

30/2-26-86 JS (1) cat

SLAC-PUB--3873
DE86 006953

OVERVIEW OF DATA FILTERING/ACQUISITION FOR A 4π DETECTOR AT THE SSC*

Summary Report of the Data Filtering/Acquisition Working Group
Subgroup A: Requirements and Solutions

DR-156

A. J. LANKFORD

Stanford Linear Accelerator Center
Stanford University, Stanford, California, 94305

G. P. DUBOIS

California Institute of Technology, Pasadena, CA 91125

1. Introduction

The task of the Data Filtering/Acquisition Working Group was to examine the feasibility of acquiring data at SSC event rates from a 4π detector with approximately three-quarters of a million electronic channels. The scope of this examination included all electronics between amplifiers on each detector element and transfer of data to off-line computer facilities. In particular, subgroups of the working group examined (1) data buffering, digitization, and reduction, (2) architectures to handle the data flow from the electronics on the detector to an online farm of processors and to mass storage, (3) strategies for reducing the trigger rate using the processor farm, and (4) general considerations such as event sizes and trigger rates.

This report provides an overview of the work on data filtering and acquisition. Section 2 reviews the assumptions made about the detector, event rates, and event sizes. Section 3 outlines the overall picture of data flow through the data acquisition system. Section 4 describes the problems of and the general approach to handling of the data during analog and higher level trigger decision periods. Section 5 sketches flow of the data to the online processor farm. Section 6 comments on software trigger strategies. Section 7 sketches aspects of the overall picture of a generic data acquisition system. Finally, Section 8 provides a summary of major issues and some needed developments. A number of separate reports, from subgroups or based on presentations made at the workshop, are included in the proceedings of the workshop and referenced by this overview.

* Work supported by the Department of Energy, contract DE-AC03-76SF00515.

Contributed to the Workshop on Triggering Data Acquisition and Computing
for High Energy/High Luminosity Hadron-Hadron Colliders,
Batavia, Illinois, November 11-14, 1985

②

MASTER

2. Data Flow Requirements of a 4π Detector

50-1

This section summarizes assumptions made by the Data Filtering/Acquisition Working Group about features of SSC performance, general 4π detectors, and triggering which impact the design of a data acquisition system. In general, these assumptions are drawn from the Proceedings of the 1984 Summer Study of the Design and Utilization of the Superconducting Super Collider.¹ In particular, please see "Detectors and Experiments for the Superconducting Super Collider"² by M.G.D. Gilchriese and "4 π Detectors"³ by G.J. Feldman, M.G.D. Gilchriese, and J. Kirkby.

2.1 INTERACTION RATE AND TRIGGERS

The SSC design luminosity of 10^{33} cm⁻² sec⁻¹, with an expected total inelastic cross-section of approximately 100 mb, will yield an interaction rate of 10^8 Hz. With 33 nsec between beam crossings, the crossing frequency will be 3×10^7 Hz, and there will be 3.3 interactions per crossing.

The trigger will proceed in three stages: an Analog Trigger, followed by a Higher Level Trigger, and finally a Software Trigger. The results of the studies of the Analog Trigger⁴ and Higher Level Trigger⁵ working groups at this workshop suggest that the Analog Trigger will provide a rejection factor of about 1000 in 1 μ sec, reducing the 3×10^7 rate to between 10^4 and 10^5 Hz, and that the Higher Level Trigger, without using any digitized information, will provide an additional rejection factor of about 100 in an additional 10 μ sec, reducing the rate to between 100 and 1000 Hz. Further reduction in rate, to about 1 Hz is left to the Software Trigger.

2.2 DETECTORS

The 4π Detectors Group at Snowmass '84 described three detector examples in Ref. 3. Most details of the detector, and hence most differences between the detector examples, are not important to the general features of the data acquisition system for the detector. Details of the type of detector and detector element occupancies will influence the details of the electronics of each channel as discussed in the report of the Hardware Subgroup of this working group.⁶ In addition, the occupancies will determine the amount of multiplexing possible at various stages in the data acquisition. In general, the extent of charged particle tracking, due to the large number of measurements necessary per track, will determine the amount of data flowing through the data acquisition system. For that reason, the model detector considered here resembles the SCD detector of Ref. 3 with its complete charged particle tracking. However, in order to consider somewhat

more severe requirements, each wire in the model detector is considered to have pulse height measurement at both ends for third coordinate reconstruction.

