VXS VITA 41 melds fabrics and FPGAs: Achieving tomorrow's backplane performance today

By Andrew Reddig

Incremental changes to VME occur routinely, though most often they are designed to add more data movement capability, either on the VMEbus itself or over dedicated backplane connections. The new VITA 41 VXS standard defines a switch architecture that increases – by an order of magnitude – the available bandwidth to the backplane. VITA 41 also is fabric-agnostic, allowing several fabric "flavors" to be used to fit the I/O situation.

Editor's note: *For more information about VITA products in this special VITA Standards magazine issue, please refer to the Product Guide on page 40.*

The venerable VME backplane standard seems to have nearly as many lives as the luckiest of cats. Designed into numerous military processing systems, a large part of its success and longevity is attributable to the way the standard has been developed through several generations of I/O bandwidth increase while maintaining backwards compatibility with previous generations of VME hardware. This is a crucial factor in military applications where systems can be in service for many years and incremental system requirements changes are the norm in an era of continuous spiral refresh.

VITA 41, known as *VME Switched Serial* or *VXS*, has provided the latest injection of adrenalin into the VME standard by adding support to the backplane for a high-speed serial, switched fabric interconnect capability. VXS increases the maximum achievable I/O bandwidth between cards by *more than an order of magnitude.* It also provides support for legacy VME cards, protecting the investment made in existing systems while providing a smooth upgrade path to address increased system demands on I/O and processing capacity.

VXS makes possible many advanced architectures. By building systems based on VXS switch cards, it is possible to create highly interconnected card slot meshes that increase the I/O bandwidth between processing cards by a greater order of magnitude. Such architectures allow systems to be built that can offer performance levels comparable to those promised both by developing and emerging standards such as VITA 46 (VPX) and AdvancedTCA. In addition, VXS supercharges high-performance, nonfabric architectures that rely on efficient, low-level protocols implemented strictly in FPGAs.

VXS basics

The VXS standard defines two types of cards: *payload* cards, which are essentially VME cards with an enhanced P0 connector, and *switch* cards, which use high-speed connectors to provide the interconnect between multiple payload cards. Each payload card supports up to 2.5 GBps full duplex bandwidth, increasing the potential throughput by an order of magnitude over previous generation switched fabrics such as RACE++. Each switch card supports up to 18 payload slots along with interswitch connections, providing up to 27.5 GBps full duplex bandwidth. Examples of VXS payload and switch cards are shown, respectively, in Figures 1 and 2. Figure 1 is a QinetiQ Neptune payload board, whereas Figure 2 is a QinetiQ Callisto switch card.

Because a VXS payload card uses the traditional VME P1 and P2 connectors, a VXS system can directly support legacy VME cards in VXS payload slots. This protects the investment made in existing systems while providing a smooth upgrade path to address increased system demands on I/O and processing capacity. Legacy cards can be VME64x cards or can include P2-based fabrics such as RACE++ or StarFabric.

Figure 1

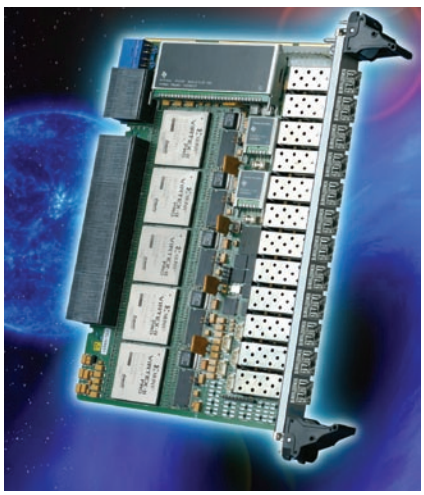

Figure 2

A key feature of the VITA 41 architecture is the layered approach that the standard uses to define payload cards, switch cards, and backplanes. The base standard, VITA 41.0, defines the mechanical, connector, power, and system management interfaces but does not define a specific switched fabric interconnect. Each possible switch fabric is defined in a protocol standard that maps the switched fabric protocol into the high-speed differential pair connections laid out in VITA 41.0. There are VITA 41 protocol standards defined for InfiniBand, Serial RapidIO, Gigabit Ethernet, PCI Express, and Advanced Switching.

The combination of the VXS layered architecture and the common physical and electrical interface used by today's high-speed serial switched fabrics allows board-level suppliers to create "fabric-agnostic" building blocks. Many of the VXS products on the market today are based on this approach, using FPGA technology to implement products that can be adapted to whichever switched fabric the user requires through the selection of specific IP cores. This approach allows users to select one hardware platform and use it in multiple applications, selecting the appropriate FPGA IP for the fabric in each system.

