High Availability Electronics Standards

R.S. Larsen[1, 2]

Stanford Linear Accelerator Center, 2575 Sand Hill Road, Menlo Park CA 94019

larsen@slac.stanford.edu

Abstract

Availability modeling of the proposed International Linear Collider (ILC) predicts unacceptably low uptime with current electronics systems designs. High Availability (HA) analysis is being used as a guideline for all major machine systems including sources, utilities, cryogenics, magnets, power supplies, instrumentation and controls. R&D teams are seeking to achieve total machine high availability with nominal impact on system cost. The focus of this paper is the investigation of commercial standard HA architectures and packaging for Accelerator Controls and Instrumentation. Application of HA design principles to power systems and detector instrumentation are also discussed.

I. INTRODUCTION

The International Linear Collider will be one of the largest and most complex accelerators ever built. The sheer number of single-point failures in multiple systems makes it imperative that a major design emphasis be placed on Availability, the amount of time the machine is operational during planned running time. An accelerator availability modeling program developed for the ILC, *Availism* [1], has shown that the ILC cannot operate with acceptable up-time unless all systems are designed with a very high tolerance to individual unit failures. 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.

The main focus of this paper is the development of HA packaging standards for Accelerator Controls and Instrumentation. Current modular standards, some invented by the physics community, achieve improved up-time by rapidly interchangeable plug-in modules. However this requires machine interruption to replace a failed module, which means lost machine productivity. In addition, current modular standards share a parallel bus backplane in which a single-point failure can disable the entire crate and interrupt the beam for hours. *Availism* shows that for ILC, reliability of all key systems must be improved significantly to meet an availability goal of >85%. This is true whether all systems are in a single tunnel and inaccessible during running, or in a parallel service tunnel with all active components accessible for immediate repair. The Baseline model for ILC is a parallel service tunnel. Even so, many sub-components need improved Mean Time Before Failure (MTBF), some by large factors that are impossible without new design strategies. 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.

The chief contributor to machine downtime currently is magnet power supply systems, followed by RF power systems, and then by instrumentation and controls. For instrumentation and controls, the key high availability strategy has been developed years ago in the form of the standard instrument module[3]. The main availability strategy for such systems is very short repair time (Mean Time to Repair, MTTR). In other words, when a module fails, identify the cause and replace it quickly. However this simple operation usually interrupts the machine or experiment to make the repair, leading to a recovery time which can be much longer than the actual replacement. In contrast, the guiding principle for ILC design is to make repairs without interrupting machine operations if at all possible, and to keep the probability of 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 use of modular instrumentation standards that had an impact on 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[4]. 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. SPECIAL CHALLENGES OF THE ILC

The question has been raised: Why is High Availability design such an issue for the ILC and not the LHC? There are several responses to the question:

[1] R.S. Larsen is Assistant Director for Electrical Systems, Operations Division at SLAC and heads the SLAC ILC Electronics R&D Program.

[2] Work supported by US Department of Energy Contract DE-AC03-76SF00515.

[3] NIM, CAMAC, FASTBUS and VME. See Ref. [2] for a brief historical review.

[4] The Advanced Telecommunications Computer Architecture, announced in June 2004, is being designed by a large industrial consortium to standardize products for the 10B$/year Telecom industry. This is a major advance in standardization for the industry. See Ref [3].

