



WP175 (v1.0) May 15, 2003

# *High-Speed Serial Interconnects*

## *Technical Advantages, IC, and System Design Strategies*

*By: Bert McComas  
InQuest Market Research*

---

Companies across a wide range of industries are witnessing a transition from parallel to high-speed serial I/O solutions to reduce system costs, simplify system design, and provide scalability to meet new bandwidth requirements. Serial solutions will ultimately be deployed in nearly every type of electronic products imaginable, from chip-to-chip interfacing, backplane connectivity and system boards, to box-to-box communications. This document focuses on the dynamics of this transition in the connectivity solutions market.

## Market Segments

**Mainstream PC** — Ethernet, USB2, and 1394 initiated the transition to serial interfaces, followed by Serial ATA, and later by PCI-Express to replace existing, purely parallel interfaces. Silicon integration and new high-speed chip-to-chip interfaces offer fast data paths for all performance-critical functions. PCI will exist for many more years as legacy expansion.

**Servers** — PCI-X is firmly entrenched and will not be replaced in the near future. Servers continue to depend on high-speed parallel chip-to-chip mezzanine-level interfaces (such as IMB, Hublink, and HyperTransport). Servers also represent a vital early market for the fastest serial interfaces, such as 1 GB and 10 GB Ethernet, FibreChannel, and Serial ATA.

**Communications & Industrial Computing** — New platforms in these sectors must maximize processing power, expansion capacity, and deliver throughput that scales accordingly. The existing market offers numerous intelligent line cards or single computers using different processors running different operating systems and applications, each handling different stages of the data or packet processing flow. With greater processing power available at the chip and board level, the challenge is to optimize the flow of data, to eliminate potential bottlenecks, and to maximize expansion. Future services such as “fiber to the curb” require performance levels higher than could be previously achieved. The interconnect standards of interest in this market include switched Ethernet, FibreChannel, RapidIO™, HyperTransport, XAUI, Serial RIO, InfiniBand, PCI-Express, Serial Attach SCSI, PCI, CompactPCI, and others.

Since the mid-1990s, PCI has been the glue that has held all platforms together. PCI bus performance has been improved by increasing speed and expanding bus width. However, this approach has also driven up costs, and limited the number of connections because of its shared bus model. PCI is not ideal for state-of-the-art electronics because of performance and expansion requirements. Designers need more aggressive solutions to achieve high-speed and flexible connectivity in their designs.

## Serial Changes Everything

During the last few years, serial interconnect technologies has matured to enable high-speed switched architectures with excellent performance and scalability, as well as low pin counts and cost. Clock speeds do jump from 33-133 MHz in the parallel connectivity world, to 2-10 GHz in the serial connectivity world. Beyond that, numerous technology advancements and capabilities are associated with today’s advanced serial interfaces. These are highlighted in the following table:

*Table 1: Comparison of Serial and Parallel Interfaces*

Serial Interface	Parallel Interface
Packetized bus transactions	Granular bus transactions
Point-to-point switched architecture	Shared bus topologies
Differential low-voltage	3V signaling
Full duplex	Half duplex
8b/10 b encoding with embedded clock	Parallel clock and data

The following are additional serial interconnect features:

- Layered protocol allowing choice of PHY for PCB, cable, or optical
- Multiple lane configurations, such as 1x, 4x, 8x, and 16x
- Isochronous capabilities
- Split-transactions
- Flow control
- Configuration and power management

While some of these attributes are also available in parallel interfaces, full serialization is the final step in the transition that includes:

- Shrinking pin count to its minimum by embedding the clock
- A high-speed analog physical transport layer suitable for long PCB traces or for external cabling (copper or optical).

## Advantages of Serialized Interfaces

One of the most important potential advantages of serial interface technologies is cost reduction. While cost premiums may be acceptable for a short time, the market requires long-term cost advantages (in addition to other strengths) before it moves forward. Even at high performance levels, serial technology helps reduce connector costs, package costs, and possibly board costs.

Additional advantages are explained below.