The model detector for data acquisition considerations is:

1. Coverage of rapidity between ± 5.5 units.
2. Charged particle tracking with 100 samples per track with drift times of 100 nsec. 175,000 wires instrumented with drift time measurement and pulse height measurement at both ends.
3. Calorimetry with 50,000 towers of 3 longitudinal segments each. Wide dynamic range pulse height measurement provided by two gain ranges. Some means of separating signals from nearby crossings.
4. Muon tracking with 100,000 wires with drift time and charge division measurements.
5. Vertex detection is to be treated as a special case and is not considered here.
6. Total electronics channel count is 850,000.
7. Some technique of waveform sampling will probably be used for signal digitization for all detector types, in order to provide the multihit capability necessary for tracking devices and the pulse shape information necessary to deconvolute calorimeter pulses. The total information content of all samples before zero suppression will be 20-80 MByte depending on detector details.

2.3 EVENT SIZE

The typical minimum bias interaction at the SSC is expected to have about six charged particles (and three neutrals) per unit of rapidity. Thus, there will be about 66 charged tracks in the tracking detectors and about 100 charged and neutral particles in the calorimeters for the model detector covering 11 units of rapidity.

For an event with a multijet trigger with a jet threshold of 500 GeV, there may be 30-40 charged particles per jet plus a background which resembles a minimum bias event. For two jets there will be about 150 charged tracks, and for four jets there will be about 200 charged tracks. Typical event size calculations use 200 charged particles plus 100 neutrals.

The overall event size from the model detector will be dominated by the tracking chambers where there will be 100 hits per track (as compared with three samples per calorimeter tower). In addition, during the 100 nsec drift time of

the tracking detector there will be signals from interactions during two or three previous beam crossings and two or three subsequent beam crossings. Thus, in addition to hits from tracks in the event of interest there will be hits from the equivalent of nine minimum bias events. The total number of hits is then equivalent to the number from 800 charged particles, or 80,000 hits. Assuming that the drift time, pulse height, and third coordinate measurements from each wire can be reduced to five bytes including a wire label, the 80,000 hits will produce 400 KBytes of charged particle information.

In the calorimeters there will be about 1200 particles assuming resolving times of about 100 nsec, and three times as many for resolving times of 300 nsec. Allowing for two hit towers per particle on average and three hits per tower, there will be 7200-22,000 hits per trigger. However, with timing information about the hits, only hits from the events from different beam crossings than the trigger can be rejected, leaving 3200 hits per trigger. At four bytes per hit, there would then be 12 KBytes per event. There may be multiple samples per signal if high occupancies demand that signals from different crossings on a single channel be deconvolved (see Ref. 6); however, we assume that this deconvolution is performed in hardware and count only one sample per event. Furthermore, it will be possible to pack the information from all three sections of a tower in six bytes with a label, leaving only 6 KBytes per event. These numbers are in any case negligible compared to the charged particle data. The quantity of data from the muon tracking system will also be negligible.

Thus, 400 KBytes of data are expected from a typical event after all reduction.

3. General Model of Data Flow

The general model of data flow from the detector to off-line processing consists of a series of stages, or levels, of processing, buffering, and filtering of the data. A level in the data acquisition system can be modelled as shown schematically in Fig. 1. An input data stream must be buffered while portions of the data are used to make a trigger decision. The buffer can be thought of most simply as a delay line into which data is placed for the length of time required by the trigger decision. The data appears at the output of the delay line as the trigger decision completes. It is filtered, i.e.: retained or discarded, according to whether the trigger decision was "accept" or "reject". In reality, the buffer may be digital or analog, or even a shift register or a physical delay line if it can preserve necessary attributes of the signal throughout the decision time. The data stream can be processed at any point, before or after buffering and filtering.

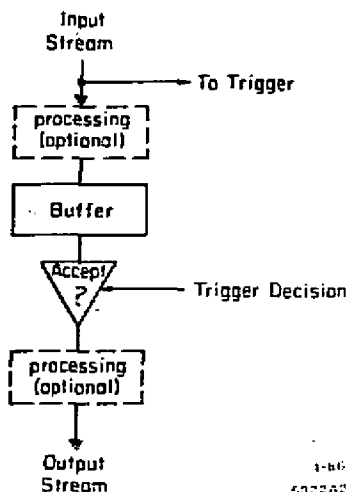


Fig. 1. Representation of a level of buffering and processing in the data acquisition system, corresponding to a level in the trigger decision process.

