

Implementing FPGA clusters and mesh architectures



By Bob Walsh

VME is the preferred choice for many deployed systems such as those in high-rel military applications. And signal processing experts agree: FPGAs, with their inherent parallel processing architectures and high-speed off-chip serial interconnects, make ideal DSP engines. Coupled into mesh architectures and installed on VME boards using the VME Switched Serial VXS standard, FPGAs can rival standalone DSPs and dedicated cluster computers.

As processing capabilities continue to expand, signal processing systems use larger amounts of sensor data – in resolution, bandwidth, and number of channels – to perform their functions. As these systems are being developed and evaluated, architectures are required that can keep up with the signal processors' capabilities. The tools and technologies that enable faster signal processing systems, including FPGA-based processing, can be used for significant processing work in clusters of tightly connected FPGAs. This approach can have significant size, weight, power, and cost advantages over general purpose processors.

VITA 41 VXS

The VITA 41 VXS standard offers useful building blocks for FPGA clusters by defining two types of boards, both of which are very similar to standard VME64 6U by 160 mm boards. The first type, called a *payload* board, is identical to a VME64 board except that it has a special P0 connector mounted between the traditional VME P1 and P2 connectors. This connector supports eight differential serial ports, each running at 3.125 Gbps. Typically, these are connected in two sets of four, so-called *4x lanes*. These 4x lanes can each transfer data at about 1.25 GBps.

The second board type is called a *switch* board. A switch board is the same size as a payload board, but it has five new high-speed serial connectors. A switch board terminates 80 of the differential serial lines, or 20 4x connections, each at 1.25 GBps or 25 GBps total.

The simplest implementation of a VXS system arranges the payload boards in a dual star configuration, as shown in Figure 1. Up to 18 payload boards are

included, along with two switch boards. Each payload card connects to each of the switch boards with one 4x lane. The remaining two 4x lanes on the switch boards are used to connect them to each other. The implication is that I/O and processing take place on the payload cards and communicate through the switch cards. Note that VXS does not specify a protocol to be used over the serial ports; several protocols are expected to be used, including PCI Express, InfiniBand, and Serial RapidIO, to name a few.

Advanced topologies

The star configuration is not mandated by the VITA specification, as that is only a suggestion. In addition, the concept of a set of nodes connected to a switch is sup-

ported by VITA 41, but it is not required. It is possible, for instance, to build a backplane that does not include any switch boards at all, just payload boards. The serial lanes can be connected in whatever way makes sense for the application. It is also legal to build a backplane that has more than the standard two switch board positions. The switch boards are defined to have very high I/O bandwidths. If the boards are populated with processing nodes instead of some protocol switch, a very powerful processor can be built.

FPGA processing boards

Current technology makes it possible to build very powerful processing boards using FPGAs instead of general purpose or DSP processors. The VXS switch




Figure 1

VITA 55/VSP

Many embedded signal processing systems just have fixed connections between the elements of the system. By design, these systems don't gain anything from using a switching fabric such as PCI Express or RapidIO because the connections never change. Therefore, the overhead involved in implementing switching protocols is wasted. For this reason, a new protocol called Virtual Streaming Protocol (VSP) has been proposed as VITA 55. VSP is a straightforward, low overhead, point-to-point protocol with no switching or fabric functions. It is designed to minimize the use of FPGA resources over static links. Conversely, by eliminating the unnecessary protocols, FPGAs linked together offer maximum performance.

board form factor provides enough I/O bandwidth to keep such a processor busy. One example of a VXS processing board is the QinetiQ/Tekmicro *Callisto* board, as shown in Figure 2.

This module is a VXS switch board with five Xilinx Virtex-II Pro FPGAs mounted on it, each one with a 128 MB external DDR RAM block associated with it. This type of board provides extended connectivity by connecting the five FPGAs with a 64-bit parallel bus that supports full-duplex data rates above 1.5 Gbps. Each FPGA has three or four 4x lanes connected to the VXS backplane, each running at 1.25 Gbps. It also has 12 front-panel fiber or copper transceivers running at 2.5 Gbps each, for connectivity to other chassis or systems.

