening of the appropriate PCI slot(s). If it is fixed by a screw, retain the screw for use in step 5.
- 4) Insert the card in the chosen PCI bus slot. Press the card firmly down so the contacts are securely seated in the connector
- 5) Secure the card bracket to the computer chassis (if appropriate, use the screw from step 3).
- 6) Replace the cover on the computer.




Figure 1 A D330 Adapter

An example of an Adapter in the PCI-to-SCI family is shown in Figure 1. Any PCI-to-SCI Adapter has two or more I/O connectors, marked with SCI IN and SCI OUT. VHDCI SCI cables are used for interconnecting such Adapters to other system components, e.g. other Adapters or Switch Ports. There are no jumpers or switches on the Adapter. All configuration is controlled from software.

2.2.2 Inserting the Cable Connectors

When inserting a cable into a connector on the PCI-SCI Adapter, take care that the cable connector is inserted straight, and make sure that no excessive force is used. It is highly recommended to use a torque screwdriver for fastening the screws, and max. torque should be 0.4 Nm. If such a tool is not available, human fingers is a good alternative. Excessive force may break the connector housing, especially if the screw on one side is fully tightened, before the other is entered.

2.2.3 Interconnecting Two Nodes

Prerequisites: 2 PCI-SCI Adapters, 2 SCI Cables. It is assumed that the 2 Adapters are installed in computers, as described in the previous section, and as an example, they're called A and B.

- 1) Insert one cable from Adapter A: SCI OUT to Adapter B: SCI IN.
- 2) Insert the other cable from Adapter B: SCI OUT to Adapter A: SCI IN.
- 3) Turn Power back on, and observe that the LED between the SCI IN/OUT connectors turn green.

This completes the hardware installation of a two-node system.

2.2.4 Interconnecting Two Nodes with Channel Bonding

2.2.4.1 Two Separate 1D PCI-to-SCI Adapters

It is possible to obtain twice the bandwidth of a single PCI-to-SCI Adapter by installing two cards (in two different PCI buses). This requires also twice as many cables, and the two cards are connected in parallel. Let's assume we have two nodes, called A and B, which have two Adapters each, called A0, A1 and B0, B1. The connections should be done as follows:

- 1) Connect one cable from Adapter A0: SCI OUT to Adapter B0: SCI IN
- 2) Connect one cable from Adapter A1: SCI OUT to Adapter B1: SCI IN
- 3) Connect one cable from Adapter B0: SCI OUT to Adapter A0: SCI IN
- 4) Connect one cable from Adapter B1: SCI OUT to Adapter A1: SCI IN

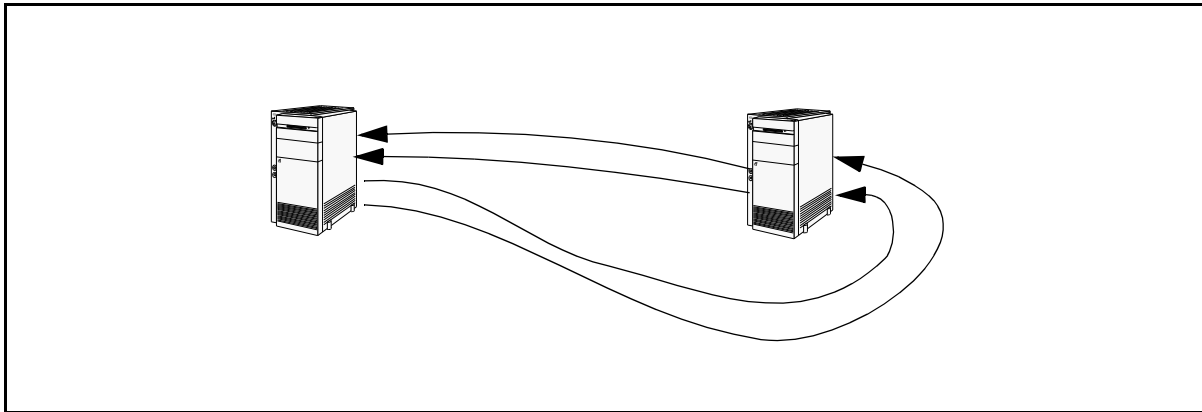


Figure 2 Interconnecting a 2 Node SCI Cluster with Channel Bonding

2.2.4.2 The D350, Dual 1D PCI-to-SCI Adapter

Dolphin has a special Dual Cycle Adapter (D350) that in one x8 PCI Express slot, can connect two virtual channels between two nodes. The interconnection scheme is exactly the same as above, except for the fact that the two channels are located on the same Adapter. Let's again assume two nodes, A and B, shall be connected together.




