

DOE DE-FG02-03ER41281

INTERNATIONAL LINEAR COLLIDER ACCELERATOR PHYSICS R&D

Senior Staff: George D. Gollin, Michael Davidsaver, Michael J. Haney, Michael Kasten
Undergraduates: Jason Chang, Perry Chodash, Will Dluger, Alex Lang (University of Wisconsin), and Yehan Liu

I. Introduction

Our ILC work has concentrated primarily on technical issues relating to the design of the accelerator. Because many of the problems to be resolved require a working knowledge of classical mechanics and electrodynamics, most of our research projects lend themselves well to the participation of undergraduate research assistants. The undergraduates in the group are scientists, not technicians, and find solutions to problems that, for example, have stumped PhD-level staff at Fermilab.

Since our ILC projects are closely related to efforts at Fermilab and Argonne, we spent a day at the end of last summer presenting our findings at both labs. Each of the students assembled a coherent description of their activities and findings into a PowerPoint presentation, which they showed during the morning at Fermilab and the afternoon at Argonne. I was very impressed with their talks, and was delighted to see the level of interest in our findings at the labs. It was clear that we were considerably further along in addressing a number of issues in high reliability controls architecture and monitoring than was the controls group at Argonne. And our ILC beam work attracted good reviews at Fermilab.

II. ILC Results in FY08

A. Damping ring kicker studies

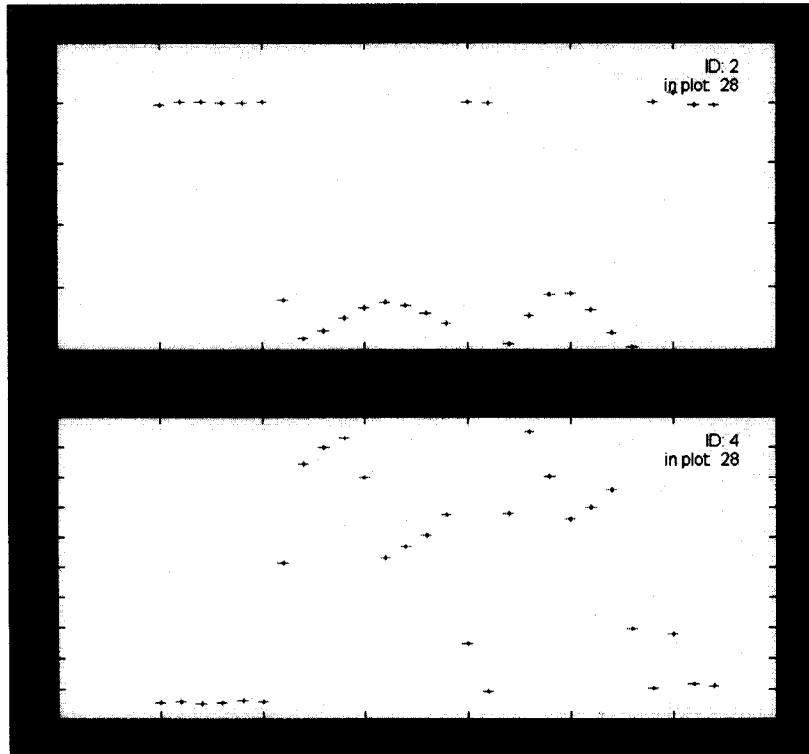
The ILC Reference Design Report calls for 6.7 km circumference damping rings (which prepare the beams for focusing) using “conventional” stripline kickers driven by fast HV pulsers. Our primary goal was to determine the suitability of the 16 MeV electron beam in the AØ region at Fermilab for precision kicker studies. The stripline kicker we designed and installed in collaboration with Fermilab is shown in the following figure.



Stripline kicker installed in the AØ 16 MeV beam at Fermilab. Beam enters from the left.

We found that the low beam energy and lack of redundancy in the beam position monitor system complicated the analysis of our data. In spite of these issues we concluded that the precision we could obtain was adequate to measure the performance and stability of a production module of an ILC kicker, namely 0.5%. I was surprised, and encouraged, by this. One of the important components in our ability to analyze our data was the beamline simulation written by REU student Alex Lang. Alex was able to determine that the beamline's analyzing magnet was being run into saturation and that the field map provided by Fermilab was inaccurate, and could not be used successfully without substantial modification.

