

SystemIO Technology Promises High-Speed Connectivity Across Multiple I/O Standards

As implemented via Virtex-II Platform FPGAs, SystemIO technology addresses both the physical interfaces and networking protocols for high bandwidth connectivity and throughput.

by Rina Raman

APG Director of Applications Engineering, Xilinx
rina.raman@xilinx.com

Moore's Law, which predicts computer processor speeds will double approximately every 18 months, has proved to be remarkably accurate for more than 35 years. The exponential increases in processor speeds have enabled ever higher bandwidth networking. The problem now, however, is that I/O busses have only doubled their frequencies every three years. Thus, the level of performance of a system at the board level is now dictated not by the speed of processor but by the limits of the I/O bus. While delivering greater bandwidth, numerous emerging I/O standards – RapidIO™, POS-PHY™ Level 4, 10 Gb/s Ethernet, XAUI, HyperTransport™, Fibre Channel, among others – provide various choices and architectural options. Xilinx Virtex®-II Platform FPGAs address this proliferation of standards with the SystemIO™ solution, which provides system interconnectivity at both the physical I/O interface and the networking protocol levels.

I/O Signaling Standards

Virtex-II devices are designed to support many signaling standards, including essential interfaces to high performance systems. For instance, memory devices operating above 180 MHz can only use SSTL (Stub Series Terminated Logic) or HSTL (High Speed Transceiver Logic) I/O standards.

Many I/O standards – including the traditional switching standards of LVCMOS (Low Voltage CMOS), memory interfaces of HSTL and SSTL, and even LVDS (Low Voltage Differential Signaling) – are not keeping pace with the increased demand for more bandwidth. Designers have tried to overcome these bandwidth limits by using more pins and/or larger bus widths. This has made the traditional I/O standards more “pin-intensive.” Significant problems arise as pin counts grow into the hundreds and thousands, creating routing congestion on printed circuit boards. Although Virtex-II Platform FPGAs can easily route inside the device, it is extremely difficult to interface with all the pins on a PCB. More layers of interconnect routing in the PCB cause dramatic increases in overall board costs. The number of pins running these standards also produces problematic electromagnetic interference.

Additionally, the most popular bus architecture is a “shared” medium where multiple entities use the same bus, each waiting in turn for its opportunity to complete its transaction. As the size of audio and video data streams increase, the waiting period gets longer. Ultimately, overall performance decreases along with performance predictability. Therefore, the industry trend is to move from a shared bus to a point-to-point link, typically configured as a switched fabric.

Beyond the Shared PCI Bus

The most popular shared bus is the PCI (Peripheral Component Interconnect) bus. It has become a general-purpose bus for personal computers and embedded systems supporting many different applications. A 33 MHz, 32-bit PCI bus can support one source and five destinations with an overall bandwidth of 1 Gb/s. At 66 MHz and 64 bits, the bandwidth rises to 4 Gb/s, but then the bus only supports one source and two destinations. The PCI shared bus structure has diminishing performance returns, is less predictable than a point-to-point solution, and is limited to internal systems.

Ideally, we would like to establish a link and then “burst” the data over a really wide data bus. This would maximize bus efficiency. Bus efficiency is highest with long bursts.

Decreased bus efficiency has many causes. Although we may have economies of scale for most applications, the shared PCI bus compromises graphics performance. For example, when using a PCI graphics card, the graphics card needs to refresh every few milliseconds. In order to do this, it needs to have immediate and frequent access to the PCI bus. As a result, other PCI cards cannot send huge data bursts. This is because the PCI arbiter is designed to ensure that no one component can send large bursts of data.

In graphics intensive applications, an AGP (Advanced Graphics Port) is a much better choice than a PCI bus. Although AGP is based on PCI technology, it is designed especially for the high throughput requirements of 3D graphics. Rather than using the shared PCI bus for graphics data transmission, the AGP introduces a dedicated point-to-point channel so that the graphics controller can directly access main memory. Virtex-II Platform FPGAs support both PCI and AGP I/O standards, as well as many others.

Another drawback of the PCI bus is that it has no termination – or rather, it is series terminated. This means that it relies on reflective wave switching. Although reflective wave switching is inexpensive and has relatively low power consumption, series termination requires the system to wait for the reflection, so we lose valuable time.


Switched Fabric Advantages

A switched fabric is more scalable for high performance and ultimately, offers a lower cost solution for high bandwidth applications. In the past, data rates weren't high enough to warrant using a switched fabric within a system, because this would require many point-to-point connections and would increase system complexity. With the dramatic increases in performance enabled by the new interface standards, however, switched fabric solutions are becoming cost-effective.

With the huge existing PCI infrastructure, bridges are needed to interconnect the various high-performance I/O standards. Virtex-II Platform FPGAs are key building blocks in the transition from shared buses to switched fabrics. The Virtex-II architec-

ture supports universal switching capabilities to these new standards, making them the logical choice for system designers.

As we mentioned earlier, the key advantage of SystemIO technology is that it



provides both physical interfaces as well as various cores that support the network protocols. With the transition from shared busses to switched fabrics comes a variety of different source-synchronous (parallel) protocols. The following are some of the leading source-synchronous standards that are emerging:

- **RapidIO** – Originally organized to support the processor and local bus markets, the RapidIO interconnect architecture has been embraced by the networking and storage markets.
- **InfiniBand™** – Founded by an industry consortium, InfiniBand targets remote storage, servers, and networking devices.
- **HyperTransport** – Formerly known as LDT (Lightning Data Transport), HyperTransport was jointly developed by AMD and API to replace PCI in high-speed computing applications. It has gained some acceptance in the networking space.
- **OIF - SPI4** (Optical Internetworking Forum - System Packet Interface Level

4) – This standard is defined for 10 Gb/s Ethernet applications and optical networking applications demanding OC-192 performance.

- **10 Gb/s Ethernet** – Using an XGMII interface, this IEEE standard is positioned to drive the convergence of LAN and WAN technologies.

In addition, there are emerging serial channel standards as well. Two of the most popular standards in this arena are:

- **Fibre Channel** – A favorite for the storage area networking (SAN) market, Fibre Channel uses optical fiber, coaxial cable, and/or twisted-pair telephone wire.

- **XAUI** (pronounced ZOW-ee) – This new standard targets 10 Gb/s serial channels by bonding four 3.125 Gb/s transceivers. XAUI targets the OC-192 and 10 Gb/s Ethernet markets for WAN and LAN routers.

Conclusion

Having a solution that supports all these various system interfaces is crucial to success in the marketplace. The Virtex-II Platform FPGAs' SystemIO solution offers exactly this – support for physical interfaces as well as cores that support the network protocols for all the common and emerging system I/O interfaces. Virtex-II FPGAs enable high performance interfaces to memories from Cypress, IDT, Micron, SiberCore, GSI Technology and others, as well as interfaces to networking ASSPs (Application Specific Standard Parts) from vendors such as AMCC, PMC Sierra, and Vitesse. Now, with Virtex-II SystemIO solution, you can pick any standard, and any vendor offering that standard in their ASSPs, and rest assured that Xilinx Virtex-II Platform FPGAs will support that standard.