Hot-plug Based Activation and Deactivation of ATCA FRU Devices

P. Predki, D. Makowski
Technical University of Lodz
Department of Microelectronics and Computer Science
Lodz, Poland

***Abstract*—One of the most important features of the Advanced Telecommunications Computing Architecture (ATCA) contributing to its exceptional reliability and availability is its hot-swap functionality.**

In order for the user to be able to add and remove the components of an ATCA shelf without the necessity of switching the power on and off the PCI Industrial Computer Manufacturers Group (PICMG) specification clearly enumerates the stages a Field Replaceable Unit (FRU) has to go through upon insertion into and extraction from the shelf. These stages form the activation and deactivation processes that occur every time an element is changed in the ATCA system.

This paper focuses on these processes placing the emphasis on the Electronic Keying (EK) implementation in the Intelligent Platform Management Controller (IPMC) software developed for the self-designed ATCA Carrier Board (CB). This CB is considered to be used in the Low Level RF (LLRF) control system of the X-Ray Free Electron Laser (XFEL). It utilizes the standard-defined PCI Express (PCIe) interface as well as introduces proprietary protocols in form of Low Latency Links (LLL).

***Index Terms*—Intelligent Platform Management Interface, Advanced Telecommunications Computing Architecture, Hot swap, Carrier Board, Shelf Manager, Advanced Mezzanine Card, Electronic Keying, Field Replaceable Unit, PCI Industrial Computer Manufacturers Group, XFEL**

I. INTRODUCTION

ATCA is considered a significant improvement over the Versa Module Eurocard (VME) architecture that had been used in the Free Electron Laser in Hamburg (FLASH) [1] and is a worthy candidate to be implemented in the LLRF control system of XFEL. Many features of the ATCA architecture, such as separation of management and payload power or voltage, current and temperature monitoring make it far more reliable and accessible than the VME architecture. In addition to that, one of the problems in the latter architecture was the multitude of connections which are all located on the front board panel. When the system operator wants to change a component of the system there is no other way of doing that other than switching the power off, unplugging all the cable, extracting the module and, upon insertion of the new module, rewiring it. This, apart from taking a lot of time, requires the system operator to know exactly how to connect the module in a proper way. Taking into consideration that a laser control system requires a lot of connections it is not hard to commit an error. A mistake in wiring the module may lead to malfunction of the whole system. Best case scenario the system will not work properly. Worst case scenario a part of the system can become physically broken and be in need of replacement. Such a solution also makes it harder to determine the source of malfunction because one needs to check all the connections which there may be a lot of.

A much better, safer, more foolproof and reliable solution is the ATCA architecture. The LLRF control system design includes no connections on the front panel. They are all moved to the backplane and implemented as Zone 2 (Z2) connectors for data transmission and Zone 3 (Z3) connectors for analog signal transmission [2], [3]. The PICMG specification clearly states what kinds of connections can be made this way and its flexibility should satisfy all the requirements of the system engineer. Along with solving the connection problem there is also no need to power down the system or its part. The separation of management and payload power allows us to take advantage of the hot-swap functionality. By pulling the hot-swap handle the user informs the Shelf Manager (ShM) a component is about to be removed from the system. The ShM goes through a process which ultimately disconnects the payload power of the module and allows the user to remove it safely from the system. This process as well as the reciprocal one, taking place when a module is inserted into the shelf is described in detail in the following sections [4]–[6].

II. CARRIER IPMC SOFTWARE

Because of the specific application of the system as well as the implementation of many custom solutions such as the RIO and LLL interfaces it was not possible to use off-the-shelf IPMC software. Thus, such an application is being developed for the CB according to the PICMG specification [6] and it can handle all the events specified in this paper.

