

## EMERGING NEW ELECTRONICS STANDARDS FOR PHYSICS\*

Raymond S. Larsen, SLAC National Accelerator Laboratory, Menlo Park CA 94025, U.S.A.

### Abstract

A unique effort is underway between industry and the international physics community to extend the Telecom industry's Advanced Telecom Computing Architecture (ATCA and MicroTCA) to meet future needs of the physics machine and detector community. New standard extensions for physics are being designed to deliver unprecedented performance and high subsystem availability for accelerator controls, instrumentation and data acquisition. A key feature is a unique out-of-band imbedded standard Intelligent Platform Management Interface (IPMI) system to manage hot-swap module replacement and hardware-software failover. An additional goal is to achieve a much higher degree of *interoperability* of both lab and industry designed hardware-software products than past generations of standards. This presentation will describe status of the hardware-software extension plans; technology advantages for machine controls and data acquisition systems; and examples of collaborative efforts to help develop an industry base of generic ATCA and MicroTCA products in an open-source environment.

### BACKGROUND OF ELECTRONICS STANDARDS FOR PHYSICS

#### *NIM, CAMAC and FASTBUS*

Modular instrument standards for physics laboratories were invented beginning in the 1960's with developments at Rutherford and CERN in Europe and at the National Bureau Standards in the US. The NIM Nuclear Instrument Module development was specifically undertaken at the behest of the US Atomic Energy Commission for reactor and high energy physics instrumentation because of the plethora of incompatible modular solutions that were springing up to serve these growing markets. In effect NIM was developed by a consortium of large physics labs in the US and Europe and quickly became the package of choice for both reactor controls and physics experiment nanosecond pulse resolution high speed counting electronics. That choice caused a shakeup in the industry as some companies with existing modular products resisted the change while at the same time new companies appeared eager to adapt to NIM. The NIM package was very simple compared with later high density modules with integrated computer data buses for inter-module and higher level data communications; it consisted mainly of a standard crate, power supply, signal levels and front panel connectors for coaxial inputs and outputs.

CAMAC, appearing roughly a decade later, was the first attempt to incorporate both a standard data bus and controls protocol into an industrially supported standard

that integrated analog and digital functions. It was quite controversial even within the labs because it boldly aimed to perform a wide range of analog-digital functions traditionally supported by the Test and Measurement market. In fact CAMAC never penetrated these markets very much outside of the labs but both NIM and CAMAC have had a remarkable survivability within the labs.

A third lab standard came in the 1980's at the specific request of experimental physicists who considered the CAMAC data bus obsolescent for experiments in terms of the data bandwidth requirements of ever larger experimental detectors and ever faster data buses and main frame host computers. Thus was born FASTBUS with a much larger form factor and another order of magnitude increase in bandwidth, wider read and write data buses compatible with the fastest computers, and a fully modern control structure, but still based on parallel transfers because high speed serial networking was still in its infancy and far too slow for data acquisition and processing. FASTBUS served well for more than a decade but much of its potential market diminished when most detector electronics became imbedded in special ASIC (Application Specific IC) designs within the detectors, with the great advantage of minimizing large expensive cable plants.

#### THE ADVENT OF ATCA

#### *Enabling Technologies for ATCA*

There has been no major movement in architecture of modular instruments in accelerators and physics for two decades, but the architectures of technology drivers used by the physics communities have changed dramatically with the advent of field programmable logic arrays with imbedded high speed serial drivers and receivers, and in the past decade, the advent of ATCA, the Advanced Telecommunications Computer Architecture, and its daughter product MicroTCA platform. These are the first architectures designed specifically for high availability of 0.99999 at the shelf (crate) level, achieved through a combination of redundancy, an Intelligent Platform Management Interface (IPMI) and true hot-swappability of failed or failing modules facilitated by routing all IO connections through a Rear Transition Module (RTM).

It is worth examining the enabling technologies that makes a new standard possible. In the case of xTCA a list would include:

1. High Density Interconnects: ZD fabric connectors with multi-gigabit performance, precision edge-card and stacked edge-card connectors with 180 connections per card in less than 3 inches length are marvels of engineering,

\*Work supported by US Department of Energy Contract DE AC03 76SF00515.

2. Multilayer buried via backplanes with 20 or more layers essentially replace a huge network farm in about one foot of width

3. Powerful FPGA Processors with over a thousand ball-grid-array pins with imbedded PCs and SERDES (Serializer-Deserializer) transceivers running at multi-gigabit speeds

4. Intelligent Platform Management standard chips and supports allowing cheap imbedded diagnostic from any remote device to a control room or internet link

5. Sophisticated redundant power components from 48V to less than 1V designed for high availability and remotely tuneable conversion systems to minimize switching noise

All these technology developments were driven largely by the computer and communications businesses where productivity and income are measured in up-time/availability. Although experimental physics can claim credit for many important technical discoveries, it is the engineering industries that bring together the host of technology improvements that enable xTCA to claim a “next-generation” title. It takes a very large highly motivated collaboration to make such a revolutionary new standard possible.

