SLAC-PUB-12208

November 2006

High Availability Instrumentation Packaging Standards for the ILC & Detectors

R. W. Downing[1] and R. S. Larsen[2], *Members, IEEE*

***Abstract*– ILC designers are exploring new packaging standards for Accelerator Controls and Instrumentation, particularly high-speed serial interconnect systems for intelligent instruments versus the existing parallel backplanes of VME, VXI and CAMAC. The High Availability Advanced Telecom Computing Architecture (ATCA) system is a new industrial open standard designed to withstand single-point hardware or software failures. The standard crate, controller, applications module and sub-modules are being investigated. All modules and sub-modules are hot–swappable. A single crate is designed for a data throughput in communications applications of 2 Tb/s and an Availability of 0.99999, which translates into a downtime of five minutes per year. The ILC is planning to develop HA architectures for controls, beam instrumentation and detector systems.**

I. Introduction

The International Linear Collider will be one of the largest and most complex accelerators ever built. The sheer number of possible single-point failures requires Design for Availability. An availability modeling program developed for the ILC, *Availism[1]*, has shown that the ILC cannot achieve acceptable up-time unless all systems are designed with high tolerance to individual unit failures. The simulation takes into account the Mean Time Before Failure (MTBF) of all components of all subsystems. The results showed that the ILC would require that the MTBF of some components would have to improve by more than an order of magnitude – deemed to be impractical. As a result ILC designers are exploring High Availability (HA) design of the entire accelerator, including major machine systems, utilities, cryogenics, magnets, power supplies, instrumentation and controls. This design philosophy does not ignore reliability but adds redundant architectural features so most failures do not shut the machine down. *Availism* shows that reliability of all systems must be improved significantly to meet an availability goal of >85%.

The main focus of this paper is the development of HA packaging standards for Accelerator Controls and Instrumentation, but the techniques are equally applicable to data acquisition and detector systems. The Telco industry has driven HA design for many years and has recently collaborated on a new open specification, the Advanced Telecom Computing Architecture, ATCA or AdvancedTCA. Adoption of a commercially available standard electronics is of prime importance for ILC to achieve a common design platform, low engineering costs and high system reliability.

The Baseline model for ILC is a 15 km long machine tunnel with a parallel service tunnel that allows access to most of the control and instrumentation electronics while the machine is running. While statistics show that controls and instrumentation of existing machines contribute less downtime than power systems, the sheer size of the ILC demands an effective reliability improvement in all subsystems by at least an order of magnitude.

For instrumentation and controls, the key high availability strategy was developed decades ago in the form of the standard instrument module. The main availability strategy for modular systems is very short repair time (Mean Time to Repair, MTTR). However this simple operation usually interrupts the machine or experiment to make the repair, leading to a recovery time typically long compared with the module swap. In contrast, the guiding principle for HA design is to make repairs without interrupting machine operations if at all possible, and to keep unavoidable interruptions to an absolute minimum.

Fortunately many new tools are now at hand to address these challenges. While the physics community in the past drove the development of modular instrumentation standards that impacted industry, it has been dormant for more than a decade in making any significant improvements. Industry, on the other hand, has forged ahead into new territories driven by their own need to build scaleable modular systems that can operate even while isolating faults and making repairs. This has resulted in a powerful new open standard instrumentation system, as well as many improvements in the associated power supply systems, both of vital importance to the ILC. Similar design strategies are being adapted to all controls, instrumentation and power electronics systems in ILC, techniques that will be invaluable to detector systems as well.

