9/19/89 89 (2)

SLAC-PUB--5042

DE89 016946

**DATA ACQUISITION AND ONLINE PROCESSING REQUIREMENTS
FOR EXPERIMENTATION AT THE SUPERCONDUCTING
SUPER COLLIDER***

A. J. LANKFORD

Stanford Linear Accelerator Center, Stanford University, Stanford, CA 94309

EDWARD BARSOTTI and IRWIN GAINES

Fermi National Accelerator Laboratory, Batavia, IL 60510

Abstract

Differences in scale between data acquisition and online processing requirements for detectors at the Superconducting Super Collider and systems for existing large detectors will require new architectures and technological advances in these systems. Emerging technologies will be employed for data transfer, processing, and recording

*Invited talk presented at the 4th Pisa Meeting on Advanced Detectors:
Frontier Detectors for Frontier Physics,
La Biodola, Isola d'Elba, Italy, May 21-26, 1989*

*Work supported in part by Department of Energy (USA) contracts DE-AC03-76SF00515 and DE-AC02-76CH03000

1. Introduction

At the Superconducting Super Collider (SSC), interactions of colliding 20 TeV protons are expected to occur at a rate of approximately 100 MHz. These interactions will be studied by high energy physics experiments consisting of more than a million electronic channels each. High performance data acquisition systems will be needed to move the large quantities of data from the detector elements through trigger processors and to mass storage. Extensive online processing will be required to filter the number of interactions and the amount of data per interaction down to the rate of approximately 10 to 1000 interesting events per second and to overall data rates which are compatible with mass storage techniques and future offline computing capacity. The architecture of the online processors, the efficient high-speed transfer of data among the processors, and the effective management of processing resources and software will be crucial.

This paper highlights some of the differences in scale associated with data acquisition and online processing for an SSC detector as compared to the large detectors of the current generation and attempts to identify some of the technological advances which will enable the new problems to be solved. It does not attempt to describe in detail, or to completely outline, a data acquisition system for an SSC detector. Efforts at such an overview have been the subject of past workshops [1,2], the most recent being held in Toronto early this year [3]. A more complete overview is also the subject of an ongoing study by the Task Force on Electronics, Triggering, and Data Acquisition at the SSC, which operates under the auspices of the SSC Central Design Group.

The distinctive features of the data acquisition system for an SSC detector are:

- front-end electronics based on custom VLSI,
- high-speed data collection and data transmission using fiber optics,

MASTER

EB

- parallel event building,
- massive processor farms,
- large amounts of data placed into mass storage, and
- intensive use of processors for triggering, calibration, data compression, and monitoring tasks.

Following a brief overview of the data acquisition system as a whole, the subsequent sections of this paper describe these features.

2. System overview

An overview of the data acquisition system for an SSC experiment is shown in figs. 1 and 2. Figure 1 emphasizes the overall architecture of the trigger and the data acquisition system; whereas, fig. 2 emphasizes the use of parallelism in implementing the architecture.

The overall architecture is determined by the architecture of the trigger. As in current experiments, the selection of interactions is accomplished by the trigger in a series of progressively more selective and more time-consuming stages. Our model in fig. 1 shows three such stages, or "levels." Levels 1 and 2 are referred to as the "prompt" triggers and consist mostly of special-purpose processors constructed from hardwired logic which can perform relatively simple decisions fast, in 1 to 2 μ sec for Level 1, and 10 to 100 μ sec for Level 2. As will be mentioned later, more general purpose processors embedded in special purpose architectures may play an increasing role at Level 2 in the SSC era. Level 3 of the trigger, the "nonprompt" or "higher-level" trigger, is performed in an online processing farm by high-level language programmable processors, perhaps preceded by additional special-purpose processors.

The number of sensing elements in an SSC detector is expected to be very roughly 10^7 with interactions occurring at a rate of 10^8 sec⁻¹. Data from the detector elements is buffered in the front-end electronics while this rate is reduced by the prompt trigger. Only the data required by each stage of the trigger is transported to the trigger processors while all the data is buffered. The data input to the Level 1 trigger is expected to be mostly analog; whereas, the data input to Level 2 is expected to be digital.

