PMC

XMCs: The popular PMC adds a fabric interface

By Andy Reddig
and Greg Novak

Traditional PMC modules have been the de facto standard for adding I/O expansion capability to embedded systems. The recently adopted standard for Processor PMCs enables the expansion of PMC module functionality to include processing as well as streaming I/O. As embedded systems transition from legacy parallel buses to high-speed switched serial fabrics, mezzanines are keeping pace through the XMC standard for switched fabric mezzanine cards. These new fabric-enabled mezzanine cards will allow the popular PMC form factor to continue its evolution, supporting the next generation of embedded systems with scalable, modular I/O and processing modules.

Introduction

The original PCI Mezzanine Card (PMC) standard emerged about 10 years ago as an outgrowth of PCI Local Bus popularity. With the wide availability of PCI silicon, PCI became the most cost-effective way of implementing modular I/O, but the traditional PCI card did not suit the embedded market's packaging or environmental requirements. The PMC standard repackaged the PCI electrical interface into a compact mechanical form factor, providing a low profile, rugged version of PCI that could facilitate I/O expansion in VME and CompactPCI systems.

The PMC form factor's modularity allowed manufacturers to build single board computers with application-specific I/O added through PMC modules. This approach also allowed OEMs to add leading edge I/O as it became available, so that refreshing the technology did not require replacing the computer.

As processors shrunk in size and power, it became possible to fit an entire processing subsystem (CPU, memory, networking, and PCI bridge) into the PMC form factor. These PMC capabilities became the genesis of the ANSI VITA 32 standard for Processor PMCs (PrPMCs). This standard maintains compatibility with legacy PMC modules but adds sideband signals to support using PrPMCs as an embedded system's primary processing resource. Leading manufacturers have now developed PrPMCs. Developers creating small, fast, and flexible computing solutions for telecom, industrial automation, medical imaging, defense, and aerospace applications are adopting PrPMCs.

Today PMCs and PrPMCs have exploded in popularity. Embedded PMCs and carriers in form factors ranging from VME, CompactPCI, and ATX to custom designs have been shipping for 10 years. The PMC ecosystem also includes more than 100 companies with modular PMC products covering all types of I/O and processing.

Evolution to XMCs

High-performance embedded computing has always pushed the envelope for computing and I/O technology, often with severe constraints on volume, power, and environment. As processor throughput and I/O bandwidths march inexorably upwards, new interconnect standards are emerging to support the insatiable demand for bandwidth.

With recent advances, parallel buses such as PCI have kept up with demand through performance enhancements. For example, VITA 39 added PCI-X capability to PMC modules. As performance requirements continue to rise, traditional bus-based products will be supplanted by switched fabric architectures using open standards such as RapidIO, PCI Express, or Infiniband. The XMC standard supports migration of legacy PMC modules to switched fabrics while maintaining backward compatibility with existing PMC modules and carriers.

The XMC standards committee started with the PMC standard and focused on adding switched fabric interconnects. Traditional PMC modules have two or three connectors (P1 through P3) for the 32- or 64-bit PCI bus interface and an optional fourth connector (P4) for user I/O. An XMC module adds one or (optionally) two new connectors (P5 and P6) to support high-speed differential signals for fabrics or user I/O. These new connectors implement the XMC switched fabric interconnect independently of the PCI bus interface; they even include power, ground, and basic control to enable an XMC card to function properly with just the fabric interface, allowing elimination of the PCI connectors. Figures 1 and 2 provide a better understanding of these updated mezzanines. Figure 3 shows an XMC layout.

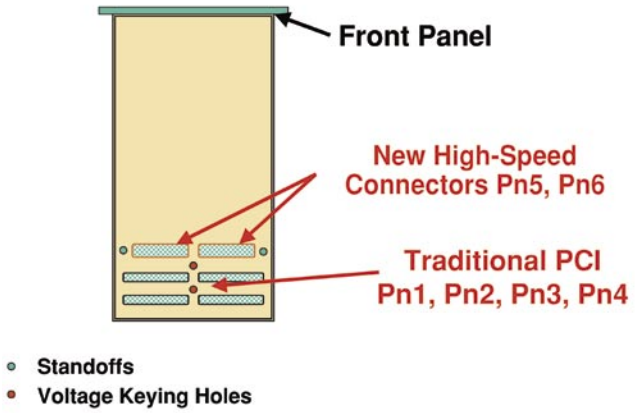

Figure 1

From the beginning, the VITA XMC standards committee has agreed upon a well-defined set of goals to focus discussions and quickly create this new standard. The overall vision calls for seamlessly adding two new high-speed connectors while maintaining backward compatibility with standard PMCs and PrPMCs, including their conduction-cooled alternatives. The challenges

