A HIGH SPEED DATA ACQUISITION SUBSYTEM FOR THE SSC

David Cutts and Richard Partridge
Brown University

Tricia Rankin
University of Colorado

Robert Downing, Michael Haney, Larry Jones
David Knapp, Jon Thaler†, and Marianne Winslett
University of Illinois

Hirofumi Fujii, Yoshihide Sakai, and Yoshiyuki Watase
KEK

Pekka Sinervo
University of Pennsylvania

Steve Quinton and Peter Sharp
Rutherford Appleton Lab

† Contact Person.
Address until July, 1990: SLAC - Bin 71
Email: DoctorJ at SLACvm.BitNet
Phone: 415-926-2553

1 January 1990 - 31 December 1990

Prepared for the U.S. Department of Energy
SSC Detector Subsystems Research and Development Program

Table of Contents

Abstract

We propose to develop a high speed data acquisition subsystem capable of meeting the needs of the high p_T solenoidal detectors being considered for the SSC. Our goal is to produce a detailed design of the data acquisition subsystem and demonstrate its feasibility through simulation and development of small scale prototypes of key components. The principal areas of proposed effort are architectural design, data collection electronics, parallel event building, and simulation of the subsystem. We also plan to establish collaborations with other groups working on triggers, high speed data links, on-line and off-line event processing, and data storage.

1. Introduction

A model architecture for SSC data acquisition systems has been discussed at a series of workshops on electronics, data acquisition, and triggering (see figure 1). The essential features of this model are:

- Data collection electronics to funnel data from the front-end electronics onto a moderate (100-1000) number of high speed data streams. The data collection electronics may also be required to provide data ordering and/or rate leveling functions.
- Trigger algorithms to recognize interesting events. A separate trigger data stream is likely to be required. The trigger must remain synchronized with and in control of the event data collection. The data rate must be reduced to a manageable amount before data movement can take place.
- A parallel event builder to assemble data from various detector elements into complete event records containing all data associated with a particular event. A moderate (100-1000) number of output streams transmit the event records to the processor farm.

- A parallel processor farm to apply sophisticated trigger filter algorithms and perform final event selection. An event is sent to one of the available processors in the farm, which then applies high-level trigger algorithms to determine if the event is accepted.
- A host computing facility for permanent recording of accepted events, the user interface, and monitoring of detector operation.
- High speed data links operating at bandwidths of at least 100 MB/sec. It is anticipated that fiber optics will be used to move data off the detector to the parallel event builder, while either fiber or copper may be used for links between the event builder, processor farm, and host computer system.

The time from initial conception to completion of construction will be longer for SSC detectors than has previously been the case. The increased latency period will have a greater impact on data acquisition design than on other detector subsystems. This impact is a result of the interaction between two effects. First, high SSC data rates will require that more event selection me made by the data acquisition subsystem than has previously been the case. Second, greater latency will increase the probability that the primary physics goals of SSC experiments will have changed substantially during the design and construction period. This means that if the event selection criteria have not been implemented with sufficient flexibility, the ability of the data acquisition to adapt to new goals will be reduced, and the ability of the detector to do physics will be compromised. Thus, the shortest possible latency period and the most flexible possible design will be crucial to a successful data acquisition system. The large cost of the data acquisition subsystem is likely to preclude any substantial retrofitting; it must be built right the first time. Our approach to the DACQ subsystem design is founded on this.

We propose to design, simulate, and prototype the components of the data acquisition system associated with moving data from the front-end electronics to the processor farm. We anticipate that the hardware trigger will reduce the event

rate from 10^8 Hz to $\approx$ 1-10 kHz, with the hardware trigger based on calorimeter energy sums, electron identification from calorimeter pattern recognition, tracks in the muon system, and possibly charged particle tracking and/or TRD information. With a typical event size of 1 MB, the data acquisition system must be able to handle a bandwidth of at least 10 GB/sec. This data rate will originate from perhaps $10^{5\pm1}$ front-end integrated circuits and must be delivered as complete events to a farm of 10^3-10^4 high-speed processors.

We plan to begin our work by focussing on the needs of a high-p_T detector at the SSC. We want to stress, however, that we expect many features of the data acquisition system to be the same for many future experiments. Accordingly, we hope to work closely with groups studying other aspects of data acquisition for SSC detectors and for future upgrades of existing collider detectors. We will use our simulation studies to identify the range of experimental conditions for which various solutions are adequate.

