SLAC-PUB-10453
November 2004

Electronics Packaging Issues for Future Accelerators and Experiments

R. S. Larsen, *Life Fellow IEEE,* and R. W. Downing, S*enior Member IEEE*

Abstract— **Standard instrument modules for physics reached their zenith of industrial development from the early 1960s through late 1980s. Started by laboratory engineering groups in Europe and North America, modular electronic standards were successfully developed and commercialized. In the late 1980's a major shift in large detector design toward custom chips mounted directly on detectors started a decline in the use of standard modules for data acquisition. With the loss of the detector module business, commercial support declined. Today the engineering communities supporting future accelerators and experiments face a new set of challenges that demand much more reliable system design. The dominant system metric is Availability. We propose (1) that future accelerator and detector systems be evaluated against a *Design for Availability (DFA)* metric; (2) that modular design and standardization applied to all electronic and controls subsystems are key to high Availability; and (3) that renewed Laboratory-Industry collaboration(s) could make an invaluable contribution to design and implementation.**

Index Terms—Instrument Modules, Data Acquisition, Accelerator Measurement and Control, Modular Design, Machine Availability.

I. DESIGN FOR AVAILABILITY

Future Accelerators and Detectors require huge investments and must perform to very high standards. In the future International Linear Collider (ILC), lost productivity in the 20+ mile machine will cost $50-100,000 per hour. The productivity metric is *Availability (A)*. *A* =1 only if the machine is delivering *usable beam*, and the detectors are recording *usable data*. Thus all systems must be simultaneously functioning most of the time to achieve high average Availability.

Current accelerators claim average *A*=0.7-0.85. The much smaller light sources have taken steps to increase Availability, due to customer demand, and most claim *A*=0.95 or more.

Fig. 1 is an Availability chart from the NLC ZDR[1] showing twelve machine subsystems, each with *A*=0.99 (a Linac is weighted as 3 subsystems due to its size), giving an overall machine *A*=0.85. This is a completely arbitrary choice. We should strive for subsystem *A* >0.99 to raise the overall to as close to *A*=1 as possible. Ultimately cost and complexity will determine the reasonable upper limit, but downtime is so expensive that a modest increase in subsystem capital cost is more than justified.

Scheduled Operating Hours: 6500			
	Weight	Availability	Unscheduled Outage (hours)
e^- Inj, Source and Linac	1	0.99	66
e^- DR and Compressor 1	1	0.99	66
e^- Booster Linac and Comp. 2	1	0.99	66
e^- Main Linac	3	0.97	195
e^- Final Focus and Dumpline	1	0.99	66
Subtotal e^- machines:	**7**	**1**	**458**
e^- Inj, Source and Linac	1	0.99	66
e^+ Source and Linac	1	0.99	66
e^+ Pre-damping Ring	1	0.99	66
e^+ DR and Compressor 1	1	0.99	66
e^+ Booster Linac and Comp. 2	1	0.99	66
e^+ Main Linac	3	0.97	66
e^+ Final Focus and Dumpline	1	0.99	66
Subtotal e^+ machines:	**9**	**1**	**589**
Totals:	**16**	**0.85**	**1047**

Table 17-1. *Availability specifications for the NLC machines.*

Fig. 1: NLC Availability

Availability of a subsystem hinges on three factors: Reliability, Mean-Time-to-Repair, and System Redundancy. Briefly, no matter how reliable a single component can be made, among many thousands of units there will be frequent failures. If these failures bring down the machine, then the MTTR will determine Availability. However, if the subsystem also contains redundant components that can be automatically substituted, then there is no resultant downtime and *A* =1 despite the failure. Accessibility to the failure determines whether it can be repaired quickly, for example by a module exchange even while running, so MTTR has no impact and Availability approaches 1. Therefore a combination of reliability, modular design for quick replacement, and redundancy *at both the component and system levels*, can result in overall system Availability that approaches *A*=1. These principles are valid for any subsystem.

For example, Fig. 2 shows a new telecommunications computing instrument standard, ATCA, designed for a system Availability of 0.99999, while Fig.3 shows a solid state GigaWatt peak power pulse modulator designed to continue operation with the loss of up to 10% of its high voltage modules.

II. MODULAR INSTRUMENT HISTORY

Instrument systems in high energy and nuclear Physics laboratories from the 1960s through the 1980s depended on a triad of modular packaging and electrical standards, namely

R. S. Larsen is an Assistant Director for Electrical Systems at the Stanford Linear Accelerator Center, Menlo Park California, 94025 USA. Work supported by US Department of Energy Contract DE-AC02-76SF00515.

