



Fermi National Accelerator Laboratory

FERMILAB-Conf-88/110

**Trieste Conference on Digital Microelectronics and
Microprocessors in Particle Physics
Summary and Concluding Remarks***

Thomas Nash

**Advanced Computer Program
Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510**

August 1988

*Presented at the Adriatico Conference on the "Impact of Digital Microelectronics and Microprocessors on Particle Physics," International Centre for Theoretical Physics, Trieste, Italy, March 28-30, 1988.



Operated by Universities Research Association Inc. under contract with the United States Department of Energy

**TRIESTE CONFERENCE ON DIGITAL MICROELECTRONICS AND
MICROPROCESSORS IN PARTICLE PHYSICS
SUMMARY AND CONCLUDING REMARKS***

Thomas Nash
Advanced Computer Program
Fermi National Accelerator Laboratory
Batavia, Illinois 60510 USA

Abstract

This paper is a written version of the Concluding Remarks presented at the International Conference on the Impact of Digital Microelectronics and Microprocessors on Particle Physics. The Conference emphasized on-line data acquisition and triggering problems in high energy physics. Among the participants there was a clearly growing consensus that as these real time systems become larger they require more attention from the beginning to overall system coherence and manageability issues. We consider what this means for SSC/LHC era detectors. Given the interesting results on pixel silicon, neural networks, and parallel microprocessor based computers presented at Trieste, we speculate on some surprisingly simple, though still very radical, ideas on systems solutions for those huge detectors.

Trieste was the sixth location for a series of conferences focussing variously on the many aspects of on-line and off-line computing in experimental high energy physics. Starting at CERN, these conferences have been marked uniformly by the intense, fast changing nature of the subject matter they cover and the spectacular interest of their sites (Padua, Guanajuato, Amsterdam, Asilomar, Trieste, and next year, Oxford). Because Asilomar covered software issues in depth, the announced emphasis at Trieste was on hardware rather than software. To a large extent, the high energy physics community interprets this distinction between software and hardware as off-line *versus* on-line computing. Asilomar was populated by those working on off-line computing problems; Trieste's participants were primarily on-line and trigger specialists.

The difference between the style and concerns of these two groups, both essential to the success of HEP experiments, is extraordinary, almost that of two cultures with little communication. The on-line culture is willing to try anything, hardware or software, if it appears to have a chance of doing a required job. They place a strong emphasis on code efficiency and hardware cost effectiveness. The amount of effort required for a specific task and the overall coherence and manageability of experiment wide real time systems have not been at the forefront of their traditional concerns. Some individuals in this culture cannot

* This written version was prepared without access to the transparencies or papers prepared by other speakers at the conference. Although it thereby differs in detail and sequence from the actual concluding remarks given in Trieste, it expresses the same general conclusions.

understand, as one of them expressed it at Trieste, why all the hundreds and hundreds of MIPS worth of frequently idle on-line processors can't be applied to the off-line computing load.

Among those who develop off-line code, concerns are dominated by the huge scale of the software packages that must be prepared. Data and code management techniques are frequently used and there is considerable attention paid to the coherence of the overall skeleton of large programs. Ease of use and reliability of computers, peripherals, and operating systems takes precedence for many over processor cost effectiveness. When the computer cycles required for off-line processing inevitably reach finite budgetary limitations, cost effective processor systems developed by on-line types are, at best, tolerated.

At Trieste, despite the absence of most off-line software developers, many of their system level management concerns were beginning to make themselves felt among those responsible for large on-line data acquisition and processing systems. The ALEPH on-line electronics system is a prime example of this, even to the extent of using SASD tools in developing its on-line software. At OPAL and HERA's H1 on-line system management is being addressed through advanced human interfacing techniques taking advantage of Macintosh Hypercard tools.

Certainly, on-line systems coherence is far from being prevalent these days, particularly among smaller experiments attempting to retrofit modern processors into existing electronics. Conetti described the extraordinary variety of ways in which Fermilab's ACP Processors are being incorporated into experiments. Although widely used in standard configurations for off-line, there is no standard on-line implementation of these systems because there is nothing approaching a standard data acquisition system. In present day large experiments, the sociology of multi institution collaborations works against data acquisition system coherence. The typical scenario at design report time is to divide up the responsibilities among the various groups allowing each to define the approach to be used in specific subsystems. A few years later the problem of interfacing these subsystems becomes a hot topic. The situation is, to some extent, due to the way approval committees operate. Expertise is compartmentalized by subdetector and institution. The overall system is based on disparate subsystems, developed on the basis of correct, strongly held individual views, that don't mesh as a system.

Most experiments attempt to standardize at some level. Nonetheless, many still end up with systems that incorporate both Fastbus and VME or otherwise mix standards. Much of this results from a plague of what might be called *standards evangelism* in our business. Individuals who have developed some degree of personal expertise in something like VME or Transputers or UNIX or ACP Systems become strong advocates. At some level, this is an understandable tendency to protect an individual's intellectual investment in the expertise. However, the increasing complexity of detectors cries out for unbiased attitudes, a *secular humanist* approach, toward standards. We need to cool off the all too common electronics and computing standards religious wars which every laboratory has encountered.