II. Control System Reliability

The ILC control system touches every part of the machine. Although one usually thinks of the control system as the

---

1. R. W. Downing, Member IEEE, is President of R. W. Downing, Inc., Tucson, AZ and consults for the ILC R & D Program at SLAC

2. R. S. Larsen, Life Fellow IEEE, is Assistant Director for Electrical Systems, Operations Directorate at SLAC and heads the SLAC ILC Electronics R & D Program

[1] *Availism* is the simulation model by T. Himel of SLAC. See: U S Linear Collider Steering Group Accelerator Subcommittee, U S Linear Collider Technology Options Study (USLCTOS), Chapter 4, Availability Design, pp 200-232, March 2004

Presented at 2006 IEEE Nuclear Science Symposium, 10/29/2006-11/3/2006, San Diego, CA, USA

Work supported in part by the US Department of Energy contract DE-AC02-76SF00515

"operator" of the linac it has another equally important use - monitoring. To insure the availability of the ILC one needs to know the "health" of all components. Monitoring is already part of the machine control feedback loop. By monitoring additional elements of the accelerator that are not in the feedback loops one can predict potential failures and schedule maintenance. Changes in voltages, currents, temperatures, thresholds, pulsed waveforms or other parameters can provide early warning of problems.

For example, in monitoring the current draw in a module an unusual increase in the current could signal an overheating component and a potential failure. Redundant hardware in the service tunnel can be auto-switched to the backup and the "ill" module replaced as soon as possible.

Hardware in the machine tunnel has a different replacement schedule. Since an access will require the ILC to be down, monitoring keeps a list of failed units to be replaced when access opportunities arise. By repairing these items during the current shutdown another later shutdown is avoided resulting in increased availability of the ILC.

Since the control system has a dual function it needs to have the availability comparable to safety systems such as Personnel (PPS) and Machine Protection (MPS); these systems can be monitored by Controls but must function independently.

III. COMMERCIAL ELECTRONICS

There is a strong commitment to use commercial equipment as much as possible in the design of the electronics systems for the ILC. Many components such as processors, switches, network modules, can be obtained from various such as VME, CompactPCI, and possibly others. These are older bus based standards and do not have a standard method of control and monitoring. Their hot swap methods are not as robust as needed. The newest modular electronics specification is AdvancedTCA that seems to have the required features for availability, hot swap, control and monitoring built into the specification. It also has wide industrial support and is a growing market. As a result it was decided to use it as the model for purposes of initial testing and costing of the ILC. It will be evaluated again as the project moves to more detailed design and construction.

One can already purchase AdvancedTCA processors, node boards, switches, hardware, and other associated components. Since the specification was driven by the Telco industry the first offerings are geared to them. Special hardware that the ILC needs such as signal processors, motor controllers, analog sensors, *etc* are not off the shelf. The volumes of such hardware for ILC are attractive to contract with industry. Some military contractors are developing AdvancedTCA sensor I/O modules. By the time the ILC needs this kind of hardware the industry will have more products available.

IV. ADVANCEDTCA

Advanced Telecom Computing Architecture specification was developed by the PCI Industrial Computer Manufactures Group as PIGMG 3.0 and released January 2003. Although it was developed for the Telco market the specification has many features that are useful when developing high reliability systems. The Telco market specifies five 9's (this is 5 minutes per year down time) as the minimum availability for their equipment.

To achieve this availability the system is designed for redundancy at many levels. This does not simply mean doubling hardware, but more generally providing N+1 redundancy, one extra unit for 3-4 modules capable of sharing the load. One example is in the power feeds, which can be dual redundant, 2/3 redundant, etc., as optimized for the design. Similarly multiple cooling fans operate at a fraction of full load so that when one fails, speed of the remaining units can be increases to compensate. These actions are performed by an independent diagnostic control known as the Shelf Manager (SM). The SM module itself can be dual with automatic failover and hot swap capable. Connections to the modules in the shelf can be via radial redundant IPMI[2] connection. IPMI communicates with shelf modules and other shelf hardware such as fans, temperature sensors, voltage monitors, *etc*. The SM is connected via Ethernet to a host computer. Some control functions can be programmed into the SM and others directly by the host. Flexibility is included in the specification so the designer can tailor the hardware to meet the requirements.

Data communications are all bit serial with a number of protocols specified. Using serial links reduces single point failures that can occur in parallel buses. These links can be configured with appropriate backplanes as Star, Dual Star or Full Mesh. The Dual Star and Full Mesh offer redundancy that improves availability. The connections between shelves are also serial and can be as redundant as the system requires for the desired level of availability.

