SLAC-PUB-11263

Instrumentation Standard Architectures for Future High Availability Control Systems

R. S. Larsen, Life Fellow IEEE
*Stanford Linear Accelerator Center, Menlo Park California USA**

Abstract— Architectures for next-generation modular instrumentation standards should aim to meet a requirement of High Availability, or robustness against system failure. This is particularly important for experiments both large and small mounted on production accelerators and light sources. New standards should be based on architectures that (1) are modular in both hardware and software for ease in repair and upgrade; (2) include inherent redundancy at internal module, module assembly and system levels; (3) include modern high speed serial inter-module communications with robust noise-immune protocols; and (4) include highly intelligent diagnostics and board-management subsystems that can predict impending failure and invoke evasive strategies. The simple design principles lead to fail-soft systems that can be applied to any type of electronics system, from modular instruments to large power supplies to pulsed power modulators to entire accelerator systems. The existing standards in use are briefly reviewed and compared against a new commercial standard which suggests a powerful model for future laboratory standard developments. The past successes of undertaking such projects through inter-laboratory engineering-physics collaborations will be briefly summarized.

Index Terms—Instrument Modules, Data Acquisition, Accelerator Measurement and Control, Modular Design, Machine Availability.
**R.S. Larsen is head of Electronics for the SLAC International Linear Collider R&D Program and a Laboratory Assistant Division Director. Work supported by US Department of Energy Contract DE-AC03-76SF00515.*

I. MOTIVATION FOR DESIGN FOR HIGH AVAILABILITY

High Availability design requires analyzing the reliability-availability of each subsystem in view of the selected overall systems availability goal and a strategy developed for implementation. This is most crucial in large expensive enterprises where the cost of lost production ("Opportunity Cost") is prohibitive[1]. The problem is more severe the larger and more complex the machine. If efficient solutions can be found for the largest machines, then the by-product subsystems can be useful and cost-effective in smaller machines and experiments as well. Therefore the cost impact of new development can be easily justified for a large machine and the benefits reaped by the experimental community as a whole. This in fact was the impetus for past standards developments. The current proposal for a large new machine, the ILC, gives the entire scientific instrument community a unique fortuitous opportunity to make a quantum leap in instrument and system design that will serve the next generation of experiments and machines both large and small. It should also offer a platform to which existing valuable instrumentation designs can be easily ported. The standards of the past grew from the bottom up as new electronic devices became available. Today, the urgent driving force needed for the next generation of standards is to take advantage of the spectacular advances of the last decade in on-chip processing, analysis, memory and high speed communications transceivers and media including fiber and wireless.

II. A SYSTEMS APPROACH TO HIGH AVILABILITY DESIGN

The principles of HA design are simple. Every unit in a system has an MTBF, Mean Time Before Failure, also called Reliability. The unit can be a resistor, a computer or a power supply or a whole subsystem. The MTBF of a machine is the product of the MTBF's of its component parts. Another simple parameter is MTTR, the mean time to repair a failed component that is interrupting operation. Availability, the parameter if interest, is the proportion of time that an entire machine, or machine plus experiment, is actually operating correctly, compared with the planned time of operation, and its maximum is 100% or A=1. Simply, Availability A= (MTBF-MTTR)/MTBF.

It is important to think of large scientific machines as production plants like an electric power plant or a refinery, with a desired Availability goal of A approaching 1. Typical high energy physics machines today operate with A's of around 50% of calendar time, while light sources (much smaller machines with commercial customers) would not stay in business without A's of 0.95 or more. None of these older machines were designed with HA as the chief goal; they were most commonly designed for lowest cost with a still acceptable up time, often with serious compromises to availability (insufficient spares and maintenance coverage, long cool-down and MTTR, for example).

Figs. 1 and 2 depict a 3-level controls and instrumentation system for a large accelerator or experiment. At Level 1 is the centralized control room or data acquisition computer complex consisting of a large farm and associated interfaces. At Level 2 are the node control elements providing processing and links to controls and data acquisition elements for a sector

[1] Larsen, R.S. and Downing, R.W., Electronics Packaging Issues for Future Accelerators and Experiments, Paper NS-24, IEEE 2004 Nuclear Science Symposium, Rome.