### *Motivation for xTCA Open Standards*

Industry drives new technology because it has proven to lead to very profitable businesses. The only reason industries drive new standards is to cut costs of development and improve profitability and market share. The fact that the \$100B per year telecom industry would develop xTCA in collaboration indicates that it is a very smart thing for business profitability. This manifests itself in three factors:

1. Quicker and cheaper development cycle
2. More funding available for company's value-added products
3. Shorter time to market profitability of those products
4. Better customer environment offering wider choices to optimize system, more competitive cost with *interoperable* modular solutions.

A common platform costs a collaboration participant little or nothing to develop but benefits of adoption can be huge. The ATCA platform in a relatively few years has exceeded \$1B in annual sales. The lessons of economy and time to market (total project readiness) applies equally well to the much smaller and typically cash-strapped physics controls and instrumentation industries.

### *Genesis of xTCA for Physics*

ATCA's superior technical features attracted teams in the US and Europe studying the requirements for the International Linear Collider (ILC) and led to the 1<sup>st</sup> xTCA for Physics Workshop in conjunction with the

2007 IEEE Real Time Conference at Fermilab. At this venue of xTCA tutorials given by Industry ATCA Founders, physics needs for extensions of IO and timing features were identified which led to an invitation to form Technical Subcommittees (Working Groups) under PICMG, which was eager to support expansion into new applications areas. In 2008 DESY made a decision to explore xTCA for possible use on the new XFEL superconducting linac LLRF and control systems, which further spurred interest in a standards collaboration.

At the 2<sup>nd</sup> xTCA for Physics Workshop in November 2008, action was recommended to form working groups under PICMG, which then was completed just before the 3<sup>rd</sup> Workshop in May 2009. In June 2009 work began with weekly hardware and software working group web meetings.

## **LAB-INDUSTRY COLLABORATION - INFRASTRUCTURE**

Between June 2009 and August 2011, a dedicated team of Lab-Industry members met weekly to hammer out hardware options and details. Industry experts were crucial to making rapid progress. The effort bore fruit and by August of 2011 formal extensions for physics have been captured in two new specifications including and one still in progress:

1. A standard MicroTCA backplane connecting AMC Advanced Mezzanine Cards
2. A robust rear IO capability for low-to-high bandwidth mixed signals
3. Standard RTM interfaces for both ATCA and MicroTCA
4. ATCA - Adding and defining fast timing and precision clock synchronization and distribution to the standard dual star and dual-dual star gigabit serial backplanes

In addition to hardware developments, a roadmap for software standards or guidelines was undertaken to simplify interfaces between different high level software platforms and real-time operating systems, common modules and imbedded firmware. These are bearing fruit and guideline releases are expected in the coming months of 2012.

Examples of the new ATCA and MicroTCA hardware are shown in Figures 1-4.

The two new standards, PICMG3.8 for the ATCA RTM, and MTCA.4 for the complete MTCA shelf with AMC-RTM, backplane timing and trigger distribution, have been released to the PICMG Committee for final approval. In addition the ATCA timing guideline under construction specifies a methodology for clock and trigger distribution using “extended option” serial lines in the standard ATCA backplane.




Figure 1: New ATCA RTM Interface PICMG3.8 (Courtesy M. Huffer, SLAC).




Figure 2: AMC-RTM-Shelf Concept MTCA.4.




Figure 3: 12-Slot Redundant Shelf MTCA.4 (Schroff).




Figure 4: 12Slot Redundant Backplane MTCA.4 (Courtesy K. Rehlich, DESY).

### Non-Physics Applications of “xTCA for Physics” Standards

These developments are by no means restricted to physics applications but are expected to be widely used for telecom, industrial control and military applications. Thus the titles of the documents deliberately do not include the word “physics”.

### LAB-INDUSTRY COLLABORATION - APPLICATIONS

#### Generic AMC Development

For applications beyond infrastructure, laboratories and industry partners are developing AMCs and RTMs based on the concept of a “generic” AMC coupled with application-specific RTMs. For example, a generic high speed digitizer AMC card can be used to instrument RF feedback systems, several different kinds of beam position monitors and toroid beam current monitors, through the use of RTM signal conditioning and calibration adapters. A generic FPGA AMC or Industry-Pack AMC Adapter with RTM interface can service a large range of laboratory and industrial control applications. Examples of applications under development are shown in Figures 5 and 6.




