Case Study: Cryptography Computing Cluster Blends FPGAs, Quad CPUs

Designed for military cryptography applications, a cluster computing design marries an array of FPGA computing resources and power control processors tied together over PCI Express.

Steve McReynolds, Technical Support Manager
Trenton Technology
Mark Hur, Director of Marketing
Pico Computing

Choosing the right architecture can make a huge difference in how complex signal processing computing applications perform. Applications such as weather predication, bomb blast analysis and financial modeling, for example, require floating-point computations. Meanwhile, other similarly complex applications such as the Smith Waterman algorithm, military cryptography and image processing do not typically require floating-point computations. For non-floating-point applications, large arrays of FPGAs can dramatically speed up the computation of the overall solution algorithm.

This case study takes a look at how an array of up to 84 FPGAs managed by a PICMG 1.3 System Host Board (SHB) communicating to the FPGA boards over PCI Express links on a PICMG 1.3 backplane, enables 10,000:1 speed improvements in typical hardware cryptography applications.

Get Connected
with companies mentioned in this article.
www.cotsjournalonline.com/getconnected

Figure 1

Unlike general-purpose processor arrays—in which a very large number of gates must switch to perform even the simplest decisions—FPGA computing modules, such as this Pico Computing E-16, use a minimal set of gates to make basic logic decisions.

FPGA Acceleration with Intel Processors

General-purpose processor arrays—including graphics processors—are effective and easy to use. But many such processors are burdened with extremely high power requirements. In such processors, a very large number of gates must switch when performing even the simplest decisions. In contrast, FPGA modules, such as the Pico Computing E-16 shown in Figure 1, use a minimal set of gates to make basic logic decisions.

The high-performance computing system described here requires approximately 500W of system power. Over half of that power is consumed by the PICMG 1.3 System Host Board (SHB) with its two Quad-Core Intel Xeon Processors E5335. The Trenton MCXT system host board shown in Figure 2 takes full advantage of the capabilities of the Quad-Core Intel Xeon Processors to manage the complex data traffic flow needed by the high-performance computing system.

In cluster computing designs, equally important as the processing elements is the flexibility of the interconnect that links processing elements together. Rarely can any one military application justify its own purpose-built hardware optimized for a particular algorithm. Typically a machine designed for one military application does not have the optimal interconnect strategy for another algorithm. With that in mind, the PCI Express (PCIe)-based interconnect strategy

used in this system allows tremendous flexibility connecting the processing elements (PE). The PEs may be connected to their nearest neighbors, to other processing elements hosted in the same box, and even across boxes with no change to the FPGA algorithms and minimal impact on speed. Unless a PE is capable of performing significant amounts of processing independently, the potential for high levels of parallelism is limited. The PCIe interconnect strategy used in the Pico Computing SC3 Super Cluster overcomes this inherent limitation.

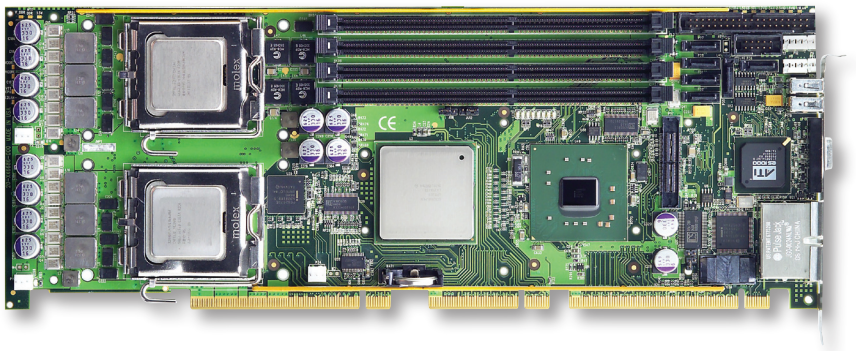

Figure 2

In the SC3 SuperCluster, Trenton's MCXT system host board shown here takes full advantage of the capabilities of the Quad-Core Intel Xeon Processors to manage the complex data traffic flow needed by the high-performance computing system.

