Off-the-shelf management solutions for AdvancedTCA boards

By Mark Overgaard

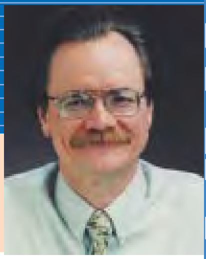

AdvancedTCA (PICMG 3.0) platforms deliver a dramatically higher level of interoperable management visibility and control than existing CompactPCI platforms supporting the CompactPCI System Management specification, PICMG 2.9. The Intelligent Platform Management Interface specification, on which all PICMG management architectures are based, is an excellent platform management framework, but lacks key functionality such as hot-swap awareness. Almost a third of the 430-page PICMG 3.0 R1.0 specification is devoted to defining the necessary extensions.

Achieving the manageability benefits of these extensions requires considerable new functionality in the Intelligent Platform Management (IPM) Controller that represents an ATCA board to the Shelf Manager and overall System Manager. The IPM Sentry Board Manager software and BMR-100 reference design for IPM Controller hardware allow developers of ATCA boards to focus on their unique added values rather than developing and validating the management subsystem on their own. IPM Sentry components are extensively tested, including successful operation with multiple independent shelf manager implementations in PICMG-sponsored interoperability workshops. (Please see "Management Building Blocks Speed AdvancedTCA Product Development" in the March, 2003 issue of CompactPCI Systems *Magazine for an introduction to ATCA management and the IPM Sentry management components.)*

Management aspects of an ATCA board

Figure 1 shows the key backplane interfaces of an ATCA board, including:

- Dual redundant -48 VDC power from the backplane. A DC/DC converter provides management power for the IPM Controller, and, when enabled, power for the rest of the board.
- Dual redundant IPMB-0, the primary in-shelf management link to the IPM Controller. An 8-bit (including parity) Hardware Address is used to establish the IPMB-0 address through which a board communicates.
- Fabric, Base, and Update Channel Interfaces providing the point-to-point backplane interconnects in a particular board and slot. The IPM Controller, through the Electronic Keying (E-Keying) process, enables the appropriate subset of the ports on those interconnects.
- Synchronization Clock Interface and Metallic Test Bus, both of which are bused on the backplane. They are subject to a separate E-Keying process to avoid conflicting uses of those buses within a shelf.

The payload portion of an ATCA board implements the main function of the board and communicates with the IPM Controller over the Payload Interface. AdvancedTCA leaves the details of the Payload Interface entirely to the board implementer. One example interaction across the Payload Interface would be coordination of graceful shutdowns of the Payload when initiated via the IPM Controller.

The IPM Controller uses non-volatile storage for various purposes, including as a repository for the FRU information that describes the Fabric, Base, and Update Channel Interfaces implemented by that board for use in E-Keying.

AdvancedTCA requires a watchdog timer external to the IPM Controller processor. If this watchdog is ever allowed to expire, the IPM Controller is disconnected from IPMB-0 and undergoes a reset operation as a recovery measure. It is crucial that such a recovery reset not disturb the operation of the payload; therefore, many of the payload-oriented controls should be persistent across such resets. Controls in this category include the Point-to-Point E-Keying and the Payload Power Enables.

In addition, ATCA boards are required to include at least one temperature sensor; when the thresholds configured on that sensor are exceeded, the Shelf Manager is notified and can take corrective action. Potentially, the Payload could also be notified of such events over the Payload Interface.

As shown in Figure 1 (where the IPM Controller portion is *definitely* not to scale), a board's IPM Controller is also responsible for managing key operator interfaces, including:

- The handle switch, which is integrated with the lower board handle to indicate whether the board handle is open or closed, enabling automatic responses to board insertion and extraction actions.