The data acquisition model studied by the Data Filtering/Acquisition Working Group consists of three levels, corresponding to the Analog Trigger decision (Level 1), the Higher Level Trigger decision (Level 2), and the Software Trigger decision (Level 3). In fact, a considerable amount of flexibility is available. More levels are possible, and may simplify processing or data flow. Furthermore, any level can consist of sublevels, as will be necessary in levels involving extensive processing such as Level 3. Processing, such as digitization, zero-suppression, correction, and compaction, can be applied at whatever point is convenient for overall data flow considerations. Section 4 and Ref. 6 discuss levels 1 and 2 of the data acquisition. Sections 5 and 6 and Refs. 7 and 8 discuss Level 3. Figure 2 outlines the three levels of the model used.

The inputs to Level 1 are the raw detector signals from the approximately 850,000 electronic channels arriving at the beam crossing rate of 3×10^7 Hz. Data is stored in the Level 1 buffer for the 1 μ sec required by the Level 1 trigger. The Level 1 buffer must be in effect dual-ported in order to avoid deadtime at the input while data is being readout from it.

Input to Level 2 buffers is at the rate of 10^4 - 10^5 Level 1 triggers per second. Each event candidate is 20-80 MByte (or equivalent analog size) before zero

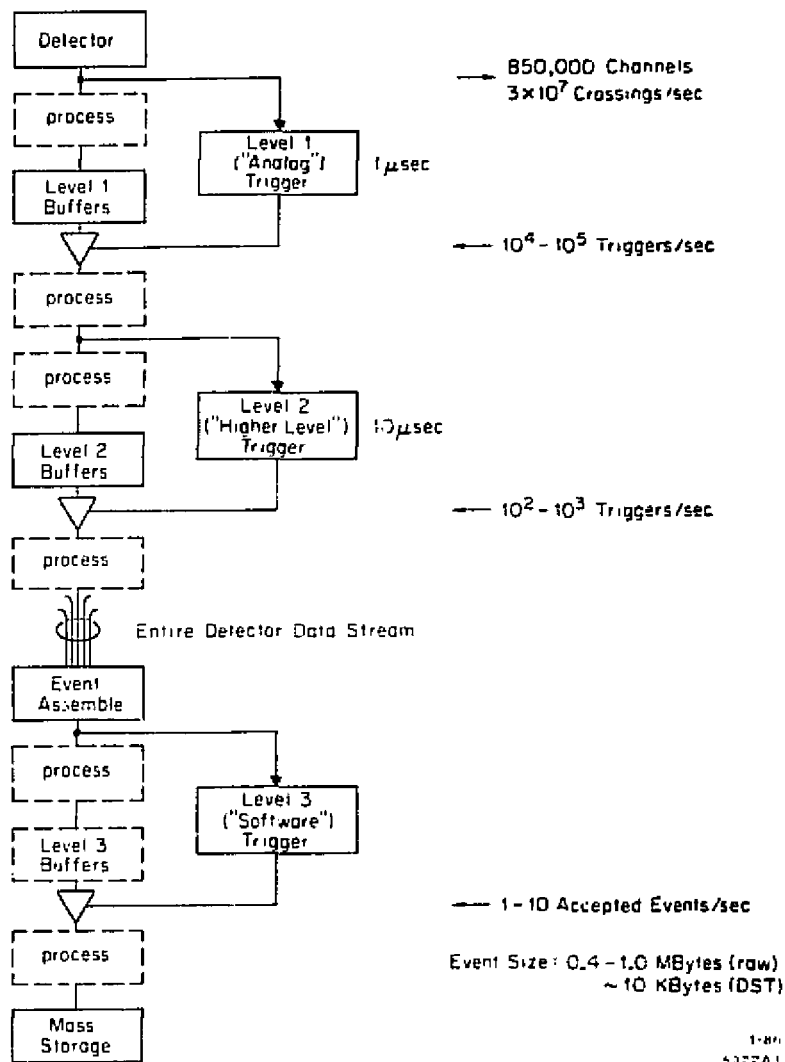


Fig. 2. General model of data flow through levels of the data acquisition.

suppression and further data reduction. The Level 2 trigger decision time and buffering time are about 10 μ sec. Some processing schemes could take advantage of shorter decision times on some events. The Level 2 trigger can use either raw data or processed data from Level 1.