Presented at RT 2005 IEEE Real Time Controls Conference, 5-10 June 2005, Stockholm Sweden
Work supported by Department of Energy contract DE-AC02-76SF00515

of a machine or a subsystem of a large detector. At Level 3 are the front-end modules, which gather and process raw machine or experimental data and receive controls and calibration instruction from the Level 1 via the node concentrators at Level 2. All communications are based on dual star topology except that additional mesh elements are easy to add as well. Fig. 1 shows this possibility.

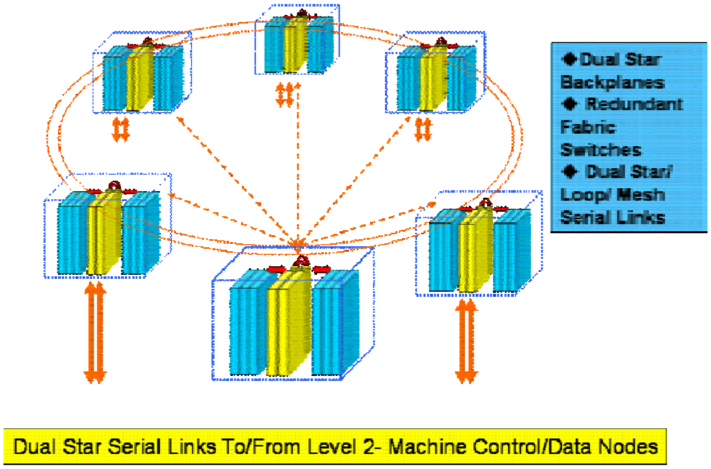
◆ Dual Star Backplanes
◆ Redundant Fabric Switches
◆ Dual Star/ Loop/ Mesh Serial Links
Dual Star Serial Links To/From Level 2- Machine Control/Data Nodes

Fig. 1. Level 1 Central Control Computer Topology

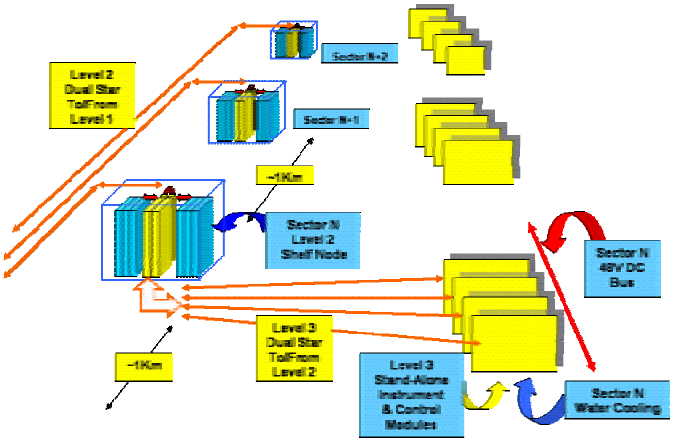
Level 2 Dual Star To/From Level 1
Sector N+2
Sector N+1
~1Km
Sector N Level 2 Shelf Node
Sector N 48V DC Bus
~1Km
Level 3 Dual Star To/From Level 2
Level 3 Stand-Alone Instrument & Control Modules
Sector N Water Cooling

Fig. 2. Level 2 Sector Nodes & Level 3 Standalone Front-End Module DAQ and Control

The most basic elements along an accelerator are standalone modules near the signal source serving a dedicated purpose, such as detecting and processing signals from a single beam position monitor, RF pickup (fast data) or a vacuum pump or magnet mover (slow data). All data are assumed to be digitized at the front end. Communication links also carry precise clock and timing information, as well as high and low-speed data for data and controls respectively. For pulsed machines timing signals with a resolution of about 10 nanoseconds can be carried with the control data, but higher precision data is carried on a dedicated temperature compensated star fiber network. The control data protocols are designed to be robust and noise-immune.

In terms of Availability, the larger the complex, the higher the risk that failures will occur so frequently that repairs cannot keep up and A→ 0. For example, with 1000 power supplies in a system each with a MTBF of 50,000 hours and replacement time averaging 2 hours, failures would occur about every 50 hours and the 2 hour repair would drop Availability from that one subsystem alone to A= 0.96. With 10 such subsystems, the average Availability of a major system (e.g. linac) would be 0.67, and with 16 such systems making up a large machine like the ILC, A (full system) would be .0016 or practically zero.