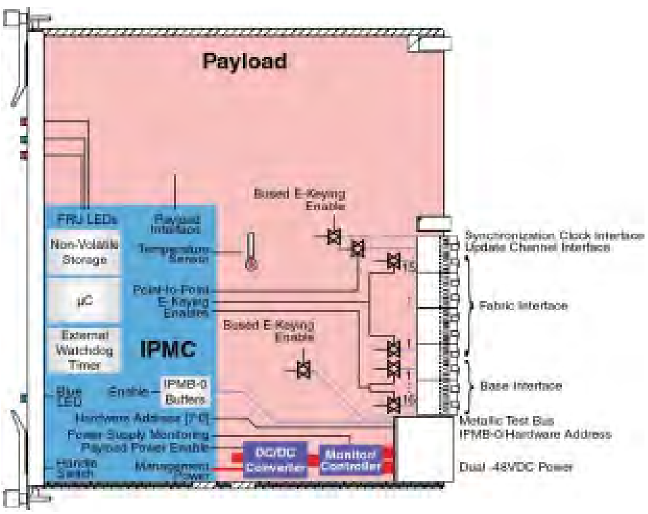

Figure 1

- The blue hot-swap LED, which is fully on when the board is safe for extraction, fully off when extraction is not safe, and blinking during transitions between those two states.
- Some number of Field Replaceable Unit (FRU) LEDs, which communicate overall status of the board and its various functions. AdvancedTCA defines interoperable mechanisms for management software to control LEDs, including blink rate, color, and lamp test functions.

IPM Sentry BMR-100 hardware architecture

The IPM Sentry BMR-100 is a reference design for the hardware portion of an AdvancedTCA IPM Controller, complete with schematics and bill of materials. By using this reference design as the basis for the required IPM Controller on their AdvancedTCA boards, developers can save time and expense compared to designing a compliant IPM Controller on their own. Even better, the companion IPM Sentry Board Manager software is already developed and tested on BMR-100-compatible hardware, saving substantial additional time.

The primary intelligence in the BMR-100 design is provided by a Texas Instruments TMS320VC5470 (C5470 below), which includes two processors on a single compact chip:

- a 47.5 MHz 32-bit ARM7
- a 100 MHz C54x Digital Signal Processor (DSP)

The DSP is dedicated to a soft implementation of the dual redundant IPMB-0, while the ARM7 handles the remaining processing responsibilities for the IPM Controller.

Given the built-in SDRAM and Flash controllers of the C5470, single chip SDRAM and Flash storage need not be more expensive than much smaller capacities of SRAM and SEEPROM storage used in other management controller designs. With BGA packages, even the board real estate requirements are comparable. Larger memory has a number of benefits, but especially ensures that the extensive functionality additions defined by ATCA can be accommodated.

As required by AdvancedTCA, an external watchdog timer ensures that if the C5470 fails, due to hardware or software issues, the IPM Controller will be detached from IPMB-0 and reset. Reconnection to IPMB-0 only happens when the Board Manager is confident of the health of its underlying hardware. As a further safety measure in the BMR-100 design, the watchdog strobe signal is routed through the same I2C-accessed GPIO block that holds the payload power enables. Therefore, when the watchdog can be strobed, the Board Manager is automatically able to control payload power as well.

The BMR-100 reference design uses I2C-accessed temperature and voltage sensors, which allows for convenient placement of the appropriate number of sensor devices wherever on the board the measurements are most effectively done. Integration of these sensors with the IPM Controller processor would have some benefits, but it is generally preferable to route the simple digital I2C bus to remote parts of a board rather than potentially tricky analog signals. The specific sensor devices shown in Figure 2 (Dallas Semiconductor DS-75 and Analog Devices ADM1024) have built-in software support in the IPM Sentry Board Manager, but

