The Next Big Thing - Moving Beyond CompactPCI

Joe Pavlat
President, PICMG

When PICMG released the first CompactPCI specification in 1995, the PCI bus was viewed as a high performance mechanism for moving and handling large quantities of data. The PCI bus could transfer data at over 130 Mbytes/second in its 32 bit, 33 MHz version and promised 500 Mbytes/second in the near future using 64 bit transfers at 66 MHz. Compared to the ISA bus and VME, with transfer rates of a few to a few dozen Mbytes/second, PCI was fast. It was supported by virtually every microprocessor and peripheral chip supplier and was truly the first universal chip-to-chip and board-to-board interconnect. CompactPCI incorporated PCI silicon, and the software to run it, in a robust and modular form factor.

As CompactPCI grew from a single idea to an entire family of interrelated specifications, seventeen at last count either ratified or under development, a funny thing happened. Starting with the addition of the H.110 Telephony Bus to CompactPCI in 1997, the PCI bus itself began to morph from a heavy duty data hauler to more of a command-and-control bus with other specialized buses added to the backplane to handle data transfer. H.110 provided a time division multiplexed, isochronous bus tuned for telephony – something for which PCI was never intended. The addition of H.110 also reflected another dynamic, that being the acceptance of the CompactPCI form factor by mainstream telecommunications equipment providers for certain classes of applications. Vendors rushed to provide communications products based on CompactPCI for the same reason Willie Sutton robbed banks – because that's where the money was.

The move to add additional data buses within the CompactPCI framework has been accelerating. The PICMG 2.16 subcommittee, led by Performance Technologies, has worked very hard over the last year to develop a comprehensive specification for incorporating a switched Ethernet fabric in the CompactPCI form factor. I expect this will have been ratified by the PICMG membership as this article goes to print and it appears that it will be widely deployed. Earlier this year, the PICMG 2.17 subcommittee was formed to develop standards for adding the Starfabric switched fabric architecture to

CompactPCI. Over 50 companies are participating in this effort, making it PICMG's biggest effort to date.

So what's going on here? Several things. First, the fundamental desire to move data about faster and faster is clashing with the laws of physics head on. Parallel buses like PCI are no longer able to increase in speed very much, as every connection to the bus, be it a chip or a board, adds capacitance, and that slows the maximum signaling rate. PCI-X, with its allowance for just one peripheral slot at 133 MHz in its desktop configuration, may be the practical end of the line for multi-drop data buses.

To go faster, designers are moving towards what are commonly known as switched fabric architectures, wherein a single source of data is connected to a single destination for data at any given moment. Switch elements control the traffic between large numbers of sources and destinations in a dynamic fashion, and are often very sophisticated chips in their own right. At last count, several dozen switched fabrics have been announced, with several, including Infiniband, Starfabric, and switched Ethernet, moving into production. Fabric-based architectures are ideal for building high availability systems, as individual faults can be relatively easily isolated.

The PICMG 2.16 and 2.17 efforts are good examples of the industry adding new fabric architectures to a proven platform. CompactPCI is continuing to evolve and improve it performance and will remain a solid solution for large classes of communications applications for many years to come. It is well understood, well accepted by customers, and there are many hundreds of suppliers worldwide.

PICMG 2.16 offers aggregate throughput on the order of 20 Gbits/sec, which is approximately twenty times higher than 32 bit 33 MHz CompactPCI systems and spans more than twice as many slots. PICMG 2.17 uses first generation StarFabric technology to provide aggregate throughput of 50 Gbits/sec, and the potential for deploying independent 64 bit 66 MHz CompactPCI and H.110 bus segments at each slot. These enhancements are available in an environment that accommodates the CompactPCI legacy, and in a form factor particularly well suited to remote deployment. It's fairly clear that CompactPCI, which is admittedly maturing, nevertheless provides a cost effective platform with considerable remaining upside potential in terms of performance.

But is CompactPCI alone suitable for all of the high end communications applications over the next ten years? I think the simple answer is no.

Market pressures for technologies like "fiber to the curb" will require platforms capable of switching terabits per second in a single shelf, a level of performance one hundred to one thousand times higher that can be achieved today with off-the shelf products. The Telecom Equipment Manufacturers (TEMs) are demanding this level of performance going forward. They also cannot afford the luxury of spending years to design proprietary equipment as they have done in the past. The time is right for a new multi-vendor, open architecture that delivers this level of performance and can address the $100 billion market that exists right now.

CompactPCI has its own set of characteristics that limit its ability to reach the terabit per second switching rate that is becoming a requirement. As higher and higher raw data rates must be supported, new connector technology must be incorporated. The IEC 2mm connectors used in CompactPCI were not designed to support the multi-gigahertz, low voltage differential (LVDS) signals of many of the new fabrics. The majority of CompactPCI products use the larger of the two specified form factors, 6U by 160 mm, but even that board size is too small for large footprint chips, memory modules, and lots of other circuitry. The amount of power that can be delivered to a CompactPCI board is insufficient for many new designs, with single chips needing 50W and single board computer designs totaling 100-150 watts per board. The 0.8 inch pitch used in CompactPCI and VME is simply too small going forward, as the amount of cooling possible with forced air is well below those numbers. Mainstream CPU chips have been announced that simply will not fit on boards spaced on 0.8 inch centers. Heat sinks and tall memory modules add to the pitch problem.

It is time to create a whole new family of specs that are specifically designed to support the new switched fabrics and address the emerging requirement for open industry standards capable of delivering terabit performance. These specifications should be fabric-agnostic, as different fabrics have different "sweet spots" in the application space. The industry needs a standards-based approach to star and full mesh fabric interconnects, 2.5 Gbits/sec LVDS signaling, larger board area and pitch, and more power and cooling per board. The industry specifically needs to better understand the needs of telecom equipment suppliers, including NEBS and ETSI compliance, fiber interconnects, redundant and robust power and system management, and international requirements. And the industry needs to develop something quickly, using existing Eurocard mechanical standards if appropriate.

Is this realistic? I think so. A small group of companies began work in mid-2001 with the goal of creating an architectural definition for a fabric-based series of specifications to address these issues. A call for participation was issued to PICMG members on September 5th, 2001 , and the work will now move to the larger PICMG membership. PICMG has the breadth of experience in its members to develop significant and useful next generation standards for switch fabrics, and has the processes to do so fairly quickly. The current goal is to have a finished base specification and at least one subsidiary fabric specification completed by late 2002.

Does this mean the end of CompactPCI? Absolutely not. CompactPCI is a proven platform for a wide variety of communications, networking, medical instrumentation, industrial control, and military applications. It continues to evolve and improve its performance and applicability. But it can no longer be considered the highest performance platform available, as switched fabrics are re-defining what the high end means. That's a good thing, as PICMG members can now take much of the knowledge learned from CompactPCI and leverage it into a new, high end platform for entirely new classes of applications. That's a win-win situation for developers, vendors, and customers.