



PCI Express enables high-end embedded computing applications

By Jim Ison

E mbedded computing has various meanings to each company designing products for the “embedded” computing market. For every definition of “embedded” there is a definition of what is considered “high-end.” One thing typically remains the same when “high-end” and “embedded” are used together... there is never enough processing power to be considered “high-end” while remaining physically or thermally small enough to be considered “embedded.” This is especially true when companies choose to adopt industry standard architectures based on widely available commercial desktop silicon. The form factor may not fit the application, or the high-end computing power may be overly constrained by the thermal requirements.

As PCI Express continues the imminent replacement of PCI as the host add-in card bus of choice, high-end embedded computing applications are poised to take advantage of this quantum leap in technology through several industry standards. Many of these standard architectures from the PCI-SIG, PICMG, and VITA organizations are pending release in the coming quarter. These standards will aid the embedded designer with a multitude of issues associated with high-performance embedded computing. In this article, Jim explains new technological advancements of the PCI Express bus that will enhance the architecture choices of high-end embedded computing designers. The standards discussed in this article include PCI Express Cable, COM Express, and CompactPCI Express with an overview of SHB Express, XMC, and MicroTCA.

PCI Express, in the most basic sense, is *packetized* PCI transmitted serially over several transmission media. The media can be traces inside a backplane, motherboard, or add-in board, or over a twisted pair cable in many standardized mechanical form factors. It is ideally suited toward high-speed chip-to-chip, board-to-board, and box-to-box applications. PCI Express uses Low Voltage Differential Signaling (LVDS) to transmit the PCI packets over, in the most basic form, a four-wire bus running at a clock speed of 2.5 GHz. This four-wire bus is referred to as a PCI Express lane. The lane provides a total available bandwidth of 5 Gbps. A single lane between two PCI Express end point devices, along with any of the optional sideband signals for enhanced features, is called a x1 (*by one*) link. Designers can place several lanes between PCI Express end points in parallel to achieve higher bandwidth links of x1, x4, x8, and x16, yielding a range of 5-80 Gbps of total bandwidth. Recent PCI Express press releases by the PCI-SIG plan on doubling the clock rate of second generation PCI Express (Gen2) to 5 GHz beginning in 2006. That would yield data rates of 10-160 Gbps late next year.

In addition to the hardware portion of the specification, PCI Express is inherently backward compatible with PCI in regards

to operating system and application software. This compatibility allows the application and driver developer to use the same software tools used to develop PCI-based software. This is in contrast to the add-in card change from ISA/EISA to PCI that required new tools and operating systems.

PCI Express Cable

The first architecture to aid in high-end embedded applications is a PCI compatible cable expansion/extension capability based on PCI Express. PCI Express Cable is a standard undertaken by the PCI-SIG to transmit the host PCI Express bus over a high-speed cable. This can be done internal to a system enclosure or external in a box-to-box type application. Using a cable as shown in Figure 1, it is possible to extend the PCI Express bus approximately six to seven meters from the host CPU complex without the need for active equalization to suppress the inherent noise.



Figure 1

This particular cable is a x8 PCI Express external cable from Molex capable of transmitting 40 Gbps of data plus the PCI-SIG defined sideband signals.

Transmitting the host bus over copper cables opens a new world to the embedded designer. The PCI Express Cable enables a high-end computing core in a cooler area of a machine to host embedded I/O subsystems in remote, thermally constrained areas of the machine. The host and I/O system can be of different form factors suited to the location or performance each system requires. For example, a high-end, dual Intel Xeon class host system could provide the computing power for an operator interface and a high-speed data link to a high-end embedded I/O subsystem based on MicroTCA, PC/104, 3U CompactPCI Express, or proprietary form factor.

A compelling application of PCI Express Cable includes an expansion system, a set of products that extends the host bus of a system an arbitrary distance from the host enclosure to an expansion enclosure. This approach enables designers to insert more add-in boards into the system than the host system was originally designed for. A simple example of an expansion system is using a host interface board, cable, and 19-slot expansion chassis to extend a 4-slot ATX motherboard host system to a 20-slot system. Expanded systems in excess of 100 add-in boards are likely possible utilizing PCI Express expansion.

PCI Express Cable has a unique advantage over other expansion systems currently on the market. With PCI Express acting as

both the host bus and the cabled expansion protocol, it does not require drivers or conversion from the host bus to the expansion protocol then back again. This eliminates a root cause of some of the throughput latency of the expansion link. PCI Express offers a level of software compatibility and performance scalability unparalleled in even the most modern generation of cabled expansion systems currently on the market.

