

VXS: Switched Serial Gigabit Fabric for VMEbus

The VMEbus, well into a third decade of widespread deployment, continues as the dominant bus structure for high-performance embedded systems. In an industry characterized by a steady succession of new device offerings with speed and density increases every few months, VMEbus has retained this leadership position not simply because it was based on a sound electrical and mechanical architecture. Indeed, the major reason for its longevity has been a series of performance and feature enhancements promoted and nurtured by a broad base of VMEbus vendors. Before we discuss VXS, the latest enhancement, we will first review some basics of switched fabrics.

Switched Fabrics

A switched fabric is a system for connecting devices together to support multiple simultaneous data transfers usually implemented with a crossbar switch, as shown in Figure 1. Data is sent in packets with information contained in the packet header for identification, routing and error detection and correction. To ensure adequate performance for any given system, the interconnecting fabric can be as simple as a point-to-point connection between two devices, or a more complicated architecture that may include switches, routers, hubs, and repeaters.




Figure 1. Switched Fabric Concept.

Because of recent advances in serial data technology, the new generation of switched fabrics uses serial links. With bit data rates now in the gigahertz range, these new serial interfaces can easily rival their parallel counterparts. In many cases the transition from parallel to serial occurs only at the lowest levels of the OSI layer model. In this way, existing protocols are maintained so legacy products with parallel interfaces can be supported with hardware adapters that convert the physical layer interface to the new serial link.

This strategy has been extremely successful in allowing the new serial technology to be inserted seamlessly. One excellent example is the migration of parallel flat cable SCSI to serial Fibre Channel as the interface of choice for the latest generation of high-performance hard disks and disk arrays.

One of the major benefits of these new serial interfaces is the reduced number of signal lines and smaller connectors and cable. This results in enhanced system density, simpler system integration, lower installation costs, and easier maintenance. Another benefit is the ability to use copper cable for low cost local connections or optical

cable for fast, long haul transmission of data. Again, the physical layer can be made completely transparent to the protocol layer.

Yet another benefit of serial links is the ability to gang together multiple serial links to boost data throughput. Since the signal in each single bit link contains embedded clock and timing information, each link can propagate on its own across the channel and transceivers at each end can handle the multiplexing and demultiplexing for 1x, 4x, 8x or 16x ganging at a relatively low hardware level.

Once the benefits of switched serial fabrics became apparent, embedded systems vendors sought ways to take advantage of this technology for a wide range of interconnection needs: boards to peripherals, boards to boards, chassis to chassis, and facility to facility. Not only are switched serial fabrics attractive alternatives for front panel interconnections, they are also extremely appropriate for backplane data traffic to augment or replace the conventional parallel backplane bus.

Popular Switched Serial Fabrics

Embedded system vendors are faced with these switched serial fabrics:

Infiniband is primarily aimed at server and storage system connectivity for box-to-box links.

StarFabric provides transparent serial links between PCI devices.

PCI Express is Intel's initiative for connecting processors and boards in personal computers and workstations.

Hyper Transport is AMD's solution for chip-to-chip and board-to-board connections in personal computers.

RapidIO is targeted for chip-to-chip and board-to-board connections for real-time COTS embedded systems.

These fabrics are all vying for position. Aside from some valid technical pros and cons for each fabric, the key issues for switched fabrics tend to be business issues. For example: which major vendors are backing each standard? How easily can these new fabrics be integrated into existing software operating system environments? What components are available for bridging to existing hardware and processors? What kind of switches are available? And finally, can the fabric achieve a high production volume to make the parts inexpensive and easily connected?

Serial RapidIO

One of the switched fabrics, Serial RapidIO, is especially well suited for real-time

VXS promises to solve even the most challenging data transfer requirements of real-time embedded systems. It will assure the continuing popularity of VMEbus by allowing legacy boards to coexist with the new VXS versions.