The algorithms which select event candidates at each level of the trigger determine both the data bandwidth required for input into the trigger processors and the data rate between stages of the data acquisition. Within a given experiment, a certain amount of flexibility is available with respect to choosing at which trigger level to deploy selection criteria. Moreover, the algorithms available to different experiments are detector and physics dependent. For instance, the final rate of interesting events is quite different for experiments studying high p_t phenomena and ones studying decays of beauty. Furthermore, the study of high p_t phenomena allows event selection based largely upon calorimetric information which is relatively easy to incorporate in a prompt trigger; whereas, efficient selection of large numbers of beauty events requires triggers which utilize tracking and vertexing information with the full resolution of the detector not normally available to prompt triggers. Thus, there is a high degree of interplay between the capabilities of the trigger system and of the data acquisition system at each level in the system.

Very roughly, a combined rejection of the prompt triggers of a factor of 10^3 to 10^5 is expected. The resulting trigger rate of 10^3 to 10^5 corresponds to a data rate output from the front-end electronics of between 0.2 and 100 Gbytes/sec if the event size is between 0.2 and 1.0 Mbytes. In order to reduce the data from the ten million detector elements to less than a megabyte, considerable processing power must be provided to suppress data from elements without signals and to compress data and filter data from hit elements.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

The data bandwidth necessary to transport data from the front-end electronics to the processor farm can be provided by a relatively small number of parallel data paths. For instance, 10 Gbytes/sec could be provided by 200 links capable of 0.5 Gbits/sec/link. On the other hand, conventional techniques of collecting the data from the entire detector in a single event builder module before transmitting it to a target processor are no longer feasible at these proposed data rates. The required data throughput is supplied by a switching network which assembles events in parallel.

Practical online processor farms may provide aggregate computing power of up to one million VAX-11/780 equivalents. For instance, 5000 processors of 200 MIPs each would provide between 10 and 1000 "VAX-seconds" of CPU time per input event candidate. It is hoped that the Level 3 trigger operating in the processing farm can reduce the trigger rate by a factor of order 100. The final trigger rate would then be 10 to 1000 events/sec. The event size may also be reduced in the processor farm; however, it is doubtful that it can be reduced below a few hundred kilobytes. On the other hand, mass storage systems exploiting parallel drives are capable of recording over 100 Mbytes/sec and large archival systems can provide storage and access to the accumulated volume of recorded data.

This data acquisition model emphasizes the upper limits of system performance. Subsequent sections of this paper will attempt to justify some of these performance goals. On the other hand, if the performance of the trigger processing is highly sophisticated and successful at identifying its physics targets, ultimate performance may not be required in all detectors. The extremes of system performance sketched here provide the capacity to absorb data rates higher than expected if necessary, at least until trigger algorithms and performance reach the required levels. In addition, the use of parallelism throughout this data acquisition architecture, as illustrated in fig. 2, enables adaptation to the scale of performance required.

3. Integrated front-end electronics

For some time now amplifying electronics have been mounted in close proximity to particle detectors for improved analog performance and for increased immunity to RF pickup. More recently, several groups working on silicon tracking devices and the SLD collaboration have pioneered large-scale integration of electronics for mounting on the detector. In the case of silicon microstrips, the density of connections and limited space for cables leading to and from amplifiers, particularly in 4π detectors, have led to amplifiers, sample-and-holds, multiplexers, and in some cases sparse-scanned readout on a single custom VLSI chip. The SLD collaboration has incorporated similar functionality into all of their electronic systems, through the use of hybrid and semicustom monolithic amplifiers and custom sampling and buffering devices, in order to achieve a more cost-effective, space-efficient, and reliable electronics system for a detector with a large number of channels.

The same principles will be applied to the electronics of SSC detectors for the same reasons. However, the large number of channels in an SSC detector and the very high interaction rate require well-integrated custom solutions in order to limit power dissipation as well. Power considerations drive SSC electronics to incorporate pipelined buffering in the front-end electronics mounted on the detector. On every electronics channel, signals must be buffered from each of approximately 100 beam-crossings which occur during the time required by the first level of trigger selection. In addition, the desire to limit dead-time in face of the high interaction rate demands that the analog front-end electronics continue to sample subsequent crossings at the same time as buffered data is read out for triggered interactions.

Consequently, the front-end electronics for SSC experiments is driven to solutions like those being worked on by several groups for various types of SSC detectors.