The input to Level 3 is 100-1000 Hz. Final hardware processing and buffering is done in preparation for the data being moved into a farm of Level 3 processors. Event sizes are 300 KByte - 1 MByte at this stage. Full events are assembled for the first time in the Level 3 processors. At Level 3 the power of full event reconstruction can be implemented as necessary to perform the final selection of events. Full Level 3 decision times may be many seconds; however, events are filtered in parallel at this stage instead of in a pipeline. The desired output rate from Level 3 is between 1 and 10 Hz to some mass storage device. The amount of data preprocessing performed on-line in real-time, as opposed to off-line, is again very flexible.

4. Data Handling During Hardware Trigger Decisions

Each electronics channel, particularly channels of tracking detectors, will provide a number of samples per measurement. The total number of samples for nearly a million channels represents a vast amount of data for each beam crossing which must be retained without deadtime for the one microsecond required for the Analog Trigger decision. The data from 30 crossings will accumulate while the Analog Trigger decision is made.

The above problem recurs during the Higher Level Trigger decision period, however with different parameters. The total amount of data to be stored can be reduced by suppressing zeroes and by processing multiple measurement samples into simpler quantities. Only data from crossings accepted by the Analog Trigger need be retained; however, accepted crossings still occur at a frequency between 10^4 and 10^5 Hz. Moreover, the period for which the data must be retained during the Higher Level Trigger decision is 10 μ sec. At the higher end of the expected Analog Trigger rate some buffering would be necessary to handle fluctuations in the interval between accepted crossings.

Before the data from crossings accepted by the Higher Level Trigger is sent to the Software Trigger, the data must be preprocessed. This preprocessing certainly involves zero suppression and data formatting. It also involves applying channel-by-channel corrections. Furthermore, to restrict data transfers to the software trigger processor farm to manageable rates, preprocessing must involve deriving drift times (or distances) and pulse heights (or third coordinates) for

tracking measurements and, most likely, deconvolving calorimeter pulse height measurements.

The basic solution to the problem of handling large amounts of data until the Analog Trigger decision is to pipeline the data in some fashion on a channel-by-channel basis. In some cases, shift registers or analog delay lines may be sufficient; however, in most cases memories, either digital or analog, which can function in a dual-ported fashion would provide pipelined buffering which could be dead-timeless. Much preprocessing can be done on a channel-by-channel basis, which would also provide the necessary pipelining during the Higher Level Trigger decision. Alternatively, data could be buffered as during the Analog Trigger decision, and then preprocessed at the lower rate of accepted events from the Higher Level Trigger. This solution offers either more preprocessing time or the opportunity to multiplex and thus save electronics. A considerable amount of flexibility is available in how data is buffered and when it is processed; however, choices depend upon the details of the detectors and possibly upon what data is needed by the various levels of trigger. In all cases, the buffering and processing for large amounts of detailed data will require new VLSI developments, such as improved analog memories, large dynamic range FADC's, specialized hardware processors, and multifunction circuits in a single package. These developments are within the realm of existing technology. The report of the Hardware Subgroup⁶ discusses in more detail the design of a possible solution to the problems of buffering and preprocessing the large quantities of data. All solutions are expected to involve much local processing and data correction. Special-purpose processors to assist in the determination of calibration constants may be useful. Diagnostic features which permit the testing of each stage of processing will be necessary.