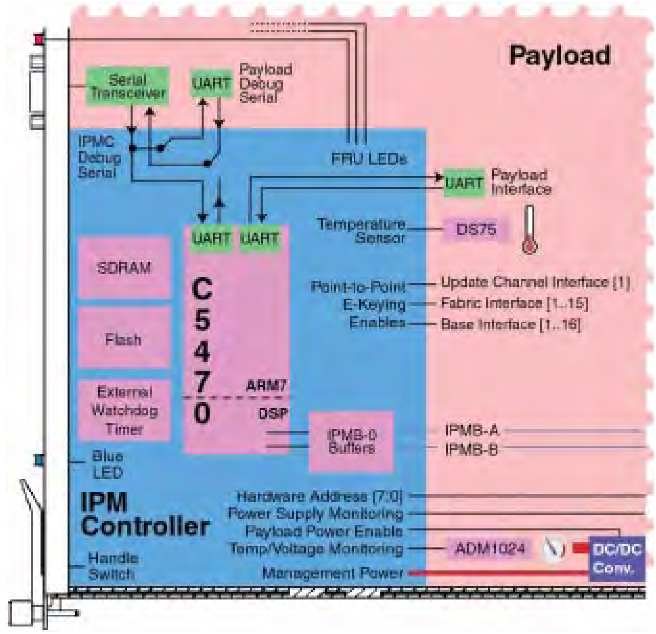

Figure 2

other devices can be substituted if the corresponding software support is added.

The BMR-100 reference design uses a UART-based Payload Interface. UARTs are either built into or easily added to almost any payload processor. Using one of the C5470's built-in dual UARTs as the Payload Interface is applicable to and cost-effective for almost any board design. The other built-in C5470 UART provides a flexible serial debug port, which allows serial access to the IPM Controller in the lab or in the field. Optionally, a switch on this interface can be used to provide serial debug access to either the IPM Controller or the payload processor.

IPM Sentry Board Manager software architecture

The IPM Sentry Board Manager implements the software aspects of an AdvancedTCA-compliant IPM Controller, running on the BMR-100-based hardware integrated into an ATCA FRU. Figure 3 shows the overall architecture of the Board Manager.

The Board Manager includes the following components or component types:

- **Intelligent Platform Management Foundation (IPMF):** Provides the core services for an IPM Controller that are applicable across AdvancedTCA board and FRU types (including FRUs running either IPM Sentry Shelf Manager or Board Manager). For each functional area (such as overall FRU state management, E-Keying, sensor management, cooling management, etc.), IPMF provides a combination of function APIs that can be called directly by its client and callbacks to the client to minimize the size and complexity of corresponding IPMC-dependent code.
- **BMR-100 Client:** Configures and invokes IPMF specifically for the BMR-100 reference hardware.
- **IPMC-Specific Configuration:** Source module(s) that customize the operation of the Board Manager to a specific IPM Controller implementation. For IPM Controllers that stick closely to the BMR-100 reference hardware design, changes in the Board Manager outside this block may be