This work is typified by the circuits being developed by the University of Pennsylvania and the Catholic University of Leuven [4,5] for readout of drift chamber systems. This front-end chip set consists of two custom multichannel chips. A preamplifier, shaping amplifier, and discriminator are implemented on a bipolar chip. A time-to-voltage converter, analog memory buffer, analog-to-digital conversion, output multiplexing, and control logic are implemented on a CMOS chip. For an SSC detector, these chips will replace both the boxes of detector-mounted amplifiers and the crates of remote FAST-BUS TDC modules found in today's large detectors, as well as the hundreds of long cable interconnections. Further detector-mounted multiplexing and data preprocessing will replace today's crate-level scanners and segment interconnects.

4. Data collection

Although each type of detector component will have custom front-end electronics appropriate to its measured quantities, the control and readout of the front-end chip sets will be sufficiently similar that a common readout scheme may be achievable for the entire detector. Data from as many as several hundred thousand front-end chips, each with data rates of very roughly tens of kilohertz, must be multiplexed onto a manageable number, perhaps 100 to 1000, high-speed data channels which provide an aggregate data rate of several to 100 Gbytes/sec. A hierarchical solution to data collection, starting with groupings of nearby detector channels and proceeding towards large groupings of all the data from one region of solid angle, is appropriate. The entire data collection process, reducing the number of data paths to the few hundred to be input to the parallel event builder, will occur within and on the detector.

The most ambitious solutions to the problem of data collection, those aimed at the highest achievable rates of data transfer, are data driven. At each step in the data

collection process, every data source is pushing data into intermediate buffers as the data becomes available. Data collectors then gather the data from the buffers at the highest possible rate and push the data into the next stage of buffers. The bandwidth of all data links can be used to full efficiency. The data is transmitted with appropriate event and channel tags; however, packets of data do not necessarily correspond to individual events. The process of event building is therefore to a large extent decoupled from the data collection and transmission. In these data-driven schemes, control is minimized as data is moved along a series of simple data transmission links. Operation of such a system should be easy to verify and trouble-shoot, since verification and fault identification will be amenable to a series of communications tests, which in fact could be performed by simple expert systems.

5. Data transmission

Transmission of data to each stage of data collection will be via links of technology appropriate to the bandwidth required at that stage. Data collected from the front-end chips, where bandwidths are low, will be transported via copper buses on detector-mounted printed circuit boards. At the other end of the data collection process, the perhaps hundreds of long links carrying the data from all parts of the detector to the parallel event builder in the control room will be high-speed fiber optic links. The speed and number of links at that stage will be determined by practical considerations, such as the cost and the size of the switching network in the event builder. The transition from high-speed copper links to fiber optic links of modest speed will occur at some intermediate stage.

The principal advantage offered by fiber optic transmission is that of high bandwidth, particularly over distances longer than several meters. Fiber optics promise performance

that makes data acquisition of gigabytes per second feasible. Fiber optic transmission also offers the important advantages of immunity to electromagnetic interference and low transmission losses. In addition, if used within the detector, they offer advantages in size and mass over copper cables. Radiation hard fibers are available to a level of some megarads and exhibit some self-annealing. Gallium arsenide electronics, which is normally used to drive and receive high-speed fiber optic systems, is intrinsically radiation hard to an even higher level.

The fiber optic needs of the computer industry are driving technology to increased performance and decreased cost for links similar to those needed for SSC data acquisition. The present cost of fiber optic links, including transmitter, receiver, electronics, and optics, expressed in terms of cost per bit transferred is smallest for links of about 100 Mbits/sec, where it is only about \$2/Mbit/sec. Whereas, 500 Mbit links now cost more than \$10/Mbit/sec, industry is currently developing integrated gallium arsenide electronics for such functions as serialization and laser drivers, detectors with integrated receivers and deserializers, cross-point switching, and multichannel devices for byte-wide transmission. Compact gallium arsenide lasers which can be mounted on the same hybrid as the integrated driver electronics are also being developed. These developments are well-advanced and all in the range of about 1 to 2 Gbits/sec/link. Gigabit per second links should be quite accessible for use in SSC experiments.

6. Parallel event builder

The parallel event builder addresses the bandwidth bottleneck arising in traditional event builders where data all passes through one path. In a parallel event builder, a number of input data paths from the detector are connected to a number of output data paths to the processors, and all the data paths can be active simultaneously to maintain

the aggregate bandwidth. The number of input and output data paths need not be equal; however, if bandwidth is nearly optimized then the numbers are naturally the same.