Of course the reason that large machines are able to run today is that critical components, instruments and systems are designed with some redundancy features to minimize MTTR. For example, control computers in critical applications are either fully redundant or 1/N redundant, where an extra processor among a group of N can carry the load if one unit should fail, or carry the load with degraded performance if the main system should fail. The RF power stations in a linac or storage ring often have 1/N redundancy so operation can continue with one (or more) units failed. In the case of the linac, there must be sufficient spare stations to maintain the beam energy; while in a storage ring the current can reduced to keep operating when a single RF station fails. In the controls and instrumentation areas, quickly interchangeable standard instrument modules have had enormous benefit in reducing MTTR and raising Availability. However, most critical power components below the systems level, with the exception of instrument modules at the crate level, have not been designed to HA principles and therefore pose a risk to overall Availability. This leaves an enormous number of components in a typical machine as sources of “Single Point Failures” in which the failure of a single element brings down the entire machine. Examples are non-redundant power supplies, modulators, controllers, timing modules and links, machine protection systems and beam sources. The goal of HA design is to effectively eliminate the impact of single point failures by a combination of imbedded intelligence quickly detect failure or impending failure, built-in unit redundancy, non-invasive replacement/ repair strategies, and system level redundancy for critical components such as klystrons which have a known low MTBF.

III. CHOOSING AN AVAILABILITY GOAL

Taking the example of the ILC, 16 Systems, each with an Availability that is the product of an assumed 10 subsystems, make up the full system, to which is assigned an Availability goal of A ≥ 0.99. The required average Subsystem Availability needed to achieve the full system Availability of A ≥ 0.99 is between four and five 9’s, as shown in Fig. 3. As will be seen, the choice of a goal is arbitrary, and once chosen, the subsystem performance requirements are dictated. Obviously these are averages; not all subsystems will behave equally.

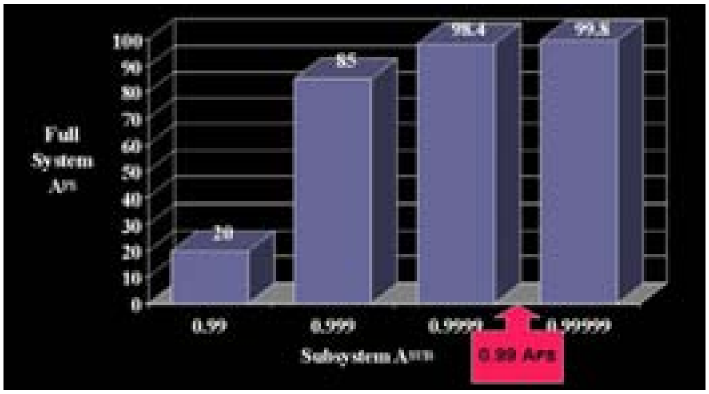

Fig. 3. A Full System vs. A Subsystem for ILC (160 Subsystems)

IV. BASIC STRATEGIES OF HIGH AVAILBILITY DESIGN

The basic strategies are summarized:

Every *Unit*, from instrument module to computer to power modulator or power supply, should include 1/N redundancy in a quickly replaceable modular design.

Every critical *data or timing link* should have an available redundant path with automatic switchover capability.

Every *Subsystem assembly* where multiple components cluster, such as an instrument crate or a multi-board solid state modulator, should have imbedded intelligence to detect and locally respond to internal failures such as

Re-route data or power flow around a failed module or channel,

Automatically power down a failed module for removal without interrupting operation of the shelf or crate, and

Automatically notify main control of any internal failure so maintenance can be prioritized.

Every *Subsystem* should have 1/N redundancy so an entire unit can be taken off-line and replaced without interrupting machine operation.

Obviously, full dual or triple redundancy of all systems would double or triple basic costs, which would be impractical. However it is unnecessary, because 1/N redundancy coupled with other strategies will be more effective in cost-performance.

V. SPECIFIC STRATEGIES FOR HA DESIGN

Referring to Figs. 1 and 2 depicting only the controls and instrumentation systems, specific strategies are summarized as follows:

The Central Computer system shown in Fig. 1 should have both a 1/N redundancy of processors plus the ability to re-route traffic if a failure occurs.