AdvancedTCA modules connect to the Shelf Manager for monitoring and control. Each module can be powered on or off in stages. After the module has come to a low power state the SM checks the electronic key for the device to insure its compatibility with the network and sees that the shelf has enough power available for the unit. Only after the system has validated the unit is it brought to full power and enabled on the network. In addition to enabling and disabling modules, the SM monitors the "health" of the device, *e.g.* voltages, currents, control registers, *etc*. Some actions can be taken by the SM autonomously; others are communicated to the host.

The rich set of features that AdvancedTCA provides has made it the choice for initial design of the ILC control system. In addition, the modules appear adaptable to many physics

[2] IPMI is an Intel control bus that uses I^2C physical layer with Intelligent Platform Management Interface protocols.

applications, and the diagnostic management functions of the SM can be applied to many other subsystems.

V. ADVANCEDTCA SHELVES

Fig. 1 and Fig. 2 show a 14 slot shelf. In Fig. 1 the SMs are at the left side. Some other manufacturers place the SMs in the lower fan area. The specification only specifies the functionality, not the placement. The figures also show the dual star backplane configuration, most likely the one the ILC will use. For most nodes the bandwidth appears adequate, and if more is needed the full mesh can achieve 2 Tb/s per shelf.

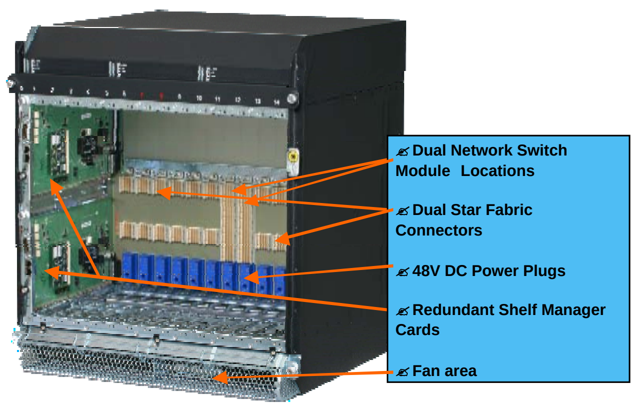

Fig. 1. Front view of AdvancedTCA shelf with a Dual Star Backplane

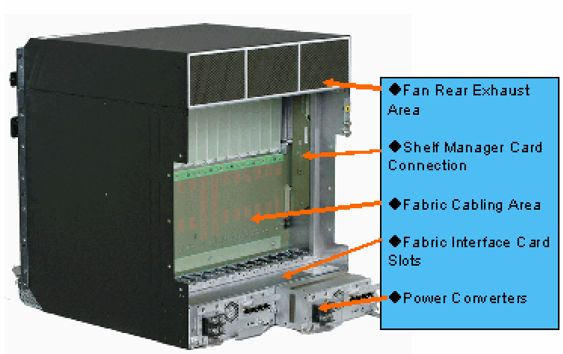

Fig. 2. Rear View of AdvancedTCA Shelf

Six variable speed fans are in the top of the shelf pulling air from the bottom plenum. The fans are hot pluggable. At the bottom rear are the power connections. These are redundantly wired through the backplane to every module for full dual redundant sources. Power modules provide filtering for the 48 volt feeds and prevent conducted noise being fed into the power system from the DC-DC converters on the module.

An area above the backplane is reserved for user connections. A Rear Transition Module (RTM) can be fixed in this area and connect to the front module. It remains in place as the front module is inserted and removed. The RTM can have circuitry on it, however, rear service may be difficult. The RTM will mainly be used as a cabling point and can be viewed as a form factor changer. In this configuration the cables remain fixed as electronic modules are serviced resulting in greatly improved reliability. Past experience has shown that plugging and unplugging cables is a major source of failure both from miss connections and damage to the cable and connectors.

Fig. 3 shows 2 and 6 horizontal slot AdvancedTCA shelves along side the 14 slot vertical shelf for comparison. Other sized are also available from several vendors. In some the modules are horizontal and fewer in number. Cost savings can be achieved if a full 14 slot shelf is not needed. These shelves have all the features of the larger shelves, only fewer slots. If one uses a 24 inch wide rack instead of the common 19 inch, a 16 slot shelf can be used.

