StarFabric: Adding flexibility to CompactPCI designs

By Ananth Jasty

CompactPCI Systems

StarFabric technology is one of the most popular switch fabric architectures today, largely due to its unique ability to seamlessly extend and enhance standard PCI, while retaining the investment in legacy systems. StarFabric's freedom from the restrictions of traditional parallel bus designs allows CompactPCI system designers the ability to develop more powerful, higher I/O systems with greater reliability and performance while avoiding expensive proprietary components.

The PCI bus has evolved from the simple I/O buses of the early x86 era. Unfortunately, with its emphasis on simplicity and low cost implementation, the flexibility and scalability available in high speed serial interconnects required by today's industrial computing systems has been lost (see Figure 1). StarFabric meets these needs by offering all the advantages of serial switched interconnects, while still keeping the overwhelming benefits provided by inexpensive, ubiquitous CompactPCI and PCI devices.

PICMG 2.17 and 3.3

Over the last year, StarFabric has begun to come into its own as a system interconnect. With the PICMG 2.17 specification approved and the PICMG 3.3 specification in development, StarFabric has answered the needs of industrial systems designers who are looking for the flexibility required in today's complex system designs. With the plethora of StarFabric, CompactPCI, PMC, and backplane designs introduced, developers and system integrators have the ability to build complex systems from the ground up, or by using larger functional building blocks. PICMG 2.17, which routes StarFabric topologies within a standard CompactPCI platform, is an extension to the PICMG 2.x family of specifications. The StarFabric interconnect provides a high-performance, reliable, and scalable point-to-point interconnect within a CompactPCI chassis. Adopters can choose to replace all or some of the existing bused slots on a CompactPCI backplane with StarFabric. PICMG 2.17 topologies can coexist with 64-bit PCI, H.110 (PICMG 2.5), and CompactPCI Packet Switching Backplane (PICMG 2.16) connections in the existing CompactPCI architecture.

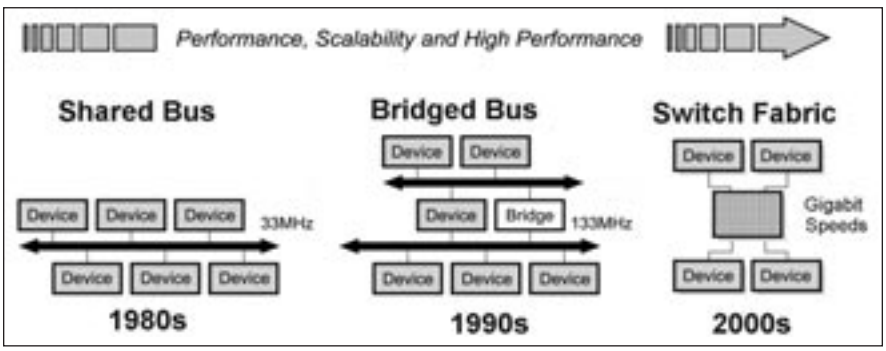

Figure 1

The PICMG 2.17 specification supports up to 21-slot backplanes, multiple enclosures forming large virtual backplanes, full backplane redundancy down to the slot level, chassis bandwidths up to 190 Gbits/sec (for a 21 slot chassis with 2 link slots), QoS/CoS support, hot-swap for fabric links and other advanced features. The StarFabric specification provides a solution for next-generation embedded equipment vendors who want to utilize open standard technology while adding features such as scalability, high availability, and quality-of-service. It balances the benefits of flexibility and the need to be compatible with existing PCI software, driver, OS, and hardware infrastructure so that adopters can maintain existing investments while solving next-generation challenges.

There is currently a set of AdvancedTCA – PICMG 3.x specifications approved and released within PICMG. The PICMG 3.x specifications define a new generation architecture for building high-end carrier grade equipment, which is complementary to PICMG 2.x family of CompactPCI specifications. PICMG 3.0 is a fabric agnostic general specification for mechanics, board dimensions, power distribution, connectors, system management, etc. Initial subsidiary specifications for fabric interfaces are:

- PICMG 3.1 (Gigabit Ethernet)
- PICMG 3.2 (InfiniBand)
- PICMG 3.3 (StarFabric)

The future of StarFabric

One recurring issue in current CompactPCI designs has been the geography of large systems. Applications such as medical imaging and industrial controls benefit from the extended distances StarFabric allows. Likewise, they benefit from the rack scalability and the freedom to use multiple rack configurations and form factors for the best use of valuable rack real estate. In addition, StarFabric's greater flexibility allows chassis to become interchangeable units. Many current designs use specialized I/O components to allow the port densities required. With StarFabric, inexpensive commodity parts can be used in bulk, lowering the system cost and allowing the use of more mature, volume-tested components. Also, there are still many specialty components only available in the PCI form factor, which can be easily integrated with StarFabric's factor agnostic design, while still benefiting from the hot swap and redundancy capabilities inherent to the fabric.