Figure 3 Dual Channel D350

- 1) Connect one cable from Adapter A Ch 0: SCI OUT to Adapter B Ch 0: SCI IN
 - 2) Connect one cable from Adapter A Ch 1: SCI OUT to Adapter B Ch 1: SCI IN
 - 3) Connect one cable from Adapter B Ch 0: SCI OUT to Adapter A Ch 0: SCI IN
 - 4) Connect one cable from Adapter B Ch 1: SCI OUT to Adapter A Ch 1: SCI IN
- And that concludes the interconnect of a two node system with Channel bonding

22.5 Interconnecting Multiple Nodes in a Ringlet

Several nodes can be connected in a ringlet, using the same scheme as above. As an example, we assume 3 Adapters called A, B and C, are to be interconnected.

- 1) Insert one cable from Adapter A: SCI OUT to Adapter B: SCI IN
- 2) Insert one cable from Adapter B: SCI OUT to Adapter C: SCI IN
- 3) Insert one cable from Adapter C: SCI OUT to Adapter A: SCI IN

In a similar manner larger ringlets can be interconnected, see Figure 4 and Figure 5. Figure 5 gives an example on interleaved cabling, which eliminates the need for long cables.




Figure 4 4 Node SCI Ring Topology.




Figure 5 4 Node SCI Ring Topology with Interleaved Connections.

22.6 Interconnecting Multiple Nodes with the D535

When interconnecting PCI-to-SCI Adapters by means of a switch, each port of the switch has an input and an output connector, just like a regular node, and one or more nodes can be connected to a switch port, following the same scheme as described above for regular ring.