5. Data Flow to the Online Processor Farm

With as many as 1000 Higher Level (Level 2) triggers per second and as much as 1 MByte per event, the required data flow capability between the detector component electronics and the processor farm must be 1 GigaByte per second. In fact, some excess capacity must be available to handle variations in intervals between triggers. These large bandwidths can be handled by a manageable number, of high-speed busses operating in parallel. For instance, twenty-five busses of 40 MByte capacity, which is similar to current FASTBUS capabilities, would provide the 1 GByte/sec necessary. Each bus would be fed from a buffer containing data from some portion of the detector. The entire event is sent to a single processor; consequently, each processor must have access to all the busses. Since the bandwidth into any one processor is similar to the bandwidth of any

one bus, many events (as many as there are busses) must be in the process of being transferred to as many processors at any moment. These considerations are discussed in more detail in Refs. 9, 10, and 11 and in the report of the Model Architectures subgroup.⁷

6. Software Trigger Strategies

The Software Trigger must reduce an input trigger rate as high as 1000 Hz from the Higher Level Trigger to a final rate of about 1 Hz. The conditions which the Software Trigger will demand are dictated by the various physics processes being studied and will naturally be the logical "or" of several conditions. Most simple requirements will already have been applied by the hardware triggers. In fact, if the input rate to the Software Trigger is as low as 100 Hz, as the Higher Level Trigger Working Group at this workshop felt was possible, the rate will be low because many of the trigger requirements will have been applied in hardware. Consequently, the tools available to the Software Trigger in reducing the trigger rate will be determined by careful tradeoff with the hardware triggers. The hardware triggers in most instances will be faster, but the Software Trigger will afford more flexibility. Information from the hardware triggers will guide the Software Trigger. For instance, rather than attempting to reconstruct all tracks, the Software Trigger can simply confirm the tracks found by an efficient hardware track finder. Even if no hardware track finder is used by the Higher Level Trigger, the Software Trigger will want such a device to increase its speed.

The Software Trigger, by offering the capability of fully reconstructing an event, affords the trigger all the tools available offline in selecting events. However, full event reconstruction is expected to take about 1000 seconds on a VAX-equivalent computer. Consequently, it will be impossible without a million VAX equivalents to fully reconstruct all the events input to the Software Trigger. To reduce the amount of processing power required the Software Trigger must consist of levels, with progressively more time-consuming analysis occurring at the higher levels. The actual bus architecture of the processor farm need not be tiered although the processing is. The Software Trigger algorithms must be carefully designed for execution speed. The Software Trigger should only confirm and pursue the conditions which caused the hardware trigger. For instance, if the hardware trigger tagged an event as a candidate for the decay of a W -pair into leptons then the Software Trigger should not investigate the event as a possible four jet event. Reference 8 describes the Level 3 filters of the CDF experiment as an example of a software trigger and draws some implications for software triggers at the SSC.

As an experiment develops at the SSC – as its physics goals and its hardware become better understood – much of the event selection initially done offline can be done online by the Software Trigger. Similarly, some data preprocessing and event reconstruction done offline can be moved online. Computing power that would have been available offline for such processing could be placed online. The division of computing power between online and offline is largely a logistical issue; however, the principal advantage of online event reconstruction, or reconstruction of portions of events such as track segments, would be to reduce the amount of data recorded per event in order to record more good physics events. For instance, the DST information on a fully reconstructed event will be two orders of magnitude smaller than the reduced data input to the online processor farm. A spectrum of choices exist between a 1 Hz trigger rate of events with raw data and a higher rate of events with track vectors only. The rate of processed events will be limited by the amount of processing available online.

7. Overall Online Computing

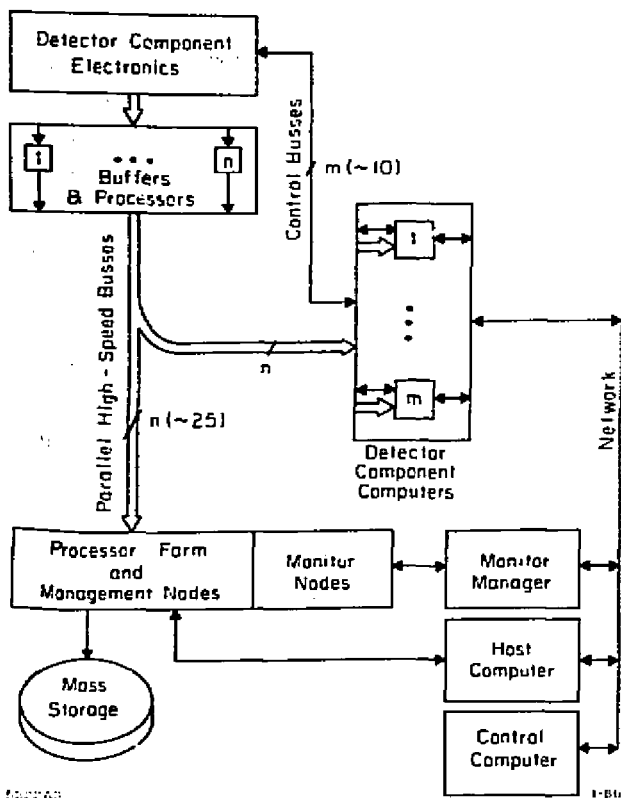
The overall organization of online computing for an SSC experiment can be fairly conventional, similar to an online system for a LEP-generation experiment. This section very briefly describes the online computing necessary to manage the flow of data and to perform other online tasks. Since a fairly conventional approach appears to be adequate, this subject was not studied in depth at the workshop.