Figure 2. VXS Payload Card


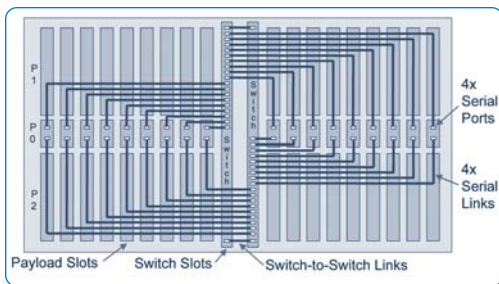


Figure 3. VXS Switch Card

Figure 4. VXS Backplane with 20 slots



embedded systems. RapidIO is promoted by the RapidIO Trade Association whose founding members include Motorola, Lucent, Mercury, IBM, Texas Instruments, et al.

Objectives of this high-performance, packet-switched, interconnect standard include fast interprocessor communication, networking of DSPs, high-speed backplane interconnect, and efficient chip-to-chip and board-to-board transfers. These speed goals are addressed with scalable serial bit rates of up to 10 GHz. Performance and efficiency are achieved through a combination of a low-overhead protocol plus hardware error detection and correction. By offloading these tasks from the processor, Serial RapidIO is well suited for real-time applications, where shared coherent memory, channel predictability and low-latency are essential.

At the physical layer, Serial RapidIO uses the same differential current mode signaling as other standards including Fibre Channel, Infiniband and 802.3 XAUI.

VXS

During the last few years, the VITA 41 committee of the VMEbus Standards Organization has been defining a switched serial backplane fabric for VMEbus called VXS. The specification defines a VXS Payload Card, a VXS Switch Card and a connector scheme for a backplane family to support VXS.

Because of the “fabric wars”, the VXS specification was defined to be fabric agnostic: there are five sub-specifications, one for each of the five fabrics described previously. The switched fabric architecture chosen to connect the boards across the backplane is a ganged 4x, full-duplex serial channel, so that each interconnect supports data flow in both directions simultaneously.

Serial bit rates are defined for frequencies up to maximum of 10 gigabits/sec, although lower frequencies are supported for the first systems. With the 4x ganging and a nominal bit frequency of 2.5 GHz, the input path and output path are both capable of moving data at 1 Gbyte/sec.

VXS Payload Card

The VXS Payload cards are processor, CPU, memory, and data converter 6U VMEbus cards with the VXS interface added. They have standard P1 and P2 connectors that implement the standard VME64x backplane interface. A new P0 backplane connector mounted between P1 and P2 handles

two 4x, full duplex switched serial ports. Figure 2 shows a simplified view of the Payload card and the new VXS P0 connector.

At a 2.5 GHz clock frequency, each VXS Payload card can move data in and out at an aggregate rate of 4 GBytes/sec, two orders of magnitude above the original VMEbus backplane specification!

VXS Switch Card

Unlike the Payload card, the VXS Switch card has a 6U VME board form factor but no P1 and P2 connectors. Instead, it has up to eighteen 4x full-duplex serial connectors and a power connector. The VXS Switch card implements the crossbar switching to connect Payload Cards together. Figure 3 shows a view of the VXS Switch card.

VXS Switch cards can have any number of crossbar switches and any number of serial ports. They may also include other interfaces to networks for communication and storage devices, as well as front panel serial ports to other VXS Switch cards in the same chassis or in adjacent racks. Optical serial ports could be used for remote high-speed data transfers.

VXS Backplane

The VXS backplane can take on many different layouts to accommodate specialized system needs, but will normally handle two to twenty Payload cards and one or more Switch cards. The standard board-to-board pitch of 0.8 inches is maintained throughout, and other VMEbus card cage mechanical hardware is compatible. The objective is to connect the two 4x serial links of each Payload card to links on the Switch card(s) to support the necessary board-to-board connectivity. Some smaller systems may require only a few Payload slots and a very simple Switch card, while others may need to use a full width backplane and multiple Switch cards to handle the required traffic.

Figure 4 shows one example of a 20-slot VXS backplane that holds 18 Payload cards divided equally in each half, and two Switch cards occupying the two center positions. One serial link from each Payload card is wired to one of the Switch cards while the second is wired to the other Switch card.

With a maximum of 18 serial link connections on each Switch card, all 18 Payload cards can be connected to each other through two redundant paths, namely through both of the two Switch cards. This dual redundancy is attractive for many applications requiring fault tolerance and high availability. VXS Switch cards also have additional serial links that join Switch cards together, providing yet another path for routing. □