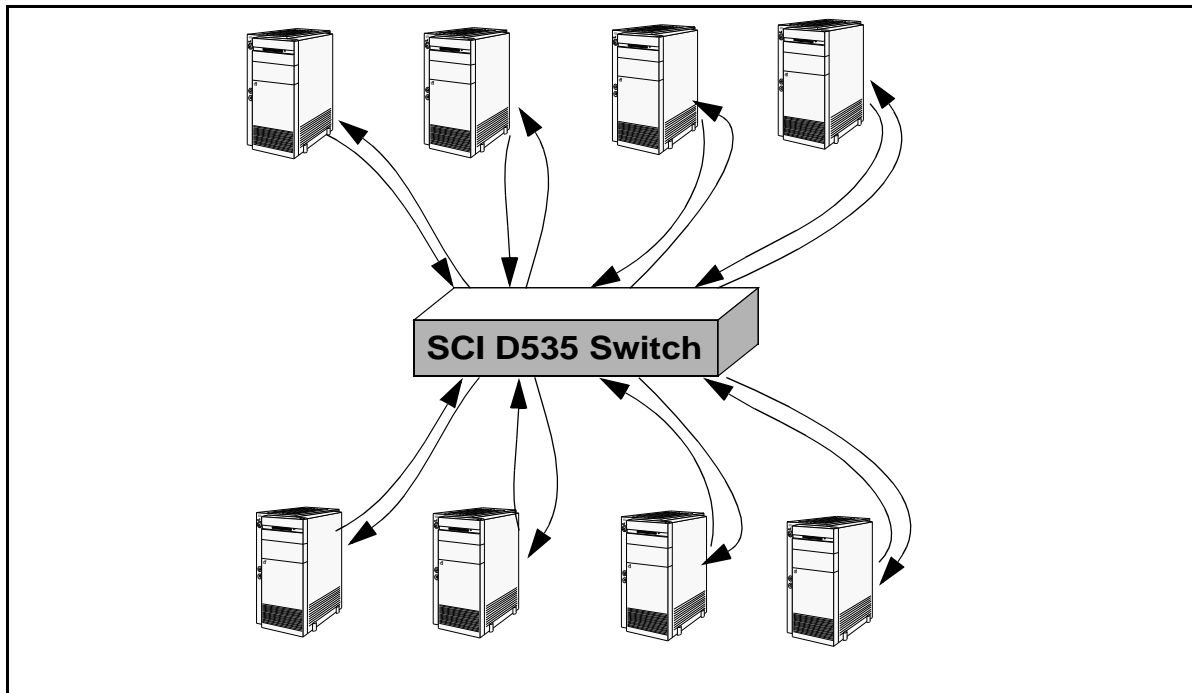


Figure 6 8 Node SCI Switch Topology

22.7 Interconnecting Multiple Nodes in an SCI Torus topology

By using 2 dimensional (2D) PCI-to-SCI Adapters, nodes can be interconnected in a 2D torus configuration. Such Adapters have two SCI-links and can be connected as shown in Figure 7 and Figure 8.

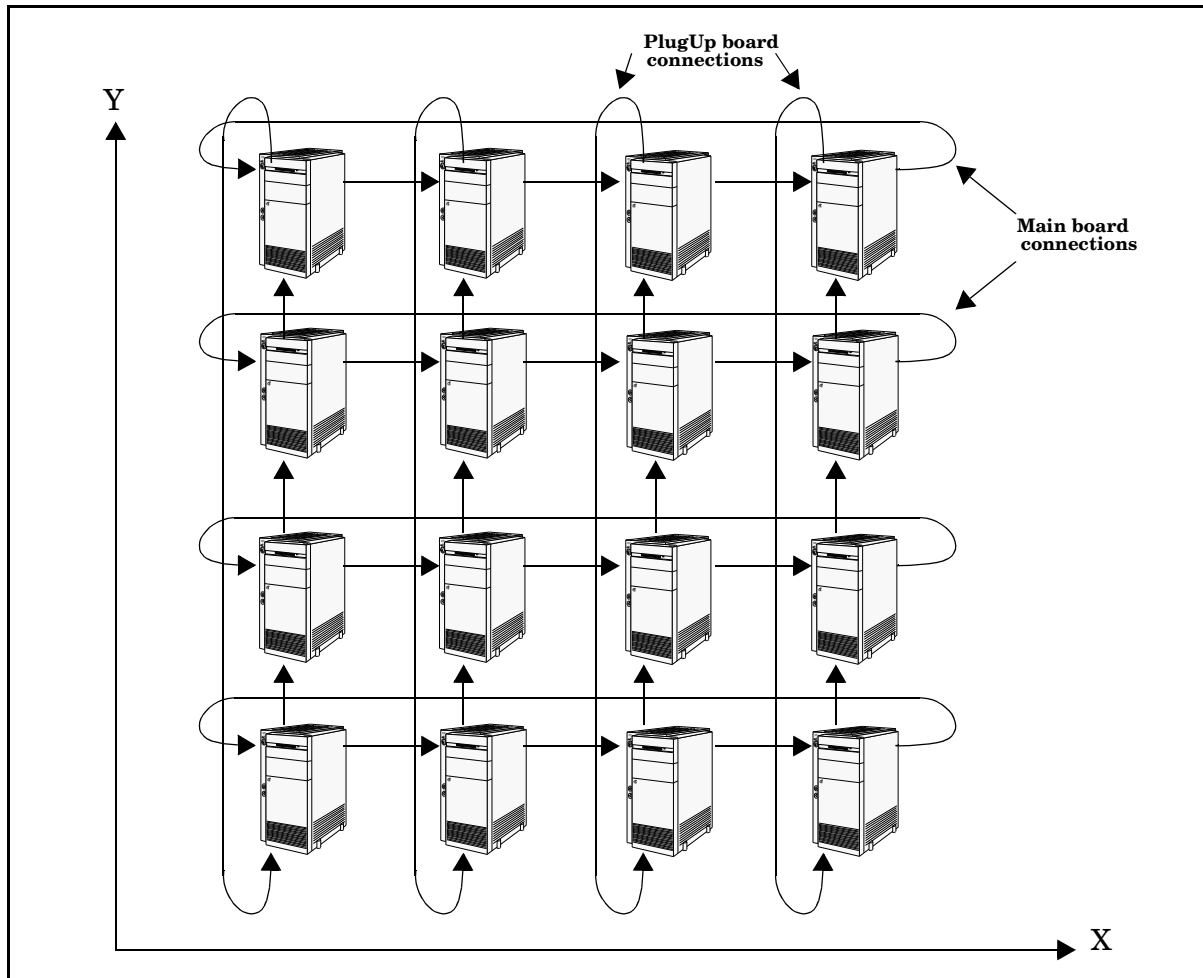


Figure 7 16 Node SCI Torus Topology (4x4)