1. Availability is a crucial concern for *all accelerators and detectors*. It is a measure of the efficiency of the machine in performing its designed tasks. A chronically inefficient billion-dollar class machine speaks to poor design and poor execution. This is not acceptable when tools are available to accomplish high performance at little incremental cost.
2. The ILC, unlike a storage ring, is a *repetitive one-shot machine* where every feature has to function perfectly on every shot. An injector for example cannot fail while the experiments continue running on the already stored beams for possibly many hours before beams disappear. This is a routine failure occurrence in storage ring systems that is masked by the stored beam.
3. The ILC is of a scale of linac never before attempted. The complexity of a single linac collider like the SLC at one-tenth the length is well known. It took a very long time to commission and to finally achieve design performance. In this case both electron and positron beam shared the same linac, in principle less likely to fail than two linacs pointing at each other.
4. Not only the complex of positron and electron injectors and their respective damping rings, but critical components such as the quad magnets in the main linac need to be available all the time, or beam emittance will be lost and the beams will fail to collide.
5. A computer program designed to estimate availability has been benchmarked using real data from operating machines. The original availability goal for the NLC and TESLA machines was A=0.85 (85% uptime); however to achieve it required each of 16 machine subsystems to all be available 99% of the time. This has never been achieved before. Present estimates from the *Availism*[5] program indicate that even with factors of improvement in component MTBF's of 20-50 times present performance, it will be very hard to achieve ILC availability above 0.85.

III. BASIC HA DESIGN PRINCIPLES

The fundamental Availability equation is:
A= (MTBF-MTTR)/MTBF →1 when MTBF→∞ or MTTR→0.
The principles of HA design are simply:

1. Maximize MTBF of basic components
2. Minimize MTTR

This is achieved by:

3. Modularizing all system basic components
4. Adding *partial redundancy* into the system in the form of extra modules, N+1 or N+n where n is much smaller than N.
5. Introducing, wherever possible, module *hot-swap* capability.

Consider the example of a power supply in *Figure 1*: Some ILC magnets are critical in that loss of a single power supply will stop the machine. This is approximately true for the linac superconducting quadrupoles. In this case instead of a single power supply with a controller, the unit is modularized into several parts: A shared bulk DC supply [3] feeding several individual supplies and loads; N modular DC-DC converters operating in parallel to supply the total current needed for each load; and separate controllers to manage each supply and the total system. Both the shared bulk supply and each controller are redundant. The entire system is designed to auto-failover: If one of the shared bulk supplies fails, the second unit carries on to support the full load while the diagnostic system notifies the control room of the failure. If a DC-DC converter fails, the system continues as above. If a controller fails, its redundant neighbor takes over. All three of these features are needed to make a power supply system for the ILC that will be available 99% of the time. For repair, it seems feasible to make the N+1 or N+n converter modules as well as the controller hot swappable or perhaps even the bulk for low power supplies. However safety regulations will dictate what can be done in practice. Note that if a unit fails, it is not necessary to replace it instantly (although the quicker it is done the better in case of another failure in the same unit); one can usually wait for a programmed maintenance time to do this. In any case, if hot-swap were permissible, the most common failures could be eliminated as a source of machine interruption and *MTTR for the system can in fact be made zero.*

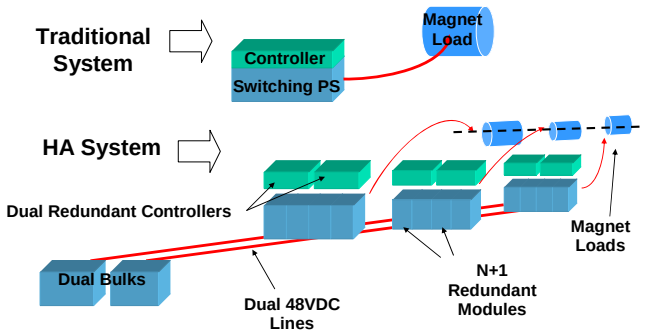

Figure 1: High Availability Power Supply System Concept

Note that in linac design, redundancy at a high level is always incorporated in that the machine is longer than needed for its design energy and this extra length represents an allowable number of simultaneous failures of RF stations with minimal operations impact. Although this is also planned for the ILC, it does not help the quadrupole magnet situation, where a failed single magnet could stop the machine.

In controls and instrumentation, similarly, the machine has redundancy of critical elements like steering magnets and beam position monitors. However these must be employed frequently enough all along the machine to allow failures of single units without interruption except for a tuning readjustment, which should be counted as part of MTTR even if it only temporarily knocks the beams out of collision.