Other embedded applications for the PCI Express Cable are found across virtually all embedded markets. For example a high-speed docking station link for a high-end handheld or portable device useful in medical services, inventory control applications, or commercial laptops could employ PCI Express Cable. Another architecture a cabled solution could address is a noncontinuous backplane. This could take the form of several small backplanes in a nonconventional configuration, such as arranged in a circle or around a corner. In more traditional applications, an internal cable can replace the riser card of a 1U server where the add-in cards are mounted perpendicular to the motherboard.

COM Express

Another important standard is COM Express, which packs powerful PCI Express computing cores in small form factors for the embedded systems designer. COM Express is a PICMG effort to standardize PCI Express implementations of Computer-On-Module technology. COM Express standardizes two separate form factors and several different pin-outs, offering a choice to embedded developers.

Important features of COM Express include:

- Processor architecture independent
- Support for Gen1 and Gen2 PCI Express with two impedance controlled connectors
- 125 mm x 95 mm x 18 mm and 155 mm x 110 mm x 18 mm form factors
- Support for up to 32 lanes of PCI Express in several configurations
- Support for hybrid modules with a combination of PCI Express/PCI pin-outs
- Support for high-speed serial I/O and legacy parallel I/O
- Up to 160 W power budget per module

These modules allow embedded system designers to focus their core competencies on a carrier card that includes only the custom I/O functions required of the application. The designer can then attach the COM Express computing core module to the carrier card to form a customized embedded single board computer. The form factor and capability of the module proves useful in designing high-end handheld devices, custom shape carrier boards, and customized I/O carriers. The computing core of the carrier can be easily scaled to the application or upgraded with a new plug-in module, protecting the design from obsolescence.

Several PICMG member companies have announced COM Express modules and road maps. Figure 2 shows PFU Systems' basic form factor COM Express module.

CompactPCI Express

For the embedded systems larger than a handheld device the CompactPCI standard has undergone a transformation to CompactPCI Express. CompactPCI Express is a PICMG standard pending release in the second or early third quarter of 2005. The base standard, named PICMG EXP.0:




Figure 2

- Improves power delivery to individual CompactPCI Express cards
- Supports Gen1 and Gen2 PCI Express bandwidth with an improved connector
- Includes provisions for CompactPCI/CompactPCI Express hybrid systems

Base features of CompactPCI, such as user I/O pins, rear I/O transition modules, support for telephony buses, and base mechanicals remain in the CompactPCI Express standard. This means for a 6U CompactPCI Express card the J3-J5 mechanical attributes remain the same as the base CompactPCI standard.

In contrast, the J1 and J2 of CompactPCI are replaced with improved connectors. Power delivery is achieved using a 7-pin Universal Power Module (UPM). The UPM is capable of delivering over 400 W of power to individual cards. The high-speed PCI Express interconnect is achieved with a 3-row Advanced Differential Fabric (ADF) connector. Two ADF connectors are used to provide up to 120 Gbps of available PCI Express bandwidth with to the backplane. A mini enriched 2 mm hard metric (eHM) functions in several capacities depending on the slot in which it is used. The eHM is a keyed connector that can provide rear I/O in 3U card form factors, power to low power (<34 W) cards, PXI trigger signals, or geographical addressing. Switch cards, used to support larger numbers of CompactPCI Express cards through fan-out, have a five-position UPM and a card edge filled with ADF connectors to maximize the fan-out to additional PCI Express slots.

Early manufacturers of embedded systems utilizing CompactPCI Express expect to leverage the wide array of available 3U and 6U legacy CompactPCI and PXI I/O cards. This is accomplished with readily available CompactPCI chassis and newly designed hybrid CompactPCI Express/CompactPCI system backplanes. The placement of switches and bridges can include direct backplane integration, rear pallet bridges, or slot loaded switch/bridge cards. Hybrid systems with CompactPCI Express and CompactPCI from One Stop Systems are now entering the market.

SHB Express

System Host Board Express is the passive backplane PICMG 1.3 specification. SHB Express defines a new PCI Express host single board computer form factor to support the passive backplane PCI/PCI Express market.