Several schemes for parallel event builders have been discussed. These schemes generally utilize a matrix of buffer/router nodes, as discussed in previous SSC data acquisition workshops [6,7], or utilize switching networks. The schemes have many similarities, particularly the need for extensive buffering to smooth out event-to-event fluctuations in amounts of data on each link and the need to balance the average data rates on each data path. These needs arise from the fact that the bandwidth will be limited by the longest event fragment of each event if the buffers are insufficient or by the slowest data path if rates are not balanced.

Parallel event builders using switching networks exploit advances in the technology of cross-point switches arising from the communications and computer industry. A generalized network would allow the interconnection of input and output data paths in any combination; however, the events can be built using simpler networks. An interesting implementation which minimizes control is that of the "barrel-shifter" suggested by Mark Bowden and Ed Barsotti of Fermilab [8]. In this scheme, each input path is sequentially connected to an output path in a cyclic fashion.

Let us examine a simple case of a barrel shifter, that of a four-input four-output switch. Consider a detector that produces four data streams, A, B, C and D, each with a series of equal size event fragments. As illustrated in fig. 3(a), data passes through the system in fixed-length packets with each input channel delayed by one packet time slot relative to the adjacent channel. With the switch control set to zero, the first data packet (1A) passes directly through the switch along with three empty packets. The switch control is then incremented by one [fig. 3(b)] and packets 1B and 2A are transmitted. During the next time slot [fig. 3(c)] packets 1C, 2B, and 3A are transmitted. After one

rotation of the switch control, the system reaches a steady-state condition as shown in figs. 3(e) and 3(f). Parallel event fragments are converted to assembled event streams with no loss of bandwidth. The same principle can be extended to any $N \times N$ switch. After one rotation of the switch control, all the fragments from event E are delivered to a single output port (E modulo N).

Figure 3 is an idealized version of the barrel shifter switch operation. Because event fragments are not all of equal length, segregating event data by packet in a real system would not be efficient. Some provision must be made which allows a single event fragment to span several packets or allows fragments of several events to occupy a single packet. To eliminate any such correlation between event and packet boundaries, the hardware actually maintains "logical" FIFO buffers for each input and output channel instead of the single buffer shown. Data placed in one logical buffer of a transmitter appears in the corresponding logical buffer of a receiver, independent of event boundaries. Data appearing in the receiver is essentially a memory image of the complete event at the detector. Some minor formatting may take place in the receiver or as a preprocessing step in the higher level trigger software.

7. Online processors

The rapid increase in processing power available in commercial integrated circuits presents the high energy physics community with important opportunities for SSC era data acquisition systems [9]. The processors that will be available in the late 1990's will allow enormous amounts of computing power to be utilized online and will permit commercial high-level language programmable devices to be used for tasks previously performed by home-brew, hard-wired, or microcoded devices.

Reduced Instruction Set Computer (RISC) microprocessors are already commercially available with processing power of 20 VAX 11/780 equivalents per chip, and a number

of different manufacturers expect chips of 100 VAX power by the early 1990's. Digital Signal Processors (DSPs) and other more specialized chips offer even greater amounts of processing power with only slightly less convenience.

RISC microprocessors have gone from an academic research project to practically a computing industry standard in a short period of time. Every leading semiconductor manufacturer and computer vendor have RISC projects underway. The current generation of RISC processors has already surpassed the more common Complex Instruction Set (CISC) architecture (typified by the Motorola 68030 and the Intel 80386) in performance, and shows signs of surpassing mainframe performance as well.

What is RISC, and why do these processors have such high performance? The principle of RISC is to keep the instruction set of the processor as simple as possible, so that all instructions can be executed in a single clock cycle. This is in contrast to the prevailing design philosophy of the 1960's and 1970's, where instruction sets were filled with an enormous variety of instruction types and addressing modes, supposedly to make it easier to write compilers for high-level languages. However, study of the instructions used by compilers indicated that only a small fraction of the complex instruction sets were frequently used, and that the complexity led to a large loss in performance even when performing the simplest instructions. Eliminating many of these instructions results in an architectural simplicity allowing the remainder of instructions to be executed in one processor clock cycle. Infrequently performed complex operations are done in software rather than in hardware, so that there will be no performance penalty for the vast majority of simpler operations.

