Optimizing Embedded Systems for Broadband 10 Gigabit Ethernet Connectivity

FPGA technology is revolutionizing broadband systems development.

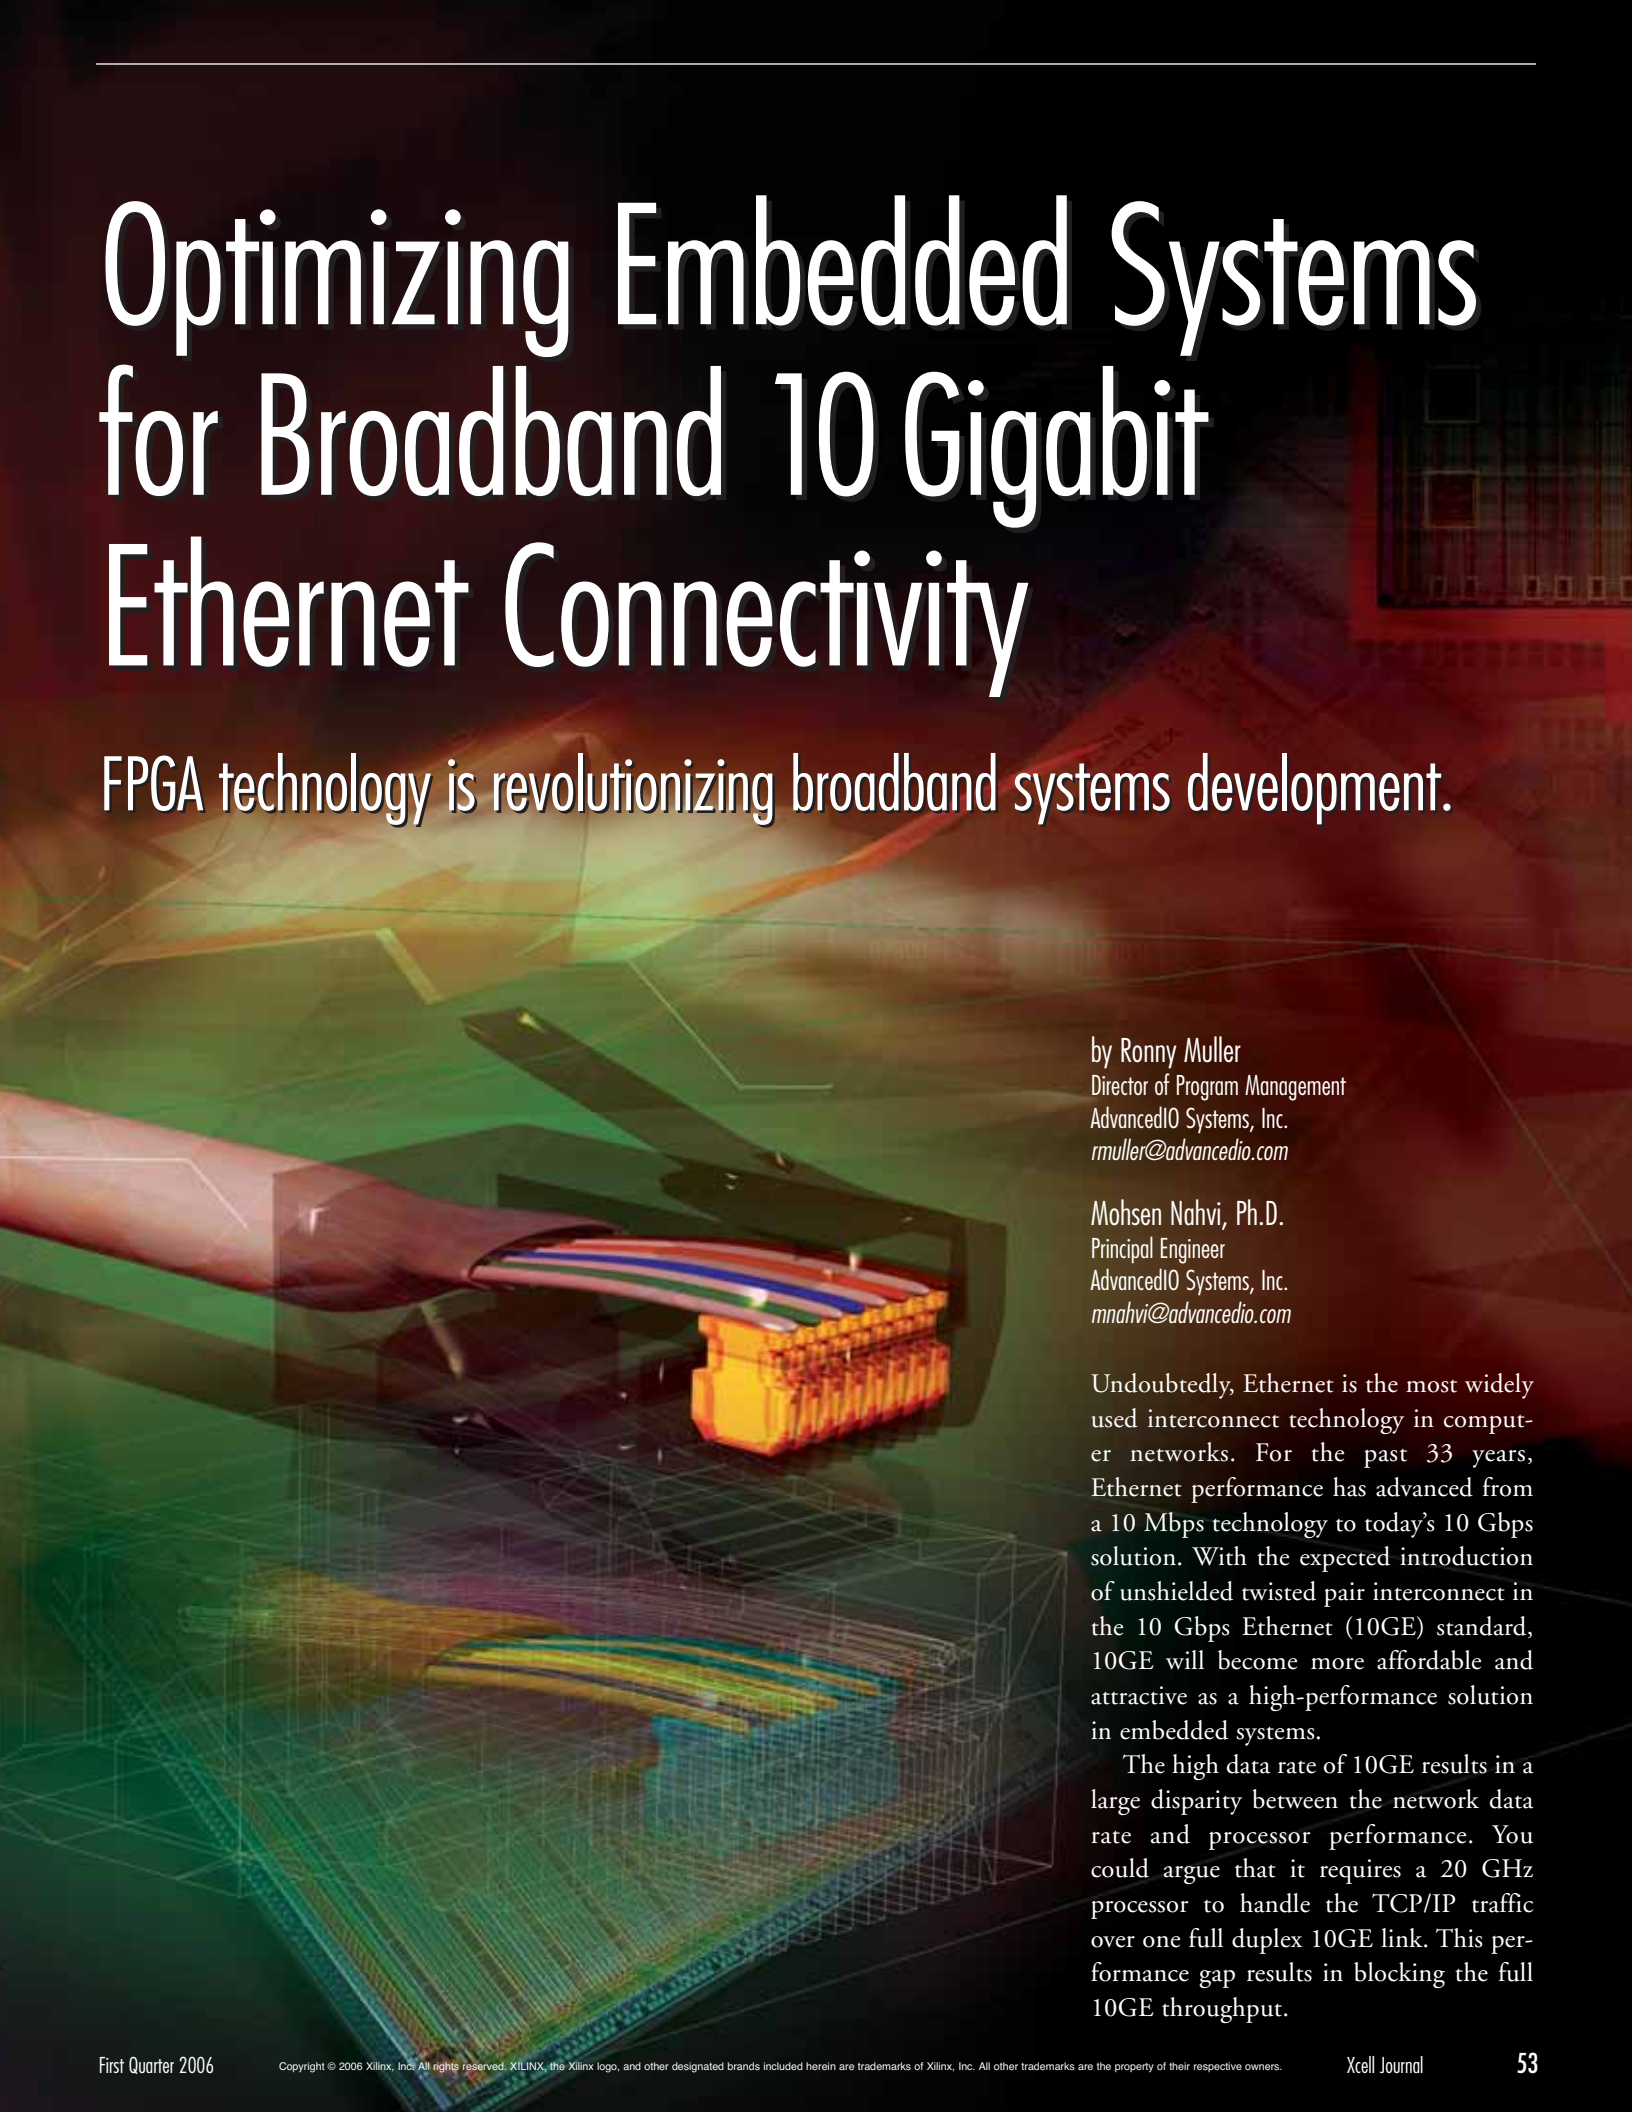

by Ronny Muller
Director of Program Management
AdvancedIO Systems, Inc.
rmuller@advancedio.com

Mohsen Nahvi, Ph.D.
Principal Engineer
AdvancedIO Systems, Inc.
mnahvi@advancedio.com

Undoubtedly, Ethernet is the most widely used interconnect technology in computer networks. For the past 33 years, Ethernet performance has advanced from a 10 Mbps technology to today's 10 Gbps solution. With the expected introduction of unshielded twisted pair interconnect in the 10 Gbps Ethernet (10GE) standard, 10GE will become more affordable and attractive as a high-performance solution in embedded systems.