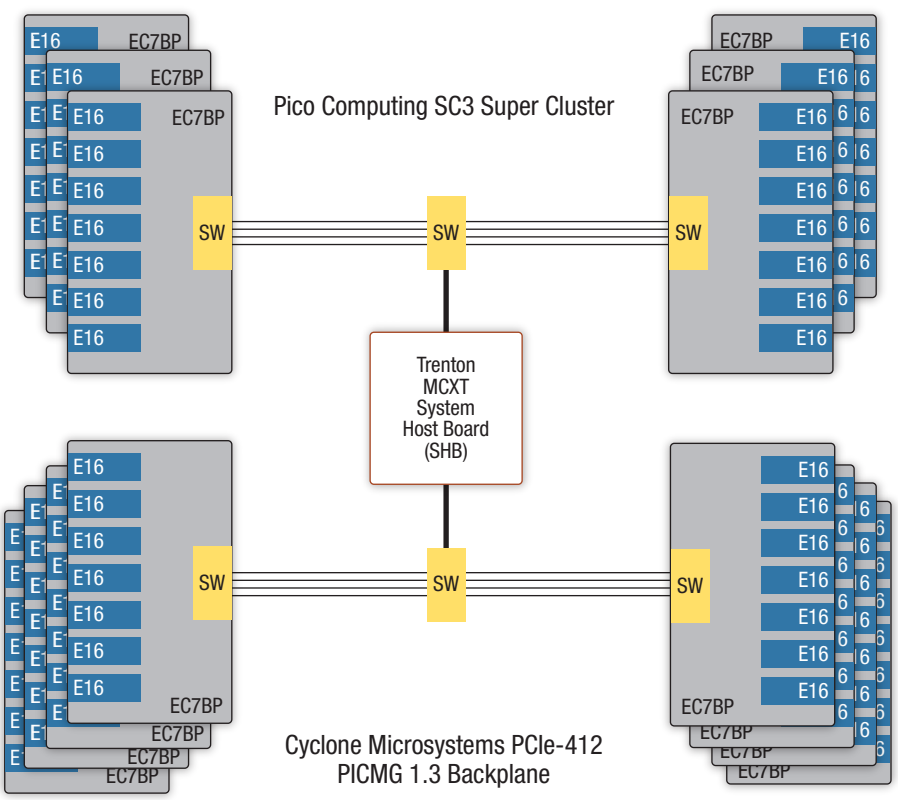

Figure 3

Illustrated here is the system architecture of the Pico Computing SC3 SuperCluster. The SC3 is the third generation of FPGA-based cluster computers from Pico Computing. The Cyclone Microsystems model PCIe-412 backplane is used and features fourteen PCI Express option card slots, one SHB slot and two PCI-X slots. A Pico Computing EC7BP PCI Express backplane plugs into the PCIe card slots on the backplane.

FPGA-Based Cluster System Architecture

The block diagram in Figure 3 illustrates the system architecture of the Pico Computing SC3 SuperCluster. The SC3 is the third generation of FPGA-based cluster computers from Pico Computing. A central element of the SC3 is the PICMG 1.3 system host board from Trenton Technology, featuring two long-life Quad-Core Intel Xeon Processors. The Cyclone Microsystems model PCIe-412 backplane is used and features fourteen PCI Express option card slots for the FPGA backplane boards, one SHB slot and two additional PCI-X option card slots. A Pico Computing EC7BP PCI Express backplane plugs into the PCIe card slots on the backplane. Figure 4 illustrates how each EC7BP is capable of supporting up to seven standard FPGA cards such as the Pico Computing E-16 card, and each E-16 card features a Virtex-5 LX50 FPGA with 46,080 Logic Cells. Using this combination of off-the-shelf components enabled the system to be developed, manufactured and tested in a very short period of time.