The global controls system includes data communications, timing and RF phase references and which must be delivered to every critical instrument and pulsed power device with the same reliability as the AC line power supply. The concept is shown in *Figure 2*. These core

[5] *Availism*, the simulation model by T. Himel of SLAC, quantifies MTBF of all components and estimates availability for various architectural choices such as one vs. two tunnels, redundancy of sources and RF, etc. See US Linear Collider Technology Options Study, USLC Steering Group, March 4, 2004,. Chapter 4: Availability Design, http://www.slac.stanford.edu/xorg/accelops/

functions of controls are planned to be fully redundant with auto-failover hardware and software. Beyond this core which reaches into nodes at every sector of the machine, as shown in *Figure 3*, individual control and monitoring functions can be tailored with HA features in accordance with need. The overall system level of redundancy, such as RF stations, determines the level of need. For example, extra RF stations allow a unit to be taken offline and serviced *while the machine is running*, impacting availability only briefly while the control system switches stations which we assume are in operational standby mode at all times.

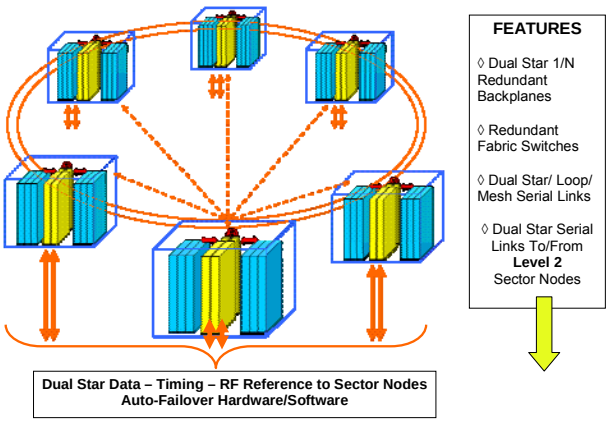

Figure 2: Main Control HA Concept

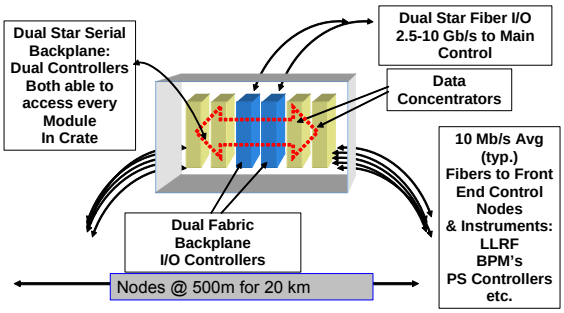

Figure 3: Linac Control Node HA Concept

Underlying all these systems is a new robust Diagnostic Layer that is essential to managing the overall HA system complex. This layer will vary in form between subsystems but its functions remain constant, namely:

1. Provide local intelligence to monitor the health of every subcomponent down to the module level
2. Help protect and isolate faulty modules by disabling local power
3. Monitor internal performance parameters and warn the control room of impending problems
4. Possibly take local action to avoid interlock trips that will interrupt the machine, such as by reducing power out of a local RF unit approaching its limits and compensating elsewhere to keep beam energy constant
5. Flag modules that need attention by hot-swap or programmed repair before multiple failures occur to cause a machine interrupt

Examples will be discussed in the following sections.

IV. HA DESIGN CONSIDERATIONS FOR DETECTORS

To a major extent the large detectors already employ high availability design architectures. The structures are inherently highly redundant in terms of layers of the various chambers, and every layer has multiple individual channels. Even a partially crippled detector that loses a segment can often keep on taking data, although at reduced efficiency. At the same time, a detector that shares a beam as in the Large Hadron Collider (LHC) must be highly reliable as a unit since a serious on-line failure could mean loss of mission for an extended period. Even if immediate repair were possible, access and repair could take weeks or months. The stakes for overall reliability/availability are high indeed. Designs tend to be unique and unproven with respect to long-term performance. It is expected that problems will be fixed with upgrades as necessary. In the execution of the front end electronics, low mass requirements together with design for radiation and magnetic field tolerance severely restrict design choices.