A second level of 1/N redundancy should be incorporated in the form of a complete identically configured crate or shelf.

The remote Shelf Nodes shown in Fig.2 could be connected with a second link to nearest neighbors for backup. However if a shelf fails it would be best to include a completely redundant shelf as well. This is not a huge expenditure as there are not that many of these sparsely located nodes. A second node would also be useful to maintain communication with critical sector safety systems for both machine and personnel.

The design of front-end modules could also include 1/n redundancy internally in terms of the number o channels. To make this effective would require a re-routing system at the inputs for the analog signals. In some cases this is simple and in some cases problematical due to compromised noise performance. However, slightly compromised noise performance in a critical application is far preferable to a loss of mission operations.

Some commercial implementations are illustrated in Figs. 4-5. For the Central Processor application, the new Advanced Telecommunications Computer Architecture commercial package, ATCA, is a strong candidate. This standard, released in June 2004, was designed specifically for High Availability with a typical system goal of A=0.99999. This is the kind of figure that is need for a typical subsystem of a large machine or detector with 10-15 subsystems to obtain an overall ILC Availability of A>/=0.99 or 99% up time (Fig. 3). The strategies include intelligent management of the crate or shelf by automatically rerouting traffic around a failed channel, powering down a failed board so it can be removed and replaced while the remaining system keeps operating, and sensing of fan airflow and temperatures and controlling fan speeds to compensate for a failed fan. Its main shortcoming is that crate temperatures can get unacceptably high for experimental applications where the hardware is not easily accessible, or is totally inaccessible during long periods of running, such as inside an accelerator tunnel or detector.

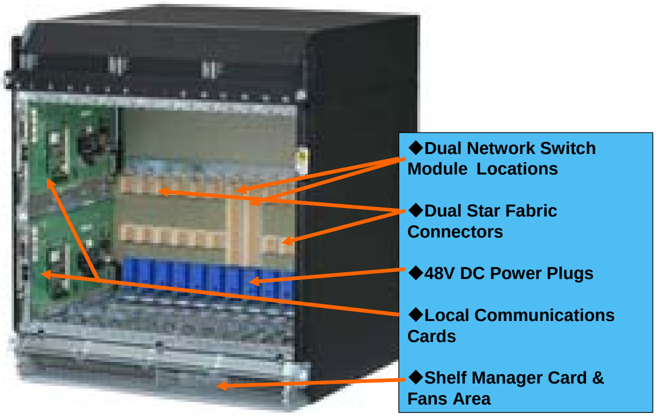
◆Dual Network Switch Module Locations
◆Dual Star Fabric Connectors
◆48V DC Power Plugs
◆Local Communications Cards
◆Shelf Manager Card & Fans Area

Fig. 4. ATCA Shelf (Crate) Features

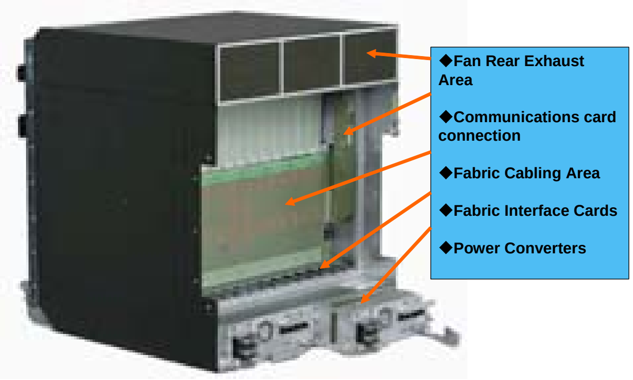

Fig. 5. ATCA Shelf Rear Features

The Front End Stand-Alone Card system concept is shown in more detail in Fig. 6. Here we visualize a dual 48V (e.g.) power bus system running the length of a machine sector, upon which cards can be placed either singly or in groups as the situation dictates. The idea is to minimize cable runs from the sensors. The card has all connectors for signals, power and cooling on the bottom side so the module can be machine-manipulated if desired for removal and replacement with the machine running. Communication is by dual serial bus, fiber if there is no radiation and wire or wireless if there is radiation that would degrade fiber. The electronics itself needs to be protected from neutrons and gammas so concrete shielding is visualized. Timing is also provided if the needed precision is 1 nanosecond or less. All timing and data go to the Sector Hub.