### Lowest Pin Count Per Connection

The new high-speed serial interfaces offer the lowest pin count per connection and maximum performance through a fixed number of connector pins. You can offer more I/O channels per device and more I/O channels per connector. As a result, the cost per I/O channels is minimized. Multiple interface standards can be implemented in a single IC or grouped on a single connector. Using only a small number of pins, serial interfaces may be added to existing backplanes to dramatically enhance performance, or new small form factor backplanes can be defined. The serial interface also allows the use of low-cost or smaller IC packages, which results in the reduction of the size or number of ICs required in a system with fewer traces, smaller board area, and/or PCB layers.

### Higher Speeds and Improved EMI

Serial interfaces achieve higher clock rates than other interfaces because of their analog asynchronous designs with embedded clock, point-to-point connection, and low-voltage differential signaling. They are less susceptible to crosstalk, noise, and skew. Particularly with pre-emphasis (signal equalization), they can tolerate longer connection distances than digital interfaces. Serial interfaces allow reduced EMI and easier FCC compliance check compared to wide high-speed interfaces. You can confirm this by spectrum analysis of radiated emissions, and testing sensitivity to injected noise.

### Dedicated Bandwidth, Extensible, Switched Architecture

Low pin count per connection allows you to design a multi-port switch IC to sustain maximum dedicated bandwidth with independent, non-blocking data paths between the ports in your system in a cost-effective manner. This is important in the following scenarios:

- bandwidth utilization is high or unpredictable
- bandwidth utilization is expected to change as system resources scale
- bandwidth data processing load scales

Dedicated bandwidth increases expansion capacity, increases cumulative system bandwidth, and keeps latencies minimized without impacting connector pin count requirements. These performance advantages can be important for real-time or isochronous applications that need deterministic bandwidth and/or latency between many connected devices. This architecture also makes it easier to implement a hot swap capability for higher reliability and better uptime.

## Reference Comparison - PCI-Express & XAUI

Although there are many differences between the parallel and serial versions of a particular interface, the most tangible difference is pin count relative to performance. In this section, PCI-Express and PCI-X are compared, followed by a comparison of XAUI and XGMII.

To understand the pin count and performance of PCI-X, its current performance extremes need to be analyzed. 32-bit PCI-X requires 61 signal pins, and operates at 66 MHz, delivering 256 Mb/s, which is shared by up to four slots. 64-bit PCI-X 2.0 QDR requires 89 signal pins, and operates at 533 MHz, delivering 4 Gb/s to a single slot. PCI-Express operates at 2.5 Gb/s delivering 250 Mb/s bandwidth in each direction, or 500 Mb/s full duplex directly to each slot over four signal wires (two differential pairs).

The next comparison is between XAUI and the 10 Gigabit Media Independent Interface (XGMII), which is a full duplex digital interface between the 10 Gb Ethernet MAC and PHY. XAUI is a serialized version of this interface.

XGMII is a full duplex interface operating at 312.5 Mb/s per wire with 32 data bits driven in parallel in each direction. In addition to its 64 data signals, it has 10 clock and control signals for a total of 74 pins. Its data rate is 1.25 Gb/s in each direction. XAUI is a four lane interface, where each lane is full duplex operating at 3.125 Gb/s. As a result, there are four differential signal pairs in each direction, totaling 16 signal pins delivering 12.5 Gb/s. Given the 8B/10B encoding/decoding overhead, the maximum bandwidth is reduced to 10 Gb/s or 1.25 Gb/s.

Because of signal count, skew, and other problems, XGMII cannot be used for routes longer than 7 cm, or to support numerous interfaces in a single chip. For this reason, chip-to-chip, board-to-board, and chip-to-optical module applications are not very practical with this interface. To address this issue, the industry defined and adopted XAUI. Automatic de-skew and other features allow the XAUI interface to be routed to distances of 50 cm on PCB. Additionally, its low pin count makes it easier to develop complex ICs for switches and other applications. XAUI is very attractive as a MAC-PHY interface, but it will also be used as a chip-to-chip interface, or for backplanes without using a PHY.