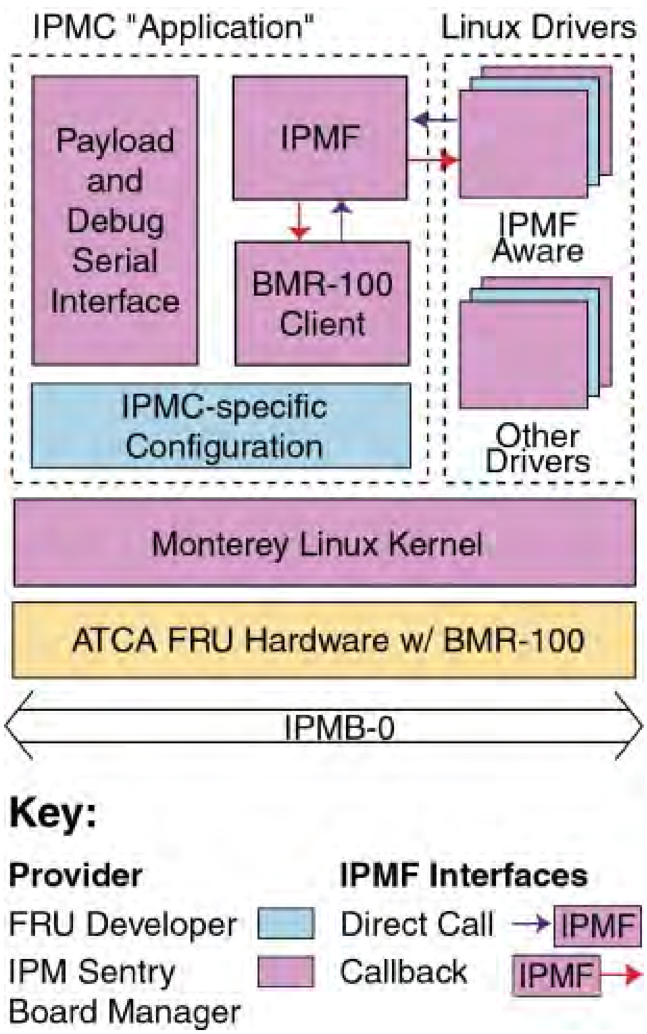

Figure 3

unnecessary.

- **Payload and Serial Debug Interface:** Implements an ASCII interface between the IPM Controller and 1) the FRU Payload and 2) a Serial Debug interface.
- **Monterey Linux kernel**: The kernel of Monterey Linux (www.montereylinux.com), tuned and optimized for the Board Manager context.
- **IPMF-aware Linux device drivers:** Linux compatible drivers that also make use of IPMF services. For instance, the driver for the ADM1024 sensor device can be configured to automatically interact with IPMF to create IPMI voltage and temperature sensor objects corresponding to the physical sensors in one or more ADM1024 devices. Board developers may create their own IPMF-aware device drivers for custom sensor devices.
- **Linux device drivers:** Other Linux-compatible device drivers. These drivers may be specific to the BMR-100 hardware (such as the collection of drivers that implement the soft IPMB-0 working with code executing on the C5470's DSP subsystem) or generic Linux drivers (such as the drivers that make up the generic parts of the Linux I2C subsystem). Alternatively, FRU developers may create their own drivers for custom devices.

This architecture for the Board Manager supports a wide range of adaptability to the specifics of the IPM Controller for a given ATCA board or other intelligent FRU. At one end of the spectrum, only simple changes in the IPMC-Specific Configuration module(s) may be necessary. At the other end of the spectrum, the Linux kernel foundation enables a rich source of existing device driver and other software that can be leveraged to address unique requirements. Furthermore, the Board Manager is delivered with a complete open source development environment (including JTAG access facilities) with one-stop support.

IPMF provides built-in facilities for key areas of IPMC functionality, but if those facilities are not suitable for a given IPMC, they can be modified or replaced entirely. For instance, there is a Simple Power Management Facility that implements the payload power management negotiations defined by ATCA for the case that the payload is either on or off (as opposed to the richer model available in ATCA, where there may be up to 20 power levels). Similarly, there is a FRU LED management facility that implements the fairly rich semantics defined by the ATCA FRU LED commands (flashing, local and shelf manager control, lamp test, color choice). Either of these facilities can be used as is, modified, or replaced entirely.

The Board Manager implements an ASCII-based Serial Interface Protocol (SIP) over the Payload Interface and the Serial Debug Interface. The SIP is based on and compatible with IPMI's Terminal Mode and supports many of the management interactions with the IPMC that are defined for ATCA shelf managers. Among other things, the payload or an operator accessing the serial debug interface can:

1) Discover the capabilities and settings of the E-Keying-governed interfaces
2) Discover the LED facilities of the board and control the local settings of those LEDs
3) Request that the IPMC provide a watchdog facility for the payload, resetting the payload if the watchdog is not periodically strobed
4) Negotiate a graceful shutdown of the payload

Though it is not represented in Figure 3, the Board Manager also supports a stand-alone mode of operation where no shelf manager is present (perhaps in a simple one-board chassis). Some ATCA board developers are choosing to design their board to facilitate such dual uses. The BMR-100 hardware is fully capable of supporting the functional extensions required by such a mode, compared to the minimum responsibilities of a generic IPM Controller.