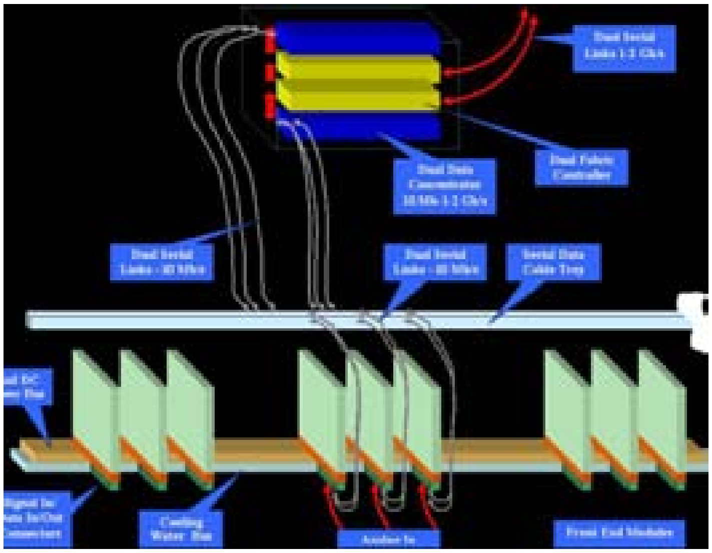

Fig. 6. Front End Stand Alone Card Concept

VI. New power systems diagnostic controller

In addition to the HA strategies embodied above, a new strategy is proposed, namely to add a special diagnostic controller into every power unit, such as a modulator or large power supply. Moreover it is proposed to use a modular architecture for all such devices in order to employ a 1/n redundancy strategy at the unit level. Like the ATCA system, the modules of critical power supplies that would interrupt the machine if they failed would be hot swappable. With this approach, it is easily shown that a power supply subsystem of 1000 supplies can obtain the needed Subsystem Availability of >0.9999 to obtain an overall machine Availability of > 0.99,as shown in Fig. 7.

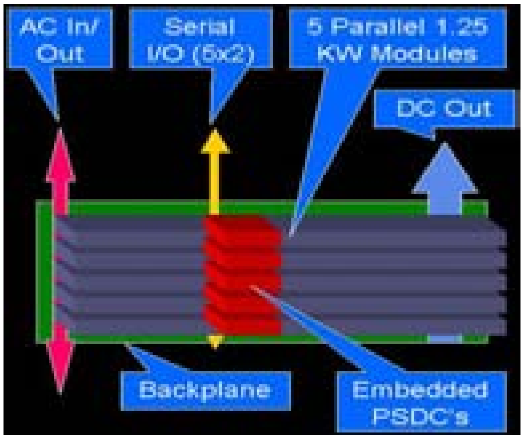

Fig. 7. Modular 4/5 Redundant HA Power Supply with Embedded Diagnostic Controller

With the basic strategies outlined, extremely high subsystem Availability can be obtained without the necessity of unrealistically high Reliability of the basic components. Choices can be made between these two when engineering the Subsystems. Another example of a very important component is shown in Figure 8. This is a new style Klystron Modulator called a Marx bank, in which identical cells are stacked to provide the total voltage and current needed for an ILC 10MW pulsed tube. The important feature of this design is that a very small number of extra cells can yield an HA design that meets all the requirements of Subsystem Availability.

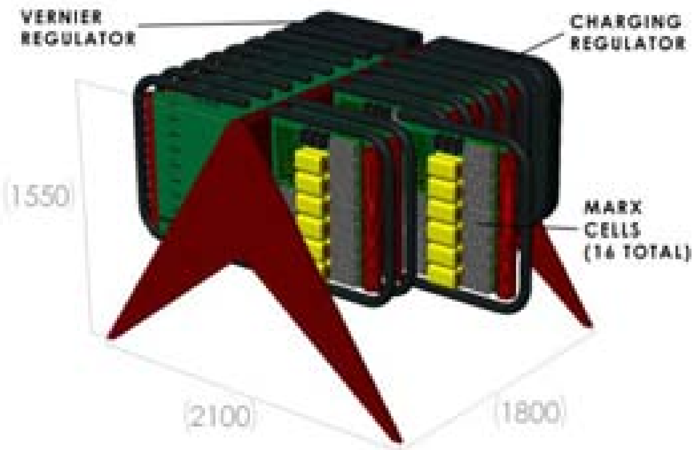

Fig. 8. HA Marx Bank Modulator for ILC

In addition to HA architecture, this transformer-less unit has the advantages of smaller size, air cooling, higher efficiency, reduced size and weight, and reduced operating cost.

