

The Advanced Telecom Computing Architecture specifies a modular approach to building complex communications equipment, housed in enclosures, called shelves, such as this one. Variations on this shelf include the number of slots, the back-plane type, cooling schemes, power-supply capacity, module orientation, rack height, and compliance with various standards (courtesy PICMG).



## PCI Express and ASI: AVOIDING CHAOS AT 2 Gbps

**AT DATA RATES AS HIGH AS 2 GBPS PER LANE, LITTLE THINGS CAN MAKE BIG DIFFERENCES IN SIGNAL INTEGRITY. EXPERTS DISAGREE ON WHETHER THINKING ABOUT WHAT'S REALLY GOING ON CAN IMPROVE PERFORMANCE OR ONLY MAKE IT WORSE.**

**P**CIEXPRESS (PCIe) is rapidly becoming the dominant standard for high-speed data transmission in PCs, notwithstanding the existence of myriad other high-speed point-to-point serial protocols, most of which predate PCIe. What's more, elements of PCIe,

particularly the physical layer, may become even more pervasive as part of ASI (Advanced Switching Interconnect), an ambitious protocol targeting applications that require performance higher than that of typical desktop and laptop PCs. The ASI SIG (Special Interest Group) names communications, storage, servers, and embedded systems as target markets but does not limit ASI's potential to those markets.

ASI appears to bring many benefits to high-speed data-transmission. For example, unlike the host-centric PCIe, ASI enables peer-to-peer communication among a system's peripheral devices—or among CPUs in systems that incorporate multiple CPUs. Because it allows multiple conversations to take place simultaneously among these devices, the peer-to-peer feature can greatly speed information flow in complex systems. Other ASI features that PCIe doesn't support include QOS (quality of service); protocol tunneling, which enables ASI-based systems to support multiple protocols; and the fabric switching that is at the heart of ASI.

Nevertheless, according to its detractors, ASI faces stiff competition from older protocols, such as RIO (RapidIO)—especially with the new RapidFabric extensions, InfiniBand, FibreChannel, and a sizable list of proprietary protocols that have established sub-

stantial bases of users, most of whom allegedly see no particular advantage in changing horses for the foreseeable future. Many in the industry regard Ethernet as a key competitor, but the ASI SIG's position is that that long-lived protocol complements and does not compete with ASI.

Current single-lane PCIe implementations transmit data at a maximum rate of 2 Gbps in each direction. The raw data rate is 25% higher—2.5 Gbps—because PCIe, like most high-speed-serial protocols, uses 8B/10B (8-bit/10-bit) coding to embed the clock within the data stream. However, both PCIe and ASI allow the use of multiple lanes, which can increase data-transfer rates by factors as large as 32. Already, advanced video subsystems in some premium-priced desktop PCs that target video gaming are using unidirectional 16-lane-wide versions of PCIe to transmit data at 32 Gbps—equivalent to 4 Gbytes/sec. Despite such mind-numbing speed, these desktop systems don't require ASI's advanced features, so a unidirectional version of PCIe is the appropriate choice.

#### RAPID DEPLOYMENT

Manufacturers of PCs and peripherals are deploying PCIe at a rate that, figuratively speaking, rivals the rates at which data travels on the bus (**Reference 1**). And these manufacturers seem undeterred by difficulties in obtaining from the PCI SIG certain detailed information that could come in handy to engineers creating new equipment designs. A key reason that gaps in the information that the SIG has so far provided have not hindered adoption of PCIe is the voluminous information available from the PICMG (PCI Industrial Computer Manufacturers' Group). PICMG's role complements the roles of protocol-specific industry groups, such as the PCI SIG and the ASI SIG. Although PICMG defines no communication protocols, it does publish system-physical-architecture specifications that enable designers to successfully implement protocols that other groups have defined. Moreover, PICMG member companies provide the hardware for these implementations. Recently, PICMG completed what it calls the largest spec-writing effort in its 11-year history. The result is a group of specifications collectively describing the modular ATCA (Advanced Telecommu-