Figure 5: Generic AMC + Application Specific RTM: Top- IF Input (A. Young, SLAC); Bottom- RF Down-Mix (K. Rehlich, DESY); Right: ADC-DAC (Strück).

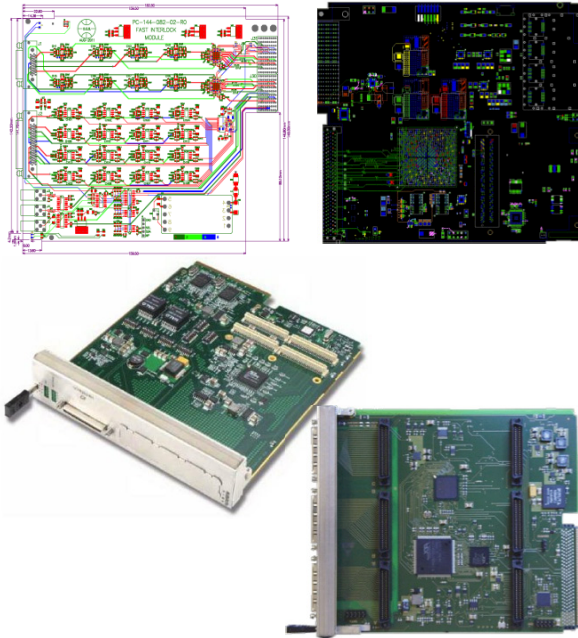


Figure 6: Top-Interlock FPGA-RTM. Bottom-PMC, IP AMC Adapters (Non-RTM Versions for Illustration) (Courtesy D. Brown, SLAC; TEWS, Vadatech)

Again the principle is of shared development costs and goals to reduce development costs, promote rapid development of key COTS components, and quicker time to market or project-readiness.

The lab-industry collaboration is unique and has resulted in a very fast development cycle. Industry is now looking to design for broader use which will expand market and reduce cost for highly cost-sensitive laboratories.

### HW-SW Interoperability Potential

*Interoperability* is a keystone of ATCA-MicroTCA and PICMG has a formal mechanism to test products for compliance. These tests are conducted in an open laboratory hosted by a vendor every 6 months or so, and test the interfaces that assure the management system functions correctly so a non-functioning module cannot damage a shelf or another module in the shelf.

However the biggest potential gains by the lab-industry collaboration stand to be made by increasing *the range of module interoperability* of new designs in hardware, firmware and software interfaces. Standards of the past have failed to make much impact at the software level because standards were basically confined to hardware and a data acquisition transmission protocol; there was no attempt to standardize timing beyond the data transfer protocol and modules did not have a management layer with a standard protocol for internal diagnostics and hot swap was never realizable on the parallel backplane. Moreover standard virtual API's and middleware interfaces designed to a standard template are for the first time within reach.

## PROGRESS TOWARD COTS SUPPORT

A major concern for new-project readiness for (1) Infrastructure and (b) Generic AMC Applications is the level of COTS support. Although in its infancy, due to the strategy of building on the most modern platform available, xTCA can claim a significant level of early support.

Tables 1 and 2 summarize current status of COTS developments as of August 2011.

Table 1: COTS Infrastructure Support MTCA.4

Infrastructure	Description	COTS Availability
Development Shelf non-redundant	6-payload shelf with PU, integral cooling fans	2 vendors
Station Node Shelf dual star redundant	12-payload slot shelf , hot-swappable fan tray(s)	3 vendors
Modular Power Supplies	12V Power Units (PUs) 300/600/900W	2-4 vendors
Hub Controller (MCH) – full featured for timing needed	MCH Controller w/ integral IPMI shelf manager, hot-swap, access to radial timing option	2 vendors Switches for radial timing lines need development
IO Controller Processor (IOC)	Generic AMC processor running Linux, EPICS	2+ vendors
Timing Module	1 or 2-wide AMC (SLAC needs EVR compatible, needs adaptation)	1 <sup>st</sup> units available (U. Stockholm), need COTS sources (2)

Table 2 also shows the variety of Application-Specific Adapters planned for the corresponding RTMs.

