By Lynn Brielmaier,
Application Engineer,
Communication Platforms Division, RadiSys Corp.

Bringing home the bandwidth: using CompactPCI for high bandwidth telephony applications

INTRO

CompactPCI, developed by PICMG (PCI Industrial Computer Manufacturers Group), a consortium of over 450 companies who collaboratively develop specifications that adapt PCI technology for use in industrial and telecommunications computing applications, is an open computing standard that combines the high-performance PCI bus with the rugged Eurocard format popularized by VersaModule-Europe (VME). Unlike many proprietary hardware solutions found in telecom applications, CompactPCI uses off-the-shelf PCI silicon, design tools, and software, thus lowering the overall development costs. This gives OEMs the ability to bring different product configurations to the telecom market sooner at a lower cost.

CompactPCI is desirable to the telecom industry because it marries VME form factors with PC architecture systems and the H.110 bus standard. Additionally, the technology has various layers of high availability technology building-blocks; hot-swap standard, system management, multi-computing, dual-host support, and failover management -- making it the ideal solution for today's embedded telephony applications.

The first layer which makes CompactPCI so desirable among telecom OEMs is the CompactPCI hot-swap standard. Equipment utilizing CompactPCI, with the hot-swap standard, has a practical advantage over the traditional AT-style systems. First, there is vertical front accessibility to the front panel for easy access. Secondly, the existence of rear I/O transition cards helps the system operator (the IT manager) of the telecom equipment have easy access to the system by allowing the removal or replacement of cards without touching the cables connected to the system.

Additionally, the CompactPCI hot-swap standard addresses three levels of hot-swap capability that a peripheral card and system may provide. First, there is basic hot swap. This standard lets the system operator safely insert and remove cards to a powered system with no software involvement. Next, full hot swap provides the signaling and protocol for coordinated insertion and removal of cards. With full hot swap, the operating system (OS) receives interrupts for hot swap events, including when the card is inserted and when it is about to be removed. The OS can then coordinate the use of the card. Finally, there is high availability hot swap. This standard offers a slight enhancement over the full hot swap standard. With this standard, network service provider can centralize the hot swap controller to force a local reset on a peripheral card. Additionally, the system operator can centralize the hot swap controller to control power on/off to a peripheral card and to monitor a "healthy" signal on a peripheral card.

SYSTEM MANAGEMENT

The second high availability infrastructure advantage that will help drive CompactPCIÒ technology into the telecom market is the adoption of IPMI (Intelligent Platform Management Interface) specifications (developed by Intel Corp., Hewlett-Packard, Dell Computer Corp. and NEC) and its system management. Using the I2C BUS, IPMI is the common interface of intelligent hardware that is used to monitor physical characteristics of the hardware (in this case switch, PBX, server, platform, etc.), including temperature, voltage, fans, power supplies and chassis. These monitoring abilities provide information that enables system management, recovery, and asset tracking that help drive down the total cost of ownership (costs outside of the initial software and hardware of maintaining a system) for the system operator and thus reducing system downtime.

The functionality of IPMI fits into a CompactPCI system by having a baseboard management controller embedded on the host single board computer (SBC) called an SBC management controller. This controller presents the IPMI standard interface to whatever management software may be running on the host processor. In addition to providing access to management data specific to the SBC, the SBC management controller also provides access to the IPMI management bus (IPMB). The IPMB bus is routed to all CompactPCI peripheral slots, thereby enabling the management of all peripheral cards. IPMB also allows for the management of other pieces of CompactPCI systems such as power supplies and cooling modules.

According to Intel Corp., " IPMI specifications give IT managers access to platform management information and control features that allow more accurate prediction of hardware failures, diagnoses of hardware problems and initiation of recovery actions."