#### AT A GLANCE

- ▶ High-speed serial PCI Express, which is rapidly becoming the dominant technology for interconnection in PCs, stands an excellent chance of much broader application because its wide use will hold down component costs.
- ▶ PICMG standards for physical implementation of many communication protocols can take much of the guesswork out of the design of systems, including those based on PCIe and the related ASI.
- ▶ Adaptive equalization makes practical ultrahigh-speed-serial-bus designs that might otherwise perform unreliably—even after tedious manual tweaking.
- ▶ Industry experts disagree not only on how best to characterize components used in implementing PCIe, but also on whether a cookbook implementation approach invites or avoids problems.

nications Computing Architecture).

Among many others, the PCI SIG's PCIe specs and the ASI SIG's ASI specs complement the ATCA specs—although PICMG prefers to call the protocol specs subsidiary to ATCA. Versions of the ATCA spec define how to implement the architecture with different communication protocols. One spec covers both the ASI and the PCIe versions. This commonality should come as no surprise because the two protocols share both their physical and data-link layers, although ASI's switch fabric will almost certainly require module-resident hardware that PCIe doesn't use. What may come as a surprise, though, is that PICMG has so far found that one card-cage design, one backplane design, and one module format are sufficient to implement versions of ATCA that support at least half a dozen communication protocols.

Module manufacturers have even developed prototype plug-in ATCA-format line-card modules that they can adapt to any of several communication protocols simply by loading the appropriate software. However, according to Gilles Garcia, director of marketing for switched-fabric products at AMCC (Applied MicroCircuits Corp), line cards built around network processors need more than software to adapt to ASI or RIO/

## DON'T LET FIXTURE PERFORMANCE INVALIDATE CONNECTOR MEASUREMENTS

By Orlando Bell, GigaTest Labs



**Figure A** You can easily measure the properties of interconnects, such as differential traces on a pc board, without introducing impedance discontinuities, such as those that can occur with SMA connectors.



**Figure B** A flexible probing system allows accurate measurement of a wide variety of interconnect hardware.



**Figure C** Instrument calibration at the probe tips uses a precise impedance-standard substrate.

The PCI SIG (Special Interest Group, [www.pci-sig.com](http://www.pci-sig.com)) has published the “PCI Express-Connector High-Speed Electrical-Test Procedure.” This document describes the methodology for measuring the performance of PCI Express connectors to 10 GHz, using a custom test fixture and a four-port VNA (vector-network analyzer). You can also use a more accurate and flexible probe-based methodology for such high-speed serial-interconnect measurements.

Fixture-based validation of connector designs is desirable for several reasons. First, the drive for industry standards works best when all interested parties avoid correlation problems by using consistent test hardware. In addition, using SMA connectors simplifies implementation of the fixturing interface (Figure A) because SMAs are the de facto standard connections for high-frequency measurements.

Unfortunately, in many cases, connectors and traces on high-speed test fixtures can make the fixtures’ performance inferior to that of the devices the fixtures are supposed to help characterize. Often, a fixture exhibits signal loss greater than that of the device under test. If the device is inherently a two-port element—that is, a device that has one input and one output—with little crosstalk between the ports, you can use T-matrix manipulation to subtract the fixture contribution. On the other hand, if the device under test has four ports that suffer from fixture-level crosstalk, de-embedding the fixture contribution becomes difficult. To overcome these fixture limitations, you must usually move the measurement reference plane as close to the device under test as possible. You can achieve this close spacing by using microwave probes (Figure B).

By using a calibration substrate (Figure C), you can de-embed the measurement system’s response all the way to the probe tips. In many

cases, you can directly probe the device; in the worst case, you can add vias to the board for use as probing points. Either way, the measured data usually requires little or no de-embedding.

