DATA FILTERING - ACQUISITION GROUP

REPORT OF THE HARDWARE SUBGROUP

A Feasibility Design for the Readout of a 4π SSC Detector

Participants

James Dunlea	Ohio State University
Henry Kasha	Yale University
Spenser Klein	Stanford Linear Accelerator Center
William M.Morse	Brookhaven National Laboratory
Leo Paffrath	Stanford Linear Accelerator Center
Marleigh Sheaff	University of Wisconsin

Peter. S. Cooper, Chairman
Yale University
Department of Physics
P.O. Box 6666
New Haven, CT. 06511

ABSTRACT

The hardware subgroup took as its charge the study of the front end electronics for a 4π SSC detector. We present here a feasibility study for such a readout including a conceptual worst case design. We conclude that within the assumptions enumerated below such a readout system is feasible with electronics and systems which are either presently available or are likely to become available within the next few years. The cost of such a design is non-trivial (~100$/channel x 10^6 channels), but for this exercise no attempt at economy was made.

ASSUMPTIONS

We took the parameters of the SSC and 4π detector as those discussed at the beginning of the workshop. In the spirit of doing a worst case design study we have included none of the improvements to these parameters generated by the other groups in the workshop. We discuss below some of the implications of relaxing these assumptions.

The SSC was taken to produce a luminosity of 10^{33} /cm^2/sec with a 33 nsec period between bunch crossings. Two levels of hardware triggering are assumed with properties given in the table below:

ASSUMED TRIGGER PROPERTIES

	Delay	Rejection	Rate [Hz]
Events			10^8
Analog Triggers	1 μsec	1000	10^5
Second Level Triggers (requires processed data)	10 μsec	100	10^3

The delay time is the delay until the trigger signal arrives at the readout electronics which require it. This clearly depends, in the case of the analog trigger, upon where the various electronic systems are located (see below). The 4π detector is assumed to be of conventional design with central tracking, longitudinally segmented calorimetry and surrounding muon chambers. The required number of readout channels for each subsystem is given in the table below:

ASSUMED DETECTOR PROPERTIES

	Elements	Channels
Central Tracking	175k wires * 2 ends	350k
Calorimeter	50k Towers * 3 segments * 2 ranges	300k
Muon chambers	100k wires * 2 ends	200k
TOTAL CHANNELS		850k

The central tracking chamber is assumed to be a drift chamber with multiple hit readout on each end of each wire. For simplicity the Muon

chamber readout is assumed identical to the drift chambers. The drift time is assumed to be 100 nsec and we have taken for desired resolutions 1 nsec on leading edge timing and 10% on end-to-end charge division. This will give ~100 μm resolution in radius (R) and 20 cm resolution along the beam line (Z) assuming 2 m long wires.

For the calorimeter we have taken as a boundary condition that we must be able to resolve and measure a single minimum ionizing track and not saturate on a 10-TeV electromagnetic shower. This requires a dynamic range of 10^6. We concluded that analog and digitizing circuits with these dynamic range requirements are not feasible so we have required two readout channels per calorimeter element. With a high and low gain channel for each element the dynamic range required per readout channel is 2×10^3 which is only somewhat beyond the current state of the art. We assume an integration time for calorimeter signals of 300 nsec in order to accommodate conventional calorimeters like liquid argon.

The design we have attempted begins with signals from suitable pre-amplifiers for each detector element and ends with zero compressed data from a single event stored in a set of memories. The third level trigger, which will consist of a large farm of microprocessors, would fetch its data from these memories.

Finally, we assumed that the maximum acceptable readout dead time is 10%.

FEASIBILITY DESIGN

In general, the readout must perform three main functions. It must provide a delay for all signals for 1 μsec until the first level (analog)

trigger arrives. It must process these signals to extract the data relevant to the trigger event, deconvolving signals when necessary to separate the data of interest for that of events near in time. Finally, it must provide a readout with data sparsification in order to reduce the amount of data which must be processed by the third level software trigger at manageable proportions. All this must be accomplished with sufficient parallelism to keep the readout dead time to less than 10%.

In addition to these minimum requirements in our thinking we applied a significant prejudice - an aversion to cables and connectors. The thought of 1 million cables was more than some of us could bear. We conceived of much of the front end electronics being mounted directly on the detector. The clear implication of this choice is the necessity of removing 50 - 500 kw of heat from the detector. We examine the relaxation of this requirement below.

The basic design is shown in the block diagram, figure 1. The particular case shown is for the calorimeter but the drift chamber case is essentially identical.

The physical layout envisions that everything from the preamp through the buffer shown in figure 1 is mounted in crates on the back of the calorimeter. Individual crates are connected by data buses which are terminated by local readout controllers which drive fiber optics data links to remote memories.

The readout is in six stages. We will discuss them in turn for the case of the calorimeter as shown in figure 1. The differences for the drift chamber readout will then be addressed.