The PCI Express serial links between the SHB slot and the individual E-16 FPGA cards are managed by a PCI Express fan out switch on each EC7BP backplane card and four 32-lane PCI Express fan out switches on the Cyclone PCIe-412 PICMG 1.3 backplane. A PCIe-to-PCI-X bridge chip on the backplane manages the link between the SHB and the two 64-bit/100 MHz PCI-X option card slots. These standard PCIe fan out switches can simultaneously coalesce multiple PCIe lanes. This provides considerable bandwidth overlap while keeping any inherent latency penalties down to a minimum. Granted, the PCIe lane architecture is not as fast as some gate-to-gate FPGA technologies that have been built. But PCIe-

based FPGA clusters can be expanded to include very large arrays where cluster computing is used to satisfy the requirements of complex applications.

The SC3 runs under Linux and naturally exploits the 64-bit architecture of that operating system. For algorithm development, any particular card may be run under Windows (XP or Vista), as would be most likely available on any laptop. Loading the FPGA image is managed entirely through the PCIe bus (in either architecture). The choice of host operating systems is, of course, not binding on the operating system that runs on the FPGA. On the SC3 there is no processor on the FPGA card so an operating system on the card was not a choice. That said, the PE can be replaced with the FX part, which includes a PPC 440 processor.

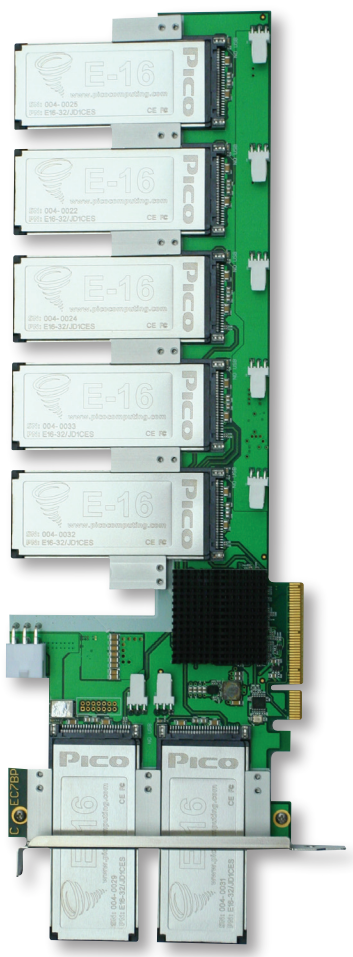

Figure 4

Each EC7BP is capable of supporting up to seven standard FPGA cards such as the Pico Computing E-16 card, and each E-16 card features a Virtex-5 LX50 FPGA with 46,080 Logic Cells.

Engineering Challenges and Solutions

As shown in Figure 5, the Pico Computing SC3 SuperCluster has five EC7BP PCI Express backplane cards installed with seven E-16 FPGA cards plugged into each board. The PCI Express architecture is built on the same logical addressing model developed years ago with the PCI bus and as such, PCIe is subject to similar limitations on the number of supported device segments. Given the architecture of a typical PICMG 1.3 passive backplane and the BIOS configuration of a system host board, the number of allowable PCI buses was exceeded in the initial system design with the maximum number of FPGA cards installed.

To correct this condition, engineers at Trenton Technology made several modifications to the MCXT board's BIOS to allow up to 256 segments on the PCIe bus. Each Pico Card requires approximately 3 1/7 segments when allowances are made for all of the supporting devices contained in the system. With 256 segments managed by the BIOS, a maximum of 84 Pico Cards can be used effectively in this SC3 SuperCluster.

Particular care was paid to cooling the SC3 SuperCluster. The standard fans were replaced with higher-volume fans; however, it was not necessary to deviate from the standard air-cooled chassis design. Each FPGA has a temperature sensor built in, and the chip will shutdown if it overheats. The same is true with the Quad-Core Intel Xeon Processors used on the Trenton MCXT system host board. Care must taken to guard against FPGA overheating since the internal algorithms are open to the user design and it is possible to drive the FPGA so hard that it will overheat. Design steps are taken to find the point where the increased FPGA clock speed will cause the FPGA to "burn up" and then the clock speed is adjusted down-