VITA 41.7 mesh backplanes

As noted above, the VITA 41 specification describes one possible topology, a *dual star* arrangement, but it does not mandate this or any other topology. Also known as *VXS Processor Mesh*, VITA 41.7 is a new configuration for VXS that offers tremendous bandwidth potential via mesh connectivity options. Companies such as Elma

Bustronic have introduced backplanes in several configurations, some with only payload cards with direct connections to each other, as well as configurations with five switch/processor slots interconnected in a mesh arrangement. These backplanes take advantage of the advanced capabilities of the VXS concept.

Some FPGA clusters use chains of FPGAs to process large amounts of data in parallel. The data flows through each chain or pipeline and is processed by each chip in turn. FIR filters, digital downconverters, or demodulators might be implemented this way. The individual pipelines don't interact with each other, except perhaps at the end of the chain. The simplest way to implement this sort of application would be to use all the FPGAs on one board to handle an entire chain, transferring data from one to the next through the parallel interconnect bus. Input and output data would use the backplane serial connections.

More complex applications might need more complicated interconnections. Matrix or image processing operations would need a partially or fully connected mesh, since the data flows in two dimen-

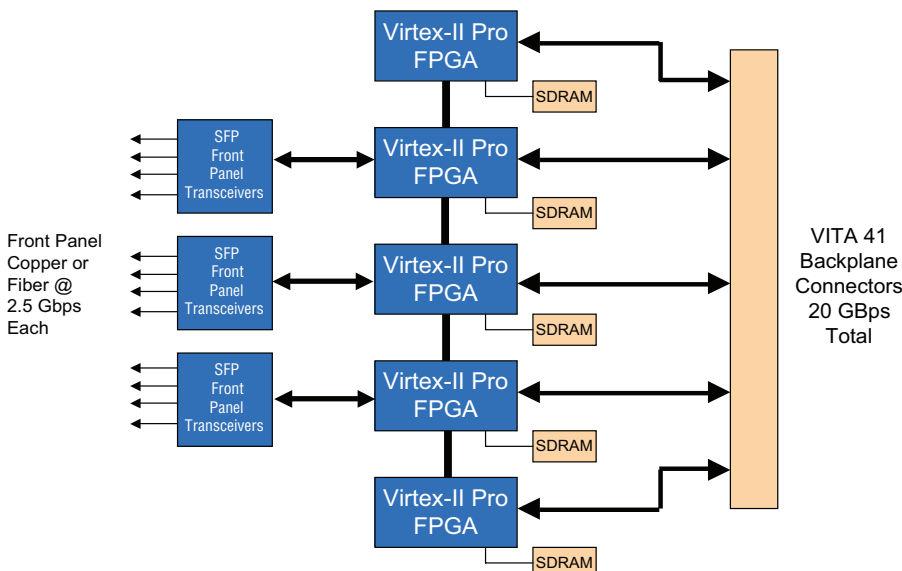


Figure 2

“The new approach using FPGA-based processing, VXS, and the mesh architecture can be used to do significant processing work by clustering tightly connected FPGAs, resulting in significant size, weight, power, and cost advantages over general purpose processors.”

sions in these sorts of problems. The Elma Bustronic processor mesh backplane would be a good fit for these kind of problems, because it provides a very rich set of interconnections.

By bringing the latest technologies together for very demanding I/O-centric signal processing applications, we have enabled the creation of systems that were previously built using general purpose processors. The new approach using FPGA-based processing, VXS, and the mesh architecture can be used to do significant processing work by clustering tightly connected FPGAs, resulting in significant size, weight, power, and cost advantages over general purpose processors. Ω

Bob Walsh is a field applications engineer for TEK Microsystems, Inc., based near Austin, Texas. Bob has extensive experience in the embedded computing industry at leading companies such as Mercury Computer Systems and SKY Computers. Previously, he was an embedded engineer at Texas Instruments. He holds an MS degree in Electrical Engineering from Southern Methodist University, as well as an MS in Management and Administration from the University of Texas at Dallas.

For more information, contact Bob at:

TEK Microsystems, Inc.
2 Elizabeth Drive
Chelmsford MA 01824
Tel: 512-869-2592
E-mail: bwalsh@tekmicro.com
Website: www.tekmicro.com