READOUT ELECTRONICS - CALORIMETER

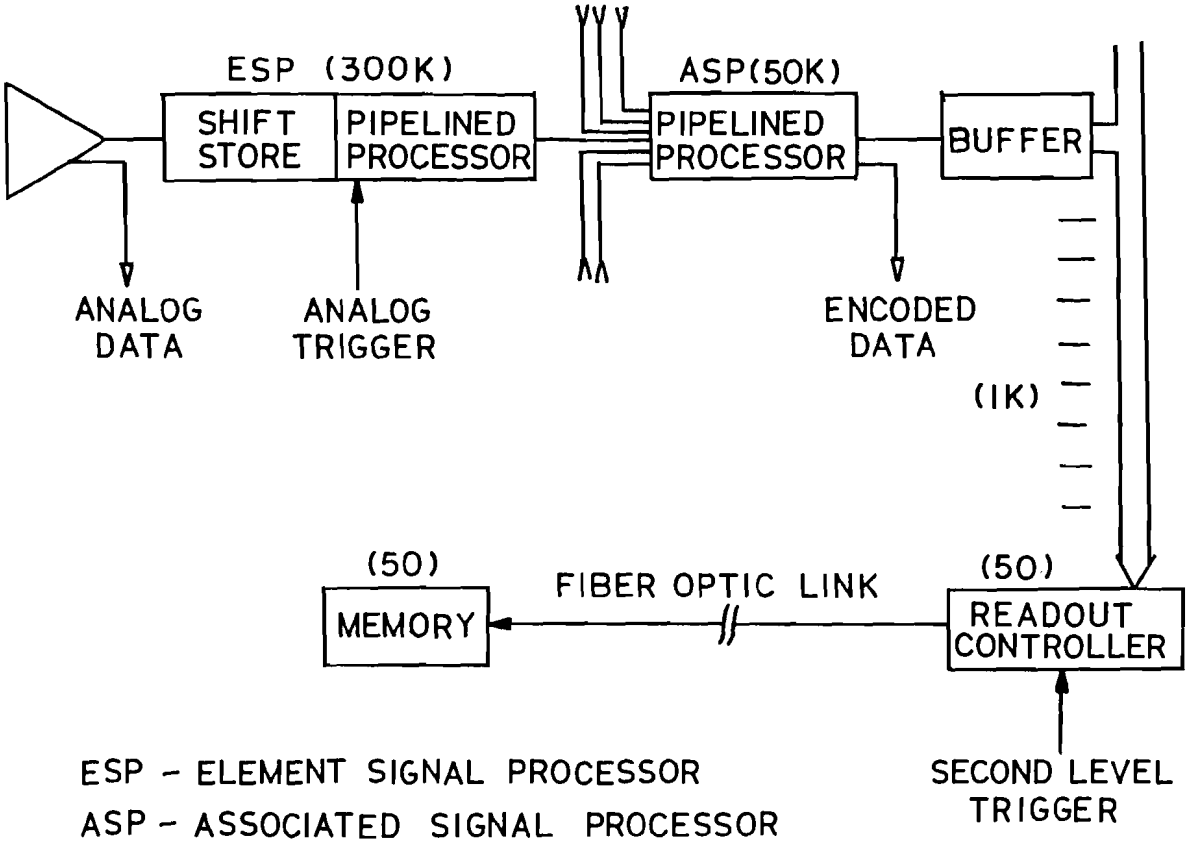

FIGURE 1

The first two stages are conceived of as a single custom designed chip with a multiplicity of one per readout channel (six per calorimeter tower). The first stage is a shift store which continuously samples the input waveform in fixed width time buckets driven by an external clock. The model we used was the SLAC AMU chip which is an analog waveform storage device with 256 storage buckets, a maximum clock rate of 120 Mhz and a cost of $4 per channel. This chip does not have the linearity or required dynamic range of 2000 to be used in this application. A significantly improved version,if technically feasible, might do this job. Another possibility is to flash ADC the input waveform and store the data

digitally. 100 Mhz flash ADCs do exist, but not with 11 bits of resolution. Some development in chip technology is required here. We believe that these developments are feasible and likely within the next several years.

The second stage of this chip, which we have named an Element Signal Processor (ESP), is a pipelined processor which,when started by the analog trigger, digitizes the waveform data in the range of times near the triggered event, if it was not already stored in digital form, and processes it to extract the relevant parameters for the event of interest. Assuming 100 nsec per pipeline step,the ESP can execute 100 operations in the 10 μsec between analog triggers. To eliminate dead time at this stage several techniques are possible. Perhaps the simplest is to have two shift stores, one for storage while the second is being readout.