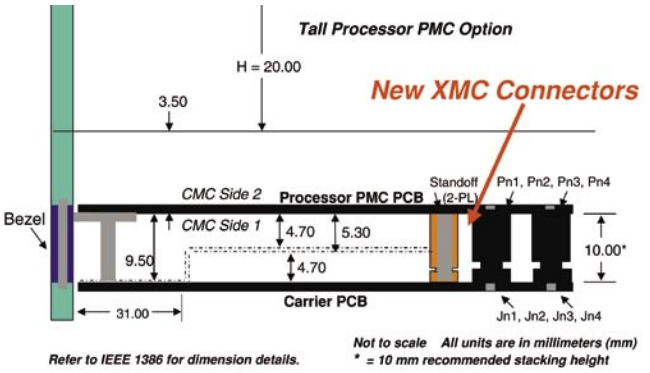

Figure 2

Figure 3

arise in making it backward compatible yet straightforward. Many standards have failed or languished due to proliferation of options. While options minimize committee debate, they confuse users, diffuse sales volume, and ultimately drive up cost. The committee is taking the time to debate issues in order to minimize signal and form factor options. Finally, in order to speed acceptance and release of this standard, benefiting from previous LVDS-compatible connector signal testing is key. A significant head start resulted from the collaboration with the RapidIO Trade Association (RTA). The extremely open discussions VITA and RTA held and the RTA organizations' giving VITA access to specifications and a detailed report summarizing connector evaluations resulted in rapid development of the initial XMC draft standard, ultimately benefiting the entire embedded community. The committee's goals are summarized as:

- Add multiple high-speed interface serial fabrics to PMC/PrPMC mezzanines.
- Leverage applicable research and specifications to speed completion of this standard.
- Utilize PMC/PrPMC form factor and tooling for 6U and custom applications.
- Continue compatibility with traditional PMC sites to control cost and speed adoption.
- Focus on compatibility with CompactPCI and VME.
- Minimize proliferation of options.
- Maintain conduction-cooled compatibility.

XMC compatibility with PMCs is key

Maximizing compatibility requires considering the effect of changes on both the XMC modules and XMC carrier cards. Initial XMC carriers will include both PCI and high-speed fabric connections, allowing carriers to seamlessly support traditional PMCs and the newer XMCs. In developing products for a new mezzanine and carrier standard, a *chicken or egg came first* dilemma occurs. PMC adoption taught that carrier sites likely come first, ultimately driving extensive mezzanine development. An XMC site will support both traditional PMCs and XMCs within the same real estate, enabling carrier vendors and OEMs to create XMC-capable carriers economically. As performance and availability dictate, legacy PMCs and PrPMCs can mix and match with high-speed XMC I/O and processor modules, preserving compatibility while providing a migration path to future switched fabric architectures. Ultimately, as the availability of XMC mezzanines increases, carriers will eliminate the PCI connectors and simply support high-speed fabric mezzanines.

The XMC standard, like its PMC relative, supports:

- I/O controllers as well as processor modules.
- Nonintelligent and intelligent carriers.
- Singlewide and doublewide modules.
- Front panel or rear I/O connections.
- Air-cooled standard and tall modules or conduction-cooled products per the VITA 20: PMC Conduction Cooling standard.

VITA 42: XMC Switched Mezzanine Card Auxiliary series of standards specify the detailed requirements of XMCs. VITA 42.0 (abbreviated XMC.0) is the base document that defines the mechanical interfaces, connector details, and common control signals. The supplemental protocol standards map specific switched fabric interconnects into the XMC framework. The first three protocol standards under development are:

- XMC.1 for Parallel RapidIO
- XMC.2 for Serial RapidIO
- XMC.3 for PCI-Express

Look for the very first modules and carriers to use the XMC.1 Parallel RapidIO embedded switched fabric offering at throughput rates from 1 to 2 Gbits/sec. XMC.2 and XMC.3 products will likely follow quite rapidly, using serial fabric technology to support data throughputs up to 8 Gbits/sec (4.0 Gbits/sec full duplex). Future protocol standards will define new switched fabric interconnects as market requirements dictate.

The schedule of any new standard helps drive deployment and adoption. The XMC standards committee is focused on rapidly completing the base XMC.0 standard (it is also a key input gating compatibility and wording of the protocol standards). XMC.0 has completed several reviews and ballots and it is quickly moving towards task group approval in the April 2004 timeframe, followed by ANSI approval in about four months. All three protocol standards have drafts available awaiting XMC.0. These three standards are expected to move almost as a group through standardization and follow the base standard approvals by about two months. XMC.3 has the added benefit of input from the PICMG-Express committee, where coordination of schedules is being coordinated to maximize XMC benefits for PICMG-Express and the entire embedded community.

Conclusion

The XMC standard is the next evolutionary step in mezzanine technology. XMC provides all the flexibility that made PMC successful, maintaining backward compatibility with legacy PMC modules. Architects are enthusiastic about this standards-focused, multi-vendor path supporting the popular switched fabric interconnects required for high-performance systems well into the future.

For further information, contact Andy or Greg at:

Motorola Computer Group
2900 South Diablo Way
Tempe, AZ 85282
Tel: 800-759-1107
Web site: www.motorola.com/computer

TEK Microsystems, Inc.
2 Elizabeth Drive • Chelmsford, MA 01824-4112
Tel: 978-244-9200
E-mail: info@tekmicro.com
Web site: www.tekmicro.com