In the following figure the mean kicker deflection and RMS width of the deflection are shown as a function of run number.

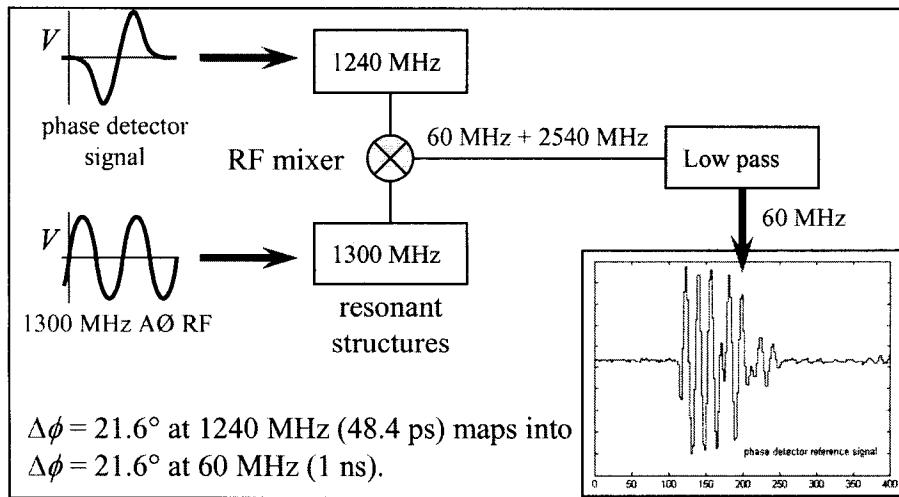


Kicker deflection (mrad) and RMS width (mrad) vs. run

We concluded that the kicker was stable to an accuracy of $\sim 2.0\%$ and that we could measure this precision to an accuracy of $\sim 0.5\%$. As a result, a low energy beam like that at AØ could be used as a rapid-turnaround facility for testing ILC production kicker modules.

B. Bunch timing work

The ILC timing precision for arrival of bunches at the collision point is required to be 0.1 picosecond or better. We studied the bunch-to-bunch timing accuracy of a “phase detector” installed in AØ in order to determine its suitability as an ILC bunch timing device. A phase detector is an RF structure excited by the passage of a bunch. Its signal is fed through a 1240 MHz high-Q resonant circuit and then down-mixed with the AØ 1300 MHz accelerator RF, as shown in the following figure.



Phase detector signal path

Because the mixer output preserves the phase difference φ between input signals $\sin(\omega_1 t)$ and $\sin(\omega_2 t + \varphi)$, a small shift in the arrival time of a bunch is mapped into a much larger shift in the time structure of the $(\omega_1 - \omega_2)$ signal.

We used a kind of autocorrelation technique to compare the phase detector signal with a reference signal obtained from the phase detector's response to an event at the beginning of the run. We determined that the device installed in our beam, which was instrumented with an 8-bit 500 MHz ADC, could measure the beam timing to an accuracy of 0.4 picoseconds.

Simulations of the device done by UIUC student Jason Chang showed that an increase in ADC clock rate to 2 GHz would improve measurement precision by the required factor of four. As a result, we felt that a device of this sort, assuming matters concerning dynamic range and long-term stability can be addressed successfully, would work at the ILC.

