VITA 41 (VXS) – Evolutionary open standard with scalable bandwidth

By Andrew Reddig and Steve Birch

The VMEbus architecture has become one of the most successful interconnect standards in the history of embedded computing. A key element of that success has been the VMEbus mandate for backwards compatibility. Each new addition to the VMEbus roadmap has maintained mechanical, electrical, and software compatibility with all pre-existing VMEbus boards, whether it was a protocol enhancement such as 2eSST, or a new extension such as StarFabric.

Although the processing throughput and functional density provided by today's 6U VME boards could hardly have been imagined by the original VMEbus architects, all of today's boards are still compatible with VMEbus boards designed 20 years ago. As VMEbus is enhanced to support the next generation of interconnect technology, the VMEbus community's commitment to backwards compatibility will continue to be critical to VMEbus system integrators and end users as they evaluate architectures for new development projects.

Serial interconnects

The next generation of scalable embedded computing solutions will be based on open interconnect standards such as:

- Infiniband
- PCI Express
- Serial RapidIO
- 1 and 10 Gigabit Ethernet

The most noticeable change from legacy bus standards is the transition from parallel busses to one or more high-speed serial links. Each serial link uses an encoding scheme that combines its own clock and data into a single serial bitstream, making it unnecessary for board and system designers to maintain low-skew signal paths between devices. This simplifies board layout and also allows serial links to be used both inside and outside the box, providing a unified architecture for board-to-board and box-to-box communication.

The encoding of the clock and data into the serial bitstream allows much higher bit rates per pin, with current fabrics supporting data rates up to 3.125 Gbps. After allowing for 8B/10B encoding overhead, this provides an effective data rate of 312.5 MBps in a single differential pair, compared with 320 MBps using 70+ pins for VMEbus 2eSST. Future protocols are expected to use more efficient 64/66b encoding, raising the effective data rate to 378 MBps using the same 3.125 Gbps link.

Each of these standards has created its own ecosystem, typically including IP solutions for endpoints along with merchant silicon products that implement endpoints and switches. As with the adoption of PCI as a local bus technology several years ago, embedded solution providers will take advantage of the available infrastructure options when selecting serial protocols.

VITA 41 switched serial standard

The VITA 41 (or VXS) standard enhances the core VMEbus standard with the addition of a new high-speed interconnect capability to the existing VMEbus interface. VXS defines two types of boards, payload boards and switch boards.

Payload boards are very similar to existing VMEbus boards. They have the traditional P1 and P2 connectors and support a standard VMEbus interface, but add a new higher-speed P0 connector with eight full-duplex serial links that provide up to 2.5 GBps of throughput in each direction. Payload boards retain complete mechanical and electrical compatibility with legacy VME boards – any VMEbus board that uses the traditional P1 and P2 connectors can be installed in a VXS payload slot.

Current VXS technology is based on serial fabrics using 2.5 and 3.125 Gbps signal rates, but the new P0 connector is designed to support fabrics up to 6 Gbps. As serial fabrics are upgraded to higher bit rates, VXS technology will also move forward to support higher system throughput within the same basic architecture.

Small VXS systems can be built using payload boards and *distributed switching* without dedicated switch boards. In this architecture, the backplane provides passive links between payload boards without a switch board in the system. VXS compatible backplanes can support ring or mesh topologies with the data rates shown in Table 1.

The throughputs shown in Table 1 assume that all links operate at 3.125 Gbps and are fully utilized in both directions. PCI Express and InfiniBand links operate at 2.5 Gbps and will therefore have throughputs 20 percent lower than the values shown.

VXS switch boards, unlike payload boards, replace the traditional P1 and P2 connectors with higher density connectors and provide only high-speed serial interconnections. Switch boards do not support direct connection to the VMEbus, although payload boards will typically provide a fabric-to-VME bridge capability that would allow a switch board to access VMEbus resources through one of the payload boards.