In view of all these risks, HA architectures should be adopted as a design strategy for future upgrades and new projects such as the ILC detectors should be designed for emergency access for critical repairs with minimum intervention and downtime. Some general HA design comments on specific issues are:

A. Front End Electronics

On multi-channel front end boards or hybrids, consider built-in redundant silicon readout structures with availability of spare sections by switches. In principle a faulty electronics channel can be isolated and replaced by program control.

B. Local Controllers

Create redundant control paths where either controller can read out the same segment independently, so in case of failure the control system executes an auto-failover.

C. Communications channels

All high end communications should be Gigabit serial with a standard protocol wherever possible. Critical controls functions will need noise immune error detection and correcting protocols. Fibers and wires including controls and front end data and controls should be redundant with auto-failover software, especially in radiation environments.

D. Front End Board Serial Gigabit Communications

In cases so demanding of real-time response or packaging constraints that a standard protocol cannot be used, provide redundant channels as well as calibration circuitry to monitor degradation of fibers.

E. Access for Repair

Critical electronics on or within a detector should be made accessible from the perimeter. This calls for a new look at segmenting into small, plug-in packages that can be replaced from outside by a human or remote handler mechanism. The challenge is to achieve this design goal with little or no sacrifice of detector hermeticity.

F. Power Systems and Management

More inadequate designs are made on detector power systems than on the much more complex custom chip functions. HA design generally calls for basic redundancy of bulk DC supplies at a high voltage such as 48V, with all loads having local on-board DC-DC hybrid size converters. A single power converter board inside a detector would feed a small group of front end boards or chips and be remotely accessible. Power loss in cables and regulation problems of very high currents at very low voltages would be minimized. A power board for a group of front end boards should be dual-fed with failover capability. Special designs on industry-standard footprints should be designed for radiation and magnetic field tolerance.

G. Module and Board Hot-swap

Major control and power boards should be designed to be hot-swappable. This means no front panel cables or connectors. Such boards would have switches and a power management control system to isolate them automatically from 48V power in case of failure.

H. Crates

All crates should use a modern standard incorporating HA features as seen in the Telecom industry ATCA system (Advanced Telecommunications Computing Architecture). Redundancy, power management and hot-swap features are standard and supported by a growing number of supplier companies. Moreover, crate manufacturers can easily design a custom backplane since many companies will specify special routing of gigabit links and user-defined transition areas.

V. ILC HA ELECTRONICS PROGRAM

Some approximate numbers of devices in the machine are as follows: 10 MW RF stations, 720; RF cavities with RF monitoring and tuning, 18,000; BPMs, 3,000; Magnets and power supplies 17,000; vacuum pumps 10,000; Instrument racks 6,000.

The Electronics R&D program aims to investigate new HA designs for the following: Controls, Beam Instrumentation, Low Level RF, Magnet Power Supplies, Modulators, Diagnostic Interlocks, Stripline Kicker Pulsers. The HA goal for each *complete system* such as controls, power supplies, RF stations etc. is an availability of at least A=0.99. The R&D programs and progress are described briefly.

A. Controls

To achieve A=0.99 Controls requires redundancy of its core communications, controls and processing systems, timing and RF reference systems down to the machine sector node level, and communications to beam instrumentation that is machine-redundant so that a single instrument, such as a BPM, can fail without beam interruption. Core system modules must be hot-swappable to avoid the unacceptably long recovery time of turning off a crate and interrupting the machine. For beamline devices where a single failure will bring down the machine, controls redundancy should be carried to the critical device.

Progress: The ATCA system (*Figure 4*) is chosen as the model for evaluation and cost estimating purposes. This features redundant power sources and shelf (crate) power management, high-density gigabit serial communications and dual communications processors, in a 14-slot crate similar in size to VXI. The management system supports full hot swap and the entire shelf of 14 modules is designed for A=0.99999 (Five 9's). For a system the size of the ILC, 0.99 system availability requires crate level availability of Four 9's. System software studies are in process at Argonne National Lab and hardware-software studies and evaluations have begun at University of Illinois Urbana Champaign (UIUC) supported by SLAC. Part of the hardware study will be how to implement fast analog connector interfaces in the rear panel user-defined space, the gigabit serial connector for analog performance, and overall crate noise performance for analog-to-digital modules. The early deliverable is an ATCA engineering test system made available to labs with development programs, with the longer term goal of a prototype segment of a full system to serve the planned Superconducting Test Facility planned for Fermilab.