When pins on which you make no measurements require termination, you encounter the sole disadvantage of using probes. Then, you must use SMT (surface-mount-technology) resistors to terminate the pins. The red plots in Figure D show a connector’s TDR (time-domain-reflectometer) performance as measured with a fixture that has SMA connectors plus a 6-in. stripline trace on an FR4 substrate. The blue plots show what happens when you make the measurements with a microwave probe. The 10 to 90% rise time of the TDR step, is 200 psec in Figure Da and 500 psec in Figure Db.

You can readily appreciate the fixture’s effect on the connector’s response: The fixture degrades the measured impedance and delay. Also, extracting an equivalent-circuit model of the fixtured connector requires extensive effort, because you must include the fixture in the response. Doing so adversely affects the model’s quality and bandwidth.



The use of fixtures is desirable to ensure compatibility and correlation, but you pay for this convenience with reduced measurement accuracy and reduced bandwidth in the device you are testing.

### AUTHOR’S BIOGRAPHY

Orlando Bell is vice president of engineering at GigaTest Labs in Sunnyvale, CA. His key responsibilities include managing the high-frequency-characterization lab and driving the research and development of advanced test hardware for the signal-integrity market. He holds bachelor’s and master’s degrees in electrical engineering from the University of Florida (Gainesville).

### Figure D

These impedance profiles of a connector measured with a fixture and a probe show that the fixture (red curves) can mask the impedance discontinuities you are trying to measure, whereas the probe (blue curves) presents a more accurate picture. Rise time is 200 psec from 10 to 90% of the step amplitude (a); rise time in (b) is 500 psec.




RapidFabric. These protocols require additional glue logic implemented in FPGAs. Garcia sees no motivation to develop new network processors that subsume this external logic because the marketplace isn't yet demanding line cards that support ASI or RapidFabric.

**PROMISING IMPLICATIONS**

Nevertheless, ATCA's implications appear promising indeed: In telecom central-office equipment, the architecture appears poised to become a unifying force that will enable competing companies to produce physically similar hardware. This unification will bring about component-level economies of scale that help to control costs and drive product differentiation to the system-software level, at which many believe it belongs. Indeed, ATCA visionaries see the architecture's transcending the telecom market space and invading such computing areas as clustered servers, producing further economies of scale.

While it was busily at work on ATCA, PICMG scarcely ignored less complex systems. Always mindful of its roots in the older, parallel versions of PCI, the or-



**A close look at the mating side of ZD connectors, which the ATCA standard uses for high-speed differential signals, such as those that the PCIe and ASI buses carry, reveals the careful shielding between adjacent differential pairs (courtesy ERNI).**

ganization has focused on updating architectures that it pioneered for smaller systems. The best known of these is CompactPCI. PICMG is nearly ready to announce a specification for a PCIe-based version of this architecture. The group expects the old and new CompactPCI architectures to coexist for many years as the PCIe-based architecture gradually supplants the parallel version in new equipment designs. PICMG is also specifying a 95x125-mm PCIe-based mezzanine-card format that not only can supplant earlier mezzanine cards, but

also may replace PC-104 boards in some embedded-system applications. Many people think of PCIe as merely a bus for interconnects on and between pc boards within a system enclosure. Also, despite the multilane configuration that high-end video subsystems use, most applications use a single-lane, bidirectional configuration comprising two signal pairs. Still, needs exist to run PCIe through cables and to place multilane versions in backplanes. For the backplane application, which will be common in ASI-based systems, designers have a choice of connectors that are mechanically and electrically acceptable. Prices vary, but some of these parts deserve to be called downright inexpensive when you consider their amazing electrical performance. At press time, however, the PCI SIG had not named one design as the "official" PCIe backplane connector, although PICMG's ATCA specs call for the ZD series, currently available from Tyco/Amp and ERNI. Though ATCA does not reference them, other backplane connectors, including Teradyne Connec-

**FOR MORE INFORMATION...**

For more information on products, services, and design standards such as those discussed in this article, contact any of the following trade organizations and suppliers directly, and please let them know you read about them in *EDN*.