The minimum number of nodes in a 2D torus cluster is 4 (2x2), and the maximum is 256 (16x16). The nodeId assigned to each node must be specified as shown in Figure 8. The general formula for these nodes are:

- $nodeId = (x+1)*4 + y*64$

For further information about configuration, please refer to the *SCICONFIG - Configuration Tool User's Guide [4]*.

Figure 10 shows cabling details for a 9 node Torus configuration. Cables in the Y-direction should be connected to the link 1, or the plugUp board connectors, while cables in the X-direction should be connected to Link 0, or the main board connectors. See also Figure 9.

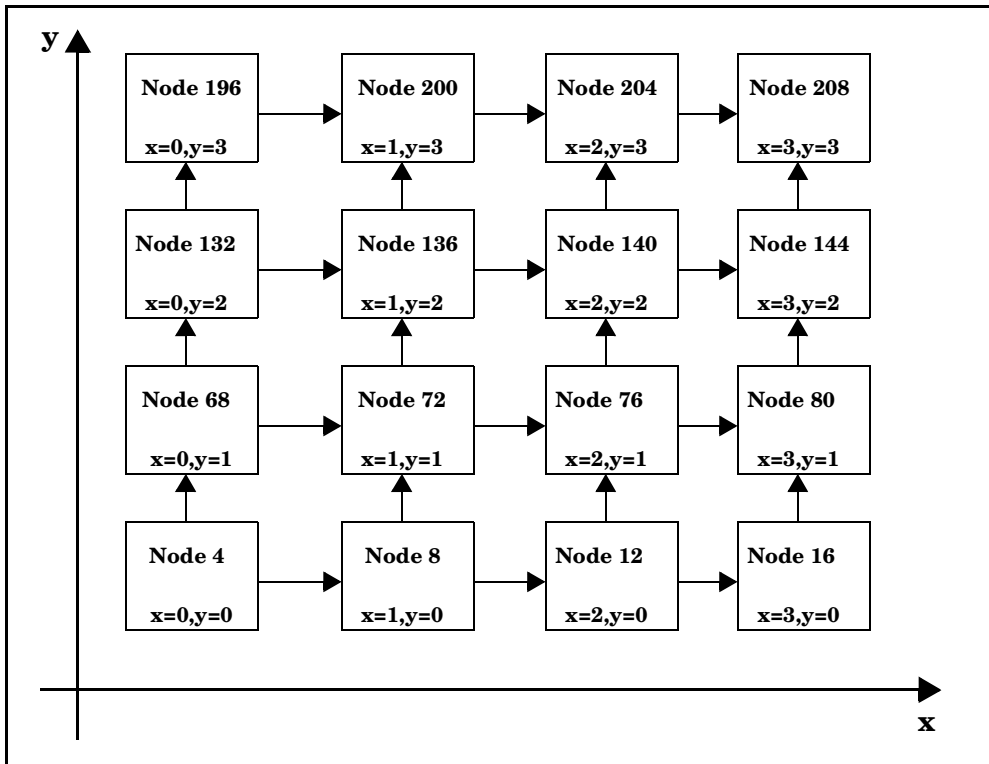


Figure 8 2D (4x4) Torus Topology (4x4) showing nodeId and x-y location.