C. High availability controls architecture studies: ATCA

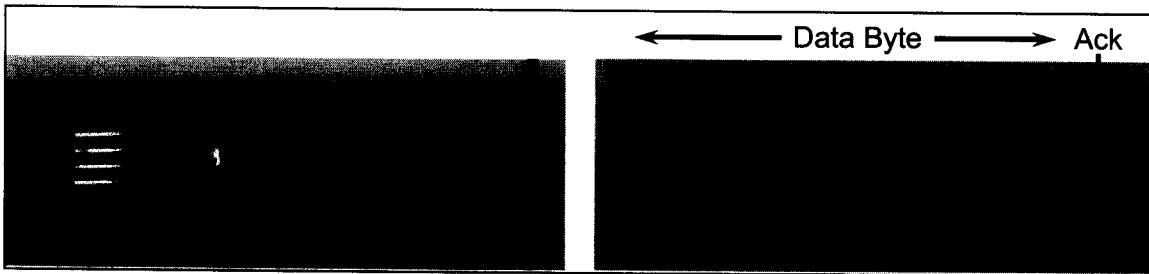
Cost effective operation of the ILC will demand highly reliable, fault tolerant and adaptive solutions for both hardware and software. The large numbers of subsystems and large multipliers associated with the modules in those subsystems will cause even a strong level of unit reliability to become an unacceptable level of system availability. An evaluation effort is underway to evaluate standards associated with high availability, and to guide ILC development with standard practices and well-supported commercial solutions.^{1,2} One area of evaluation involves the Advanced Telecom Computing Architecture (ATCA) hardware and software. An (unfair) analogy for ATCA is that of a VME (or CAMAC, etc.) electronics crate, except with dual processors, multiple/redundant communications paths across the backplane, integral availability management for predictive/automated failover, hot-swap capability (boards and power

¹ www.saforum.org

² www.linux-ha.org

supplies), and “five nines” (99.999%) availability. Understanding the scope of ATCA and adapting it to the broad spectrum of detector and accelerator electronics will require both hardware and software study and development.

UIUC physics majors Perry Chodash and Yehan Liu, along with HEP group engineer Mike Kasten, worked with an ATCA crate, processor monitors, and a small amount of ATCA circuit boards in order to develop a backplane “spy” board that would let us watch the ATCA backplane communications and pursue development of an inexpensive processor monitor that could be used as a physics-driven component of the crate-level controls system. We made good progress, and felt that we had determined a productive direction to extend this work after last summer. The backplane “spy” and a signal trace from an ATCA crate are shown in the following figure.



ATCA backplane spy and backplane data.

We felt that we had learned enough to begin designing a workable processor monitor chip if there were to be sufficient interest in ATCA shown by the ILC community.

D. High availability controls architecture studies: OpenClovis supervisory software

Fault recognition is a challenging issue in the crafting a high reliability controls system. With tens of thousands of independent processors running hundreds of thousands of critical processes, how can the system identify that a problem has arisen and determine the appropriate steps to take to correct, or compensate, for the failure? One possible solution might come through the use of the OpenClovis supervisory system, which runs on Linux processors and allows a select set of processors to monitor the behavior of individual processes and processors in a large, distributed controls network.

There is almost no expertise in the Midwest engineering community in the use of OpenClovis, except for one Argonne staff scientist. We worked with him to obtain an OpenClovis system that UIUC physics major Jason Chang installed on a small cluster of Linux machines. A schematic of Jason’s arrangement is shown in the following figure. His initial studies only used the machines inside the oval superimposed on the figure.

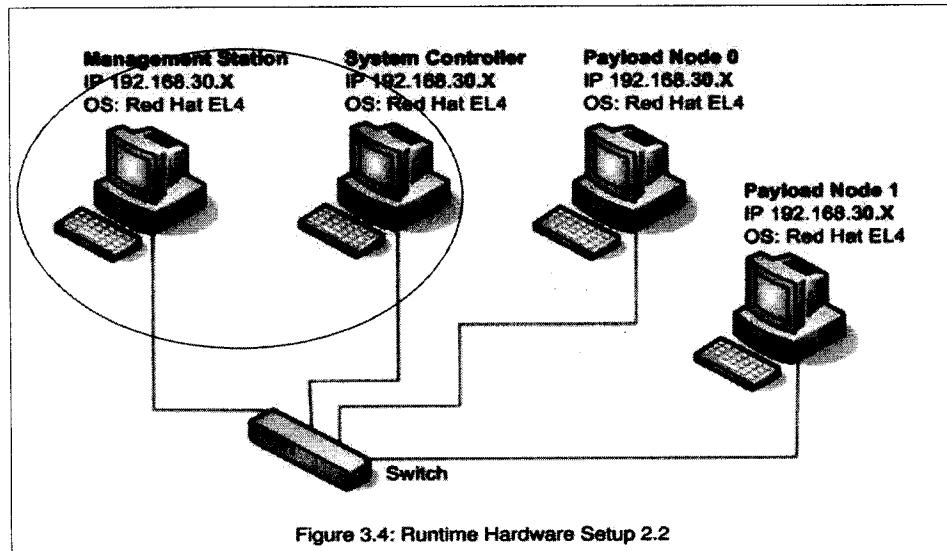


Figure 3.4: Runtime Hardware Setup 2.2

OpenClovis evaluation system.