Table 2: COTS Generic AMC Support MTCA.4

Generic AMC	COTS Availability	RTM Adapters
10/2 Ch ADC/DAC 16 bit 125 MSPS	1 vendor available 2 <sup>nd</sup> vendor due end FY11	<ul style="list-style-type: none"> <li>RF-IF down-mixers</li> <li>BPM adapter</li> <li>Photodiodes</li> </ul>
4 Ch ADC AMC 14-16 Bit 125 -500 MSPS	1-2 vendors in development	<ul style="list-style-type: none"> <li>BPM single bunch</li> <li>BPM multi-bunch</li> <li>Beam intensity Toroid</li> <li>Beam Length</li> </ul>
FPGA Virtex/Spartan, FMC optional	1 vendor in development	<ul style="list-style-type: none"> <li>Interlocks ADCs 12 bits, 8 ch @60 MSps, 16ch@2KSps</li> <li>Wire scanner interface</li> </ul>
AMC Industry Pack Adapter (2-3 IPs)	2 vendors in development	<ul style="list-style-type: none"> <li>Stepping motor control</li> <li>Vacuum control-monitoring</li> <li>Temp control-monitoring</li> </ul>
AMC PMC Adapters	2 vendors	<ul style="list-style-type: none"> <li>Timing Rx adapter</li> <li>Frame grabber adapter</li> </ul>

### Timing Module Status

Note in Table 1 the one critical module not yet available commercially is the Timing module. A second generation prototype is operational and design of the MTCA.4 RTM compatible version is in progress. The existing 1-Wide AMC unit and the 2-Wide concept is shown in Figure 7.

Note that both versions are fully MTCA.4 compatible; the RTM version simply provides more space and rear access for auxiliary IO timing ports.




Figure 7: Timing Module University of Stockholm-DESY XFEL Collaboration.

## FUTURE STANDARDS CHALLENGES

### *Commitments to Standards*

Among the many challenges facing the introduction of any new standard, a major one today appears to be whether the laboratory engineering physics community has the motivation to adopt a discipline of collaboration among themselves and with industry to advance the standard in a form that benefits all parties economically, as well as achieving the major goal of building controls and data acquisition systems with potentially negligible downtime.

The main goal of such collaboration is robust, optimized infrastructure design and a platform that will serve for the next 1-2 decades through a range of adaptations that will be jointly developed and supported by industry. A closely related goal is to reduce both engineering and production costs of state-of-the-art products for all partners, the same motivation that drove the highly competitive telecom industry to develop and adopt a common platform that offers the customer a high degree of interoperability of competing products.

The physics communities have a new opportunity to take full advantage of this potential to counteract increasing machine costs and decreasing budgets.

### *Risks*

Introduction of any new standard poses risks. The main factors driving the need for standards today are new projects and needs for upgrade of obsolescent systems of 3-4 decades in age. Often for both there is urge to change but not an urge to assume cost and schedule risk. Older systems are hampered by older risk-averse management and support and are less prone to leap to a new system.

Risk is real and the benefits need to be worth it. People available to do the job of introducing a new system, or lack of them, are part of the risk package.

A new system needs to offer compelling benefits and have sufficient champions to drive the development, commissioning and operation.

### *Risk Mitigation*

New platforms come at least a decade apart but in the life of a laboratory that is a short time. Risks associated with new technologies can be mitigated greatly by experimenting with them and trying them on small projects to reduce risk well before a large project shows up with a need. This feeds back into not just maintaining a once every few decades interest but engaging in new developments on a continuing basis, along with colleagues in other laboratories and industry. This low cost investment can clear the way for future successes in keeping laboratory systems in top form all the time, and learning how to manage the complexities of new systems with extensive intelligent platform management and auto-failover hardware-software features. Low level, low cost collaborations should focus on key problems continually and keep a clear view of where they want to take new systems in future.

### *PICMG Collaboration*

Clearly lab-industry collaborations can reap high benefits at very modest cost. PICMG is a unique window into the leading-edge technology world that has shown not just willingness but enthusiasm to support lab initiatives. The investment cost of a corporate membership for a lab is extremely low and basically gives lab engineers and scientists a free ticket to collaborate on technical issues of interest. Once established the relationships can be kept active by interests groups such as an xTCA Interest Group set up by CERN to share information on open source designs. A continuing relationship with PICMG will maintain state-of-the-art currency in technologies needed for future “world class” accelerators.

## ACKNOWLEDGMENTS

Thanks to PICMG Technical Subcommittee leaders Robert Downing (Hardware) and Stefan Simrock (Software), Secretaries Vince Pavlicek and Augustus Lowell. Key technical contributors were Dick Somes, Kay Rehlich, Mike Thompson, Dietmar Mann, Gus Lowell, Zhen’An Liu, Dariusz Makowski, Patrick Gessler, Eike Waltz, Jorges Sousa, Bruno Goncalves, Pete Hawkins, Joe Pavlat, Tomasz Jezynski and Atilla Hidvégi. Many others from the committees attended faithfully to the discussions.