## Serial, Can You Choose Just One?

Using a serialized interface for certain applications may seem like an obvious necessity, but choosing which standard to use is a more complex issue. There is no clear winner in terms of cost, market access, performance, expansion capacity, flexibility, time-to-market, infrastructure compatibility, design complexity, design debugging and validation support, future standards migration, off-the-shelf component availability, software, and so on. None of the new serial standards are expected to out perform all of the others. It may be necessary to support many different standards for years to come.

This puts pressure on system and silicon companies to design and deliver a broader range of chips, boards, and systems than ever before. They must find solutions to acquire IP, and then quickly design, simulate, and verify mixed signal systems at the IC and board level. Concurrently, manufacturers must offer the flexibility to adapt to different combinations of interface standards for different market segments, and follow future speed migrations.

This dynamic may force designers to reevaluate how they develop and productize their platforms, how they deal with the technical complexities of new standards, and how they react to changes to market trends. Implementation strategies can make or break a product. Important decisions must be made, for example:

- Should a full custom mixed signal IC design or an advanced FPGAs be used at the silicon level?
- At the board level, should cookbook design guidelines be used or should the design tools needed for end-to-end signal system level simulation, layout, and verification be acquired?

## Development Strategies and Resources

Sizing up design complexity is often a normal part of the standards selection process. Not surprisingly, almost everyone interviewed specified their chosen standard as doable, while the competing standards were branded as challenging. Clearly, it is important to find someone who implements more than one of these standards without bias and who uses available tools and resources. Although no one has implemented all of these standards yet, the FPGA vendors have implemented quite a few and are fighting both sides of the interface wars simultaneously. If there is an interface problem, the FPGA vendors want to solve it, while solving as many other problems they can in the process.

A Xilinx FPGA development platform is currently available that includes four Gigabit Ethernet fiber ports, two Serial ATA, and two InfiniBand connectors on one board, all driven directly from the FPGA. It is a clean board, not overpopulated with external circuitry. It also supports PCI, DDR, Ethernet, serial, parallel, PS2, I2C, debug ports, AC97 plus TFT, and a touch screen among other interfaces. Multiple PowerPC processors and many programmable analog physical layer transceivers that can be configured to drive the high-speed serial interfaces on this board or other interfaces such as XAUI and PCI-Express are integrated into the FPGA. Xilinx used this technology to produce the first available PCI-Express prototypes. Serial RapidIO, serial HyperTransport as it evolves, then Serial Attached SCSI are the next interfaces to be prototyped in this manner.

While the FPGA vendors often provide basic IP blocks to get designs off the ground, designers must often rely on third-party IP vendors to acquire logic cores for sophisticated new interfaces. Consequently, the reference board is not considered a completely functional and validated platform, but rather a rapid low-cost prototyping vehicle for designers developing or acquiring the necessary IP to execute one or several of the newest high-speed serial interfaces. The economic justification of using such an FPGA for silicon prototyping is strong considering a full mask set in 0.13 micron today can cost in excess of a million dollars.

The availability of this type of FPGA solution with multiple integrated configurable PHY transceivers disputes many assumptions about significant time-to-market differences between the competing high-speed serial interfaces.

Rapidly changing specifications and market requirements often make FPGA-based solutions the ideal chip interface, and in some applications, the only choice. The high

costs associated with developing state-of-the-art chips using the newest serial interfaces further tips design advantages to these latest FPGA devices. Responding to this need in the high-speed serial world, Xilinx has its new Virtex-II Pro family, and Altera has its Stratix-GX family.

## Interface Design Validation

Designing to an industry standard specification sounds simple, but there are different ways to interpret specifications. Usually, difficult judgment calls must be made about whether to implement certain features or operating modes that may be optional, non-essential, or unclear. One interpretation of the specification may not be wrong, but if another implementation of the same specification differs too much, interoperability issues may arise.