R. W. Downing is a Consultant with Fermi National Laboratory, Batavia Illinois, 60510-0500, and Chair of the VITA34 Task Group, Advanced Packaging and Cooling for Modular Electronics, VME Standards Organization (VSO).

[1] Next Linear Collider Zeroth Order Design Report, SLAC 474.

Contributed to the 2004 IEEE Nuclear Science Symposium and Medical Imaging Conference, Rome, Italy, 10/16/2004 - 10/22/2004

Fig. 2. ATCA 14, 5 & 2 Slot Crates

Fig. 3. GigaWatt Peak Power 500kV Solid State Modulator

NIM[2] (Nuclear Instrument Module), developed in the US; CAMAC (Computer Automated Measurement and Control), started in Europe; and FASTBUS, started in the US. Physics modular instrument efforts began in 1962 at Harwell[3] and CERN[4]. The US NIM Committee was formed under the National Bureau of Standards (now NIST) in March 1964 and published the NIM specification in July 1964[5]. By December 1965, an estimated 30-60% of all nuclear instruments in the US had converted to NIM; a year later the figure was 80-90%[6]. NIM quickly found applications in other fields such as nuclear science, medicine and reactor control systems. The European Standards on Nuclear Electronics Committee (ESONE under EURATOM) first began CAMAC and then proposed a partnership with the US NIM Committee to work toward a common solution.

Many instrument manufacturers in the US initially opposed the lab-developed standards. However, the labs and cooperative industrial partners prevailed and common standards were ultimately adopted in the US and Europe, formalized under IEEE and ANSI in the US and ESONE in Europe and the IEC internationally. Despite a decline in usage in large detector data acquisition, a large installed base of NIM, CAMAC and FASTBUS continues to serve accelerator laboratories, the nuclear medical and power industries and a number of industrial controls applications.

III. WHY STANDARD S?

Any significant technical community can benefit from continued discovery of standard solutions to challenging new problems.

1. Speed of design: A common misconception in the labs is that packaging standards inhibit progress among adventuresome engineers and scientists whose job is to strive for "cutting-edge" solutions. Ironically their purpose is the exact opposite, namely to provide *stable platforms* that bring order to otherwise chaotic cutting-edge solutions. By controlling the electrical-mechanical *perimeters* of modules, standards actually speed development while lowering costs, especially critical to the multi-hundred-million dollar/euro detectors of the colliding beams era. The labs' impressive innovations in detector sensors and packaging should not blind us to the fact that instrument standards have a vital role.

2. *Cost Savings:* Standards benefit users, industry and, by cost savings, the governments that support us. Governments have a fiduciary responsibility to the public to reduce costs, and to not compete with private industry. Lab engineers are responsible to help execute that responsibility, and promoting standards helps labs to efficiently transfer technology spin-offs to the private sector. In the case of the NIM system, Louis Costrell of NIST wrote in 1994:

"An economics study by the AEC in 1982 concluded that the savings (to government) resulting from NIM was at least 1.7 Billion dollars in 1982 dollars. Of course that figure would be considerably higher in current dollars and it does not include the savings in the years since 1982."[7]

3. User Benefits: Here is a short history of benefits from past standards efforts:

a. Engineers and instrumentation physicists from different labs collaborated with one another and with industry partners to produce a complete quality design that no single lab had the capacity to accomplish on its own. The "buy in" by all partners was achieved during this collaboration process.

b. End users collaborated with engineering and industry to prove the new products in actual experiments.

c. The Labs organized Electronics Pools to manage and share standard modules, as well as other lab equipment, keeping a maximum level of usage of the inventories in-hand. Maintenance groups became highly efficient as the plethora of custom designs gave way to common standards.

d. Experiment up time improved as pool-stocked spares were accessible around the clock.

e. In-house designs were compatible with commercial designs, and some migrated to industrial suppliers. Electronics engineers concentrated on circuit rather than mechanical design.

[2]NIM remained an open US AEC/ Department of Energy standard while CAMAC and FASTBUS modular systems were formally standardized under IEEE, ANSI and the IEC.

[3] United Kingdom Atomic Energy Research Establishment. Harwell.

[4] Centre Européene pour les Recherches Nucléaire.

[5] AEC Report TID-20893, US NIM Committee, L. Costrell, Chair.

[6] Development and Current Status of the Standard Nuclear Instrument Module (NIM) System, Louis Costrell, NBS Technical Note 556, October 1970.