Reprinted from *VMEbus Systems* / February 2005

Number of Boards	Mesh	Ring
Three	1.25 GBps per x4 path 6 paths per system 7.5 GBps total	1.25 GBps per x4 path 6 paths per system 7.5 GBps total
Five	625 MBps per x2 path 20 paths per system 12.5 GBps total	1.25 GBps per x4 path 10 paths per system 12.5 GBps total
Nine	312.5 MBps per x1 path 72 paths per system 22.5 GBps total	1.25 GBps per x4 path 18 paths per system 22.5 GBps total

Table 1

Each VXS switch board can support up to 18 payload boards, with four full-duplex serial links between each payload board and one of two redundant switch boards. A fully configured VXS system supports two switch boards and 18 payload boards for a total aggregate throughput of up to 45 GBps, again assuming that all links are fully utilized in both directions.

One advantage of the VXS architecture is that the backplanes are completely passive, providing high-speed interconnect between boards without any active components on the backplane itself. This implies that the backplane *is not aware* of the actual fabric being implemented, so the same backplane can be used for PCI Express, Serial RapidIO, or InfiniBand without modification. It also offers an advantage over previous generation switched fabric interconnects such as RACE++ and SKYchannel which required active interlink modules mounted behind the backplane to provide the necessary crossbar switches between boards.

VXS payload and switch boards can be designed to support a specific fabric, or they can be designed as *fabric agnostic.* Because the switched fabrics supported by VXS for the most part share common electrical characteristics, it is possible to design VXS boards that use FPGA technology and soft IP cores to support multiple serial protocols using the same basic hardware design. In some applications, the flexibility inherent in FPGA-based architectures will offset the higher power and cost when compared to hardwired solutions based on merchant silicon.

VXS Ecosystem

The number of companies supporting VXS continues to grow. As of the date of this article, several companies have either announced or committed to products as shown in Table 2.

Product	Company
VXS Payload Boards	TEK Microsystems (I/O, FPGA, multi-PowerPC) QinetiQ (I/O, FPGA) Motorola (Single Board Computer) Pentek (I/O, FPGA) SBS (Single Board Computer) VMETRO / Transtech DSP (multi-PowerPC with FPGA)
VXS Switch Boards	QinetiQ (FPGA processor) SBS (InfiniBand switch)
VXS Backplanes	Bustronic Hybricon

Table 2

System Example

The system example (Figure 1) illustrates the use of several vendor's products to create a high performance data recording system. This system example consists of the following components:

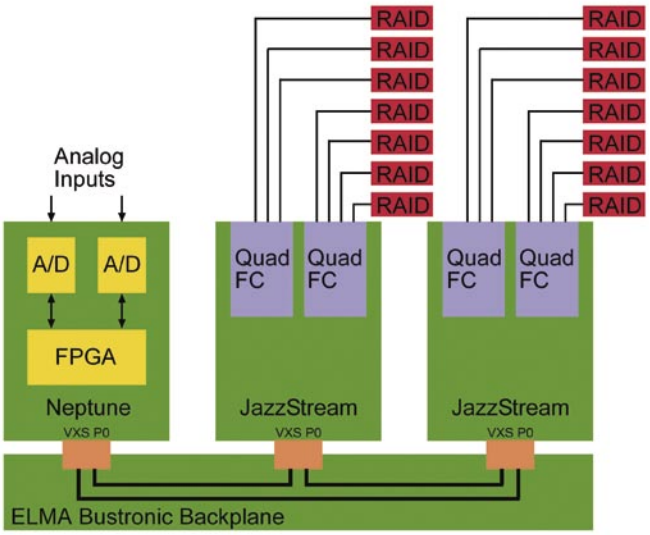

Figure 1

- A QinetiQ Neptune dual-channel high-speed digitizer
- Up to two Tekmicro JazzStream controller boards with quad Fibre Channel XMC modules

The VXS payload boards are interconnected using a Elma Bustronic VXS backplane providing three VXS payload slots and two traditional VME64x slots in a ring topology.