More specifically, the time to perform any given computing task depends on the product of three factors: the number of instructions needed to do the task; the number of clock cycles needed for each instruction, and the amount of real time required for each

clock cycle. RISC processors win by giving an enormous reduction in the number of cycles per instruction. While a VAX 11/780 averages 10.6 cycles per instruction and the Motorola 68020 averages 6.3 cycles, a modern RISC processor like the MIPS R3000 requires only 1.25 cycles per instruction. Future RISC processors will be capable of executing multiple instructions per clock cycle. Moreover, the simplicity of the architecture allows the RISC processors to run at higher clock speeds than CISCs, as well as allowing implementation in faster technologies like ECL or gallium arsenide that are not suitable for CISC architectures. This gain is accompanied by an increase in the required number of instructions, but only by 20-50%, leaving the RISCs with a large overall performance gain.

A variety of architectural techniques allow RISC processors to achieve their goal of single cycle instruction execution. Typically, there is a relatively small number of instructions and addressing modes, and a fixed instruction format. The RISC designs are often a load/store architecture, with large register sets and no memory-to-memory instructions. Control logic is usually hard-wired, with none of the microcoded control typical of minicomputers and CISC processors. Much more of a burden is put on the compilers, with sophisticated optimizations required to achieve full performance in high-level languages. The currently leading RISC implementations are those where significant effort was put into compilers right from the very beginning. In a sense, the RISC philosophy shares the complexity of processor design between the chip architects and the software writers, rather than putting all the burden on the hardware design as in a CISC processor.

The advantages of RISC architecture are by now widely recognized, and there are a large number of RISC processors already commercially available. Furthermore, all of the leading RISC manufacturers have announced plans for higher speed versions of

their chips. It is not overly optimistic to expect individual processors with between 50 and 100 VAX equivalents in performance well before the turn-on of the SSC.

Mention should also be made of digital signal processors (DSPs). These chips have traditionally only been used for special purpose applications such as dedicated trigger processors, and have not been suitable for more general purpose use due to limited memory space, no floating point, and lack of high-level languages and good program development tools. The current generation of DSPs remedies most of these deficiencies. These new chips typically have full IEEE floating point, large memory address space, and speeds of up to 100 MFLOPS. They are supported by operating systems offering high-level languages and good program development tools. The high-level software tools should make these "special purpose" processors much more accessible to the average physicist than in the past, when their coding was normally restricted to a small group of experts. These DSPs are highly suited for front end processing and triggering applications. On the other hand, they are not yet appropriate for general purpose processing farms. The RISC processors have the important advantage of allowing code development on workstations or minicomputers using the same chip set and software tools as used on the processor farms.

Industry will provide us with extremely powerful high-level language processors for both filtering (Level 3) and triggering (Level 1-2) applications. We should resist as much as possible the temptation to build hard-wired, nonprogrammable, or microcoded devices. *The real challenge will not be in providing the processing power, but rather in insuring that the extraordinarily powerful arrays of processors are actually doing what we want them to do. It is not too early to start developing tools for program specification and verification. Without these tools, we will be unable to enjoy the full benefits that processor technology can supply.*

clock cycle. RISC processors win by giving an enormous reduction in the number of cycles per instruction. While a VAX 11/780 averages 10.6 cycles per instruction and the Motorola 68020 averages 6.3 cycles, a modern RISC processor like the MIPS R3000 requires only 1.25 cycles per instruction. Future RISC processors will be capable of executing multiple instructions per clock cycle. Moreover, the simplicity of the architecture allows the RISC processors to run at higher clock speeds than CISCs, as well as allowing implementation in faster technologies like ECL or gallium arsenide that are not suitable for CISC architectures. This gain is accompanied by an increase in the required number of instructions, but only by 20-50%, leaving the RISCs with a large overall performance gain.

A variety of architectural techniques allow RISC processors to achieve their goal of single cycle instruction execution. Typically, there is a relatively small number of instructions and addressing modes, and a fixed instruction format. The RISC designs are often a load/store architecture, with large register sets and no memory-to-memory instructions. Control logic is usually hard-wired, with none of the microcoded control typical of minicomputers and CISC processors. Much more of a burden is put on the compilers, with sophisticated optimizations required to achieve full performance in high-level languages. The currently leading RISC implementations are those where significant effort was put into compilers right from the very beginning. In a sense, the RISC philosophy shares the complexity of processor design between the chip architects and the software writers, rather than putting all the burden on the hardware design as in a CISC processor.