[7] NIST Memorandum from L. Costrell to R. Schafer, October 28, 1994.

f. The instant commercial success of NIM greatly aided later step-ups to CAMAC and FASTBUS

The electronics industry itself is proof of the value of standards: Industry, not the physics labs, has persistently broken the barriers of functionality, speed and cost. Without standards at the chip, board and chassis levels to permit product interchangeability, the dynamic electronics industry we know today would not exist.

IV. TECTONIC SHIFT: ELECTRONICS 1985-PRESENT

Few industries in history have evolved as swiftly as electronics. The past decade-and-a-half witnessed the emergence of custom integrated circuits within reach of any laboratory or university; personal computers with the power of mainframes; fiber optic and wireless telephony; and the worldwide web, to name just the more dramatic changes.

By the late 1980s the landscape of detectors had completely changed from many small fixed target experiments to huge colliding beam 4π detectors. As custom chips took hold, standardization of front-end electronics was abandoned as a community activity[8]. On-board channel counts of hundreds of thousands became feasible. Investment cost was considerable but more than justified by the resulting greater cost savings and by improved performance. Readout was by serial fiber-optic link to rear mezzanine cards on FASTBUS[9], which reduced rack space dramatically from prior generations. CAMAC was used mainly for detector control functions.

At this juncture the mass of front-end electronics in large experiments has migrated into custom designs inside detectors. The market for multichannel modules in all three standards has declined and vendors have left the market. VME[10] and in some cases VXI[11] has supplanted FASTBUS due to better commercial support.

NIM appeared 40 years ago and FASTBUS 20 years ago. Meanwhile, driven by relentless demand for growth in speed and functionality in the telecom and computer industries, the electronics revolution continues unabated.

Modular Instrument systems for data acquisition and controls are using architectures that have not changed in two decades. In light of the significant recent progress toward realizing an ILC, new accelerator and detector electronics design standards based on the most advanced technologies deserve a high priority from the Laboratory Engineering community.

V. NEW MODULAR INSTRUMENT REQUIREMENTS

Accelerators need a module architecture that supports a remote stand-alone module or module cluster located a long way from its neighbors. They also need much higher data throughput than previous generations for diagnostic purposes and for discovering subtle trends through mining all stored data regarding beams. Potentially every beam bunch data should be able to be captured and stored.

Detectors need standards for custom shaped modules buried deep inside, as well as for very dense crate clusters for fast track-finding in primary, secondary and tertiary triggering. Density of data throughput is far higher than for Accelerators.

Both need high-speed serial IO that can support redundant links, from any module to any other, parallel channels for maximum throughput, and redundancy for system reliability. The parallel data bus backplanes of existing standards have too low bandwidth and reliability.

Both need liquid cooling options to accommodate new generations of high power density high-speed computing and logic array chips.

Power, grounding and shielding systems are increasingly important as chip speeds continue to increase, especially in instrumentation with extremely low-noise front-end systems such as calorimeters, or high-speed low-noise beam position monitors.

Both need on-board power systems solutions using DC-DC converters with noise performance suited to the application. External bussed systems with varying primary voltages cannot keep pace with changes in chip voltages.

Both applications need designs for radiation environments for instrumentation that for noise reasons has to reside very close to the beams.

Both need modules, controllers, processors, storage and software that can equally well support small and portable systems for bench testing, field maintenance, test beams and small experiments.

Both need sophisticated diagnostics and utilities management to improve overall reliability and availability of systems.

A primary goal is to discover new standard modular solutions that can meet the bulk of needs for both Accelerators and Detectors.

VI. NEW INDUSTRY INSTRUMENT STANDARDS

Two industry groups are designing new standards incorporating many of our requirements. The groups are the PCI Industrial Computer Manufacturer's Group (PICMG), which has developed the *Advanced Telecommunications and Computing Architecture (ATCA);* and the VME Industrial Trade Association (VITA), with its standards organization VSO working on serial packaging standards. PICMG represents telecom and computing markets with partners such as Lucent, Motorola, HP, Sun Microsystems and Intel to name a few; while VITA serves a smaller and more diverse marketplace, including military. The ATCA standard was released in June 2004 with a number of companies already offering basic crate, module and

[8] The NIM and ESONE Committees last met at the Lyon Nuclear Science Symposium in 2000. ESONE was later discontinued and NIM has been dormant.

[9] In the next generation SLAC detector, BaBar, FASTBUS was abandoned for D-Size VME (not VMEp).