The principal tasks to be performed by online computing include:

1. managing the flow of data from detector components to the processor farm and then to mass storage,
2. monitoring and controlling detector components,
3. determining and downloading detector constants,
4. providing support for the processor farm,
5. verifying the quality of the recorded data,
6. supporting detector and software development.

A block diagram of an organization of the processors necessary to perform these tasks is shown in Fig. 3. The blocks in this diagram represent tasks to be performed; however, in most cases they also represent separate processors or groups of processors to perform the tasks.

Sections 4 and 5 of this report have discussed the large blocks entitled "Detector Component Electronics" and "Processor Farm", the parallel high-speed



1-50

Fig. 3. Block diagram of processing tasks in an online computing system.

busses between them, and the processors which they contain. At the output of the detector component electronics, where data is buffered before being passed to the farm, there is the option of performing further data processing with either specialized processors or mini-farms before sending the data to the farm. In fact, these mini-farms may be a physical part of the farm while performing this logically separate task.

The accepted data flows directly from the farm to mass storage, without passing through the host computer. The mass storage is most likely local to the experiment; however, it could be located at a central computing facility

where it would be fed by a very high-speed optical link. On the other hand, offline computing power could also be local to the experiment, where it could be directly fed and maintained by the experiment and possibly flexibly allocated between online (i.e.: real-time) and offline processing. Some considerations of the relationship between online and offline computing are discussed in the report of the Off-line Computing and Networking Group.¹²

The detector is segmented into major detector components which are envisaged to be of a given detector type and to cover a certain region of solid angle. Each of these detector components will feed one of the twenty-five or so parallel high-speed busses along which the data flow to the processor farm. A detector component computer also resides on each of these busses. This computer, which may be of about the same power as present online computers, will be responsible for all tasks relating to the attached detector component which can be done without data from other detector components. For instance, it will be used in the development and testing of the detector component and will perform monitor and control of the component. It will perform calibration of the detector component and download constants to the processors in the electronics of the component. It will monitor raw data at various stages in the data acquisition preceding the farm and will control tests of all stages of data preprocessing. It will also act as a host for work relating to the detector component but demanding data from other components, such as specialized monitoring of events selected by the software trigger and graphic display of component performance. During development and calibration, the detector component computer will control the high-speed bus. During data acquisition, it will receive events in parallel as they are sent to the farm. A separate bus, of perhaps lower bandwidth, between the computer and the detector component will probably be used to download and control processors in the electronics and for detector monitor and control.

A portion of the large processor farm will be devoted to managing the triggering and preprocessing functions of the farm. Between ten and twenty VAX equivalents will be needed for this task. An additional function of the farm will be monitoring data at various levels in the software trigger decision, both to verify the trigger and to inspect the quality of the data. Although this task could be performed by the general farm processors, it will probably be done by a set of monitor nodes which spy on the data flow through the farm. The monitor nodes, however, will be similar to the general processing nodes. The monitoring task in the farm will be managed by a separate computer which will also collect and dispatch data from the monitor nodes to the detector component computers and to the main online host.

The role of the control computer is overall coordination and control of data

acquisition. It will also provide centralized diagnostic and status reporting on the performance of the detector during data taking. A separate host computer manages the processor farm, serves software development, and performs I/O. A large degree of flexibility is available in the division of tasks among the control computer, the host, the monitor manager, and the detector component computers. In fact, a separate host computer may not be necessary. A local area network will interconnect these computers, along with any additional graphics devices or other peripherals. Uniform software tools, throughout all stages of the development, implementation, and operation of detector components and throughout all levels of the data acquisition system, will be crucial to the operation of the complex system of online processors.

