

PCI Express brings the express lane to the information autobahn

By Tom Skrobacz

In the world of real-time systems, standard bus architectures including PCI, VME, PXI, and CompactPCI are fundamental to addressing a broad spectrum of high performance applications with modular instruments. Data movement within a real-time system is throttled by the slowest link in the overall architecture. While performance has been widely addressed for CPU, disk drive, RAM, and network architectures, the data bus available in standard motherboards continues to rank as the 25 MPH zone of the information super highway. As our physical highways evolved, engineers developed clover leaves, tunnels, and overpasses to eliminate bottlenecks on the road. Similarly, the future, with emerging standards like PCI Express and PXI Express, offers a breath of fresh air to the petabyte world that is quickly unfolding. In this article Tom delves into the mechanics of PCI Express and PXI Express. Time to put the Model T's in the museums – the information autobahn has arrived.

Digital data recording and playback promises blazing speeds when instruments, sensors and storage media catch up with the new standard

The PCI bus has long been the standard for expanding general purpose PCs to enable a broad range of I/O, storage, and networking capabilities. At its inception, the PCI bus provided an impressive 132 MBps performance in a 32-bit architecture. This

performance, however, is only a theoretical limit. In practice, the PCI bus is designed to be shared among multiple devices. Just as traffic lights regulate the flow of interconnected roadways, a bus manager, or arbitrator, ensures equitable use of the PCI bus by one or more competing resources. The arbitrator, by itself, results in a 20-25% overhead and results in performance degradation. The practical limit on the 32-bit bus is therefore significantly slower, or less than 100 MBps.

Over the years, chip designers have been able to compensate for the overhead on a shared PCI bus by developing new offerings like 64-bit boards and higher clock rate versions of the PCI bus. The latest PCI-X standards offer an impressive 4.3 GBps top end theoretical performance. While each successive standard provides higher aggregate bandwidths, system architects still have struggled to overcome the inherent complexities of shared bus architectures. In real-time systems where data buffering can cause disastrous data loss, the inability to guarantee a sustainable throughput forces designers to significantly over-design solutions. Still, the mighty bus arbiter has the final say in who gets the available resources.

PCI Express offers to provide an HOV lane on the information superhighway by dedicating *lanes* for high-speed traffic within a computing environment. The concept is to set up point-to-point, dedicated data lanes. The arbiter in the PCI world is replaced by a root complex that is designed to establish true point-to-point data links, and then remove itself from the equation. The results are an astounding 20 GBps data bus and zero overhead. PCI Express paves the way for Autobahn performance for the most demanding real-time applications.

Real world, real-time improvements

In real-time environments, PCI Express affords a clean pipeline for data with its point-to-point architecture. With dedicated, uninterrupted high-speed bandwidth to each individual device, or lane, and zero requirements for bridging, applications can extend over significant distances without suffering performance degradation.

The PCI Express specification will preserve software investments, and standard PCI and PCI Express cards can operate in the same

system. Module manufacturers can supply instrument and I/O controllers that plug into PXI, PCI Express, or both buses.

Hybrid backplanes in the PXI environment provide a convenient migration path for mixed technology environments. Chassis' with hybrid backplanes allow both current PXI and PXI-Express boards to co-exist. With the software interface largely unchanged, this architecture opens the doors to significant expansion while preserving current PXI technology investments.

PCI Express includes impressive expansion possibilities. The existing PCI bus can be expanded from one chassis to another with cables that can be fibre-linked. Systems with 150 I/O boards can be built, and additional chassis can be connected with 4-wire bus cabling at distances of up to 7 meters. This 4-wire bus is referred to as a PCI Express Lane. One lane connected between two PCI Express termination devices is called a *by one*, or written as an x1 link. More rapid interconnect speeds may be achieved by utilizing multiple lanes in parallel between PCI Express and termination devices such as sensors and recorders. Links can be increased from x4 to x16, which leads to performance ranges from 5 to 80 GBps total bandwidth.