III. AMC MODULE INSERTION

A good example of a FRU device activation process is the Advanced Mezzanine Card (AMC) insertion into a ATCA CB. A default type B+ AMC bay connector on the CB consists of 170 pins, most of which can be used as I/O pins [5]. Some of them, however, play a crucial role in the management of the AMC module. Among them we can distinguish the PS0# and PS1# pins, whose task is to inform the CB that the AMC module has become present. They provide confirmation that all pins of the module have mated as they are the last ones to be inserted into the connector. The active PS1# signal is used by the CB to detect the presence of the AMC module. On the other side, the PS1# signal is fed back to the module as the ENABLE# signal. Only when this signal becomes active is the Module Management Controller (MMC) on the AMC board allowed to use the IPMB_L bus (SDA_L, SCL_L pins) or read the Geographic Address (GA) inputs (GA2:0 pins) [5]. When the PS1# signal goes active the AMC module is transitioned to state M1 and its Management Power is enabled (see M-State transition diagram [5]). All state transitions are reported by the Carrier to the ShM on behalf of the AMC modules. The CB waits now for the hot-swap trigger to continue with the activation process.

Pushing the module hot swap by the user generates an event message that is sent from the AMC module to the CB. The board then sends a "Set FRU LED State" command back to the AMC with a request to perform long blinks of the Blue LED. This LED is used as a visual indicator of the state of the activated or deactivated component. The Carrier also reads Module Current Requirements record of the AMC in order to verify that it can support the component as far as power consumption is concerned. This involves checking that the maximum module current is greater than current draw and that maximum internal current is greater than current draw of all modules activated up to this point plus the current draw of the module negotiating at the moment. Another record that is of interest to the CB is the AdvancedMC Point-to-Point Connectivity Record used in Electronic Keying (EK) at a later stage [7]. This process is described in more detail in the next section.

If the Carrier can provide the necessary Payload Power it transitions the Module it to state M2. If not, a Set FRU LED State command is sent again this time requesting the Blue LED to stay on. The Module remains in M1 state. While in M2 the Carrier awaits permission from a higher level management (e.g. ShM) to proceed with the activation. The Carrier receives the Set FRU Activation (Activate FRU) command and issues a Set FRU LED State command to the Module with a request to turn the Blue LED off. It also transitions the Module further to state M4. At this point the ShM begins power negotiation by sending a Get Power Level command to the Carrier IPMC. In response, the Carrier provides the information about the necessary power level to power the Module. This conversation completes with a Set Power Level command from the Shelf Manager. At this point the Carrier enables the Payload Power for the Module and uses the electronic keying information gathered previously to enable all the compatible connections between the Modules and on-Carrier devices. When all this is done the Carrier sends an M3 to M4 transition event message to the ShM informing it that the activation process has completed and the AMC Module is now active.

IV. ELECTRONIC KEYING

EK is a mechanism by which the mandatory AMC.0 Management infrastructure is used to dynamically satisfy the needs that had traditionally been satisfied by various mechanical connector keying solutions. As mentioned in the introduction, thanks to this solution it is no longer necessary to manually connect multitudes of cables between different components of the shelf. All these connections are transparent to the user and are hidden in the backplane or, in case of AMC modules, between the AMC bay connectors and the on-Carrier devices. The system verifies fabric compatibility before enabling the connection which greatly prevents mis-operation and makes the usage of the equipment much less complicated and much safer. Still, analog links can be made using direct connections or using Rear Transition Modules (RTMs).

Continuing with the AMC module activation example, two types of topologies can be distinguished. One of them is direct Module-to-Module connection where the Carrier simply provides the connection between the Modules. In the second case there is an on-Carrier device contained by the Carrier. It can indirectly connect AMC Modules together but it can also be linked to the Zone 2 connector, providing external connection to the Modules. The Carrier IPMC is responsible for EK between the Modules and their connections to the Carrier resources.

The AMC Base Specification distinguishes several records that hold information about the various interfaces supported by the AMC Modules. Some of these records can also be present on the CB. All of them are used during EK and the Carrier IPMC decides which interfaces and which ports to enable based on that information and examining the compatibility between the Modules and on-Carrier devices. These records, including the subrecord hierarchy, are listed below.