The high data rate of 10GE results in a large disparity between the network data rate and processor performance. You could argue that it requires a 20 GHz processor to handle the TCP/IP traffic over one full duplex 10GE link. This performance gap results in blocking the full 10GE throughput.

. . . AdvancedIO provides an FPGA framework in which these high-speed interfaces are mapped, placed and routed, and guaranteed to meet timing.

The problem is even worse in high-performance embedded systems. These embedded systems typically rely on high-density distributed processing elements (PE), where each element is optimized to perform specific DSP functions and has no free cycles to perform the complex task of TCP/IP traffic processing.

To address the huge performance gap, you need TCP/IP offload engines (TOE) and protocol acceleration schemes at the network interface; in effect, smart 10GE solutions that go beyond the simple task of a network adaptor. A few ASIC-based solutions on the market do offload part of the TCP/IP stack. However, because of the huge gap between the network and the processor performance, a robust solution is essential that will not only help with offloading the TCP/IP protocol, but also enable you to address other challenges specific to applications like real-time data encryption/decryption and packet classification.

In this context, FPGA-based solutions offer an ideal solution because of their flexibility, configurability, and shorter time-to-market capabilities. An FPGA-based network adaptor allows you to offload the TCP/IP protocol processing from the host system, enabling you to implement other functionality into these intelligent adaptors to improve overall embedded system performance.