IPMI gives network service providers the ability to determine the health of their equipment, whether it is running normally or it is in a non-operational mode. Platform management based on IPMI uses intelligent, autonomous hardware that remains operational even when the host processor is down, so the platform management information is always available. The IPMI interface enables telecom equipment to be accessed

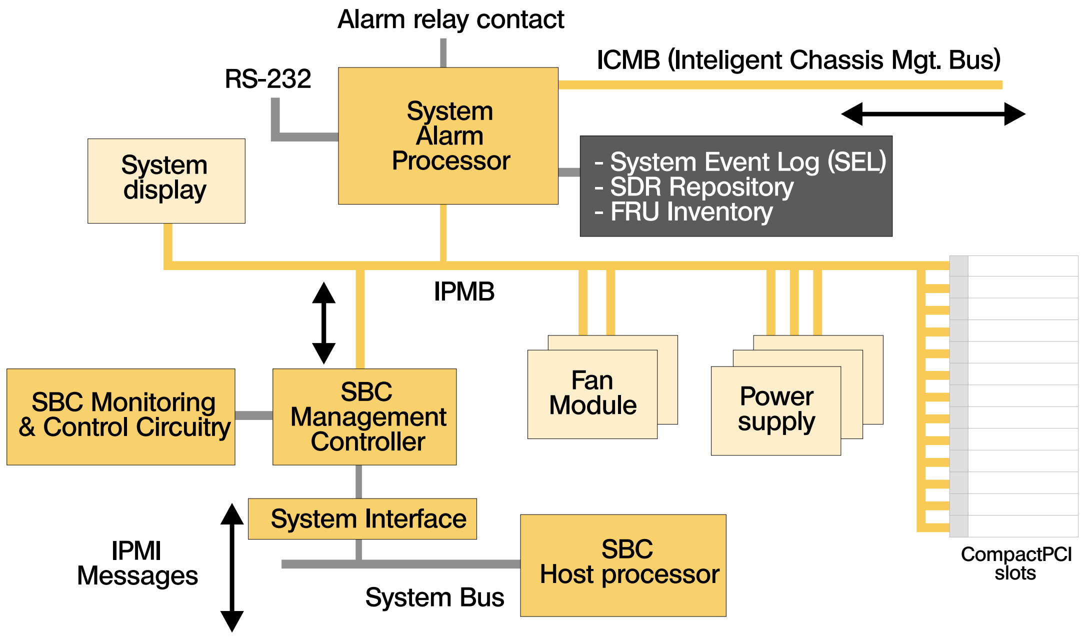

This typical system of basic managed modules found in an enterprise server class system contains an SBC, its processor modules, CompactPCI peripheral cards, two fan modules, three power supplies and a service processor (or alarm board).
The service processor has a connection to an enterprise-wide maintenance network. Each of the processor modules monitors and reports status across the Intelligent Platform Management Interface management bus to the SBC or to the alarm processor.

Figure 1. CompactPCI Platform Management Technology.

not only by management software, but also by third party emergency management add-in cards. This gives network service providers flexible and interoperable access to vital platform management information. System-to-system monitoring or management via a connected server is becoming increasingly important as system operators deploy complex system topologies such as clusters and rack-mounted configurations. In addition, the scalable nature of IPMI enables the architecture to be deployed across a server product line, from entry to high-end servers, and gives system operators a consistent base of platform management functionality upon which to effectively manage their equipment.

CompactPCI enhanced with IPMI thus gives the system operator proactive management of their equipment -- the ability to predict equipment failure problems (fault prediction) and correct them (fault prevention) before they occur, thereby increasing availability (system uptime).

MULTI-COMPUTING

Another advance in CompactPCI technology is the coming standardization of multi-computing. While a typical computer system's host processor is fully in charge of everything, with other cards in the system functioning only as slaves to the I/O controller, a multi-computing platform has several SBC modules, each of which may or may not have I/O devices, and each of which operates somewhat independently. In a CompactPCI system, the host SBC may be providing some type of special, supervisory function, or it may operate as just another peer processor.

Multi-computing adds to the high-availability found in CompactPCI technology because it enables the creation of environments with distributed processing power configured in a redundant manner. If any of the non-system-slot SBCs fail, the work can be redistributed to the surviving processors, so that it is not necessary to rely solely on the host processor.