8. Summary

The difficult or new problems for data acquisition and filtering posed by a large detector at the SSC include:

1. buffering large amounts of data for large numbers of event candidates during trigger processing, even during the fastest possible analog trigger decisions,
2. preprocessing and reducing the complex waveform samples that will be required to separate signals of events from different beam crossings,
3. transferring large amounts of data,
4. effectively reducing the trigger rate by orders of magnitude using software filters,
5. managing very large arrays of processors, both in a processor farm and embedded in detector electronics,
6. reducing, rather than increasing, the large offline computing load presented by the vast amount of data.

The solutions to these problems seem feasible; however, developments in a number of areas are necessary to realize the solutions. These developments include:

1. further development of custom chips for front-end electronics, chips with large dynamic range, storage of samples covering a microsecond, deadtimeless readout, and manageable calibration,
2. development of custom chips for data preprocessing and reduction,
3. experience managing large data transfer rates into processor farms,
4. study of trigger criteria in their physics context and study of the division of criteria between hardware and software triggers,

5. experience managing large processor farms, as well as large online computing systems,
6. integrated design of detectors and their data acquisition and study of the division of processing between online and offline computing.

Much of this development will naturally arise as part of the experimental program at accelerators currently in use or under construction, such as the $p\bar{p}$ and e^+e^- colliders. Extrapolation from the scale of detectors of that generation to SSC detectors will be believable. Workshops and electronics R&D programs can also continue to address these issues, particularly as approaches to SSC detector design continue to develop.

References

1. Proceedings of the 1984 Summer Study on the Design and Utilization of the Superconducting Super Collider, June 23-July 13, 1984, Snowmass, Colorado (Ed.: R. Donaldson and J.G. Morfin).
2. M.G.D. Gilchriese, "Detectors and Experiments for the Superconducting Super Collider", Proceedings of the 1984 Summer Study on the Design and Utilization of the Superconducting Super Collider, June 23-July 13, 1984, Snowmass, Colorado, pp. 491-522.
3. G.J. Feldman, M.G.D. Gilchriese, and J. Kirkby, " 4π Detectors", Proceedings of the 1984 Summer Study on the Design and Utilization of the Superconducting Super Collider, June 23-July 13, 1984, Snowmass, Colorado, pp. 623-634.
4. Paolo Franzini, report of the Analog Triggers Working Group, submitted to the Workshop on Triggering, Data Acquisition and Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Fermilab, November 11-14, 1985.
5. M.A. Abolins and M.J. Shochet, report of the Higher Level Triggers Working Group, submitted to the Workshop on Triggering, Data Acquisition and Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Fermilab, November 11-14, 1985.
6. Peter S. Cooper, "A Feasibility Design for the Readout of a 4π SSC Detector; Data Filtering - Acquisition Group, Report of the Hardware Subgroup", submitted to the Workshop on Triggering, Data Acquisition and Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Fermilab, November 11-14, 1985.

7. David Cutts, report of the Model Architectures Subgroup of the Data Filtering/Acquisition Working Group, submitted to the Workshop on Triggering, Data Acquisition and Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Fermilab, November 11-14, 1985.
8. J.T. Carroll, "Level 3 Filters at CDF", submitted to the Workshop on Triggering, Data Acquisition and Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Fermilab, November 11-14, 1985.
9. D.J. Linglin, "UA-1 Software and Computing", submitted to the Workshop on Triggering, Data Acquisition and Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Fermilab, November 11-14, 1985.
10. T.J. Devlin, "Data Rates for Event Builders and Farms", submitted to the Workshop on Triggering, Data Acquisition and Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Fermilab, November 11-14, 1985.
11. L.R. Fortney, report on models of event building in processor farms, submitted to the Workshop on Triggering, Data Acquisition and Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Fermilab, November 11-14, 1985.
12. M.G.D. Gilchriese and S.C. Loken, reports of the Off-Line Computing and Networking Working Group, submitted to the Workshop on Triggering, Data Acquisition and Computing for High Energy/High Luminosity Hadron-Hadron Colliders, Fermilab, November 11-14, 1985.

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.