Life cycle of an AdvancedTCA field replaceable unit

Among the most important ATCA extensions to IPMI are the provisions for FRUs that dynamically enter and leave a shelf. IPMI implicitly assumes that the FRU population of a system is static, or at least that any adjustments of the management information associated with FRU population changes are outside the scope of the specification. ATCA provides the first solution to these challenges that preserves interoperability among independent implementations. The CompactPCI specifications don't address this challenge, so CompactPCI platform vendors have implemented multiple incompatible solutions.

ATCA defines operational states for FRUs and transitions among them. Each IPM Controller manages a separate state machine for each of the FRUs it represents to the Shelf Manager. For instance, a single IPM Controller could represent several fan trays and a pair of power entry modules, with the state of each FRU logically represented by a distinct state machine maintained by the parent

IPM Controller. ATCA also defines a special hot-swap sensor type by which transitions in these state machines are reported to the Shelf Manager.

Figure 4 (simplified somewhat from the specification) shows the ATCA FRU state machine. The states are sketched briefly below to provide a feel for the interoperable solution defined by ATCA:

- **M0: Not Installed.** If the FRU physically integrates the IPM Controller, this is a virtual state, since the IPM Controller is not operational. If the IPM Controller represents separately removable FRUs, this can be a real state for those FRUs. All power to the FRU is off.
- **M1: FRU Inactive.** The FRU is installed; the IPM Controller is operational and waits in this state until the insertion criteria (such as closed ejector handles) are met. The payload is not powered and its inputs and outputs are disabled. If the FRU is shutting down (coming from M2 or M6), resources are reclaimed and E-Keying governed interfaces for boards connected to this FRU are disabled. The blue hot-swap LED is on, signaling that the FRU can be extracted safely.
- **M2: FRU Activation Request.** The IPM Controller waits for activation via a Set Activation State (Activate FRU) command from the Shelf Manager over IPMB-0. The blue LED signals this state with a long blink; extraction is not safe.
- **M3: Activation in Progress.** The FRU's IPM Controller negotiates with the Shelf Manager for a power budget; once that is granted, the FRU payload begins powering up. If the FRU is an ATCA board, it negotiates the state of its E-Keying governed interfaces with the Shelf Manager and receives the resulting enables and disables for those interfaces. The FRU autonomously transitions to M4 when activation is complete. The blue LED is off to indicate that extraction is not safe.
- **M4: FRU Active.** This is the normal operational state. The FRU is powered and cannot be extracted safely. (The blue LED is off to indicate that.) The FRU's payload may be reset, either autonomously or at the direction of the Shelf Manager, without affecting this operational state. The FRU's IPM Controller can request extraction (that is, move it to M5) if the extraction criteria are met (such as an operator opening the ejector handles). Alternatively, the Shelf Manager can direct the FRU to deactivate (that is, move to M6) with a Set FRU Activation State (Deactivate FRU) command over IPMB-0.
- **M5: FRU Deactivation Request.** The FRU's IPM Controller is requesting deactivation permission from the Shelf Manager and the blue LED shows a short blink pattern to indicate that deactivation is a possibility. No impact on or notification of the payload occurs until the request has been granted.
- **M6: FRU Deactivation in Progress.** The FRU is shutting down and I/O connections are being deactivated; the blue LED continues the short blink pattern. The FRU transitions to M1 autonomously when deactivation is complete and extraction is safe. Power to the payload is turned off and any E-Keying enabled interfaces are disabled before that transition.
- **M7: Communication Lost.** This is an abnormal state that is entered for a FRU when the Shelf Manager unexpectedly loses contact with its IPM Controller or an IPM Controller loses contact with one of its managed FRUs. Loss of contact may be due to an IPMB-0 or other communication problem or even to an unexpected extraction from the shelf. This state can be entered from any of the other states (except M0), with a return to the former state when (and if) communication is regained.