Figure 9 2D PCI-to-SCI Adapter

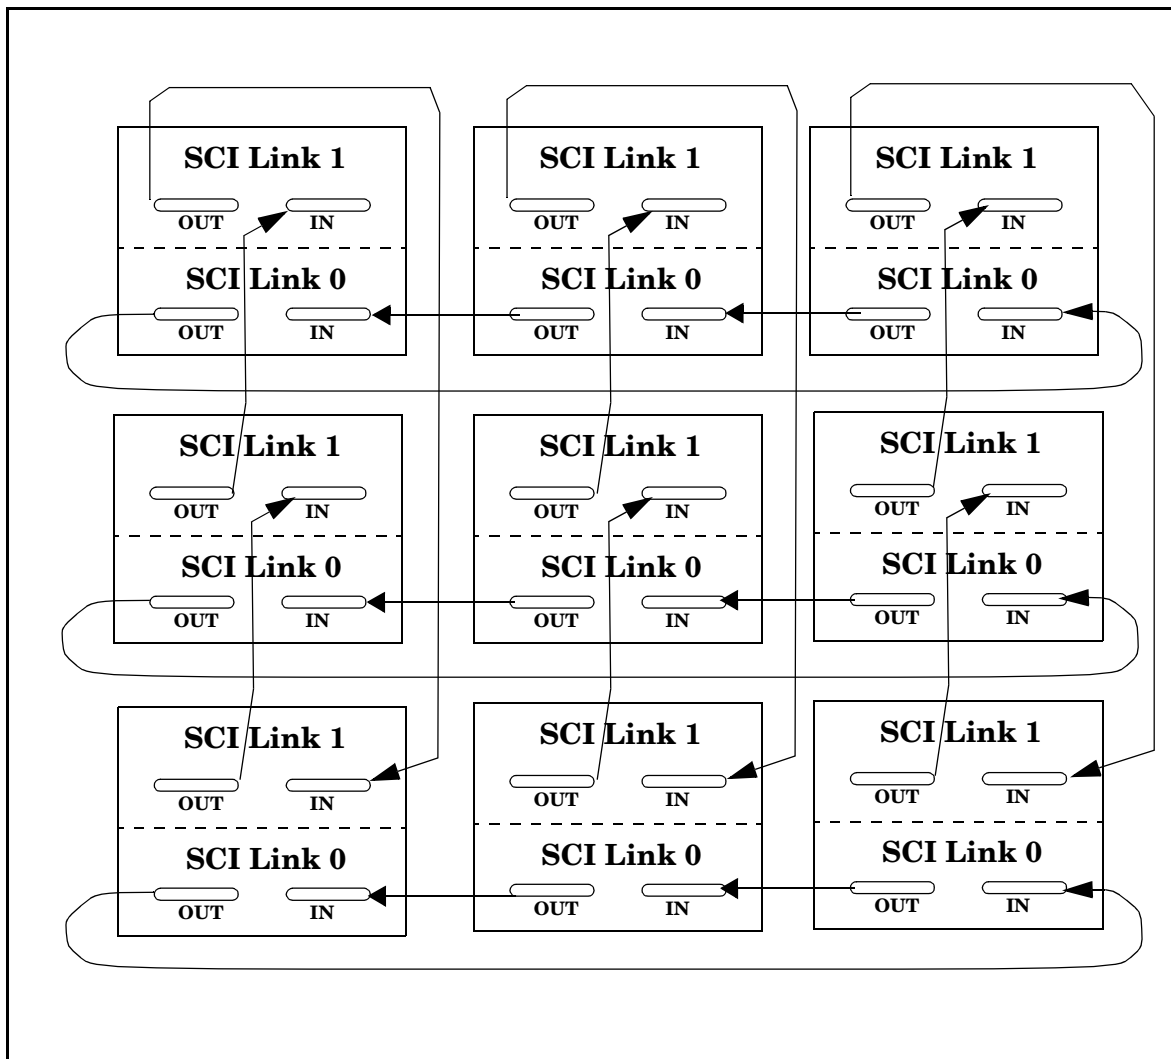


Figure 10 9 Node 2D SCI Torus Topology (3x3), Cabling Details.

For large clusters it is recommended that you avoid the need for long cables by interleaving the connections between nodes. See Figure 5 as an example.

2.2.8 Other supported topologies

2.2.8.1 A Non-Symmetrical System

It is possible to interconnect different Adapters in the same SCI cluster.

Figure 11 shows a combination of 1D and 2D Adapters in the same SCI cluster. Note that the 1D Adapters interconnected in a 2D-torus cluster, must have fixed nodeIds as indicated in Figure 11.