B. Beam Instrumentation

The degree of required instrumentation redundancy depends on the device. BPM's do not necessarily have to be redundant if they are placed frequently enough. However since they often share a controls crate with other devices they should be hot-swappable. The goal is to evaluate whether BPM's can be effectively packaged in the ATCA system or one of its standard options.

Progress: A board area comparison with current designs in CAMAC shows this is possible, perhaps including a dual channel per board for redundancy. BPM spacing in the linac is sparse, one BPM of four channels per 36m, so the unit must be a single thin standalone rackmount or a single ATCA board sharing a slot in one of the LLRF crates. Space and connector issues will be studied further by the UIUC group.

C. Low Level RF

The LLRF is the largest subsystem in the main linac, estimated at one ATCA crate for each 10 MW station plus optical transceiver connection areas, tuner driver motors, piezo drivers and a large accompanying cable plant. The LLRF equipment will fill two water-cooled racks per 36m of linac.

Each LLRF unit takes power and phase signals from each of 36 superconducting cavities to control the drive to the klystron and maintain its output power amplitude and phase in stable relationship to the beam. In addition to RF drive the system also manages the cavity tuners which are of two types, a mechanical tuner that is slow and piezoelectric tuners which are fast to dynamically compensate for fast detuning caused by the shock of beam loading.

Progress: Technically the LLRF does not have to be a full HA design because there are redundant RF stations at the system level; however, since every interruption of crate power causes serious recovery time on its own, it is highly desirable that modules be hot swappable. The ATCA is being used as the packaging model for evaluation and cost estimating purposes. Meanwhile actual design of prototypes is proceeding at both DESY and FNAL. Collaborators at KEK, FNAL and DESY are leading the LLRF studies.

Figure 4: ATCA "A=5 Nines" Basic Components

D. *Magnet Power Supplies*

. Quad and corrector magnets in the main linac are superconducting while in the injection systems, damping rings and beam delivery all quads, correctors and bends are normal conducting. For critical loads where loss of the supply would stop operations the goal is hot-swappable redundancy at the bulk, DC-DC converter and control levels. Solutions must be cost-effective. The superconducting magnet supplies are low voltage current sources that require quench protection circuitry.

Progress: An evaluation program of N+1 topology as shown schematically in *Figure 1* is underway for a 5 kW system typical of a warm quad, including 3 of 4 and 4 of 5 modular converters, dual 48V bulks and dual controllers. DC-DC converter failover has been successfully demonstrated by disabling a supply by switching off its reference and studying transient recovery. A dual controller prototype for low voltage power systems is being designed. Demonstration of a complete prototype is planned for FY07. In addition a 40-supply demonstration system has been started for the KEK ATF2 collaboration for delivery in early FY08. Dual redundant bulks will typically be sized to drive several converters, e.g. 10-20 kW units, and hot-swap requires the addition of active isolation circuitry to make the operation safe. Worst case, for safety reasons a failed bulk may be exchanged only during scheduled maintenance. If the second unit fails and interrupts the machine, both units can be changed during a short machine interruption that minimizes MTTR.

E. *Modulators*

Klystron modulators are a major cost component of the machine. The present baseline unit for the ILC is a 10-year old design that has recently been upgraded with redundant IGBT switches. However an alternate design is underway at SLAC of an all solid state Marx generator that minimizes size and eliminates the heavy transformer, producing 120kV at 140A that can run over a short cable directly to the 10MW klystron. This design emphasizes high availability through 4 of 5 redundant IGBT switches for both charging and pulse forming in each of ten 12 kV cells in the stack. In addition it includes 2 extra cells that can be automatically switched in if a cell fails. At the system level, the linacs will have 3% extra RF stations to maintain energy in case of station failure.