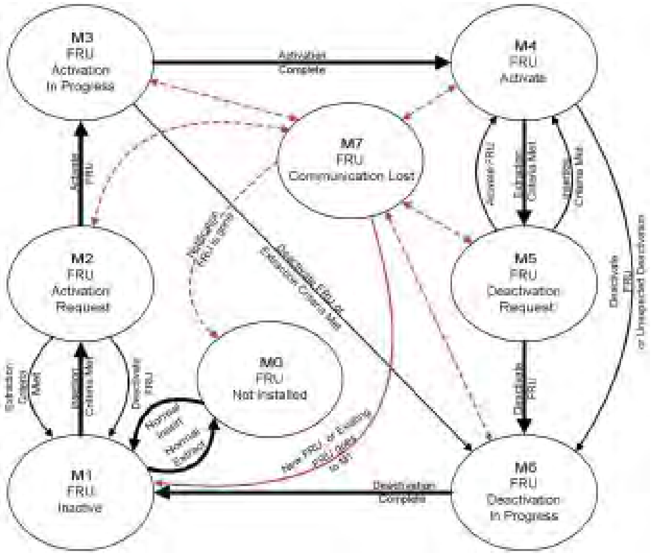

Figure 4

If the BMR-100 reference design is followed closely on an ATCA board, support for the above state machine may require little or no custom modifications of Board Manager software. Thus, the board developer can take advantage of the extensive validation that has already been done on the Board Manager executing on BMR-100-compatible hardware.

One aspect of that validation has included exercising the IPM Sentry Board Manager (as well as the IPM Sentry Shelf Manager) in several PICMG-sponsored Alpha Interoperability Workshops (AIWs). As of this writing, three of these workshops have been held, each with ten or more companies bringing their ATCA products for interoperability testing; a fourth workshop is scheduled for May, 2003. Functional areas relating to particular parts of the above state machine that have been tested thus far include:

- Basic insertion and extraction scenarios
- E-Keying for point-to-point interfaces (which mainly affects states M3 and M6)
- Power on sequencing and power budget negotiation (also primarily affecting states M3 and M6)
- Entry to and exit from state M7, in which communication with a FRU or IPM Controller is lost and potentially regained

Also tested have been several functional areas that apply to almost any of the ATCA FRU states:

- Generation and handling of temperature exceptions (which can be sent to the shelf manager at any time, based on configurable thresholds for FRU temperature sensors)
- IPMB-0 management, including exercising the dual redundancy as well as detecting and recovering from IPMB hangs

Conclusion

AdvancedTCA defines extensive additions to IPMI to create the first interoperable platform management solution for advanced modular architectures. The IPM Sentry Board Manager and BMR-100 IPM Controller hardware reference design provide off-the-shelf software and hardware components that already comply with the 130-page management section of PICMG 3.0. Furthermore, these components have already been tested for interoperability with multiple ATCA shelf managers. Use of these off-the-shelf components can dramatically simplify and accelerate the development of management subsystems for AdvancedTCA boards and other FRUs.

Mark Overgaard *is president of Pigeon Point Systems, which he founded in 1997 to focus on products and services supporting the adoption of open modular platforms to replace proprietary architectures, with an initial focus on the telecommunications market and CompactPCI. He is a leader in the technical subcommittees of PICMG (including Software Interoperability chairman and most recently in the management aspects of AdvancedTCA).*

For more information, contact Mark at:

Mark Overgaard
Pigeon Point Systems
P.O. Box 66989
Scotts Valley CA 95067
Tel: 831-438-1565
E-mail: mark@pigeonpoint.com
Web site: www.pigeonpoint.com

