StarFabric PCI bridging: Letting PCI OUT OF THE BOX

By Frank Fitzgerald

The history of PCI Bridging begins with the Peripheral Component Interconnect (PCI) specification. The first specification released by the PCI Special Interest Group in 1993 created PCI. PCI strove to bring system bus speed and register access closer to those of the CPU. Expansion boards on the open architecture bus need to allow on-the-fly base address configuration. Without this flexibility, today's PC plug-and-play features would be impossible. Expansion board flexibility, along with high speed and reflected wave switching, has given PCI significantly higher bandwidth and flexibility than the Industry Standard Architecture and Video Electronics Standards Association buses. The PCI Industrial Computer Manufacturers Group (PICMG) recognized PCI advantages when it became the basis for the CompactPCI specification.

However the designers who enhanced PCI capabilities also introduced some limitations. A 33-MHz PCI bus has a maximum practical load of 10 PCI devices. The PCB edge connector and onboard PCI device each count as a load, thus imposing a four- or five-slot per bus limit. The 66-MHz PCI bus, introduced in 1995, required even tighter design constraints. The expansion possibilities at that speed are limited to one or two slots at best. Designers have stretched systems beyond these limits by carefully controlling motherboard trace propagation characteristics. Also, the variability of expansion cards that can be inserted into each PCI system slot makes overcoming these limits difficult.

Early transparent PCI bridges such as the Digital Equipment Corporation (DEC) 21150 were critical to high-device count and high-slot count PCI system design. These devices allowed designers to create an additional PCI *hop* and create four more slots. The system controller can detect PCI devices across the bridge in a transparent manner. However, the bridge introduces latency, a key trade-off. As PCI evolved with the advent of the 66-MHz extension to the original specification, bridge devices kept pace. In 1996, Digital Equipment Corporation introduced the DEC 21154 with 64-bit bridging capability and followed up with a 66-MHz version in 1998. PCI-X, PCI-X 2.0, and PCI Express offer even higher performance and clearly demonstrate PCI will continue to be the bus interconnect of choice. Few bus architectures have shown the same degree of flexibility and longevity as PCI.

Bringing the system expansion bus closer to the CPU becomes unwieldy when a high bandwidth connection needs to propagate *beyond the box*. PCI box-to-box connections are possible. Aurora Technologies and other manufacturers have been producing these products for several years. Early PCI expansion chassis used the PCI bridge at either end of a cable to carry all bus signals from the main chassis to an expansion chassis. Limitations include cable length, cable cost, and overall expansion topology complexity. Using Ethernet to connect several systems together is an alternative. This requires intelligence within each box, and the remote systems will not maintain the efficient CPU bus connection that PCI provides.

Designers have made numerous attempts to overcome PCI's limitations. These attempts range from addressing only the interconnect scheme to completely displacing PCI. PLX Technology's GigaBridge offered a high-speed serialized interconnect and maintained PCI backward compatibility. Unfortunately, this method required an expensive multiconductor cable for box-to-box bridging. InfiniBand, initially considered a replacement for PCI, did not provide a seamless migration path or backward compatibility with the large number of existing PCI-based products. Eventually PCI-to-InfiniBand bridge products began to appear, but by this time the industry had recognized StarFabric as a solution that balanced PCI compatibility, low cost, and enhanced capability.

StarFabric has several key features that make it the optimal solution for PCI expansion. The StarFabric PCI legacy mode offers 100 percent compatibility with existing PCI driver software. PCI's parallel bus transactions are serialized and formatted into StarFabric data frames. Using StarFabric's high-speed serial Low Voltage Differential Signaling (LVDS), designers can link multiple parallel buses using inexpensive CAT5 cables. Without StarFabric LVDS, it would require at least an 84 conductor cable to bridge a 64-bit PCI bus in parallel. It takes only 32 conductors (four CAT5s) for a 556-Mbytes/sec StarFabric connection. (556 Mbytes/sec is based on a 5 Gbits/sec StarFabric connection with 89 percent data transfer efficiency). The parallel cable costs approximately $60 to $70, and electrical limitations must be less than four feet long. CAT5 cables, combined with StarFabric, allow 10-meter connections for a fraction of the cable cost. The theoretical maximum number of possible PCI bus segments is 256. In common practice, three to four buses per system is the norm. With StarFabric's simplified cabling, larger expansion topologies become more practical. A parallel bus imposes many timing limitations, and PCI is no exception. PCI buses of 66-MHz limit data setup time to 3ns. Maintaining these timing windows within a parallel cable (or even a PCB) is difficult. Designers can attain much longer interconnects by serializing the bus and embedding the clock along with the data. LVDS noise cancellation and high bandwidth can sustain 66-MHz, 64-bit PCI buses.

Figure 1 depicts StarFabric PCI bridging. In this example, a CompactPCI chassis is bridged to two PCI expansion chassis. This provides six PCI expansion slots. Showing the PCI chassis in Figure 1 illustrates cross-platform capability. This feature enables those using CompactPCI systems to take advantage of the many inexpensive PCI peripherals on the market. This *triangle* configuration allows the three systems to be interconnected using one slot in the CompactPCI chassis. PCI initialization software detects and enumerates PCI devices in the expansion topology. The StarFabric protocol also makes switching data frames possible. Switching data frames yields huge benefits. Why? Recall the limitations of parallel PCI-to-PCI bridges. How would five separate chassis be connected?