At the top end, PCI Express can allow for performance at 4 to 5 GBps in shared bandwidth or hybrid applications (4 to 8 lanes) and up to 20 GBps (first generation, 16 lanes) for dedicated applications (or a 500 percent performance improvement). That is 4 to 5 times faster than today's peripheral technology supports. There's enormous opportunity for growth, as PCI Express replaces shared bus architectures with point-to-point functionality. Furthermore, PCI Express continues to capitalize on the high performance and low cost of PC technology. Generation 2 will double the clock rate to 5 GHz and magnify performance to the next level.

PCI Express is a highly scalable technology as it readies the market to take advantage of Field Programmable Gate Arrays (FPGA), which can bridge the distance between legacy buses and PCI Express. Hardware such as disk drives, controllers, sensors, and frame grabbers can easily capitalize on low cost 64- and 128-bit FPGAs with faster clock speeds of 132+ MHz as we enter the petabyte world of massive storage potential.

Very Long Baseline Interferometry (VLBI), used in radio astronomy applications to triangulate data from outer space has an ever-increasing requirement by the scientific community to record at faster speeds. Currently at 1 GBps with up to 72 hours of continuous digital data recording, VLBI can reach a new plateau when sensors operate at GBps and do so for more extended periods.

Standard bus architectures including PCI, VME, PXI, and CompactPCI, in relative terms, have held data transfer rates at the 25 MPH zone of the information super highway. Shared bus arbitration has restrained the PCI bus to practical limits of about 100 MBps. Today, PCI Express, with a 64-bit bus and higher clock rates, offers an impressive theoretical threshold of 4.3 GBps top end. Tomorrow, with the new bus architectures, speeds will approach 20 GBps.




Figure 1




Figure 2




Figure 3

When perpendicular recording emerges in 2006, the higher densities of 230 GB/square inch will result in terabyte disk drives. Then, companies such as Conduant Corporation, can move toward implementing 1 GBps read/write capabilities by harnessing to the PCI Express lane. Their StreamStor direct-to-disk recording technology will be able to serve a myriad of new data acquisition applications in homeland security, defense, instrumentation, and scientific research.

In 2001, when terrorists ploughed through the Pentagon with a commercial airliner, much potential data was lost because slow security camera imaging was maxed out at 30 frames per second. Faster frame rates at a higher resolution would have caught valuable data between the plane approaching the structure and the fireball. Today, Conduant's partners are reaching for 2,000 frame/sec performance. Data transfer rates of 20 GBps for imaging, defined by the standard committee, are becoming possible because proposed disk controllers, frame grabbers, and storage devices will be married to the PCI Express bus with its point-to-point architecture.

Conduant is involved in designing and building real-time, long duration data acquisition systems for applications requiring ultra high-speeds. Core expertise for the company is moving data streams onto disk drives for infinite duration, and as fast as the physics will allow.

PCIExpress broadens the application set so StreamStor technology can perform to its full potential. Its most recent Amazon disk controller has a dedicated daughter card that doubles what can be done in a shared environment. The configurations (shown in Figures 1-3) show improvements from 160 MBps to 300 MBps with configurations based on PCI, LVDS, and PCI Express. Coupled with the LVDS protocol, which is a low noise, low power, low amplitude method for high-speed data transmission over copper wire, the potential is limited only by the performance of storage subsystems.

Many industry segments are poised at the onramp to the PCI Expressway, anxious for the speed limits to disappear. New applications will emerge as the dedicated bandwidth of PCI Express allows the petabyte world to open up. As we see subsystems riding this emerging bus standard, new solutions and capabilities will unfold. The information autobahn has truly arrived. **PXI**

Thomas Skrobacz, vice president of business development for Conduant, is currently responsible for worldwide sales and marketing. Before joining Conduant, Thomas served at IBM for a 16-year period. While with IBM, he pioneered software strategies to deliver distributed mainframe systems software worldwide and developed security standards for migration to the Internet. Thomas developed and deployed technologies, policies, and procedures to conduct the single largest systems integration test ever performed at IBM. He is a certified project management professional and an award-winning entrepreneur. Thomas received his B.S. in computer science from the University of Maryland at College Park.



Conduant Corporation

1501 So. Sunset Street, Suite C • Longmont, CO 80501
 Phone: 303-485-2721
 Website: www.conduant.com • Email: sales@conduant.com