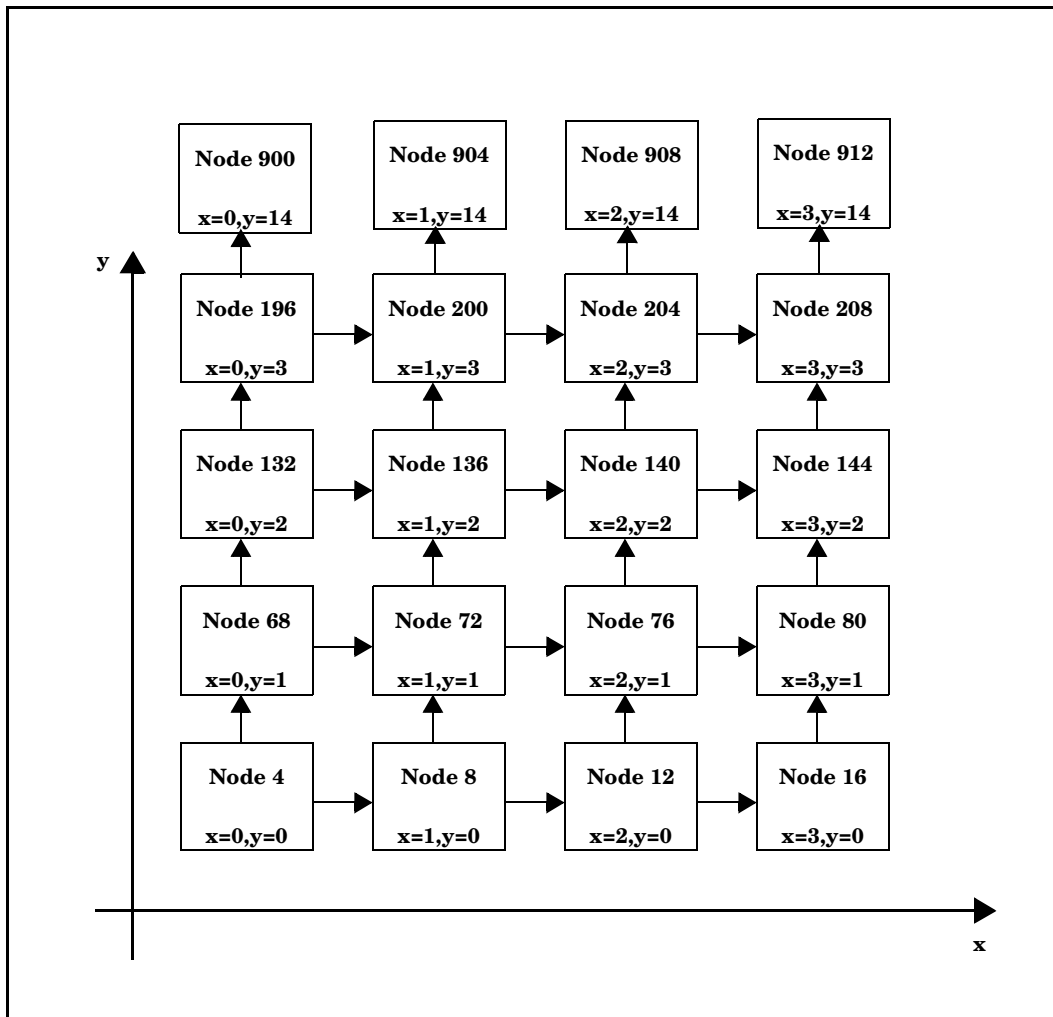


Figure 11 NodeIds for a mixed 1D/2D cluster

2.2.8.2 A Dual Fabric Configuration

It is also possible to configure a 2D cluster with dual fabric to achieve both channel bonding and redundancy. Such a configuration requires 2 ea. 2D Adapters in each machine. Figure 12 gives an example of a 2x2 cluster.

Normally such clusters will consist of the same type of machines, and if so, it is advisable (but not mandatory) to define e.g. Adapter 0 in the same slot in all machines, and that all Adapter 0s are interconnected in one fabric (fabric 0). See Figure 12. This will ease both cabling and debugging if problems occur.