DUAL-HOST SUPPORT

Another important function of the CompactPCI high-availability infrastructure is its dual-host support capability. This refers to the ability to switch out the CompactPCI host SBC without system shutdown. When the host SBC fails, the standby host takes over the responsibility for controlling the CompactPCI bus, and assumes the rest of the processing that was done by the failed SBC. Currently the PICMG Dual-Subcommittee is working on standards to define the interoperation between SBCs, backplanes and peripheral cards maintained in dual-host environments.

HOST FAILOVER MANAGEMENT

The top of the CompactPCI high availability infrastructure centers on its host failover management. This is the process of managing the transferal of data from a failed host processor to a standby processor.

There are basically four ways to design a standby processor to take over system operations once the primary processor has failed. The first two methods, lock-step processing and application-directed checkpointing, have been used for a number of years in various systems.

In lockstep processing, the standby processor processes the same data as the host processor. Once the host processor fails, the standby processor continues running with the same data. This method involves two processors running identical applications. This type of failover requires a lot of proprietary hardware and is getting more difficult to do because requires lots of pipelining and internal caches.

In application-directed checkpointing, applications running on the active system periodically send data to a standby system in order to keep the standby system up-to-date. When a failover occurs, the standby system uses the latest information it has received to keep the application running. One of the biggest drawbacks of using application-directed checkpointing for telecom applications is that failover time is long and data may be lost. This type of failure method minimizes a system failure rather than keeping a system running with no loss of service.

The next two failover approaches -- clustering and system-directed checkpointing, both of which can reside within CompactPCI technology -- are newer and hold promise for making failover transparent to application software, an important feature in telecom applications.

In clustering, large, complex applications are typically made up of individual software processes. The clustering system is an extension to the OS that allows these processes to be distributed to multiple hardware nodes usually connected to a high-speed Ethernet network. When one node fails, the processes running on that node are transferred to other operational nodes. For this type of failover recovery, applications must be written in such a way that any of the processes may fail and restart without causing the overall application to fail.

In a CompactPCI system, a host SBC and several other non-system-slot SBCs can each run as an independent node in a clustering system, creating, in effect, a cluster-in-a-box. The clustering software typically expects nodes to be connected over a LAN, so the CompactPCI multi-computing standard, which simulates a LAN, is a natural high-speed interconnect.

System directed checkpointing, a technology patented by RadiSys (formerly Texas Micro), provides a method for including a stand-by host processor which can continue operations less than 250 milliseconds after a system failure. This technology is utterly transparent to application software as well as to much of the OS. Programs do not even need to be recompiled to take full advantage of the platform's fault-tolerance. System directed checkpointing moves the entire contents of the main memory from the active host to a standby host twenty or more times per second without any application software knowledge or help. The "failover time" of 250 milliseconds makes it suitable for telecom and Internet applications where the outages of 30 seconds or more often experienced with "traditional" high availability systems are simply too long.

FUTURE OF COMPACTPCI IN THE TELEPHONY ARENA

The future growth of CompactPCI technology within the telephony market looks bright for several reasons. First, as the telecom, datacomm, Internet and wireless markets converge into one market, the demand for equipment based on open system standards will increase dramatically. This trend is most evident in the Open Source for Open Telecom initiative endorsed by PICMG. Open Source for Open Telecom is software for computer-based telecom equipment and services. Supporters of the initiative are trying to leverage the hardware components and the software environments of the mass-market computer industry to change the way telecom equipment and services are designed and delivered. One of the goals of the initiative is to provide hot-swap and telecom circuit switching software for telecommunications equipment built using the CompactPCI computer telephony specification.

Second, time to market pressures have forced OEMs to re-think the utilization of traditional telecom equipment based on proprietary and legacy systems.

High availability technologies, specifically CompactPCI and the hot swap standards are enabling the migration of the telecom infrastructure equipment from proprietary hardware and software architectures to mass market PC hardware and industry-standard software like Windows NT®. This along with the various layers of high availability technology building-blocks will help CompactPCI technology find its place in the telephony market ■

Lynn Brielmaier is an applications engineer with the Communication Platforms Division, RadiSys Corp. He has a BSEET degree from Devry Institute of Technology. He can be contacted at lynn,brielmaier@radisys.com.