The advantages of RISC architecture are by now widely recognized, and there are a large number of RISC processors already commercially available. Furthermore, all of the leading RISC manufacturers have announced plans for higher speed versions of

8. Online processing farm

Even with the large amount of processing power projected per CPU, an online processing farm consisting of quite large numbers of processors will be required. Current projections of total CPU power required of an online SSC farm are from 10^5 to 10^6 VAX-11/780 equivalents. These estimates are based as much upon what appears feasible as upon what the real requirements will be. This aggregate amount of computing power would allow roughly 10 to 100 VAX-seconds per event on average to perform final event selection and filtering. Such computing power could be provided by approximately 1000 to 5000 CPU's.

One proposed architecture for the processor farm consists of multiple processors, perhaps four, mounted on each processor board. The boards would be plugged into a standard bus, perhaps Futurebus+. Data would be delivered to the processor memories either over that bus or via a high-performance external bus from the parallel event builder. Thus, a natural architecture would include as many data links between the event builder and the farm as there are links into the event builder, all with comparable bandwidth. These links would then each feed a bank of processors contained in one or more crates. Processors in a crate could share a boot node and perhaps a mass storage device. A separate control and monitoring network would link all processors.

Computer manufacturers are becoming increasingly interested in massive parallel computing on a scale similar to the one discussed here. This scale also leads them to quite similar, although somewhat more general, solutions. At a recent workshop, participants from Intel and IBM both discussed proposals to build very large parallel processors for scientific computing with between 10^3 and 10^4 nodes using loosely-coupled RISC processors with message passing. These proposals differ from the farms that we discuss in that the commercial projects incorporate more general interconnection than needed

in the online environment. For instance, the Intel project is based on a two-dimensional hypercube in which interconnection is provided by a mesh of custom VLSI routers, each connected to four nearest neighbor routers and to a processor node. Nonetheless, the more general interconnection scheme may satisfy our bandwidth needs, in which case its cost could be compared to the advantages of commercial or collaborative support from industry.

One of the major questions concerning the implementation of processor farms with more than a thousand nodes is that of managing the data flow and the processing of such a large number of CPU's. The system design must allow continued operation as individual processors hang, fail, and are brought back online. It must allow *in situ* debugging of production code, as well as offering a development environment for new code. It must also offer facilities for verifying the operation of processors and of code. In fact, the processor farm must offer an environment with which we can feel as comfortable as we presently do with our online minicomputers and which offers all the operating system tools of the popular minicomputers in the parallel processing environment.

9. Online mass storage

The expected data bandwidth for recording on mass storage is from 10 to 100 Mbytes/sec, although rates an order of magnitude larger or smaller are possible depending on the physics goals of each experiment and on the effectiveness of online triggers and data filters. The range of 10 to 100 Mbytes/sec expected corresponds to 10 events/sec of 1 Mbyte each and to 1000 events/sec of 100 Kbytes each, which assumes that the data recorded per event can be reduced if it is desirable to record higher event rates. These data bandwidths could be provided either by high-speed recording devices now being developed or by systems of parallel recorders. Large volumes of recorded data, 100 to 1000 Tbytes/yr per experiment, must be handled in any case.

Optical tape and disk drives will probably provide the highest bandwidths and data densities available in the future. For instance, E-Systems, Inc. is developing an optical tape drive capable of recording between 3 and 12 Gbytes/sec for use in the mid-1990's. However, the most cost effective solutions will be systems utilizing components with wide commercial application. E-Systems estimates that the key recorder technology of the early-1990's will be helical scan recording on 19-mm magnetic tape cassettes. Motivated by the commercial broadcast industry, recorders are being developed for 98.8 Gbyte storage per cassette recorded at 30 Mbytes/sec and for 190 Gbytes storage per cassette recorded at 15 Mbytes/sec. Between one and seven of these drives operating in parallel could provide the bandwidth expected. The parallel drives can be interfaced to the processor farm through cross-point switches which tie the drives to the buses on which the processors reside.

Although the pentabyte of data which might be recorded in a year would require more than six million nine-track tapes or about five million 3480-type tape cartridges, it will occupy only about 5000 large 19-mm cassettes using the helical scan technology, or about 12,000 more cost-effective smaller cassettes. The libraries of these cassettes can be accessed by robotics. In fact, the system of drives, tape storage, and robotics could be conveniently located at a central offline computing facility with a high-speed fiber optic data link from the online system, as is now done at KEK. The cost of such a system depends upon the required access time to any stored data set; however, based upon large systems which they have provided in the past and on systems which they are developing for use in the early-1990's, E-Systems estimates that a system with bandwidth in excess of our requirements, with capacity of many pentabytes, and with about 15-second access time would provide storage at a media cost of about two dollars per gigabyte and a system cost of about five dollars per gigabyte.

