Selecting the Right Interconnect

Because interconnect characteristics vary, define your requirements clearly before selecting the appropriate interconnect.

by Jag Bolaria
Sr. Analyst
The Linley Group
jag@linleygroup.com

Interconnects have evolved from parallel to serial and increased in complexity to enable communications with greater efficiency and less congestion or hot points. In addition to connecting endpoints, modern interconnects define comprehensive protocols for moving data efficiently across a network of endpoints.

Thus, the networks and endpoints that must be interconnected often drive the requirements for an interconnect. These requirements include data rate, latency, lossy or lossless links, scalability, and redundancy. These requirements then drive the selection of the appropriate interconnects for a specific network.

The existing ecosystem for an interconnect technology is another important factor in selection. A good ecosystem can help reduce development cost and time to market. In this article, I'll look at some of the leading interconnects and position those for specific applications or market segments. The Linley Group's report on High-Speed Interconnects, available at *www.linleygroup.com*, provides more details on various interconnects and the leading products for each.

PCIe and Ethernet

Our research shows that the leading interconnects are driven by large-volume platforms. The economies of scale from larger volume platforms ensure low-cost building blocks and broad availability. Additionally, large deployments lead to field-proven technologies that can be applied in other platforms with minimal risk.

Two of the largest platforms are PCs and networking equipment. The PC platform drives PCI Express (PCIe) and Ethernet, while networking equipment drives only Ethernet. But because these interconnects were developed for specific applications, they are not a natural fit in many other markets. The semiconductor industry and system vendors are evolving these interconnects to meet requirements for new applications.

For example, PCIe scalability has evolved to support greater data rates and more lanes. With IOV (I/O virtualization), PCIe is evolving to support virtualization, which enables its deployment in storage systems and blade servers.

With IEEE802.3ar and BCN (backward congestion notification), Ethernet enhancements include better flow control, congestion management, and attempts to address its inherently lossy nature. These enhancements will strengthen Ethernet applicability for storage systems, data centers, and backplanes.

Although Ethernet and PCIe are now suitable for more applications, they still fall short in meeting the technical and business requirements for all systems. Blade servers, for example, use a combination of Ethernet and Fibre Channel (FC). Although OEMs may want to consolidate these fabrics, end users have a large investment in FC and want support for that now and in the future.

PCIe and Ethernet also fall short in meeting the scalability, latency, and lossless requirements of high-performance computing (HPC) applications. HPC uses a specialized interconnect such as InfiniBand, which provides better latency and scalability. In this case, OEMs will need flexible interconnect solutions to enable common platforms and thus service different user requirements.

Both dominant and specialized interconnects will continue to evolve to support greater data rates, reduced latency, and better scalability.

Everyone Else

Endpoints and specific system requirements often drive the development of specialized interconnects. RapidIO is one such example. Steered by system and chip vendors, RapidIO has evolved to address the unique requirements of the wireless infrastructure. It enables distributed computing on line cards and networking/wireless infrastructure systems better than most competing interconnects. RapidIO is also integrated on DSPs from Texas Instruments and PowerPC CPUs from Freescale.

Because base stations use farms of DSPs, it is an easy decision to use RapidIO as an interconnect in these applications. Over time, we expect RapidIO to expand to other platforms that perform digital signal processing on multiple data streams.

Examples of other specialized interconnects include XFI, SFI, XAUI, SPAUI, Interlaken, SPI-S, and KR. These interconnects were developed to address the very specific low-level requirements of each application. Although addressing all interconnects is beyond the scope of this article, let's look at a few to highlight the problems each solves and its impact in systems.

XFI and SFI are used to connect optical modules at 10 Gbps. At these data rates, the major challenge is signal conditioning, including electronic dispersion compensation for the fiber and equalization for the board traces and connectors. These requirements drive specialized components designed specifically for the characteristics of the channel through which the signal travels.

Because data at these rates may be channelized – that is, include multiple streams on a single physical link – it becomes important to add traffic management. Specifications such as Interlaken, SPI-S, and SPAUI address high data rates as well as traffic management. Because no single standard exists, we believe system designers need to design in solutions that provide flexibility in meeting current and future requirements.

The combination of 10 Gbps rates on the network and multiport line cards drives the need for greater bandwidth and therefore greater data rates over the backplane. The IEEE 802.3ap addresses this with its 10GBase-KR specification, which defines 10-Gbps serial links. In addition to equalization and pre-emphasis, it may be necessary to include forward error correction for acceptable performance over a couple of connectors and up to 40-inch traces common in backplanes. Additionally, these systems may need compatibility to older line cards, driving the need for backplane operation at 1 Gbps or 3.125 Gbps. Again, a flexible solution is critical to meet the system requirements.

Conclusion

There are many different applications for interconnects and many interconnect choices for system designers. We expect PCIe and Ethernet to be the dominant interconnects. These will be used in servers, networking, storage systems, wireless networks, and many other systems. There is, however, no one interconnect (or two) that can meet the requirements of all systems. Therefore, the industry has developed and will continue to support specialized interconnects for different applications.

Both dominant and specialized interconnects will continue to evolve to support greater data rates, reduced latency, and better scalability. Additionally, systems will need to support legacy cards.

We recommend that system designers select the best interconnect and design in flexibility to cover different interconnects and the evolving changes in each. FPGAs play a critical role in offering system designers this flexibility and supporting the broad interconnect landscape.

Xilinx Events

Xilinx participates in numerous trade shows and events throughout the year. This is a perfect opportunity to meet our silicon and software experts, ask questions, see demonstrations of new products and technologies, and hear other customers' success stories with Xilinx products.

For more information and the most up-to-date schedule, visit *www.xilinx.com/events.*

North America

July 23-27
Nuclear and Space Radiation Effects Conference 2007
Honolulu, HI

August 7-9
NI Week 2007
Austin, TX

September 18-20
Intel Developer Forum
San Francisco, CA

October 29-31
Military Communications Conference 2007
Orlando, FL

November 5-9
Software Defined Radio Forum 2007
Denver, CO

Europe, Middle East, and Africa

August 27-29
International Conference on Field Programmable Logic and Applications
Amsterdam, Netherlands

September 6-11
IBC
Amsterdam, Netherlands

Asia Pacific

August 7-10
Agilent Digital Measurement Forum
Taiwan