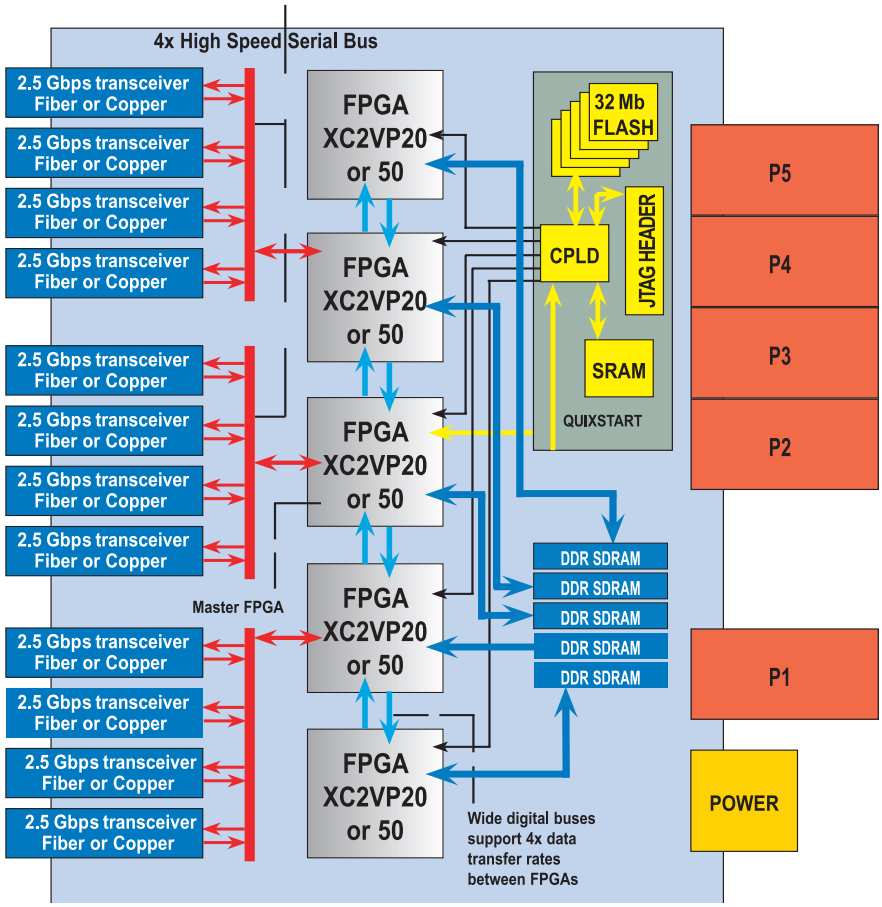

Figure 3

VXS topologies

The VXS base standard document describes one specific example of a VXS system using a combination of payload and switch cards to implement a dual-star fabric topology. The example has up to 18 payload slots and 2 switch slots, using the switch slots to provide the interconnection between the payload slots. The use of two switch slots supports redundant data paths between all cards, supporting high availability and failure recovery in many applications. A block diagram of the dual-star topology is shown in Figure 3.

Although the dual-star topology is the only architecture shown as an example in the VXS standard document, recent development effort at Tekmicro and QinetiQ has focused on using VXS components with other topologies. By implementing alternate topologies and backplane solutions, VXS payload and switch modules can be used as building blocks to create solutions optimized for a wide range of systems. Small-scale systems can be implemented using payload-only architectures without dedicated switch cards. This connects payload cards together in mesh or ring topologies to meet aggressive size, weight, and power constraints such as those needed for UAV implementations. These implementations are supported today with three- and five-slot backplanes that implement payload-only systems.

At the other end of the spectrum, large-scale systems can be implemented using payload cards for the I/O portion of a system and a dense mesh of switch cards for processing. This architecture combines traditional VXS payload cards with switch/processor cards that implement extremely high throughput to the backplane, along with integrated signal processing capability. This architecture requires a new class of switch card, combining the interconnect bandwidth of a "traditional" VXS switch card with resources for embedded signal processing. An FPGA-based switch card is a good fit for this architecture, with each FPGA node combining interconnect and processing resources in a single device.