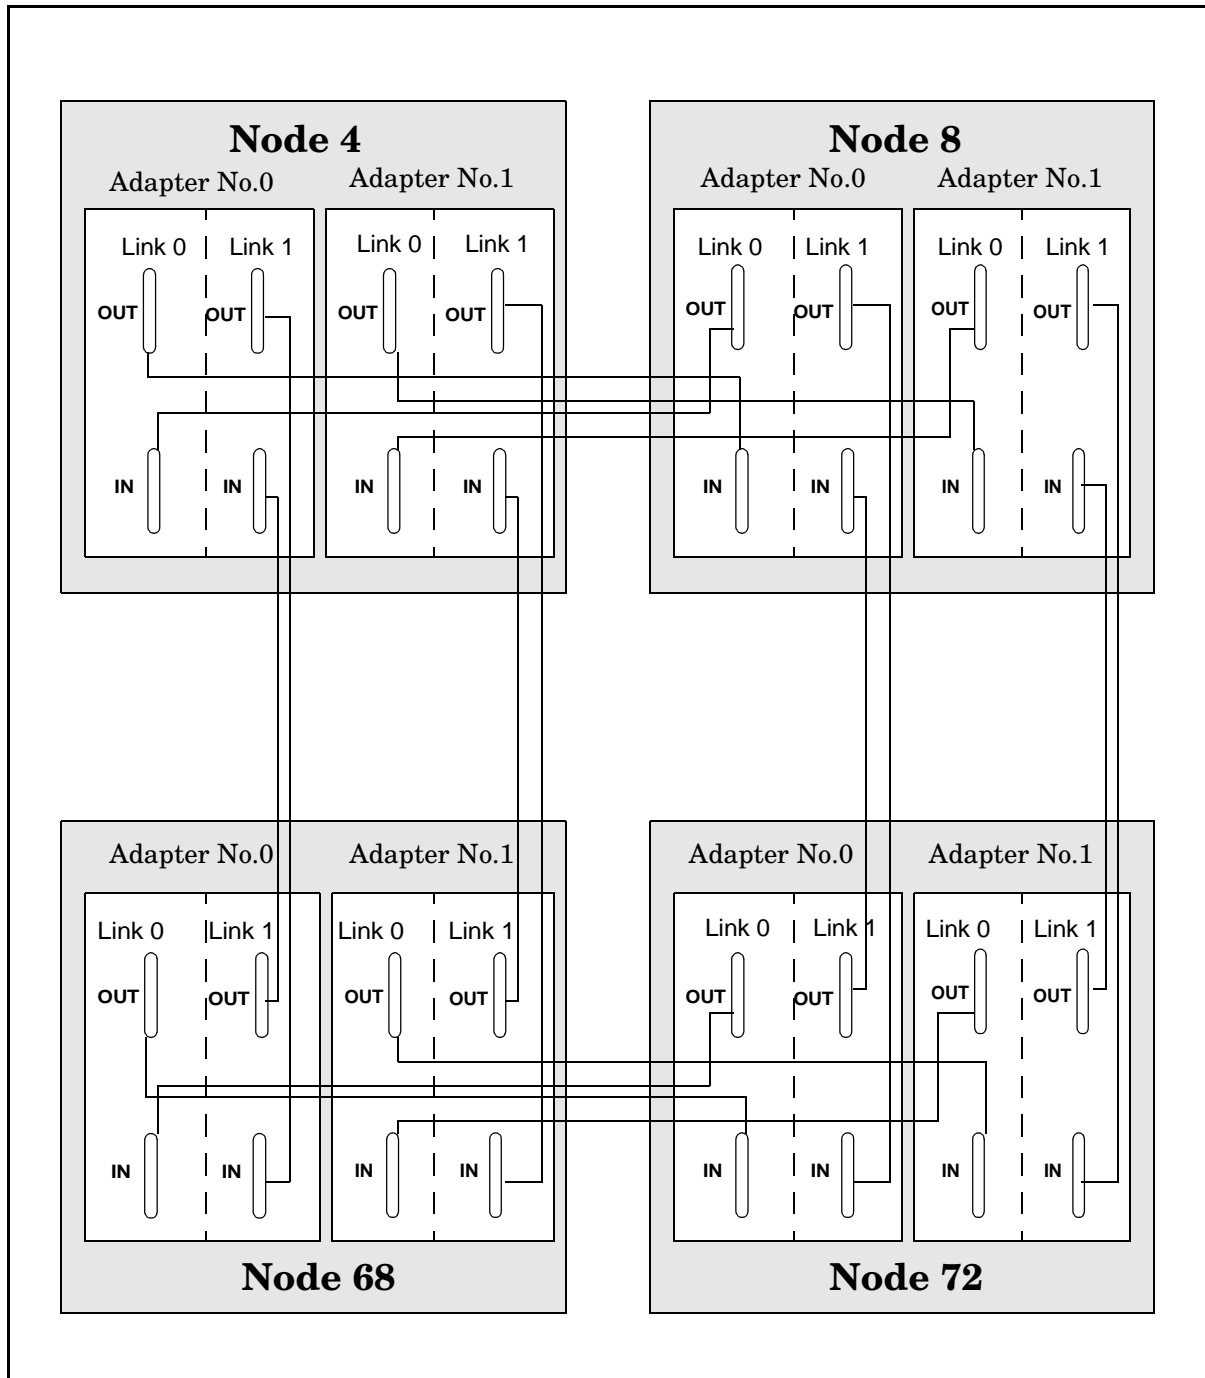


Figure 12 Connecting a 4 Node 2D (2x2) SCI Cluster with Dual Fabric

3 Troubleshooting

3.1 LED Indications

The LED(s) visible through the I/O brackets, can give some hints about the condition of the board and its operation, as indicated in Table 1.

LED Color	LED Behavior	Interpretation
YELLOW	Steady	Power on PCI-to-SCI Adapter, but * no cable installed correctly * remote computer not powered * erroneous cable.
GREEN	Steady	Cable installed and OK.
YELLOW	Short period	Self Test running after * power on * after installation of cable * software reset.
GREEN	blinking	Cable installed and OK, and traffic on B-Link.

Table 1 LED Behavior

If the LED remains YELLOW after installing the Adapters and cables, check the following:

- 1) That cables are installed in the right sequence (in to out) and that they are properly seated.
- 2) If still not GREEN, try each cable in loop-back, i.e. connect the cable between the SCI IN/OUT connectors on one node. In this way a faulty cable or a faulty Adapter could be isolated.
- 3) If all cables fail on one Adapter, it is probably the Adapter which is erroneous and must be replaced. The same exercise could be done on other Adapters to make sure that they are ok, and that the cables are ok.
- 4) If one cable fails, replace the faulty cable.
- 5) For further assistance, consult the support pages on Dolphin's web server: <http://www.dolphinics.com>. where you can find the description of the scidiag program [5]

3.2 Power On/Off