This issue is being addressed by Denali, a company that has solved a similar problem in the DRAM world. Denali produces, collects, organizes, and publishes functional models that accurately represent the functionality of the interface standard, as well as additional models that represent the precise functionality of each silicon vendor's implementation. This "Verification IP" is used to speed pre-silicon design and verification of your bus interfaces for compliance to the specification, and to evaluate interoperability with third-party implementations in its database.

The Denali functional modeling and design verification methodology applies to the digital controller function exclusively. It does not comprehend analog signal integrity and board level challenges that exist from the PHY onward, which is a completely different issue.

## System Level Mixed Signal Design & Simulation

While the IC logic design and validation process has its challenges, some logic designers are not afraid to admit that high-speed serial interfaces may introduce analog design complexity at the board and system levels. The transition from slow, wide synchronous parallel buses, to numerous lanes of 2–10 GHz asynchronous serial channels introduces new physical and electrical design complexities that have not been addressed or even understood before.

The strategy in the digital days was to sacrifice any time-to-market advantage, while waiting for a market-enabling vendor to provide reference designs, characterization data, and a strict cookbook of design guidelines. You would not be the first to market, but it was likely that your product would work.

Moving to serial allows you accomplish projects that were not possible before. Complex design variables can multiply in terms of technology, device and interconnect density, interconnect distance, simultaneous support of multiple standards, coexistence with fast parallel buses, routing problems, and the challenge of ensuring interoperability within a continuously expanding list of silicon suppliers. It is hard to imagine how a cookbook can solve these issues. Designers are moving to high-speed serial interfaces for performance, flexibility, and to push the design limits.

The traditional starting point is to review HSPICE and work within a standard eye diagram, keeping a firm handle on FR4 and connector parasitics. This only takes basic signal integrity from pin to pin in the analog domain into consideration. Perhaps this is all that can be done in a cabled network environment because the precise characteristics of the physical transport cannot be controlled. However, when multiple of multi-GHz serial I/O channels are used on boards and backplane, the physical characteristics must be known and controlled to ensure flawless reliability, the best performance, and future interoperability.



This environment requires end-to-end mixed signal simulation at the system level. It must comprehend the entire data path starting from the digital domain on one chip, through the serializer, PHY, through the package, board traces, connectors, and then back through another PHY to the digital domain on a different chip. Designers must be able to plug in different PHY models, different FR4 and connector characteristics, and simulate the transmission of a heavy stream of real data packets, staying within fixed analog timing budgets with minimum bit error rates. Then, using the same methodologies, they can synthetically route their boards within rules and limits, and re-verify.

Recognizing this need, Mentor Graphics has developed a tool that is more than an integration of HSPICE and IBIS into the front end of a digital design environment. This tool is an entirely new and comprehensive system-level development toolset based on a combination of knowledge and experience with digital, mixed signal, analog, RF and electro-magnetic design, simulation, and verification tools. The toolset comprehends mechanical, thermal, and other physical dynamics equally well. However, these are not immediately applicable to high-speed serial interface design.

Many companies are cautious about large-scale system design using high-speed analog interfaces, fearing they may repeat the costly mistakes made by others in the past. This concern requires the best available tools and resources to ensure that designs are correct, end-to-end.

## Conclusion

High-speed I/O and connectivity allow digital systems to achieve breakthrough performance processing. Serial interfaces will result in better performance scalability, interconnect density, pin count reduction, cost containment, and greater capacity.

Parallel interfaces will remain essential to the core of traditional monolithic uniprocessor or multiprocessor platforms. HyperTransport, RapidIO, IMB, and other interfaces are peers to the CPU bus, DRAM buses, etc., and will dominate in this area. However, in time, nearly all the other I/O grade interfaces will become serial.

In market segments without a pressing performance requirement, these transitions could take many years. However, for applications that demand performance and flexibility, serial interfaces will be deployed quickly, in complex configurations, while achieving system performance levels that have not been previously imagined.

---

---

## Revision History

The following table shows the revision history for this document.

Date	Version	Revision
05/15/03	1.0	Initial Xilinx release.