Summary of Working Group V, Data Acquisition and Filtering Subgroups B (Models) and C (Software)

Reported by Dave Cutts*

Brown University

Providence, Rhode Island

Contributors[1]:

Yasao Arai, Terry Carroll, Dave Cutts, Tom Devlin,

Lloyd Fortney, Dave Hedin, Cathy van Ingen,

Paul Kunz, Bill Sippach

OVERVIEW

As you have heard, others in Working Group V have considered specific applications to the problems of data acquisition for the SSC, either from the hardware direction, coping with the data at the point where it enters the system, or from the system point of view. In contrast to this specific focus, we addressed generic questions of hardware architecture and software techniques. In this summary I will repeat three themes that surfaced in our discussions which seem to me to be especially fundamental. First, we agreed that the microprocessor "farm" was the appropriate structure for the high level data acquisition event filter. Second, it was clear that to make effective use of the processing power in this farm, the

1) Those who made presentations; many others participated in various discussions.

software filter strategy must be carefully planned. Finally, a number of us felt that modelling of the entire detector data flow, including the acquisition system, was an activity vital to the success of a SSC experiment.

MICROPROCESSOR FARMS

The general problem of data acquisition for the SSC is easy to state : How does one cope with the enormous flood of data and extract the interesting physics? Given that some aspects of data (like energy sums) can be inspected quickly in hardware, one expects low level hardware-based filters to reduce this flow somewhat. Still, after this stage (or stages) of reduction the algorithms remaining to be applied are complex and difficult to hardwire, requiring correlations between different areas of the detector. At this point the data flow is perhaps a few thousand events/second, each roughly 1 MByte in size. To handle this data, and to reduce it further to a few per second of real physics interest, is the province of the high level software-based filters. The basic architecture for implementing this filter, adopted by all the new collider detectors and described already at various conferences and at other workshops, consists of an array of microprocessors, plus a mechanism for routing an entire event from the detector to a particular node within this array or "farm". Compared to a single processor, one obtains an enhancement in filtering power, or total number of events analyzed per second, that is less than but comparable to the number of separate processors working in parallel -

provided that the deadtime due to the farm's input and output is minimal. One technique for effecting very high speed transfer of data from the detector readout that would meet the requirements of SSC experiments would be to associate with each processor node or small group of nodes a number of dual-port memories, with each channel's external port capable of receiving data at high rate directly fron the detector. As an example, each of D0's 50 MicroVAX-II nodes will have 8 channels of dual-ports, each with 40 MByte/second external input; loaded in parallel from the detector this gives a maximum capacity of 320 MBytes/sec. Schemes were discussed with larger numbers of readout channels or differently arranged processor farms (see papers by Tom Devlin or Lloyd Fortney) but the basic notion was generally accepted : given proper attention to data flow problems, arrays of microprocessors was an appropriate, practical, and essentially proven solution to the need for software-filtering engines for SSC data acquisition.

SOFTWARE FILTER STRATEGY

We heard a discussion about the software filters for the TeV-I collider experiments from Terry Carroll (CDF) and Dave Hedin (D0). The basic strategy presented in both cases is the same : do first that which is easiest and gives the biggest rejection. This means sharpening in software the initial calculations done in hardware by incorporating more exact data and by making correlations not possible earlier. Beyond these steps are progressively more compute-intensive calculations such as those involving tracking which require more complete data reduction and analysis. Given the real-time demand for fast filtering decisions, this software must be highly efficient.

Beside the similiar outlines shown by the two groups there was a very important stress placed on planning, particularly by the D0 group (whose hardware filter is much less extensive than CDF's, and which correspondingly is depending more on the farm's rejection). With the complexity of the event analysis and the crucial role that this analysis has on the outcome of the experiment, a well designed, flexible, and fully understood software filter is essential. It was clear that a careful design of the software using modern design tools was important, and for the filter program particularly vital. Beyond the generic discussion of individual steps in a filter, this lesson of planning should not be forgotten for the SSC.

SYSTEM MODELLING

An extension of the above thought about the designing of the filter software is to consider modelling the entire acquisition system. The inter-relation of different elements of the detector that provide, readout, and select data must be considered in a global way. At the SSC the data reduction is a massive effort with an enormous reduction factor. To handle the flow and select the few good events,one needs a design that is based on a real understanding of the performance of each element in the system. Cathy Van Ingen described some of the efforts at using the software package RESQ (RESearch Queueing) to understand CDF's data acquisition, and to study alternative structures for their microprocessor farm. We presented a simulation of a microprocessor farm appropriate for the SSC. This simple model predicted that the performance of a farm with of order 1000 processors would be at the level necessary to handle SSC's requirements. Even this primitive model was able to give useful insights - such as the importance of excess input bandwidth into the farm, and of excess compute power within the farm.

From these discussions I conclude that the study of the filter and beyond, of the entire system, including all elements (from the detector, electronics, triggers, acquisition, and recording) is feasible and practical. It is also very important. This motherhood statement, though not original, bears repeating because it entails a fundamental change in the way detector groups work.

* Supported in part by the U.S. Department of Energy.