Since newer technology, with its shrinking geometries, requires “new”, lower voltages, it is necessary to generate these voltages onboard, since the PCIe connector [3] offers only 3.3V and 12V directly, and the bulk of the power has to be drawn from 12V. Consequently linear regulators are not longer and option, since the loss would be huge. So regular DC/DC converters are required to generate these voltages, and such regulators require a certain power off time, to discharge capacitors before a new smooth power-on.

Consequently, for all D350 (and higher) series of Dolphin Adapters, it is necessary to wait minimum 5 seconds between power off and power on.

If this power off period is too short, the symptom is normally that the LEDs stay yellow, and if trying to start the operating system, the Dolphin Adapter seems to have vanished. If this is the case, try to power cycle once more with adequate power-off-time.

Appendix F References

- [1] *IEEE Std 1596-1992*
IEEE Standard for Scalable Coherent Interface (SCI)
The Institute of Electrical and Electronic Engineers, Inc., 345 East 45th Street,
New York, NY 10017, USA
- [2] *PCI Local Bus Specification, revision 2.2*
PCI Special Interest Group, 5440 SW Westgate Dr. Suite 217, Portland, OR 97221,
USA
<http://www.pcisig.com>
- [3] *PCI Express Base Specification revision 1.0a*
PCI Special Interest Group, 5440 SW Westgate Dr. #217, Portland, OR 97221
- [4] *SCICONFIG - Configuration Tool User's Guide*
Dolphin Interconnect Solutions ASA, P.O. Box 15, Oppsal, N-0619 Oslo, Norway
<http://www.dolphinics.com>
- [5] *scidiag - SCI diagnostic tool User's Guide*
Dolphin Interconnect Solutions ASA, P.O. Box 15, Oppsal, N-0619 Oslo, Norway
<http://www.dolphinics.com>