In the case of the calorimeter, the ESP clock would run at 20 Mhz giving 50 nsec per time division and 2.56 μsec of total delay assuming 256 cells of storage. The pipelined processor of the ESP has to first determine if a pulse is present in time with the analog trigger. Since we are assuming relatively slow calorimeter signals which are likely to have risetimes longer than 20 nsec, this requires some extrapolation of the leading edge in time to determine a start time relative to the analog trigger. A major consideration for the calorimeter is the overlap of different events in time. Given the 10^8 interaction rate,and assuming that an average minimum bias event has particles hitting 1% of the calorimeter towers, the probability of an extra event within the 300 nsec calorimeter integration time in a tower which has data from a triggered event is 30%(see figure 2)! The ESP must then integrate the signal, recognizing and

The net effect of applying an analog and second level trigger to this readout scheme is an entry in 50 memories of energies (in GeV) with tower identification tags. This naturally segments the detector into 50 regions of solid angle. It is clearly possible to exploit this division for certain classes of triggers by processing these data in parallel before building the entire event in a microprocessor farm.

The drift chamber readout is essentially similar to that described above for the calorimeter. The differences are in the programs coded in the ESPs and ASPs. We considered that the ASP would associate the two ends of a single wire. This is a major cabling problem but has the advantage that the outputs are R and Z coordinates of hits. The clock rate for the drift chamber shift store will have to be 100 Mhz. 10 nsec time buckets give only 3 nsec time resolution,so extrapolation techniques in the ESP are required to give the desired 1 nsec time resolution.

RELAXATION OF ASSUMPTIONS

In light of the above worst case design, we considered what happens if some of the more restrictive assumptions are relaxed. The obvious conclusion is that things get easier and cheaper.

The higher level triggers group concluded that they required only analog data in order to generate the second level trigger. Assuming this to be the case, the readout is not required to generate any outputs on the 10 μsec time scale. This would allow multiplexing of the pipelined processors. The configuration is shown in figure 3 below. The processor in the ESP would be replaced by a FIFO analogous in function to the buffer FIFO. The function of the ESP and ASP would be combined in a single processor started with the second level trigger. Each processor would be

connected to 8 readout channels (48 calorimeter signals). It would fetch and processes in turn the data for the triggered event from the eight FIFOs. Since there is 1 msec per second level trigger the processor would have 125 μsec to complete the processing of a single channel. This could be accomplished with the pipelined architecture described above for the ESP or, perhaps, with a more conventional stored program microprocessor executing one instruction per microsecond.

ANALOG ONLY SECOND LEVEL TRIGGER

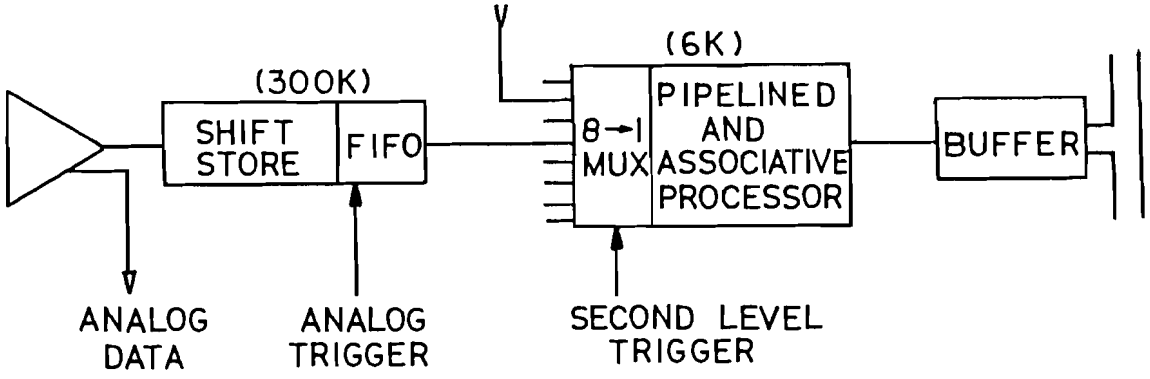

FIGURE 3

If we relax the assumption of putting all the front end electronics on the detector we solve a major engineering problem in layout and cooling. The cost is a line driver, connector, cable, connector, line receiver set for each of the 850k channels. This choice is dictated, fundamentally, by cost, not performance. If the cost of the complications of putting all the electronics on the detector is greater that the cost of the cable and connector solution then the choice is made.

CONCLUSIONS

As the name indicates we have attempted in this feasibility design to indicate what is possible under a conservative set of assumptions. The

hope is to focus on what portions on such a system are relatively in hand and where further developments both in technology and design are required. This effort is certainly not the only feasible solution to the problem and undoubtedly not nearly the best.

Two conclusions can be drawn from this exercise. First, that such a system is feasible with the electronic technology likely to be available when SSC detectors are built. Also, that unoptimized cost estimates of 100 $/channel for 10^6 channels (100 M$) are not out of scale for detectors which are likely to be five times the size of present-day collider detectors and thus five times the cost (250 M$).

It is worth noting that a project of this size naturally allows exploitation of economies of scale to significantly reduce costs. Custom designed chips are certainly required for much of the chip count of an SSC detector readout. With a million channels the cost of development and engineering of such custom designs is easily offset by the volume ordered (a 1 M$ development cost is still only 1$/chip).