We are aware of the DACQ development work at Fermilab and other institutuions. The Fermilab group planned an SSC proposal which would have included several of the members of this proposal but it has been delayed. We have worked closely with Fermilab and other institutions in the past, and it is our strong desire to join forces with them at the earliest possible date. Our group does not yet encompass the entire range of expertise needed to actually construct an SSC data acquisition system. It is important that there be a comprehensive data acquisition collaboration to insure proper systems engineering across the range of SSC detectors.

We believe it is important to begin development of the data acquisition subsystem at this time. Electronics development and test beam studies of major detector systems will require data acquisition at an early stage of the experiment. For example, the Brown University group began serious development of the DØ data acquisition architecture 5 years before expected DØ operation and has installed six data acquisition systems for use in testing various parts of the DØ experiment dur-

ing the last three years. An SSC data acquisition subsystem will be significantly more complex than the one used by DØ and development needs to begin at this time.

2. Scope of Our Proposed Work

The scope of this proposal is limited to architectural design, to the development of some of the key components of the DACQ subsystem, and to the simulation of the subsystem as a whole and of its components. We will design in detail the parallel event builder, but we will only model the data collection electronics, the trigger, the high speed data links, and the online processor farm at the "black box" (functional) level. The development of these last four components will be done by other other SSC R&D groups with whom we hope to establish collaborative efforts. At the beginning we will work to define the overall architecture of these components and to understand the constraints that data acquisition needs will put on their design. We plan to develop protocols for interfacing the different parts of the data acquisition subsystem in order to create a coherent design environment.

We present here a brief summary of each of our projects. They are described in more detail in the next section.

Data Acquisition Architecture (the collaboration)

- Define and compare alternative architectures.
- Study event data flow and control signal flow issues.
- Study the relation between the front end, event builder, and the processor farm.

Triggers (Colorado and Rutherford)

- Study the relation between the trigger and DACQ data and control signal flow.

- Study the effect of trigger processing time and background rejection on data rates and DACQ architecture.

Simulation (Illinois and Penn first, then others)

- Construct "black box" models of various DACQ alternatives and simulate their behaviors.
- Construct detailed models of the front end, event builder, and event trigger components. Simulate them within the context of the "black box" models.
- Coordinate the design of the entire DACQ system.

Parallel Event Builder (Brown)

- Compare dual-port memory and high-speed switch architectures.
- Design a dual-port memory event builder, build prototypes.

On-line Processor Farm (In collaboration with Fermilab and Penn)

- Specify the input and output data rate requirements and processing speed needs.
- Study alternative data flow architectures to and from the processors.

Data collection front end (Penn, and collaboration with Fermilab and LBL)

- Specify the behavioral requirements in the context of the entire DACQ subsystem.
- Isolate specific components that will require specialized electronics design.

Data Links (KEK, and collaboration with Fermilab and LBL)

- Specify the required rate capabilities for various DACQ designs.
- Isolate specific components that will require specialized design.

3. Components of a Data Acquisition Subsystem

The sections below give further details of the proposed R&D program.

3.1. Data Acquisition Architecture

We propose to define a data acquisition architecture which solves the problems of acquiring data at the high SSC interaction rates and to study the performance of this system via extensive behavioral simulation. The important viewpoint to be taken here is the optimization of system performance, not of each component separately. Our architecture will strive for a modularity which will enable us to upgrade components without having to redesign the subsystem. A properly designed architecture will also simplify monitoring, debugging, and diagnostic tasks during construction and operation.

The departure point for this architectural study will be the model data acquisition system examined at workshops of the Task Force on Electronics, Data Acquisition, and Triggering at the SSC during the past year. The chief features of this model[1] are outlined in the introduction and focus on the highest attainable data rates. This study will produce a more refined and complete definition of the system architecture, clarifying the requirements that each DACQ component must satisfy.

The architectural study will be organized in small subgroups of the collaboration, each addressing the central issues of one data acquisition component. Efforts will initially focus on the specification of the required functionality and performance. The collaboration wil meet periodically to report on progress and to maintain a coordinated effort. The examples below illustrate the nature of issues to be considered.

Front-End Electronics

The architecture of the data acquisition system begins with the buffering of data in the front-end electronics during the first and second level trigger decisions

and with the collection of data from these front-end buffers. The depth and nature of buffers, *i.e.* whether they are logical pipelines, FIFO's, zero-suppressed, *etc.*, affects the data rate capability, the nature of processing, and the nature of data collection of the architecture as a whole. The nature of the data collection, *e.g.* data driven or command driven, similarly affects the architecture and its performance. The front end is where high data rates stress the data acquisition most severely, so a good understanding of the interaction of the data collection electronics with the first level trigger and the data links is particularly important.