**TRADE ASSOCIATIONS AND INDUSTRY GROUPS**

**ASI SIG (Advanced Switching Interconnect Special-Interest Group)**  
www.asi-sig.org

**ATCA (Advanced Telecom Computing Architecture)**  
www.advancedtca.org

**InfiniBand Trade Association**  
www.infinibandta.org

**PCI SIG (Peripheral-Component Interconnect Special-Interest Group)**  
www.pcisig.com

**PICMG (PCI Industrial Computer Manufacturers' Group)**  
www.picmg.org

**RapidIO Trade Association**  
www.rapidio.org

**SCSITA (Small Computer Systems Interface Trade Association)**  
www.scsita.org

**VITA (VMEBus International Trade Associations)**  
www.vita.com

**SUPPLIERS OF ICs AND INTELLECTUAL PROPERTY FOR EMBEDDING IN ICs**

**AMCC (Applied Micro Circuits Corp)**  
www.amcc.com

**Intel Corp**  
www.intel.com

**StarGen Inc**  
www.stargen.com

**Synopsys Inc**  
www.synopsys.com

**Xilinx Inc**  
www.xilinx.com  
**SUPPLIERS OF CONNECTORS, CABLES, AND RELATED COMPONENTS**

**ERNI Electronics Inc**  
www.erni.com

**FCI**  
www.fciconnect.com

**ITT Cannon**  
www.ittcannon.com

**Molex**  
www.molex.com

**Teradyne Connection Systems**  
www.teradyne.com

**Tyco/AMP**  
www.amp.com

**WL Gore and Associates**  
www.gore.com  
**SUPPLIERS OF TEST EQUIPMENT, DEVELOPMENT TOOLS, OR CONSULTING SERVICES**

**Agilent Technologies**  
www.agilent.com

**CATC (Computer-Access Technology Corp—subsidiary of LeCroy Corp)**  
www.catc.com

**GigaTest Labs**  
www.gigatest.com

**NESA (North East Systems Associates)**  
www.nesa.com

**TDA Systems Inc**  
www.tdasystems.com

**Tektronix Inc**  
www.tektronix.com

Some companies offer products or services in more than one of the listed categories; the indicated category is the one most closely related to the material in this article.

tion Systems' GbX series, are electrically suitable.

And, although rumors exist that the PCI SIG has made an internal choice of a cable connector, no official announcement has emerged, and the vendor whose design the SIG is rumored to have selected refuses to make any comment—or even to provide information about the design. Purportedly, the need to resolve licensing issues has caused the announcement delay. Nevertheless, Rajeev Kumar of Intel, chairman of the ASI SIG, says that he has observed good results with routing of PCIe signals over cables as long as 7m that were designed for other standards. He mentions Molex cables that were designed for SAS (serial-attached SCSI) and Gore CX4 cables that were designed for 10-Gbps Ethernet. He adds that, in some cases, you can obtain electrically and mechanically equivalent cables from several sources.

#### A BIT EASIER THAN YOU MAY THINK

Just the idea of transmitting data error-free at 2 Gbps and more can be intimidating. Nevertheless, the task, though far from trivial and not to be approached lightly, may be easier than many people suspect. Not too many years ago, the only engineers who would even consider dealing with such high frequencies were microwave engineers. However, whereas the frequencies in PCIe data streams and microwave circuits are comparable, the technologies have less in common than you might think.

The reason is that the serial buses carry information in digital form, whereas classical microwave designs are traditionally analog. Despite the millivolt signal levels in the differential pairs that carry information on buses such as PCIe, you generally need to know the voltage levels with only about 20% accuracy; the voltage differences between the one and the zero levels are usually greater than the absolute voltage on either line of the differential pair. In other words, the buses work because they were designed to tolerate errors.