However, many of StarFabric's most advanced capabilities come from the integration of support and management software. Fully automated failover, which allows for the transparent replacement of failed components with on-hand spares, and

resource rebalancing, which adjusts the system load until human intervention is possible, are both desired applications in high reliability environments. Programmable QoS with both high priority and isochronous traffic patterns permit constant data throughput and allow for the functionality of real-time buses in the same bus framework, while high traffic devices can have their data rerouted to optimize the network for run-time conditions.

The added ability to integrate H.110 traffic into the StarFabric bus enables a single system interconnect handling all inter-device traffic on a single redundant, flexible data transport system without the design constraints of either PCI or H.110.

With StarFabric systems in active development, the PICMG 3.3 specification for AdvancedTCA factor in committee, and the projected PCI Express and Advanced Switching migration path outlined, the future of StarFabric looks to be as eventful as its present.

Ananth Jasty *is a software architect at Aurora Technologies, Inc., focusing on StarFabric development. Aurora Technologies, Inc. is a founding member of the StarFabric Trade Association.*

For more information, contact Ananth at:

Ananth Jasty
Aurora Technologies
10 Mupac Drive, Brockton, MA 02301
Tel: 508-588-6110
E-mail: ananth@auroratech.com • Web site: www.auroratech.com

StarFabric switching

A fabric card for StarFabric switching

StarFabric is largely targeted towards CompactPCI applications, owing its robust CompactPCI support and elegant CompactPCI backplane and switch implementations to its attention to industrial-class systems (see Figure 1). The PICMG 2.17 backplane-switch combination in a centralized topology allows for simpler node cards at the cost of a single slot. With Aurora's CP-SFX8 Fabric Card, the switch comes with increased bus throughput, a low latency design, and an onboard PMC site. The addition of a second CP-SFX8 Fabric Card allows for a fully redundant StarFabric network. In legacy mode, the system would not even require software changes as the fabric would appear as PCI-to-PCI bridges, while offering far more flexibility in implementation.

Figure 1

Designers utilizing the CP-SFX8 Fabric Card can get the advantages of a high-bandwidth, cost-effective solution for switching StarFabric in high-availability platforms – for greater power and flexibility. The card allows users to implement a centralized fabric topology in a CompactPCI chassis. Available in a 6U design, the CP-SFX8 features eight links to a PICMG 2.17-compliant CompactPCI backplane and supports up to eight node cards within the chassis (see Figure 2). The card also features the additional ability to connect to a second chassis for fully redundant fault tolerance, with two front panel links. Its internal switch devices interconnect so data travels no more than two hops for low latency. Using StarFabric, the card provides a port bandwidth of 5 Gbits/sec and an aggregate bandwidth of 40 Gbits/sec. Some typical applications for Aurora's StarFabric-based products include:

- Automated test equipment
- Clustering applications
- Communications (media gateways, blade servers)
- Industrial control/automation
- Medical imaging
- Military applications (radar, imaging)
- PCI expansion
- Storage (NAS, SAN)
- Video/image processing

As shown here, the CP-SFX8 Fabric Card lets you deploy a centralized StarFabric topology for PICMG 2.17 compliant CompactPCI chassis. It supports up to eight node cards in a single chassis and provides two links on the front panel for connectivity to external StarFabric devices.

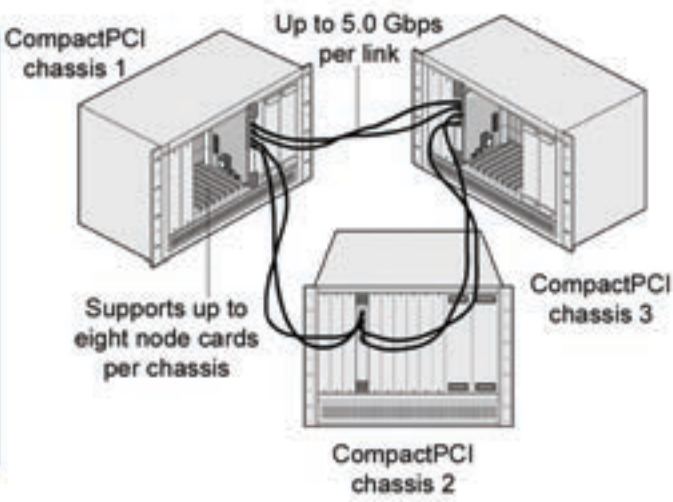

Figure 2