Progress: The first prototype Marx 12 kV cell has been successfully tested at SLAC and the full stack is under construction and scheduled for demonstration over the next 2-3 months. See *Figure 5*. If successful, it will become the new baseline for the ILC, and several units will be built for the 10MW klystron driven test programs at SLAC and possibly at other future test facilities. Industrial designs of Marx prototypes are also proceeding.

F. *Diagnostic Interlocks*

Special attention is being paid to developing much better diagnostics than in previous generations of equipment, especially high power modulators and power supply systems. The vision is for a small board-level unit to reside in each cell of a Marx or each module of an N+1 power supply system to monitor operation, including temperature of the environment and critical components, voltage and current levels in relation to equipment safety interlock limits, and critical waveforms to

show variations in performance that indicate impending problems.

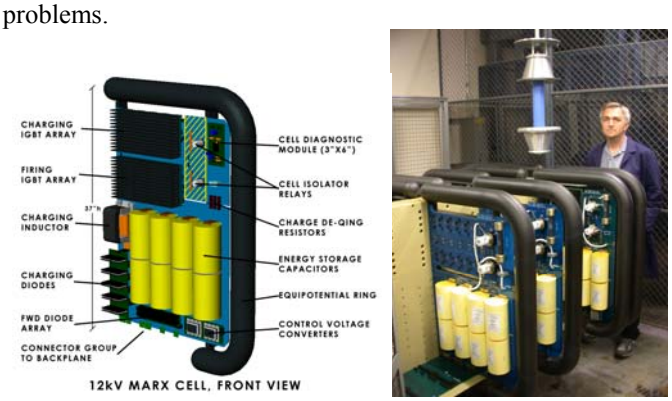

Figure 5: Marx Cell and Prototype Assembly

This information can be used to manage equipment and systems to *avoid* interlock trips, as well as to gracefully remove equipment from service for maintenance before it disrupts operations. The Diagnostic unit could also be taught to take certain actions independently rather than in present systems where every machine trip is a surprise that sends operators and maintenance personnel scurrying for answers.

Progress: A program was started in 2005 in collaboration with Pohang Light Source, Korea, to develop the first prototype of a unit for modulator and large power supply applications. This unit is complete and is undergoing evaluation. Features include fast trigger and trigger time delay management; fast and slow 8 bit and slow 16 bit converters for waveform and DC monitoring respectively; waveform memory for reconstructing trips in post-trigger operation; and serial communications to a control system IO controller via Ethernet. In addition, a special low power unit with reduced feature set has been designed for the Marx. See *Figure 6*. Its low power is necessary to operate from capacitors charged by the high voltage system, as each cell in the stack increases in voltage from ground by 12 kV up to 120kV. This unit has operated successfully on a stacked pair of Marx cells in the first prototype under assembly. See *Figure 6.*

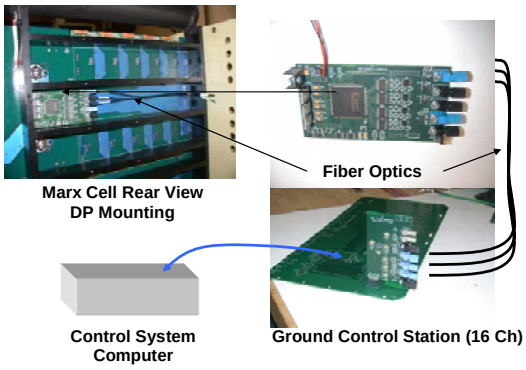

Figure 6: Marx Diagnostic Processor Prototype

G. *Stripline Kicker Pulsers*

A major challenge for damping rings is a 3 MHz kicker pulser producing +/-10kV into 50 Ohms with a baseline of 5 nsec or less and 10^{-4} or less coupling to the next pulse. If the kicker is feasible it reduces the ring from 17km to 6km. The R&D program is to produce a pulse that meets the specifications and engineer the design into a practical kicker unit. The full system will require 10-20 units in tandem to apply the total kick required. In this design as in the Marx, the system engineering aspects are most important: Inclusion of cell-level diagnostics, amplitude calibration and very precise timing control and dynamic calibration capability for the assemblage of 10-20 kickers.