As a result, a reasonable amount of ringing after a one-to-zero or zero-to-one transition is not usually fatal. Still, discontinuities and reflections in transmission paths can and do cause errors, so you must design pc boards, connectors, backplanes, and cables to minimize crosstalk and to maintain a constant-im-

pedance environment with minimal discontinuities, and you must also minimize delay differences that cause skew between signals that are supposed to reach receivers simultaneously.

In addition, when you work with signals at several gigahertz, the dielectric losses in pc substrates and cables become appreciable, attenuating signal swings and slowing (lowpass-filtering) signal edges. In all probability, without the introduction of pre-emphasis into the circuits that drive high-speed serial buses, the design problem would be orders of magnitude more complicated, and data reliability that designers now take for granted would be unachievable.

#### ADAPTIVE EQUALIZATION

The object of pre-emphasis is to shape the edges of signals that leave the driver so that the distortion (lowpass-filtering) introduced by the dielectric losses produces nearly ideal waveforms at the receiver input. But, because the amount of distortion depends on material properties and the length of signal paths, the tuning problem could be intractable were it not for *adaptive* equalization. You can find adaptive equalization in some IP (intellectual-property) cores for ultrahigh-speed serial buses. At 2.5 Gbps, the current generation of PCIe is too slow to require the technique, but if the data rate reaches 5 Gbps, future generations may.

Adaptive equalization simplifies the task of designing pc boards, backplanes, and cables that work reliably at multigigabit-per-second data rates. And it can do so without requiring tedious manual tuning or the use of pc substrates and cable dielectrics made of low-loss materials whose costs greatly exceed those of familiar materials. However, even with the use of adaptive equalization, ultrahigh-speed designs still require extreme care. For example, unless the component-level circuit models accurately describe the component behavior, no simulation of a multigigahertz design can accurately predict the performance of the complete circuit. In the area of how best to characterize and model components such as connectors, backplanes, and cables, the industry's signal-integrity experts seem to be embroiled in a controversy (see **sidebar** "Don't let fixture performance invalidate connector measurements").

To quickly and economically get new designs to market, designers need rules

that they can apply easily and without a lot of thought. To this end, the PCI SIG and its member companies have gone to unprecedented lengths and great expense to provide component characterization and design-rule information that engineers must consider in their designs.

Nevertheless, some signal-integrity consultants, such as Edward P Sayre, PhD, president of NESAs (North East Systems Associates Inc), say they are concerned that a few designers may interpret the SIG's approach as, "We've done all the thinking and encapsulated the results in rules; all designers need to do is follow the rules without thinking." Says Sayre, "That's not what the SIG intends, but if that's what some engineers believe, the ones who haven't given their designs adequate consideration are going to wish that they had taken a little more time to understand why their circuits and layouts behave as they do." Sayre describes the SIG's work as impressive, but he says that years in the field have taught him that, although design rules are invaluable, they can't replace an understanding of the underlying circuit behavior. Moreover, the most troublesome problems seem to occur when designers fervently, albeit erroneously, believe that they have followed all of the rules but haven't developed an understanding of the rationale behind those rules. □

#### REFERENCE

1. Bhatt, Ajay V, "Creating a PCI Express Interconnect," [www.pcisig.com/specifications/pciexpress/resources/PCI\\_Express\\_White\\_Paper.pdf](http://www.pcisig.com/specifications/pciexpress/resources/PCI_Express_White_Paper.pdf).

#### AUTHOR'S BIOGRAPHY

*Contributing technical editor Dan Strassberg, who has covered test and measurement for EDN since 1987, occasionally ventures into other product areas. For example, this article started out as a passive-components story. However, the larger issues of how to properly think about high-speed serial buses soon grabbed the spotlight. Strassberg holds a bachelor's degree in electrical engineering from Rensselaer Polytechnic Institute (Troy, NY) and a master's in electrical engineering from the Massachusetts Institute of Technology (Cambridge).*

#### TALK TO US

Post comments via TalkBack at the online version of this article at [www.edn.com](http://www.edn.com).