- Carrier Point-to-Point Connectivity Record
  - Point-to-Point AMC Resource Descriptor
    - Point-to-Point Descriptor
- AdvancedMC Point-to-Point Connectivity Record
  - AMC Channel Descriptors
  - AMC Link Descriptors
    - AMC Link Type Extension
    - AMC Link Type
    - AMC Link Designator

The Carrier Point-to-Point Connectivity Record is included in the Carrier FRU Information and describes the point-to-point connections implemented on the Carrier. These connections consist of an arbitrary number of Ports and are defined for all the possible links between the AMC Bays and the on-Carrier devices. The connection definitions do not need to be reciprocated. Thus, when a connection from AMC Bay A to AMC Bay B is already defined there is no need to indicate a reverse connection form AMC Bay B to AMC Bay A.

The Advanced Point-to-Point Connectivity Records are included in the AMC FRU Information and describe the Channel and Link connectivity that is implemented on the AMC Module. One or more such records can also exist in the Carrier FRU Information for each on-Carrier device supported on the platform. In these case the connections are defined in terms of Channels, each of which can be a concatenation of one to four Ports, and Links, which include the information about the interfaces supported by the Module. The AMC Link Descriptor provides information about Asymmetric Matching (required by some interfaces, e.g. PCIe [8]), Link Grouping (Ports of one Channel can operate with Ports of another Channel), Link Type (one of AMC.0 subsidiaries or OEM specific) and Port inclusion (not all Ports of a Channel have to be active in a given Link). Such a solution gives an enormous flexibility when it comes to designing complex platforms supporting various interfaces.

When a Carrier IPMC is determining the EK match the AMC Link Designator/Link Type/Link Type Extension fields are compared. For AMC Link Type governed by one of the AMC.0 subsidiary specifications there must be an exact match of those fields and compatibility in the AMC Asymmetric Match fields. For OEM specific solutions OEM-defined identification numbers (OEM GUID) are compared instead of AMC Link Type fields. A match between these fields is necessary for sending a Set AMC Port State (Enable) command by the Carrier IPMC. It is, however, not sufficient. The Carrier connectivities have to be taken into account as well as Channel Grouping IDs for multi-channel interfaces. Once the aforementioned command is sent to the Module the Links are active and communication over them may commence.

In case of the LLRF CB [9] there are three AMC Bays (numbered 5, 6, 7 as defined by the standard) connected to an on-Carrier device, which is a PCI Express switch [10]. Each AMC Bay has its own Point-to-Point AMC Resource Descriptor List in the Carrier Point-to-Point Connectivity record. Two of them define only one AMC Port (number 4) and one of them defines four AMC Ports (from 4 to 7). The Port numbers that can be used for PCI Express Link, or any other Fabric Channel Link, are defined by the PICMG and they start from the number 4. There is no need to define another Point-to-Point AMC Resource Descriptor List for the on-Carrier device since all the possible connections have already been specified. The AMC Point-to-Point Connectivity Record defines three Channels with PCI Express Link Type. Two of them enable only one Lane, which corresponds to only one Port available for AMC Bays 5 and 7. The third one, however, enable all four Lanes which also corresponds to the Port availability of AMC Bay 6 (Fig. 1).

The AMC Module that was tested with the LLRF CB was TAMC900, which supports up to PCIe x8 link (on AMC Ports 4-11) which is used to transmit data from on-board ADCs to the CPU. The current Carrier configuration allows only up to PCIe x4 (AMC Ports 4-7) connection in AMC Bay 6 so the multi-channel Link could not have been tested. However, both the x1 links in AMC Bays 5 and 7 as well as x4 link in AMC Bay 6 have been matched successfully and allowed data flow via the AMC Module and the PCIe switch to a desktop PC.

It is also worth mentioning that when the user opens the hot-swap handle the Carrier IPMC sends a Set AMC Port State (Disable) command in order to disable all the connections before the Module is extracted.