Fig. 3. Other AdvancedTCA shelf configurations

VI. STANDARDS AND LIFETIME

Semiconductor lifetimes in the market are less than 5 years for complex functions such as processors, programmable logic, *etc*. Therefore planning for upgrades for continued maintainability becomes a fact of life. Standards can help mitigate this problem. Generally the I/O interface has a long lifetime so new and old modules can communicate. This is a requirement for a standard. The main carrier boards and interfaces should have a long life. Placing the short lifetime hardware on pluggable daughter allows upgrades without changing the mother board.

Standards committees are very reluctant to alter the interface specification because of backward compatibility. VME has been around for 30 years and one can still use many old modules in new VME systems. The old card will not have all the functionality of the new, but, it will communicate with the VME bus so a new processor can talk to it. One can expect the same scenario to unfold for AdvancedTCA.

PICMG AMC.0 is a specification for daughter cards, the AdvancedTCA Mezzanine Card (AMC, see Fig. 4). AMC's have most of the features specified for the AdvancedTCA mother- board and connect to the SM for control and monitoring. AMC's are hot swappable. They can also be configured in a mini-shelf, micro-TCA, without the motherboard. However, micro-TCA does not support rear I/O and therefore is less likely to be used where rapid replacement of failed electronics modules is a requirement. Having to remove and replace cables somewhat compromises the high availability criteria since it poses another failure point in the system.

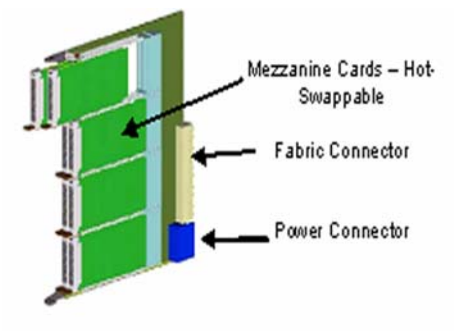

Fig. 4. AdvancedTCA Carrier Module with Eight AMC Moules

A standard motherboard can be personalized with special function AMC's such as A to D and D to A converters to lower design costs and improve flexibility.

AdvancedTCA specifies only 48 volt DC power. Each module contains converters for the voltages required. This system has many advantages and has been adopted by the Telecom industry. A major reason is that semiconductor voltages are no longer standard and power on processor chips is very high, sometimes greater than 100 watts. Supplying very low voltages at the backplane is not practical from both the number required and the currents needed. The on board conversion from 48 volts to 2 volts or less minimizes voltage drop and power loss from high currents. This also means that as voltage demands change in future the base shelf system stays in place and only the new modules have to provide the new voltage. The 48 volt power system thus insures long lifetime and compatibility.

VII. Proposed Control Systems

The proposed ILC control system is using AdvancedTCA for the prototype. Many features of AdvancedTCA are also used in special, non-AdvancedTCA hardware. The hot swap, monitoring, control are being look at in power systems. This section will only discuss the data and control architecture.

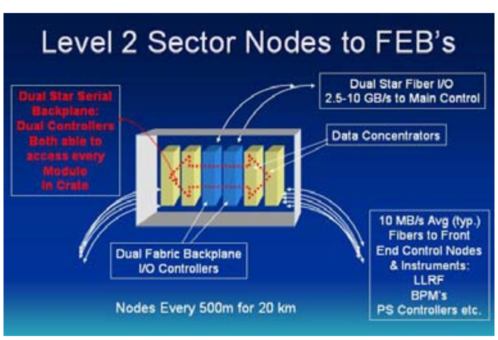

Fig. 5. Detail of AdvancedTCA Node Shelf

The heart of the system is an AdvancedTCA shelf with a dual star backplane, redundant SM's, redundant power, and redundant fiber optic links to the sectors switches and main control. Links from the front-end sensors are a mix of copper and fiber optics as required. Some of these are redundant where availability is paramount.