10. Concluding remarks

The bandwidths and processing power required in data acquisition systems for experiments at the SSC are unprecedented in comparison with the current generation of experiments. Nonetheless, by exploiting advances in technology, particularly in the area of custom VLSI and in areas driven by commercial applications, the system features sketched in the sections above will make SSC data acquisition feasible. A common thread throughout this view of SSC data acquisition is the extensive use of parallelism to provide the requisite level of system performance with practical and cost-effective components. The use of parallelism also will allow the performance of the system to be scaled to the requirements of individual experiments while utilizing the same architecture and components. The high level of system performance achievable offers flexibility in the trade-off between trigger and data acquisition performance during the life of any particular experiment.

The architecture sketched here has additional interesting and important features. The data flow and control paths have been separated, which contributes to the simplicity of the design, as well as eliminating a bottleneck in data bandwidth often encountered in the past. The overall data flow in this architecture is designed to simplify control of the flow.

In addition to the processors in the farm, the system will include large numbers of specialized processors, digital signal processors, and RISC processors operating in parallel and embedded in electronics mounted on the detectors. The heavy demands placed upon software by a data acquisition system of this scale and performance have not been addressed in this paper, and cannot be overemphasized.

The architecture sketched in this paper illustrates a direction rather than a design for SSC data acquisition. Much development work is of course required prior to a timely

implementation of a high-performance and reliable system. An important feature of the system design will be an implementation of the architecture which allows the eventual system to exploit advances in technology which occur during the development of the system, or even after the system is commissioned. The need for this evolutionary aspect of the system design is clear in the case of processors in the farm, but exists as well for other elements in the data path and for trigger processors.

Acknowledgments

This paper draws upon the work of several past workshops on SSC experimentation, particularly the Workshop on Triggering and Data Acquisition for Experiments at the Superconducting Super Collider held in Toronto, Canada, January 16-19, 1989. It also draws upon research and development currently in progress under the auspices of: SSC Generic Detector R&D Program, including work at the University of Pennsylvania and Fermilab; the Fermilab Advanced Computer Program; and the Beauty Collider Detector collaboration. It is based upon the work of the Task Force on Electronics, Triggering, and Data Acquisition at the SSC. The steering committee of this task force, which includes E. Barsotti (FNAL), R. K. Bock (CERN), R. Downing (Illinois), F. Kirsten (LBL), A. J. Lankford (SLAC), J. Thaler (Illinois), and Y. Watase (KEK), would like to thank the large number of participants in the task force work, particularly those from industry.

References

- [1] Proc. Workshop on Triggering, Data Acquisition and Offline Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Batavia (1985) eds. B. Cox, R. Fenner and P. Hale.
- [2] T. J. Devlin, A. Lankford and H.H. Williams, Proc. Summer Study on the Physics of the Superconducting Super Collider, Snowmass (1986), eds. R. Donaldson and J. Marx, p. 439.
- [3] Proc. Workshop on Triggering and Data Acquisition for Experiments at the Superconducting Super Collider, Toronto (1989).
- [4] L. Callewaert *et al.*, IEEE Trans. Nucl. Sci. NS-36 (1989) 446.
- [5] A. E. Stevens *et al.*, IEEE Trans. Nucl. Sci. NS-36 (1989) 517.
- [6] Thomas Devlin, Proc. Workshop on Triggering, Data Acquisition and Offline Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Batavia (1985), eds. B. Cox, R. Fenner and P. Hale, p. 244.
- [7] L. R. Fortney, Proc. of the Workshop on Triggering, Data Acquisition and Offline Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Batavia (1985), eds. B. Cox, R. Fenner and P. Hale, p. 232.
- [8] E. Barsotti, M. Bowden and C. Swoboda, Proc. Workshop on Triggering and Data Acquisition for Experiments at the Superconducting Super Collider, Toronto (1989).
- [9] Irwin Gaines, Proc. Workshop on Triggering and Data Acquisition for Experiments at the Superconducting Super Collider, Toronto (1989).

Figure captions

Fig. 1. Architectural requirements of the trigger and data acquisition system for an SSC detector.

Fig. 2. Overview of the architecture of an SSC data acquisition system.

Fig. 3. Schematic operation of a barrel shifting switch implemented as a parallel event builder.

4.40
0378A1

Fig. 1

Fig. 2

Fig. 3