The principle of EK is exactly the same when it comes to establishing connections between ATCA Boards in one shelf via the backplane. Similar records need to be provided in the Board FRU Information which inform the ShM of all the available interfaces the Board supports. As far as Fabric Interfaces are concerned the PICMG 3.0 specification supports only five such links (Base Interface 10/100/1000 BASE-T, Ethernet, Infiniband, StarFabric and PCI Express). However, it also allows the user to define a proprietary interface as was the case with AMC Modules. The LLRF CB introduces two such Fabric Interfaces with the following 16-byte OEM GUIDs [11]:

- Rocket IO
    - 52 6f 63 6b 65 74 20 49 4f 20 49 6e 74 66 63 65 (“Rocket IO Intfce)
- Low Latency Link
    - 4c 6f 77 20 4c 61 74 65 6e 63 79 20 4c 69 6e 6b (“Low Latecy Link)

Whenever two ATCA Boards are inserted into a shelf and they share the same OEM GUID and other parameters are compatible the ShM sends a “Set Port State (Enable) command to both Boards informing them that a physical connection is available and operational. The Boards can then exchange data using this link.

The ATCA standard allows various topologies to be used on backplanes such as dual star, dual-dual star or full mesh. The Elma backplane used for development is a 5-slot backplane

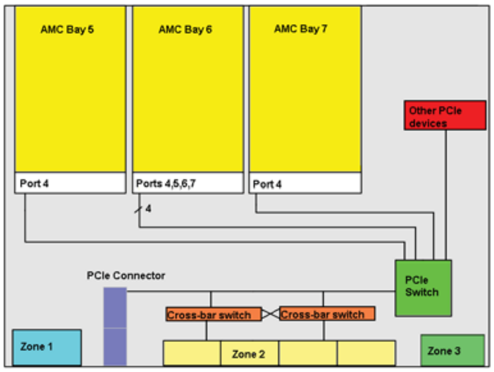

Fig. 1. PCI Express connections on the LLRF carrier board

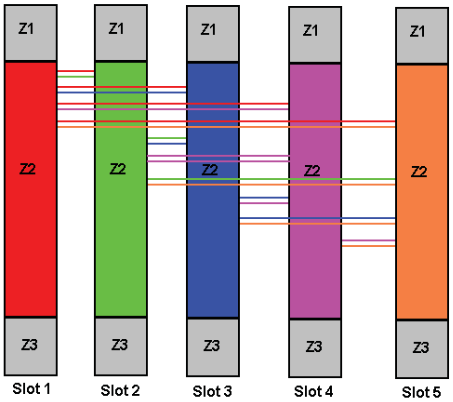

Fig. 2. Backplane Zone2 fabric connections between ATCA boards

implementing a dual full mesh topology (Fig. 2). Due to that fact each ATCA Board can be physically connected to any other board on the backplane by up to two channels provided that the channels are compatible.

V. AMC MODULE EXTRACTION

The extraction process is very similar to the insertion one with all the events occurring in reverse order. First, however, the user needs to open the hot-swap handle, informing the Carrier IPMC about the intention to remove the Module from the AMC Bay. The Carrier IPMC transitions the Module from state M4 to M5 and sends a “Set FRU LED State requesting that the Blue LED blinks at a short rate. This indicates that the Module is awaiting deactivation. At this point the Carrier awaits permission from a higher level management (e.g. ShM) to proceed with the deactivation. When the Carrier IPMC receives the “Set FRU Activation (Deactivate FRU) command it transitions the Module further to state M6 and next disable all the Ports using the “Set AMC Port State (Disable) command to the Module. After the transition to state M6 the Carrier sends the “FRU Control (Quiesce) command to the Module and awaits a Module Hot Swap (Quiesced) event from the MMC. After all this is completed Payload Power is disconnected from the Module and transition to state M1 takes place. At this point another “Set FRU LED State” command is issued requesting that the Blue LED be turned on. This is an indication to the operator that the Module can be safely extracted from the Bay. The PS1# signal going inactive indicates to the CB that the Module has indeed been removed and a transition to state M0 occurs and the Management Power is disabled.