Digitizer Board

The QinetiQ Neptune digitizer board is one member of a family of VXS fabric only payload boards that combine high-speed analog and/or digital sensor I/O functionality with an FPGA processing and communications backend. The Neptune digitizer provides a pair of 10-bit analog to digital converters (ADC) with sample rates up to 2 Gsps (Figure 2).

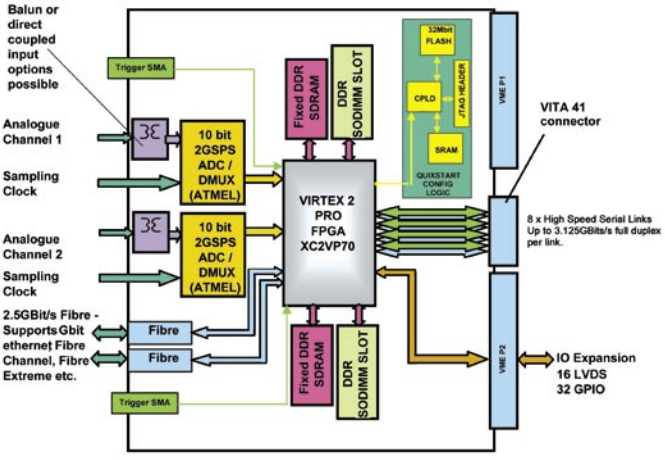

Figure 2

A Xilinx Virtex-II Pro P70 device at the heart of the board provides the following functionality:

- Interfacing of the on-board memory and ADC resources
- Reconfigurable processing resource for intensive front end DSP
- Flexible high-speed interconnect to VXS backplane

With suitable FPGA cores, it is feasible to process the very high bandwidth of digital data generated by the ADCs to achieve some initial real time data processing and bandwidth reduction prior to sending the data over the VXS backplane.

The FPGA also provides the physical support required to implement the physical layer of a number of high-speed protocols that can be implemented over the VXS backplane. FPGA cores that implement the higher protocol layers of a range of switched fabric standards allow the Neptune board to be easily reconfigured to operate with a variety of fabrics within a VXS environment.

The FPGA has sufficient spare resources to provide two additional high-speed serial links via front panel Small Form Pluggable (SFP) modules. These modules provide physical support for connection to equipment via protocols such as Gigabit Ethernet, Fibre Channel, and Serial FPDP.

Storage controller board

A similar FPGA-based approach is used for the Tekmicro JazzStream VP I/O board (Figure 3). This board uses a PLX Technology PEX 8532 switch to provide an onboard PCI Express fabric that interconnects two PowerPC processors, two PMC/XMC sites, and the VXS P0 backplane interface.

The gateway FPGA is located between the PLX switch and the VXS P0 interface. In some applications, this FPGA is configured to simply forward PCI Express traffic between the onboard and offboard serial links, creating a homogeneous PCI Express based interconnect throughout the system.

However, the use of an FPGA bridge allows the board to support interoperability with other fabric protocols either on other payload boards or on a switch board. The FPGA-based bridge may be used to hide the internal PCI Express fabric behind an endpoint, allowing the board to participate with other vendors' products in a heterogeneous system.

The software environment supplied with the board transparently manages access to onboard and offboard resources for any supported fabric, allowing application developers to manage interprocessor communications and data movement without needing the details of the underlying fabric architecture.

The gateway FPGA also provides access to bulk FLASH and DDR SDRAM memory resources. The FLASH memory is used for program and FPGA bitstream storage, and the DDR memory is used to buffer high-speed data before it is forwarded to the disk.

The disk interface itself is implemented using a Tekmicro JazzFiber XMC module with four Fibre Channel ports. This module uses 2 Gbps RAID technology to support sustained throughputs up to 760 MBps to four RAID units with real-time file system support. This allows the controller board to support an aggregate data rate of 1.52 GBps to an array of eight RAID units. Future Fibre Channel storage devices will support 4.25 Gbps links, raising the aggregate data rate to the full 2.5 GBps supported by the VXS interconnect.