In this article, we'll discuss a configurable 10GE solution that enables you to address your high-performance network needs and solve the 10GE challenges in your system.

The AdvancedIO Architecture

AdvancedIO Systems is focused on solving problems of high-speed packet switching in broadband applications using programmable logic. The company offers a line of I/O products and system solutions based on a common architecture.

At the heart of AdvancedIO's architecture is a Xilinx® FPGA, which implements the 10GE physical coding sub-layer (PCS), the 10GE media access controller (MAC), and the RapidIO (or PCIe) switched fabric interface. These standard interfaces are available as optimized logic cores from a variety of vendors.

As shown in Figure 1, AdvancedIO provides an FPGA framework in which these high-speed interfaces are mapped, placed and routed, and guaranteed to meet timing. Consequently, AdvancedIO greatly reduces the engineering effort required for designing new applications. All interfaces are fully controllable from the on-chip PowerPC™ 405 processors, allowing low-data-rate verification of broadband applications.

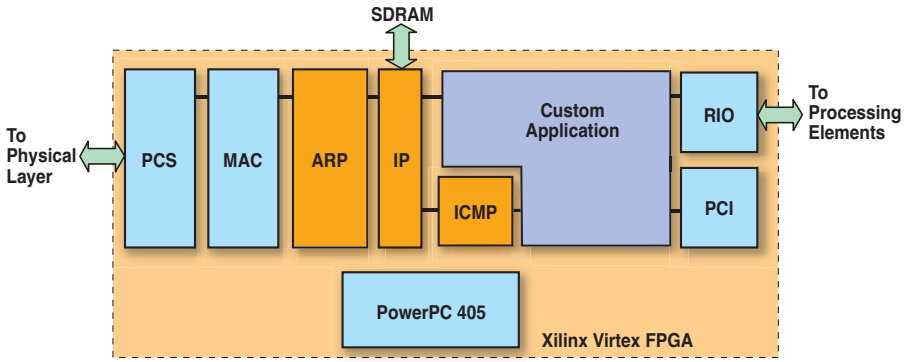

Figure 1 – AdvancedIO 10GE architecture

The V1010 is the industry's first configurable 10GE module in an XMC form factor (VITA 42), with the design based on AdvancedIO's architecture for embedded applications. One particular configuration of the module bridges between RapidIO and 10GE provides the ability to perform 10GE layer 2 and 3 packet processing at wire speed.

Three additional logic blocks facilitate the implementation of a 10GE network endpoint within the V1010 module:

- ARP. This block takes incoming IP frames and converts them into Ethernet frames by appending the Ethernet destination and source MAC addresses. It implements the network address to hardware address request and response protocol and maintains an ARP table in hardware.

Figure 2 – AdvancedIO V1010 configurable 10GE XMC module

- IP. This block terminates IP. It implements IP fragmentation and de-fragmentation at wire speed by buffering fragmented datagrams in SDRAM until the full datagram has been received. Moreover, it checks and generates IP checksums and also performs IP routing. The IP routing tables are configured by the on-chip PowerPC 405 processor.
- ICMP. This block implements the required ICMP protocol. It responds to ping/trace-route commands, for example, and reports and counts errors.

Figure 2 is a photograph of the module currently shipping.

Distributed Processing Made Easier

As shown in Figure 3, the AdvancedIO architecture allows full-speed access to switched fabrics such as RapidIO or PCIe, and can send and receive full IP datagrams to and from the 10GE IP network. Using this architecture, you can easily implement any protocol processing, from the very simple to the very complicated.

As an example, multiprocessor embedded systems with RapidIO connectivity can now seamlessly communicate with each other over an IP 10GE network, as shown in Figure 4. To achieve this, the V1010 module encapsulates RapidIO packets into UDP packets, keeping track of lost and out-of-order packets and reporting these errors to the processing elements.