Fig. 5 shows a typical node with its fiber optic connections. These nodes are of two types, one for the in-band data that controls the linacs and the other out-of-band for control and monitoring functions. The SMs would hook to the out-of-band system.

Fig. 6 is a representation of the connections of these nodes to the front end and the control computers. The lines labled FO are the redundant fiber optic links. The connections to/from the Technical Equipment (TE) are a mix of fiber optics and copper. The switch in the Technical Equipment section could be AdvancedTCA depending on economics.

The AdvancedTCA shelf in the Controls area has processors, and sensor modules. The BPM is shown as attaching here but more sensors could be connected. The connection to the Switch Shelf from the TE switch contains data that is used by the Management data path. This is indicated by the arrow connecting to the SM part of the shelf. The approximate numbers of shelves are shown at each level.

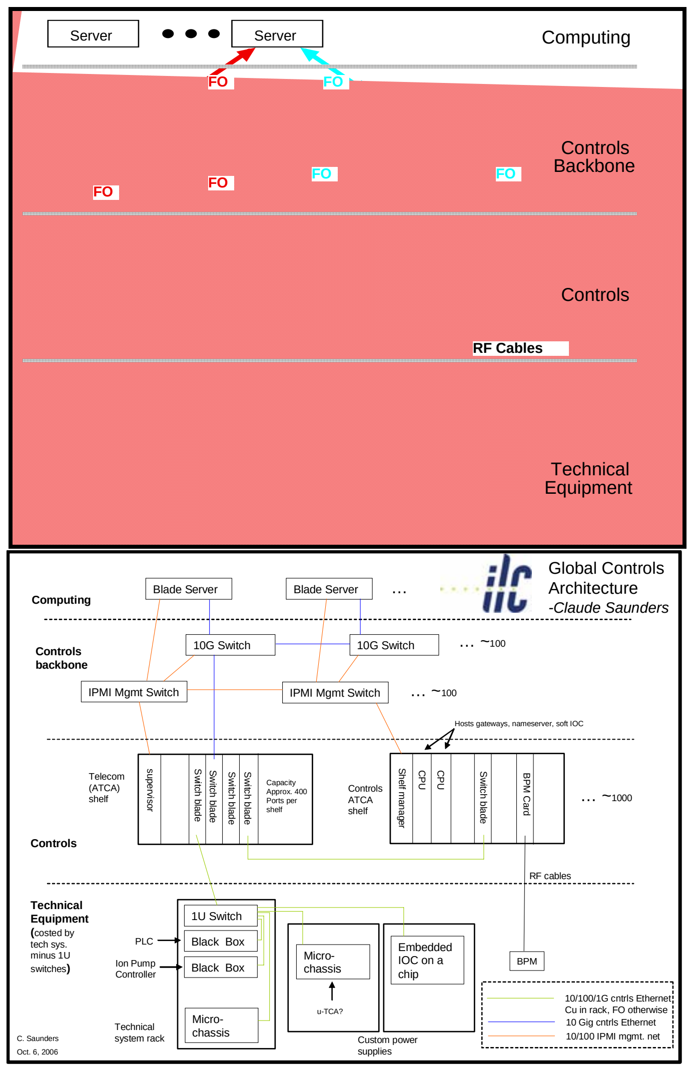

Fig. 6. General ILC Controls Architecture

New HA Power Supply systems are of special interest not only for controls and instrument modules but also for low power systems used in detector applications and intermediate supplies used in accelerator magnet systems. Industry has been active in producing very sophisticated power products for board level systems, with an intelligent I^2C management network down to the smallest converters at the chip level. The architecture is shown in Fig. 7. The standard 48 V supply used in ATCA is first reduced on the carrier board to a level of typically 12 volts for further down-conversion by smaller POL (Point of Load) chip-size hybrid converters. The latter are current sources that can be paralleled to tailor current capacity to the load. The I^2C management network allows such operations as adjusting clocks to minimize coherent noise for fast processors and logic chips, as well as for analog to digital devices; and to isolate a faulted unit to prevent taking down the entire board.