Throughput

With one storage controller board and a Bustronic 3-slot ring backplane, the system uses two board slots to support continuous recording at 1.25 GBps to 7 RAID units. By adding a second controller board, the system can support 2.5 GBps to 14 RAID

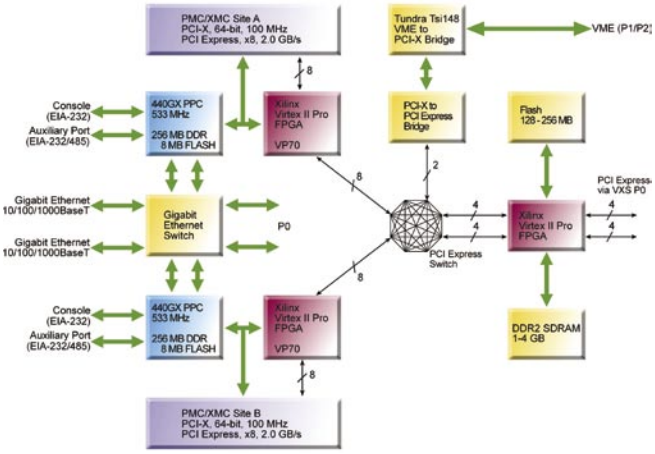

Figure 3

units while using only three board slots. Using 3U dual-channel RAID units, a complete data acquisition and storage subsystem with 22.4 TB of storage (2.5 hours at 2.5 GBps) occupies only 24U of rack space.

These data rates are well beyond the capabilities of standard bus-based architectures and require a system with switched fabric interconnect. While these data rates can be implemented using existing interconnects such as RACE++, a comparable system would need 5 to 6 board slots for the storage interfaces alone, and it would have great difficulty supporting a single 2 Gsps channel due to the board slot limitation of 533 MBps.

Conclusion

VXS offers system integrators an evolutionary advance with revolutionary capabilities. VXS systems can support up to five times the bandwidth to a single 6U slot as the fastest existing switched fabric solutions, while maintaining complete backwards compatibility with the full range of existing VME products.

Within the VXS architecture, users can build systems ranging from 2 to 20 boards, with distributed payload-only topologies in smaller systems, or dedicated switch boards for higher total bandwidth. Minimal systems support throughputs of up to 5.0 GBps, or more than 15 times the capability of 2eSST VMEbus. Fully configured 20 slot systems scale up to 45 GBps, or more than 140 times the capability of 2eSST VMEbus.

Just as VME has supported technology insertions over the last 20 years, the VXS architecture will continue the same philosophy, allowing vendors to develop compatible products and offer users a wide range of choices with evolutionary improvements as technology advances.

The good news is that VME systems integrators can start taking advantage of off-the-shelf VXS solutions today, maintaining complete compatibility with legacy VMEbus boards. In addition, they can accomplish this while they incrementally add throughput and performance, and still maintain a clear path to future technology enhancements. Ω

Andrew Reddig *is President/CTO and cofounder of TEK Microsystems, Inc. TEK was founded in the 1980s as an engineering services organization while Andrew was a student at M.I.T. Today, TEK Microsystems is a leading supplier of advanced streaming I/O products.*

For more information, contact Andrew at:

TEK Microsystems, Inc.
2 Elizabeth Drive
Chelmsford, MA 01824
Tel: 978-244-9200
Fax: 978-244-1078
E-mail: andyr@tekmicro.com
Website: www.tekmicro.com

Steve Birch *is Technical Manager of the Real Time Embedded Systems Group of QinetiQ. He was a founding member of the QinetiQ Real Time Embedded Systems group in 1999, and has been responsible for defining and bringing to market the FPGA based processing solutions now offered by the group. Steve holds a Bachelors Degree in electronic engineering from Imperial College, London University.*

For more information, contact Steve at:

QinetiQ Ltd.
Malvern Technology Centre
Malvern
Worcestershire, UK WR14 3PS
Tel: +44 (0) 1684-89-4000
Fax: +44 (0) 1684-54-3762
E-mail: swbirch@qinetiq.com
Website: www.qinetiq.com