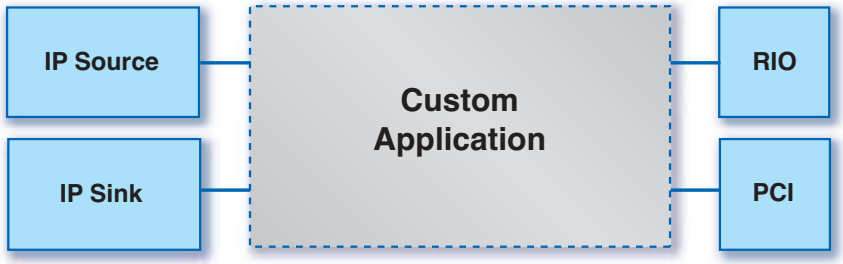

Figure 3 – AdvancedIO processing model for broadband applications

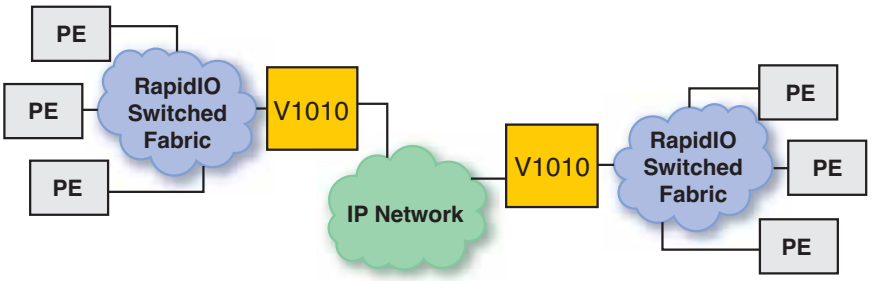

Figure 4 – Distributed remote processing over RapidIO and 10 Gigabit Ethernet

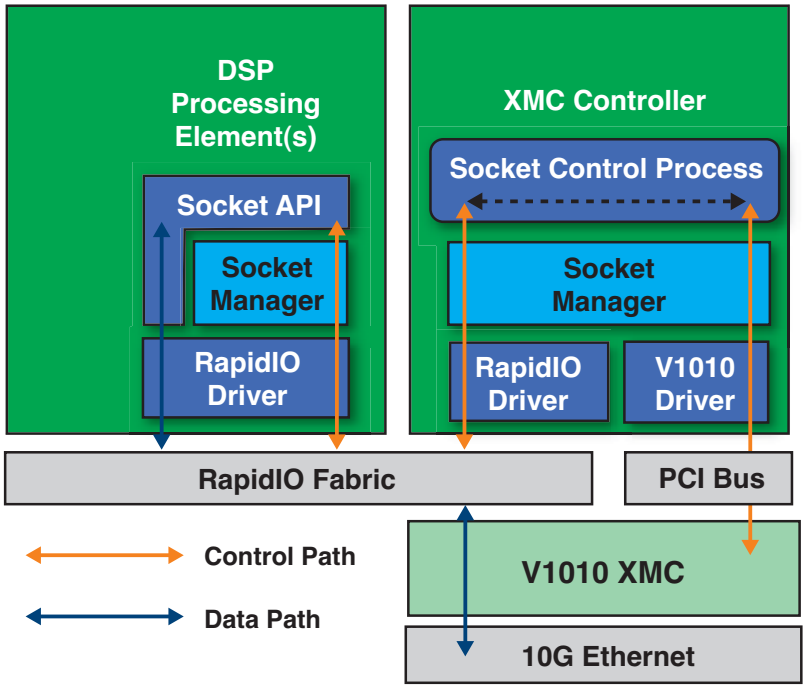

Figure 5 – AdvancedIO embedded software framework

The above tunneling system allows for inter-chassis communication over an IP network and is completely transparent to the processing elements. As all legal RapidIO packets in both message and I/O specifications are transferred over the network, you now have the capability of utilizing remote direct memory access (RDMA) over an IP network, which enables and greatly simplifies the development of complex distributed embedded applications.

The programming interface is a Unix-like socket application programming interface (API), so an application can set up and control the connectivity to the 10GE gateway from any processor element in a RapidIO network. As shown in Figure 5, the AdvancedIO software framework abstracts the complexity of the underlying functions and the 10GE data path. In a typical application, the socket manager is called, which in turn sets up the control path between a processing element and one or more 10GE gateways. Once the control path has been established, an application can fully utilize the 10GE data path.

Conclusion

With the advent of 10 Gigabit Ethernet and its foray into multiprocessor embedded systems, you need intelligent network solutions that offer more than mere TCP/IP offloading. The programmability of the AdvancedIO architecture provides the right ingredient for rapid system development and boosts the performance of embedded systems to a higher level.

The AdvancedIO architecture offers a new and unique solution to embedded computing developers, providing an efficient way for applications to take advantage of all that the 10 Gigabit Ethernet technology has to offer.

For more information on AdvancedIO's 10 Gigabit Ethernet products and services, visit *www.advancedio.com* or e-mail *contactus@advancedio.com*.