An example of a backplane for this architecture has been developed in collaboration with ELMA Bustronic and is shown in Figure 4. This hybrid backplane implements two VME64x slots, three VME64x/VXS payload slots, and six VXS switch slots. Each switch slot implements 20 x4 links for a total of 25 GBps per switch slot. The system architecture supports up to 7.5 GBps of throughput between the I/O front end and the processing mesh and a total of 112.5 GBps of aggregate throughput within the processing mesh itself. By comparison, a VITA 46.3 Serial RapidIO mesh slot supports 4 x4 links for a total of 5 GBps per slot, and an AdvancedTCA full mesh slot supports 15 x4 links for a total of 18.75 GBps per slot. Data-intensive applications such as Synthetic Aperture Radar (SAR) can utilize the additional bandwidth provided by this architecture to scale to a higher number of channels, higher bandwidth, or both.

When used with an FPGA processor such as QinetiQ's Callisto switch card (shown previously in Figure 2), the switch mesh allows the user to implement a dense interconnect of multiple FPGA processing nodes with extremely high aggregate bandwidth. The resulting Callisto FPGA

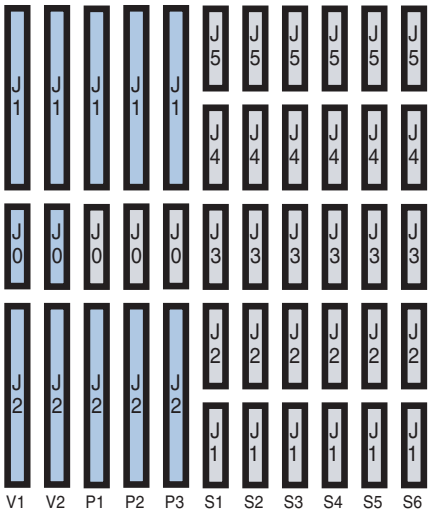

Figure 4

processing mesh interconnect topology is shown in Figure 5, which uses five switch/processor cards to implement a 25 FPGA processing cluster. Each card also has 12 fiber optic transceivers at the front panel, providing a total of 15 GBps bandwidth to other processing systems. By using the front panel I/O capability, the system can scale up to *multiple* VXS chassis for even higher embedded processing requirements. Alternatively, the front panel I/O can be used to provide open standard connections to workstation-based back-end processing, leveraging cost-effective Intel-based servers for non-real-time processing, display, and analysis functions.

Switched fabric or point-to-point?

Each FPGA processing node in the cluster shown in Figure 5 has four point-to-point connections to either VXS payload cards or to other FPGA nodes. While these links can be implemented using switched fabric FPGA IP cores, the resulting FPGA designs use a large fraction of the available FPGA logic just for communication. Unless a switched fabric is required, the FPGA resources can be more effectively utilized for signal processing functions instead of a Serial RapidIO or PCI Express core.

When a mesh topology is being used, the use of a switched fabric interconnect is often overkill for the application because most of the data flows will be point-to-point by design. In many deployed applications supported by Tekmicro and QinetiQ, FPGA-based end points have been used to create low-level, point-to-point interconnects between VXS cards, typically based on Xilinx's Aurora protocol in lieu of a full switched fabric. This configuration provides high throughput between nodes while reserving most of the FPGA's logic and memory resources for processing functions.

Because of the increasing use of point-to-point links in the systems we are building, Tekmicro and QinetiQ have recently co-sponsored a new standard called Virtual Streaming Protocol (VSP), also designated *VITA 55* by the VME International Trade Association (VITA). VSP implements a low overhead point-to-point link between FPGA-based end points and can be used on a variety of physical layers. One implementation is being defined as *VITA 41.5* and provides point-to-point interconnect between VXS payload and switch cards using Xilinx's Aurora protocol.

Another implementation will be defined to support parallel bus links between FPGA devices on a board. The range of physical implementations will allow the VSP protocol to be utilized throughout the type of FPGA processing cluster shown in Figure 5, providing a common method of interconnect for the user – whether the interface is to an adjacent FPGA on the same card or to an FPGA on another switch card or even in another chassis.