VII. SUMMARY OF HA MACHINE DESIGN

To obtain High Availability of an accelerator, all subsystems must adopt an HA architecture featuring modular unit design, redundancy, and hot swap capability for any single unit that is mission-critical to the machine. With foresight and investment in the basic components, it is reasonable to expect performance near A=1 from a machine of any size, insofar as its electronics and power systems are concerned. In the case of accelerators, the enormous cost of a crippled production machine easily justifies HA design strategies at the beginning of a new project.

VIII. *SMART STANDARDS* AND HA ACCELERATOR SYSTEM DESIGN

Instrumentation standards are an extremely important factor in engineering large projects. Custom designs of basic components such as modules, crates or power systems should be considered only when unavoidable. The long-term cost of maintaining individually engineered units where unnecessary leads to inefficient operation and loss of production. Instrumentation standards have serve well in physics research in the past but have been allowed to become obsolete as technology moved on.[2] It is now time to reconsider a set of standards that support HA design in large systems like the ILC and its detectors. At the same time, a broad collaboration should undertake this work so that industrial support is viable and so that many research users can benefit. This requires careful choices.

A common error in standardization is to become over-prescriptive and thus lock in technologies that are bound for obsolescence. This can be avoided by recognizing a primary and a secondary level of standard, on quite prescriptive and the other allowing options so the standard can flow with technology and extend its lifetime accordingly. SMART standards need to be developed with the following goals:

1. A professionally engineered, commercially available, inexpensive platform and toolkit
2. Able to accommodate technology change without losing functionality
3. Not over- or under-specified to inhibit designers
4. Primary Needs
   a. Power system, connectors & communications architecture
5. Secondary Needs
   a. Form factor and protocols

The areas we have identified for standards in accelerators may be summarized as follows:

1. An HA ATCA-like redundant Backbone for Central Farm, Sector hubs/ concentrators

2. An ATCA-like redundant communications, fast timing and mechanics for stand-alone FE cards. Fast data and timing need wire or fiber; slow data may be wireless.

3. A Redundant power system at crate level and a similar distributed system in tunnels

4. Protocols, with type and speeds chosen to suit the application

IX. STANDARDS IN LARGE DETECTOR DESIGN

Modern large detector design appears to defy standardization but that is because this aspect has been neglected for the last decade or more. There are two primary areas *inside* detectors where standards can and should be applied:

1.An ATCA-like redundant communications backbone, minimized in size as needed for application.

2. A. small intelligent card containing FPGA subset that can drop into custom hubs deep inside detector. This may have to be radiation hardened. The link may be wire, fiber or wireless.

3. A Primary dual power distribution system, e.g. 48V for lower currents, and

4. A Secondary dual power converter system to serve internal nodes for groups of Front End cards, accessible for replacement without opening the detector. Hot swappable components may also be a consideration even though detectors are not normally accessible while running.

X. CONCUDING REMARKS: CCOLLABORATION

Both the ILC project, should it become reality, and the development of standards will need strong international collaboration Not everyone will participate in the ILC which, although an important initiative that helps define the issues, .is only one part of the standardization needs being discussed. The standardization goals are more imminent and should be pursued by a broad community that will derive enormous mutual benefit from their efforts.

A revived standards program should be pursued vigorously as a community goal worthy of investment irrespective of the fate of any single large project.

XI. ACKNOWLEDGMENT

The following colleagues are acknowledged for their helpful discussions and generous sharing of information: E.J. Siskind, NYCB Real Time Computing; R. Downing, Downing Inc.; R. Chestnut, R. Humphrey, G. Leyh, R. Cassel of SLAC; C. Brooksby, LLNL/BN, and H. Shoaee, LANL.

[2] Ibid, R. Larsen & R. Downing.