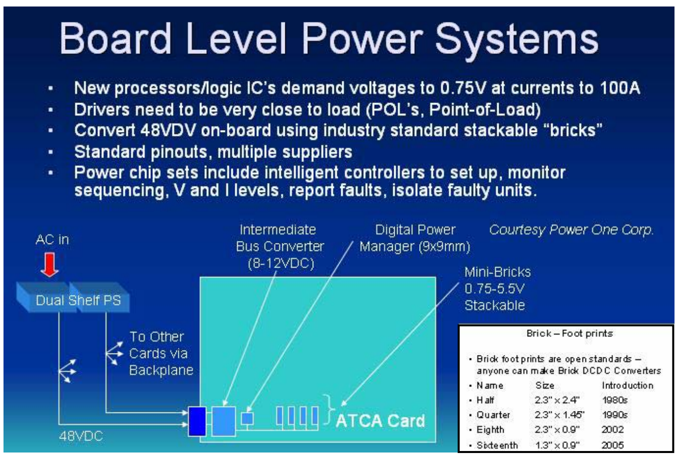

Fig. 7 ATCA Board Level Intelligent Power System

VIII. Projects

Work is in progress in testing and understanding the AdvancedTCA at several institutions. A collaboration between SLAC and UIUC (University of Illinois at Urbana-Champaign) aims to understand the AdvancedTCA system hardware and shelf management and acquisition software. Testing is being done on the SM and network boards. As soon as this is finished an adapter board will be designed for VME and/or CompactPCI. Since there are many modules of use in data acquisition test systems available in these standards and not in AdvancedTCA this is a shortcut to evaluating practical systems needed for ILC testing. Later these may be converted to true AdvancedTCA boards and/or Mezzanine boards.

Fermilab Accelerator Instruments group are developing an ATCA fast four channel A/D converter board with 14 bits resolution and a DSP for signal processing. It will be also be used to evaluate any noise issues with this type of hardware residing in the ATCA shelf.

DESY is studying porting an existing Low Level RF application design for the new XFEL accelerator to an ATCA board. The target is a unit called SimCon 3.1 with 10 A/D converters, 4 D/A converters and a Virtex II FPGA, a PowerPC™ processor and several digital I/O ports. Also the DESY Controls group is actively evaluating ATCA options and investigating products from several vendors.

KEK in Japan is reportedly studying AdvancedTCA board designs, considering both new designs and conversion of existing designs into Advanced TCA format.

ANL has been modeling the ILC Controls Architecture based on ATCA and researching the system software needed to operate the AdvancedTCA control system and associated hardware. The ATCA platform is being used as the cost model for both controls and beam instrumentation standard modules. However, considerable work is needed to qualify the system for instrumentation, since the Telco industry does not use sensitive analog instrumentation as found in the front ends of accelerators and detectors. As with past standards such as CAMAC, FASTBUS, VME and VXI, any new platform has to be newly qualified for analog amplifier, ADC and DAC performance.

IX. Conclusions

The ATCA industry open specification has the desired features for a next generation accelerator controls and instrumentation system, especially for expensive mega-machines where machine performance must considerably exceed today's accepted standards. The task of the ILC R&D groups is to test its features and progress toward strongly integrated solutions for both hardware and software. Many of the system features will also solve major issues of reliability and availability for future detector modular instrumentation and imbedded systems.

X. Acknowledgments

The current program relies on ongoing discussions and joint planning with the Global Controls, LLRF and Availability Technical Groups; as well as on working collaborations with LLNL, ANL, FNAL, KEK, DESY and UIUC. The special contributions of J. Carwardine and C. Saunders of ANL, B. Chase and M. Votava, of FNAL, M. Haney and M. Kasten of UIUC, S. Michizono and T. Naito of KEK, S. Simrock and K. Rehlich of DESY, many dedicated ILC colleagues at SLAC and all the named institutions are gratefully acknowledged.

XI. Reference

[2]. Larsen, R.S. and Downing, R.W., *Electronics Packaging Issues for Future Accelerators and Experiments*, Proceedings of the 2005 Nuclear Science Symposium, November 2004, Rome