The fact that it may prove possible to operate large present day experiments, handicapped as they are by the kind of institutionally mandated bottom up design described above, does not mean that this will continue to be possible for SSC/LHC era detectors. The order of magnitude increase in complexity of those behemoth detectors of the future simply *requires* a new and structured approach to on-line data acquisition and processor systems. It is not surprising, and very pleasing, how much agreement there was on this fundamental conclusion among participants at Trieste. We all agree that we need to get together and work out the rules for what might be called (borrowing the terminology from software) a *Structured*

Analysis Structured Design (ASD) approach to data acquisition/trigger systems. We need to work out the rules, but here are some ideas that found easy agreement:

1. **Plan the overall system first**, at the same time as the physics is planned. Worry about protocols later.
2. Require **modularity** with a well defined protocol between modules.
3. The system should be as **homogeneous** as possible with a minimum of domain boundaries between different types of electronics. There is some feeling that there is at least one fundamental domain boundary between the type of protocols and electronics required for data acquisition and the type required and supported commercially for computing systems. At the very least, this boundary occurs at the point where data is recorded, but some of us feel it may be appropriate to use such commercial computer standards as far upstream as possible.
4. **Keep it simple stupid (KISS)!** Here this famous designer's rule implies such things as keeping control and data signals separate and minimizing sexy but hard to program hardware. Hardwired, specialized data driven, and esoteric parallel approaches must be reserved for where they are absolutely required, perhaps in lowest level triggers.
5. **SASD** for on-line software and electronic CAD/CAE design and simulation tools extended to the intermodule and system level.
6. The system must be **assemblable**. Related to modularity this is an issue of being able to bring large pieces of the system together without out having to redesign the data routing. Some radical ideas have been expressed at the Fermilab ACP regarding automatic data routing where data is sent to classes of available processors or memory rather than specific addresses. This approach for large data acquisition systems is similar to that used in specialized data driven trigger systems such as the Fermilab ECL-CAMAC (part of LeCroy's ECLine) and the Nevis data driven trigger processor.

Outlandish speculation is tolerated and encouraged in concluding remarks like these. With that license, let us see how we might apply the philosophy of these rules to designing an SSC data acquisition/trigger system that is simple and manageable, though of the enormous scale required. It is expected that SSC experiments will require something approaching a trigger reduction of 10^8 and will have to cope with such complications as several events in a crossing. A major portion of this trigger reduction will necessarily be in processors which use software written in high level languages because this is the only way to deal with such a level of trigger sophistication.

Three papers presented at this conference give a hint of what the future may bring. Irwin Gaines described the new generation of ACP Multiprocessors being developed at Fermilab with individual processors 10 times faster (we have since learned they will be 20 times faster) than the present generation. What he didn't say is that a DARPA funded project at Texas Instruments is underway to develop a Gallium Arsenide version of the same instruction set RISC processors as the ACP is using. They are expecting 200 MHz (about 100 VAX 11/780 power per chip set) by the end of next year. It is no longer unreasonable, therefore, to anticipate single board processor "nodes" available in 1995 with 100 VAX per node at the usual module cost of around \$2500.

Spieler described his exciting work at LBL on pixel silicon detectors and the possibility of incorporating digitizing electronics directly on-silicon. Particularly interesting is the technique of bump bonding which allows connecting wafers full of integrated circuits that match the pixel spacing directly under the detector silicon. Listening to Bruce Denby's talk on neural network and cellular automaton applications to track and cluster finding, it doesn't take too much imagination to conceive of coupling VLSI implementations of these ideas (already being developed at Cal Tech and Bell Labs) to Spieler's pixel silicon.

A front end trigger based on VLSI could thus reasonably be expected to provide a trigger reduction of say $10^{2\pm 1}$ leaving $10^{6\pm 1}$ for higher level triggers. Can all of this remainder be handled with high level programmable processors along the lines of GaAs versions of ACP nodes? If so, we could avoid the difficulties of hardwired or other low level trigger techniques. How much processing power could we reasonably afford with 100 VAX power GaAs processing nodes? \$5 million is hardly exorbitant for such a system on the scale of an SSC detector, and that would buy 200,000 VAX equivalents. The 2000 nodes could be arranged in some sensible tree structured, self routing architecture. Using a totally unjustified hand waving argument, we can scale up from CDF's planned 100 to 1 reduction using about 50 VAXes to 200,000, and we find that a reduction of 10^6 or so with such a system in 1995 doesn't seem at all beyond the realm of possibility.

It almost seems that all we need to get a hold of the extraordinary demands of SSC/LHC type detectors are strong R&D efforts in such areas as neural networks, multiprocessors, data acquisition systems, and silicon and other detectors. At Trieste it was preaching to the converted to conclude that perhaps what we need is a few less experiments in particle physics these days -- and a few more R&D projects .