[10] VME bus (Versa Module Europa) was introduced by Motorola, Phillips, Thompson, and Mostek in 1981 and later became IEEE Standard 1014-1987. The VME community, driven mostly by non-accelerator users, has grown to support a robust and diverse laboratory, commercial and military user community.

[11] VXI was developed by a consortium of prominent instrument companies to address what were felt to be mechanical, RF grounding and shielding and functional shortcomings in VME for lab instruments. The VXI test and measurement instrument market in 2004 is estimated at $800M. See http://www.vxi.org/.

power system hardware and software; while the VITA groups are still prototyping and working toward a formal Draft Specification. Some key features of each system are summarized:

A. *Modules – PICMG (ATCA)* (See Figs. 2, 4, 5.)

a. Standard 8U motherboard
b. Single input voltage of 48/60V with on-board DC-DC converters for all local voltages
c. Standard interface section of board
d. Scaleable Mezzanine cards allow multi-functions or multi-channels on a board
e. Crate backplane interface uses high-density 5GHz connector for standard serial wire pairs.
f. No backplane parallel bus; crate provides bulk power and serial star or mesh connectors only
g. Serial standards supported: PCI Express, Ethernet, InfiniBand, Star-Fabric, S-RapidIO.
h. Glitch-free hot swapping of modules
i. Robust alignment and keying block
j. RF shielding, grounding, ESD control

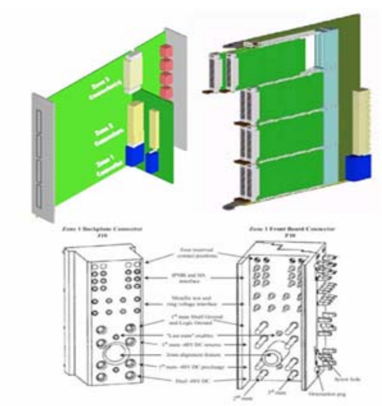

Fig. 4. ATCA Modules & Connectors

B. Crates – PICMG (ATCA)

a. Flexible number, height of modules
b. Bulk 48V DC power supplies, varying capacities
c. Power per board 200W for air-cooled version[12]
d. Intelligent crate bulk power management system delivers prescribed power to module type
e. Smooth power down/up when hot swapping
f. Backplane multiple serial connections completely user-defined star or mesh
g. Redundant circuit paths optional for reliability
h. Cooling: forced air, multiple blowers and pullers, fault-tolerant fan failure.
i. Airflow extensively computer-modeled.
j. System capacity (telecom): 2.5Tb/s
k. System Availability (up-time): 99.999%

[12] Allowable component temperature in this design is 100°C at 45°C ambient.