Progress: In 2005 a prototype designed by LLNL in collaboration with SLAC, consisting of 15 MOSFET array driven cells in an induction stack configuration, was tested in the ATF at KEK. See *Figure 7*. The pulse was considerably

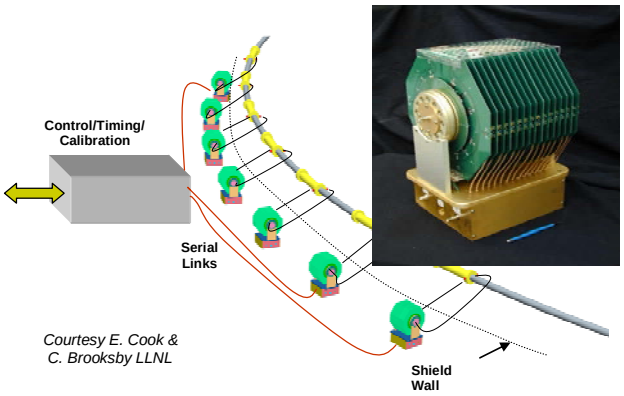

Figure 7: Prototype Damping Ring Kicker & Array

Wider than desired, the voltage was +/-5kV and the unit could only run in short burst mode not CW. Since then the unit was expanded to produce +/-8-9kV, close to the design goal, and in parallel work on new drivers carried out. At the time of writing a new push-pull pulse driver has produced a pulse with a 5.1 nsec base operating in burst mode at 3 MHz. A second prototype breadboard from a SLAC SBIR collaboration has reported a very fast rise and fall time pulse from a DSRD (Delayed Step-Recovery Diode) technique that is being investigated further as a candidate pulser. Other ILC kicker work is occurring at other laboratories as well but not reported here. The SLAC-LLNL goal is a full power and pulse rate prototype in FY07 and demonstration of tandem operation at the ATF in FY08.

VI. SUMMARY AND FUTURE PLANS

High Availability electronics design techniques are becoming standard practice in industries such as Telecom and Computing. HA design techniques can be applied to any system, not just electronics. Both Accelerator and Detector systems can both benefit if HA design is undertaken at inception or for later upgrades. This cost of N+1 redundancy is generally small and overall cost effective due to enhanced availability of the production machines and lower maintenance costs. HA design has been installed as an important performance metric for the ILC.

The ILC R&D Collaboration is investigating a range of approaches to achieving high total machine availability with nominal impact on system cost. Short term goals are to evaluate multi-gigabit serial interconnect systems for intelligent modules, demonstrate basic high availability

software failover management, and port representative core controls, serial networks and beam instrumentation designs onto new platforms. Evaluations are currently focused on the Advanced Telecommunications Computing Architecture (ATCA) system, a new industrial open standard designed for a data throughput of 2 Tb/s and A= 0.99999. Applications of HA to detector architectures is also a future goal.

VII. ACKNOWLEDGMENTS

The SLAC R&D electronics program is collaboration with the international ILC Global Design Effort (GDE) and specifically its Area and Technical groups which can be viewed at *linearcollider.org*. 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, UIUC, Pohang PLS and R.W. Downing Inc. The special contributions of E. Cook and C. Brooksby of LLNL, J. Carwardine of ANL, B. Chase and M. Votava, of FNAL, R. W. Downing, M. Haney of UIUC, S. Michizono and T. Naito of KEK, S. Simrock and K. Rehlich of DESY, S. Nam of PLS, and the many dedicated colleagues at SLAC and all the named institutions are gratefully acknowledged.

VIII. REFERENCES

[1] US Linear Collider Steering Group Accelerator Subcommittee, *US Linear Collider Technology Options Study (USLCTOS)*, Chapter 4, Availability Design, pp 200-232, March 2004.

[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.