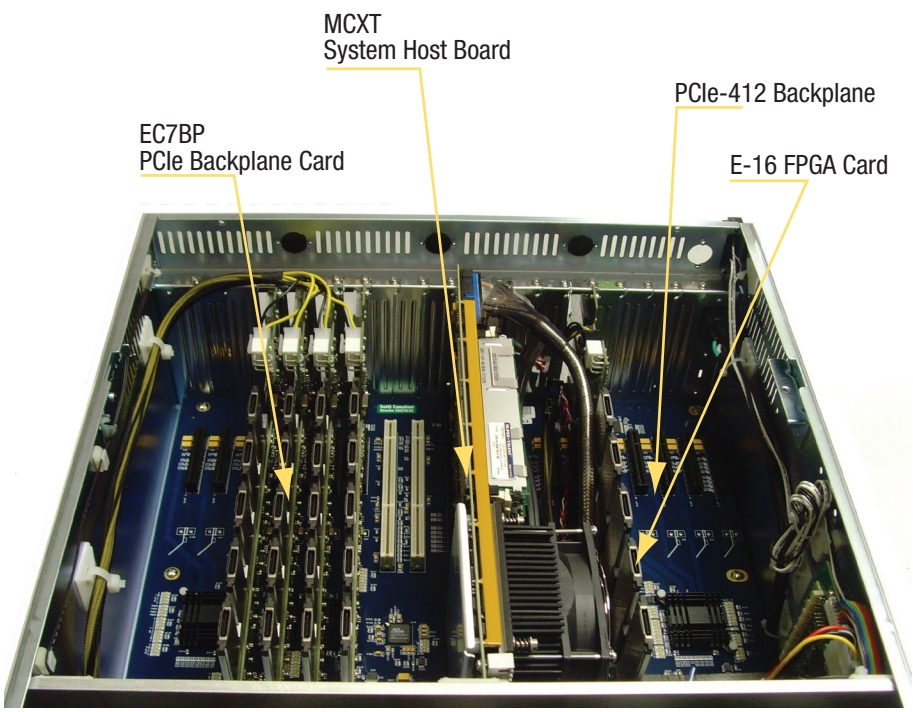

Figure 5

The SC3 SuperCluster has five EC7BP PCI Express backplane cards installed with seven E-16 FPGA cards plugged into each board.

ward to avoid this condition. While this additional step in the design process must be taken, it does illustrate the inherent FPGA flexibility and acceleration capability to fine-tune the FPGA clock speed to maximize processing efficiency of a particular algorithm.

Future Development Directions

Having the FPGA card as a separate unit with its own power regulation, memory, clocks and PCIe interface, opens up the possibility of using different processing elements. Among these possibilities under active development are: an FPGA with a PPC processor, a larger FPGA, and special coprocessors such as the MathStar chip. Single instances, or clusters, of such components can be integrated with no change to the overall system cluster architecture. The Pico E-16 card is a 34 mm wide card. The same footprint will accommodate a card with a 54 mm width, which will in turn permit a larger FPGA to be mounted on the card such as the Virtex-5 LX85 or possibly the LX110.

Both FPGAs and general-purpose processors have an impressive breadth of available development tools and a long and established track record of delivering superior system performance. General-purpose processors represent 50 years of maturity, and FPGAs have been used in advanced computing applications for over 20 years. Tools, algorithms and applications for both of these technologies have made tremendous strides. Needless to say, there are high levels of engineering and development activity at work advancing the capabilities of both of these product technologies.

An FPGA is an intrinsically parallel device and implements well with the kind of algorithm that can be divided into relatively watertight sub-processes. FPGA architectures can be expanded, more or less at will, to incorporate many diverse processing elements. This capability enables very flexible FPGA solutions in a wide variety of computing applications. The SC3 SuperCluster is a great example of how complete off-the-shelf FPGA and quad-core processor technology can be merged together by virtue of PCI Express and the PICMG 1.3 architecture, to provide the robust, high-performance computing platform needed for military cryptography and applications like it. ■■

Pico Computing
Seattle, WA.
(206) 283-2178.
[www.picocomputing.com].

Trenton Technology
Gainesville, GA.
(770) 287-3100.
[www.trentontechnology.com].