Features of SHB Express include:

- 20 lanes of PCI Express and a PCI-X bus on the card edge connector

- A dedicated connector for USB, Ethernet, and Serial ATA routing to the backplane to reduce cables to the SHB host
- Increased power capability to the host board to support higher performance processors

Embedded systems based on SHB Express range from “shoebox” style systems to 1U servers less than 17 inches deep. These form factors prove useful in embedded machine control, SCADA systems, computer telephony, and military communication applications. The processing power of such systems currently reaches dual Xeon capability from One Stop Systems and other PICMG members.

XMC

Several other small form factor PCI Express architectures will prove useful to the high-end embedded system designer. A joint effort between PICMG and VITA is underway to upgrade the PCI Mezzanine Card (PMC) standard to handle PCI Express as well as other high-speed fabric signaling. The base standard is known as PICMG XMC.0 or VITA 42 in the respective organizations. Collectively referred to as XMC, the standard defines a small form factor for processors and I/O boards that follows the exact mechanical footprint of the PMC standard with the addition of a high-speed fabric connector. The board footprint remains the same as the PMC card at 74 mm x 149 mm for a single width card. A sample XMC is shown in Figure 3.

A combination of XMC, PMC, and processor-enabled versions of these standards delivers benefits to the high-end embedded designer similar to those

produced by CompactPCI Express or COM Express. XMCs are designed to enhance a CompactPCI Express system by adding functionality to a baseboard that is connected to an embedded backplane. Designers also utilize XMCs as standalone modules connected to a custom carrier card in an embedded system. Like COM Express, this carrier card can be application specific due to functionality or mechanical requirements. At 20 Gbps available bandwidth per XMC connector and with power consumption ratings from 7 W to 20 W, the XMC standard gives embedded designers a powerful tool.

MicroTCA

MicroTCA is a specification under investigation in the PICMG aimed at aiding the embedded designer. This specification is an extension to the PICMG Advanced Mezzanine Card (AMC) standard. AMCs are the mezzanine form factor of choice for the

AdvancedTCA specification due to advanced features such as high-speed switched fabric, hot-swap, and IPMI system management support. AMCs accommodate both processor and I/O functionality. The board area of an AMC is roughly the same as a 3U CompactPCI Express card but has several choices of interconnect fabric including PCI Express, RapidIO, or Ethernet. MicroTCA aims to adapt the AMC mezzanine standard into a standalone, embedded architecture with a high-speed serial fabric interconnect.

Conclusion

PCI Express will become a valuable tool for the high-end embedded systems designer as the standards begin to release over the next few months. CompactPCI Express and MicroTCA embedded backplane based solutions offer a standard, modular, front plug form factor design option for high-end processors in small areas. XMC and COM Express offer mezzanine/carrier form factors for flexible baseboard design. PCI Express Cable reopens a chapter on cabled serial buses with higher performance than was possible with RS-232/422/485 or USB. In addition, PCI Express Cable can be combined with any (or several) of the other form factors to add an extra dimension to the architecture of the high-end embedded system.

The rewards of increased performance and flexibility coupled with the abundance of form factors available in PCI Express comes at the cost of some added complexity. With XMC and MicroTCA, manufacturers have the option of choosing serial fabrics other than PCI Express. Compatibility between modules with different fabrics must be considered. Also, several form factors have similar features and size that make for challenging architecture choices.

Systems manufacturers certainly must accept a more consultative role in overall system design with many new PCI Express architectures to choose from. The availability of off-the-shelf development systems that are application-ready, integration services based on standards based building blocks, and fast system turnaround times become critical factors in choosing a manufacturing partner.

Jim Ison is the product marketing manager for One Stop Systems and has more than 10 years' experience in the bus-board marketplace. Prior to One Stop Systems Jim has held various sales and marketing management positions centered on industrial and converged communications accounts for Ziatech Corporation and Rittal Corporation. More recently, he has held the global positions of CompactPCI product manager and director of OEM business development with I-Bus. Jim holds a bachelor's degree in Aeronautical Engineering from California State Polytechnic University at San Luis Obispo.

For further information, contact Jim at:

One Stop Systems

2235 Enterprise St. #110

Escondido, CA 92029

Tel: 760-745-9883 x1647

Fax: 760-745-9824

E-mail: jison@onestopsystems.com

Website: www.onestopsystems.com

Ihr kompetenter Ansprechpartner:



BRESSNER Technology GmbH

Breslauer Str. 34

D - 82194 Gröbenzell

Tel. 08142 / 47284-0

Fax 08142 / 47284-77

E-Mail: info@bressner.de