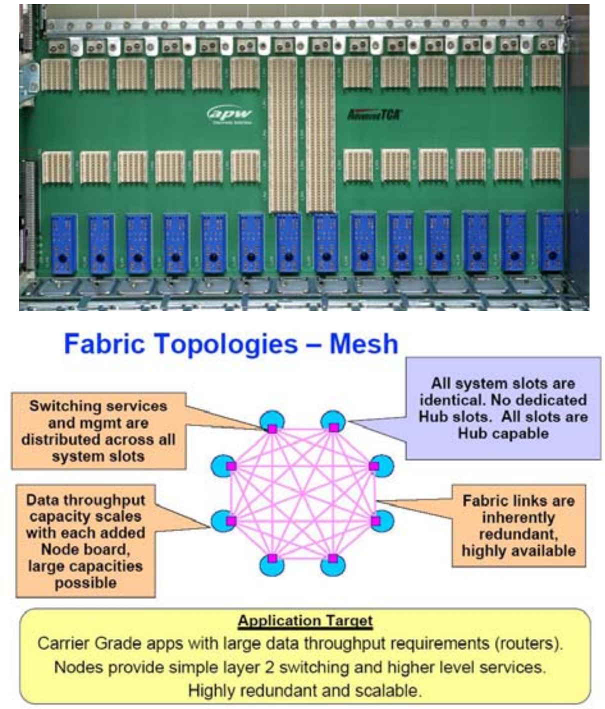

Fig. 5. ATCA Backplane &Mesh Topology

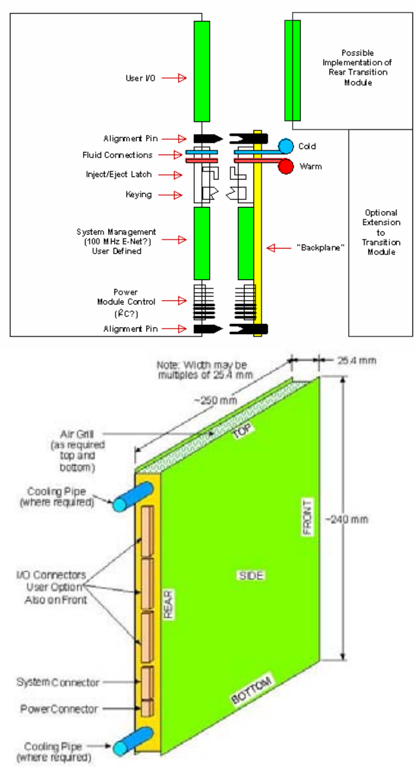

Fig. 6. VSO Liquid Cooled Module Concept

C. *Modules – VSO* (See Fig. 6.)

a. Scaleable Eurocard based Module
b. Unique connector system Tyco "MultiGig" rated up to 3 GHz serial transmission

c. Optional liquid single-board and multi-board cooling solutions, of interest to military
d. Candidates: Heat pipes, negative pressure water, refrigerant
e. Require well-designed component-to-heat sink interface for good heat transfer.

D. *Crates – VSO*

a. Crates with ATCA power management features to but also liquid cooling option from central manifold system.
b. Liquid will cost more but superior to air cooling.

VII. SUMMARY DISCUSSION

Clearly the ATCA and VITA developments are candidates as a base for future accelerator and detector applications. ATCA is an ideal architecture for high-density detector data collection, accelerator control hubs and processor farms, and distributed modules, all communicating via redundant serial links. Software management tools developed for the telecom industry may serve these applications directly.

The ATCA backplane or Shelf holds a Shelf Manager for controlling power and cooling utilities, fault diagnostics and remedial action for the local node. Diagnostic and configuration software are part of the standard product. The Shelf Manager concept is clearly the wave of the future for intelligent high availability systems.

Custom modules deep inside detectors, in addition to using standard form-factors and high-density connectors, can use the standard Shelf Manager functions regardless of topology. Standard features across all detector and accelerator subsystems would have enormous benefits project-wide.

Small or portable experimental and test beam systems must be configured and brought up quickly. A standard set of hardware and software on an ATCA-like platform, with data acquisition and control module mezzanine cards tailored to the application, could serve this need admirably. An imbedded Power PC controller could deliver data to storage by standard serial link . Reconfiguration for the next user would require reprogramming a relatively small set of hardware, using standard off-the-shelf tools.

VITA is aiming for a reliable cooling connection on normal insertion, as well as normal fan cooling. Personal Computers already use heat pipes coupled to fan heat exchangers to remove heat from the hottest chips. Liquid cooled standard modules and crates should see wide usage inside accelerators and large detectors; a number of detectors already use negative atmospheric pressure water systems.

A new Laboratory Collaboration should partner with industry on developing standards to serve the labs for the next decade or more. A combination of ATCA architecture and VITA cooling and connector design seems very promising. By inter-lab agreement, a common feature set could become a highly adaptable *de facto* laboratory standard.

VIII. CONCLUSION

Recent industry developments have opened a door of opportunity for laboratories to replace their aged instrument standards that have served so well but can no longer meet new needs. National Laboratories worldwide should reaffirm the role and importance of standards and commit the necessary engineering and financial resources for modernization.

The principles of Design for Availability should guide the next generation of both instrument standards and electronic subsystems for new accelerators. High Reliability, Low Mean-Time-To-Repair, and component and system redundancy must be balanced against cost to achieve highest possible Availability of entire machines and detectors.

Well designed standards can provide stability through several evolutions of chip design. Industry itself finds standards vital to commercializing state-of-the-art products.

Standards eliminate drudgery allowing engineers to concentrate on challenging R&D design; standards enable industrial partners to bring new products to market more quickly at less cost and risk; and standards help major projects ramp-up more quickly at least cost. Ultimately, standards benefit the end users who must learn to efficiently program, operate and maintain the machines we build.

IX. ACKNOWLEDGMENT

The authors are indebted to the following: Louis Costrell of NIST, Chair of the NIM Committee which in partnership with ESONE developed the CAMAC, FASTBUS and VMEp standards, for providing historical information and critiquing this manuscript; Kenneth Dawson of TRIUMF, member of the NIM Committee software development group who edited both CAMAC and FASTBUS standards, for critiquing this manuscript; to members of the ATCA and VITA organizations who kindly made available graphics of their developmental products; and to all former colleagues in NIM and ESONE who set the stage for the success of standard modular instruments worldwide.