We found that OpenClovis exhibited an irritating amount of sensitivity to the exact version of the Linux kernel running on the processors, and that it was poorly equipped to help us sort through problems that arose through conflicts so deep in the operating systems of the processors. But once this issue was addressed, we found that it performed as expected, recognizing crashes and process (and processor) failures. Tests of OpenClovis, in which the process monitor was surveilling several independent tasks, is shown in the following figure.

Putting OpenClovis through its paces.

Jason is one of the brightest students I have ever had in my lab, and he was the only student with enough intellectual horsepower to resolve the problems we found with OpenClovis. It was an impressive piece of work, and we now know that OpenClovis' "twitchiness" will be an issue in a very large system where one cannot necessarily trust that all processors are running a particular kernel version.

E. University ILC R&D program coordination

I am a co-principal investigator on the International Linear Collider University-based Linear Collider Detector R&D grant, along with Jim Brau (Oregon) and Mark Oreglia (Chicago). This is a joint DOE-NSF funded program that provides support for the entire U.S. university-based R&D effort for ILC detector. In spite of the problems in ILC funding, this program continues, receiving support to distribute among university and small national laboratory groups. With the delays in the ILC's schedule brought about by the U.S. budget problems, I feel this should be a lower priority effort than a university-based accelerator R&D effort. However, the Department of Energy has chosen to support some detector work at universities while eliminating funding for work on the accelerator, in spite of this.

Along with Dan Amidei at the University of Michigan I had been serving as the initiator and coordinator for the U.S. university-based ILC R&D program. This effort grew into a national program of 366 physicists distributed over 51 universities, 8 national and industrial laboratories, and 25 foreign institutions. Scientists proposed 72 different projects, to be funded jointly by the DOE and NSF. Half the projects concerned machine physics for the accelerator. Funding constraints caused the agencies to split the effort into separate detector and accelerator projects; after this happened I focused primarily on the accelerator side. Our funding was eliminated after the summer of 2007, again as a cost-saving measure. But the agencies reconsidered, and asked me to begin reorganizing a new effort that might be able to accept support beginning early in 2008.

Initially the DOE and NSF felt that a level of support of approximately \$100k per participating group would be appropriate, coming to a total of \$2.5 million to be shared equally by the two agencies. I scheduled a meeting with representatives from both agencies in November 2007. They agreed to a sensible management and oversight structure, with funds to be handled in a style similar to that for a detector project: there would be a project manager who would apportion resources based on project needs, rather than having funds arrive through a series of independent proposals and supplemental funding requests submitted by individual investigators. At the anticipated national ILC funding level of \$60 million this would be possible so I began assembling expressions of interest from my colleagues.

Last December's omnibus spending bill eliminated funding for U.S. ILC work, killing all possibility of restarting a university-based ILC program. Next year's U.S. funding for ILC work is estimated to come to \$35 million, all of which is likely to be absorbed by the national laboratories. There will be no realistic possibility of restarting a U.S. university-based ILC accelerator program next year.

The University of Illinois has been actively supportive of the ILC. The Vice Chancellor for Research provided funds for an economic impact study of the effects of the creation of a commercial superconducting accelerator industry in Illinois; we did this in conjunction with a UIUC College of Commerce consulting group. In addition, an Associate Vice Chancellor for Research was tasked to work with me on ILC issues that might involve interaction with the state's Department for Commerce and Economic Opportunities. That office also provided funds for a non-degree graduate student (Michael Davidsaver, who is now on loan to Fermilab).

The impossibility of carrying forward our ILC work is intensely frustrating.