Another advantage of an FPGA-based protocol implementation is that the user has a lot of flexibility in the selection of interconnect implementation within the system. If the application required additional off-the-shelf cards that used

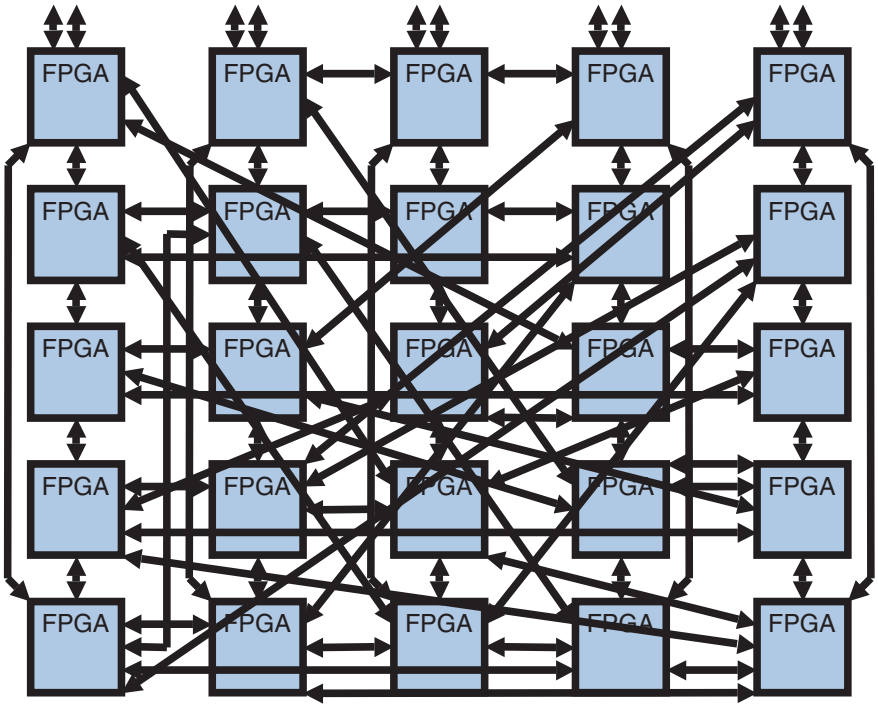

Figure 5

APPLICATIONS

VXS applications: Fabric or FPGA

The combination of dense FPGA processing nodes with a mesh fabric provides a platform for a wide range of high-performance embedded computing applications, particularly when combined with the scalability of a VXS-based I/O front end. Applications where this architecture is proving particularly useful include:

- **Adaptive beamforming**
- **Synthetic Aperture Radar (SAR)**
- **Image processing**
- **Target recognition**
- **Multispectral data fusion**

All of these applications require simultaneous parallel processing of huge volumes of data.

The advantages of a VXS architecture can also be used in SIGINT applications to dramatically increase either the input bandwidth, the number of simultaneous Digital Down Converters (DDCs), or both. By using parallel FPGA processors to apply 5-10x the processing resources to the problem, the user can select appropriate FPGA IP cores to increase the sample rate, filter size, channelizer count, or a combination thereof. The bandwidth between the FPGA nodes provides the flexibility necessary to easily share data between multiple FPGA processing nodes without fabric congestion.

switched fabrics such as RapidIO or PCI Express, the switch card could implement VSP end points for some links and a switched fabric end point for others. This allows the user to mix and match the building blocks they need for their application, potentially combining high-density FPGA processing with PowerPC or Intel general-purpose processors at the back end.

System viewpoint

The VXS switch/processor mesh architecture clearly demonstrates both the longevity and expandability of the VME platform for high-performance embedded computing. In one system, a user could be simultaneously using four stages of interconnect technology evolution.

- VME64x at 60-80 MBps across the whole chassis
- Legacy RACE++ slots with 533 MBps throughput for each slot
- VXS payload cards with 10x higher throughput at 2.5 GBps full duplex per slot
- VXS switch slots with 10x higher still at 25 GBps per slot

With such a wide range of performance in an evolutionary architecture with backwards compatibility, VXS offers the user a unique combination of leading-edge performance without compromising existing hardware and software investments or the ability to deploy incremental upgrades.

Andrew Reddig *is President/CTO and cofounder of TEK Microsystems, Inc., which was founded in the 1980s while he was a student at MIT. A 25-year embedded industry veteran and active in industry standards initiatives, Andrew is sponsor/cochair for VITA 41.5, 42.0, and 55. He also holds two degrees from MIT: a BS in Computer Science and a BS in Electrical Engineering.*

To learn more, contact Andrew at:

TEK Microsystems Inc.
Two Elizabeth Drive
Chelmsford, MA 01824
Tel: 978-244-9200
E-mail: andyr@tekmicro.com
Website: www.tekmicro.com