VI. CONCLUSIONS

Both the hot-swap functionality and the Electronic Keying support make the installation and operation of a system based on ATCA very easy and intuitive. What is more, it is much more difficult to make a mistake when adding or removing modules from the system that could prove fatal for its operation. All these factors greatly increase the reliability and availability of the ATCA-based platforms and facilitate their diagnostics and management all of which are very desirable features of complex systems such as the X-FEL control system. Taking all this into consideration Advanced Telecommunications Computing Architecture presents a significant improvement over the older Versa Module Eurocard standard.

ACKNOWLEDGEMENT

The research leading to these results has received funding from the European Commission under the EuCARD FP7 Research Infrastructures grant agreement no. 227579 and Polish National Science Council Grant 642/N-TESLA-XFEL/09/2010/0. The authors are scholarship holders of project entitled “Innovative education ...” supported by European Social Fund.

REFERENCES

[1] S. Simrock, L. Butkowski, M. Grecki, T. Jezynski, W. Koprek, G. Jablonski, W. Jalmuzna, D. Makowski, A. Piotrowski, and K. Czuba, “Evaluation of an ATCA based LLRF system at FLASH,” in *Mixed Design of Integrated Circuits Systems, 2009. MIXDES ’09. MIXDES-16th International Conference*, June 2009, pp. 111 –114.

[2] D. Makowski, W. Koprek, T. Jezynski, A. Piotrowski, G. Jablonski, W. Jalmuzna, and S. Simrock, “Interfaces and communication protocols in ATCA-based LLRF control systems,” *Nuclear Science, IEEE Transactions on*, vol. 56, no. 5, pp. 2814 –2820, Oct. 2009.

[3] V. Ricchiuti, A. Orlandi, and G. Antonini, “High speed serial links characterization for ATCA backplanes,” in *Signal Propagation on Interconnects, 2008. SPI 2008. 12th IEEE Workshop on*, May 2008, pp. 1 –4.

[4] IPMI v2.0 rev. 1.0 specification. [Online]. Available: http://www.intel.com/design/servers/ipmi/spec.htm/

[5] PICMG AMC.0 AdvancedMC Mezzanine Module R2.0. [Online]. Available: http://www.picmg.org/v2internal/specifications.htm/

[6] PICMG 3.0 advancedTCA Base R3.0. [Online]. Available: http://www.picmg.org/v2internal/specifications.htm/

[7] T. Kastner. An Integrator’s Guide to MicroTCA. [Online]. Available: http://www.ecauk.com/news/whitepapers/pdfs/Whitepaper_An Integrator’s Guide to MicroTCA Part II_ How a FRU starts up.pdf/

[8] T. Kucharski, A. Piotrowski, D. Makowski, and G. Jablonski, “PCIExpress communication layer for ATCA-based linear accelerator control system,” in *Mixed Design of Integrated Circuits Systems, 2009. MIXDES ’09. MIXDES-16th International Conference*, June 2009, pp. 140 –144.

[9] A. Zawada, D. Makowski, T. Jezynski, S. Simrock, and A. Napieralski, “ATCA carrier board with ipmi supervisory circuit,” in *Mixed Design of Integrated Circuits and Systems, 2008. MIXDES 2008. 15th International Conference on*, June 2008, pp. 101 –105.

[10] S. Szachowalow, G. Jablonski, D. Makowski, L. Butkowski, W. Koprek, and S. Simrock, “Software for data acquisition AMC module with PCI Express interface,” in *Mixed Design of Integrated Circuits Systems, 2009. MIXDES ’09. MIXDES-16th International Conference*, June 2009, pp. 145 –148.

[11] D. Makowski, A. Piotrowski, and A. Napieralski, “Universal communication module based on AMC standard,” in *Mixed Design of Integrated Circuits and Systems, 2008. MIXDES 2008. 15th International Conference on*, June 2008, pp. 139 –143.