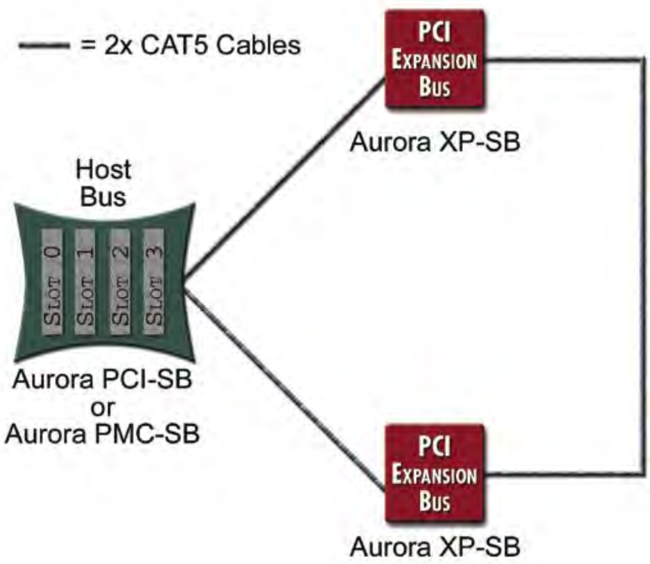

Figure 1

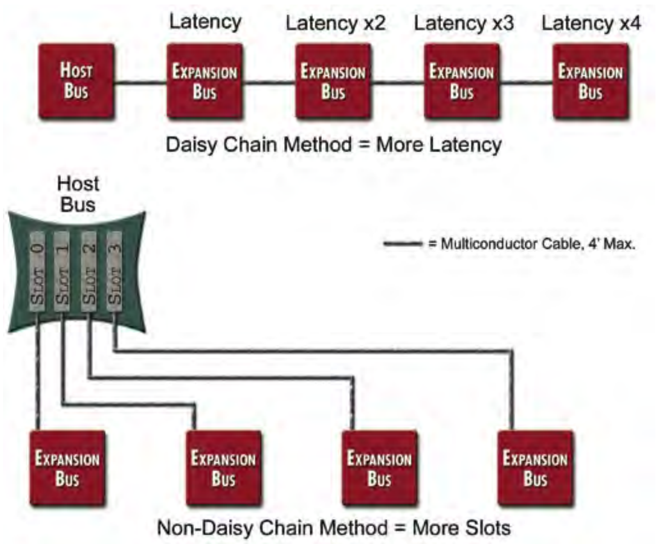

Figure 2

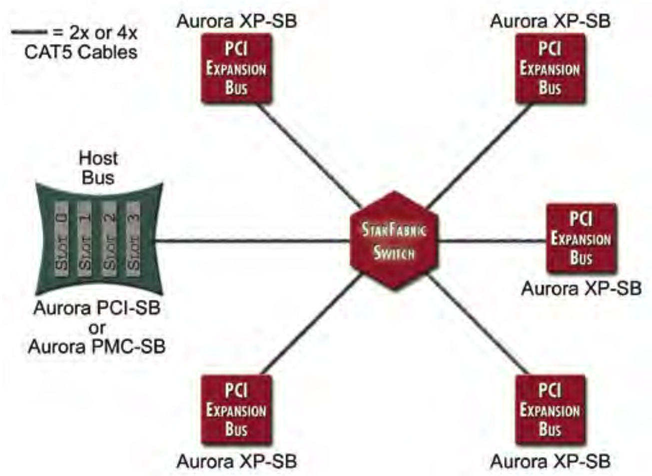

Figure 3

Figure 2 illustrates two options using parallel interconnects. The designer must choose between higher latency and slot usage. Compare these design limitations to the flexibility fabric implementation offers. Figure 3 depicts StarFabric's powerful switching feature. The StarFabric switch device forwards serialized PCI data to the endpoint designated by its PCI address, to make connecting up to six PCI expansion buses per switch a reality. Consider cable cost, complexity, and reliability. What if the application required physically isolating the system controller from the expansion systems or required testing several cards in a hostile environment? With cabling at a four-foot maximum, these requirements would be difficult, if not impossible, to reach. Interconnect simplicity helps when connecting systems with different form factors. The StarFabric interconnect typically uses two RJ-45 connectors for a 2.5-Gbit connection or four RJ-45s for a 5-Gbit connection. All Aurora PCI and CompactPCI StarFabric products utilize this interconnect.

The current trend toward 1U servers and motherboards with integrated video, USB, and other interfaces makes PCI expansion slots increasingly scarce. The need to expand CompactPCI systems to meet custom application requirements presents technical and financial challenges. For reliable, efficient, and economical PCI connections outside the box, StarFabric is the technology of choice.

Aurora Technologies offers a range of innovative StarFabric products. The PMC-SB StarFabric PCI Mezzanine Card (PMC-SB) features a PCI-to-StarFabric bridge in a versatile PMC form factor. This card can be used for CompactPCI/PCI bridging and StarFabric bridging on VME cards with a PMC site. In addition, the PMC-SB card is available in a PICMG 2.17-compliant version. The Aurora PCI-SB provides PCI/StarFabric bridging in a PCI short card form factor. Aurora's 6U CompactPCI CP-SFX8 fabric card provides StarFabric links to support eight 2.17 compliant node cards in a centralized topology. The card also features two front panel StarFabric links, thus allowing interconnection of multiple 2.17 chassis. A PMC site is available for use with a processor PMC card. PCI expansion products include XP2-SB, a PCI 2.2 compliant two slot 64-bit/66-MHz 1U chassis. XP3-SB is a 1U, three slot, 64-bit/33-MHz expansion chassis.

Frank Fitzgerald *is the director of engineering at Aurora Technologies, Inc. and the chief technology officer of the StarFabric Trade Association.*

For further information, contact Frank at:

Aurora Technologies, Inc.
10 Mupac Drive
Brockton, MA 02301
Tel: 508-588-6110
Fax: 508-588-0498
E-mail: frank@auroratech.com
Web site: www.auroratech.com