Several groups are working on the design and prototyping of the front end electronics for SSC applications. Given the close interaction between this and the rest of the DACQ system, it is important that the efforts be coordinated. We are planning to form liasons with the groups at Fermilab, LBL, and Pennsylvania, so that areas of common concern can be addressed jointly, ensuring that designs will evolve in a mutually consistent manner.

One such project has been the development at Pennsylvania of a wire drift chamber readout system that incorporates a time-to-voltage converter, an analog Level 1 storage pipeline, a analog-to-digitial converter, and a Level 2 storage buffer, which has been funded by the Generic SSC R&D program. An effort is now underway to develop a Verilog simulation of this system that would eventually become a part of the DACQ system simulations discussed above. One of the researchers directly involved in that effort is also a participant in this proposal, and will be a liason between the two groups.

Data Paths

Architectures with separate control and data paths are easier to design than those in which they are mixed. Nevertheless, the communication links must have high bandwidth and must be readily verified and diagnosed. We will define the necessary functionality to accomplish these ends in a design which takes optimal advantage of advances of communication links in the computer and communication industry. We will design these links, and work closely with other designers of data

transmission components in order to optimize the performance and behavior of the data acquisition as a whole. One of our goals will be the incorporation of error detection and recovery into the design.

The Event Builder

Higher level triggers must process complete events; therefore, an event builder is needed to assemble the events. The speed, functionality and architecture of this the event builder are critical elements of the DACQ design. As a result of the asynchronous arrival of event data to the event builder it appears that a parallel architecture is desirable. We will study alternative solutions to the problem of building events in parallel. These efforts will be directed towards identifying the parameters which determine system performance, such as bandwidth of data links and sizes of data buffers, and to defining system requirements with respect to bandwidth, control, flexibility, and robustness.

Distribution of Processing Tasks

The processor farm exploits event parallelism in order to participate in the trigger at a high level by selecting events based on assembled data. However, many processing chores, including some at the trigger level, are best performed by exploiting the parallelism *within* each event which arises from the local nature of much of the particle information within the detector. In addition, tasks such as calibration and monitoring of detector parameters and performance will probably be performed by processors distributed throughout the subsystem. We will define the role of such processors in the architecture for initialization and control as well as for event data processing.

Data bandwidth into the processor farm must be matched to the computational performance of the processors. We will study bus and network solutions for data flow from the parallel event builder into processors which optimize bandwidth and control.

Within the processor farm, limited processor intercommunications will be re-

quired for diagnostics, monitoring of processor and detector performance, initialization, calibration, and other tasks. We will define the requirements that processor farm architectures must meet, and will work closely with groups doing research on development of processor farms for trigger purposes.

Prompt Triggers

The front end trigger is likely to require a data path distinct from the event data path, over which a subset of detector information will flow. If this is the case, than the coordination of the readout (for example, the distribution of control signals) is made more complicated. We plan to study the impact of proposed DACQ architectures on the relationship between the trigger and the event readout.

The data reduction acheived by each level of the trigger affects bandwidth, buffer size, and processing power requirements. We will coordinate closely with groups studying the event features that characterize various physics processes in order to understand the range of expected trigger rates and to study the effect these rates have on the scale of the data acquisition system. We will also study the control mechanisms between the trigger and other data acquisition components.

System Considerations

A number of considerations pervade the development of the subsystem architecture. For instance, the implementation must be modular in order to allow the evolution of system performance with technological advances and with changes in experimental needs. A properly modular architecture also enables the scaling of the system to required performance levels ranging from the needs for tests of large prototype detector components to the operation of the full experiment. Facilities for system verification and maintenance and fault diagnosis must be developed. System reliability will be one of the most important criteria by which the final design choice will be made.

3.2. TRIGGER STUDIES

The trigger system can be considered as the interface between the physics goals of a detector and the data acquisition requirements. Consideration must be given both to what distinguishes interesting physics from the background and to which of the possible signals should be incorporated at the primary, secondary and later triggering stages. The goals of this group include the understanding of: the signal to noise ratio in in various components of the detector, the speed at which information can be made available to the trigger, and how one would structure the overall trigger logic in order to reject background events early in the data aqcuisition and analysis chain while retaining as much of the potential physics as one can. For example, one can study how possible or worthwhile it is to do segment finding for an inner tracking detector at the crate level when information is available only for part of the detector.

In particular we propose to study the requirements imposed on the data acquisition system by potential trigger algorithms (which may be implemented either in software or hardware). We will evaluate the data reduction rate required by a proposed trigger. We will also study the impact of any latency imposed by the hardware (or analysis time) required to implement these algorithms, in particular on pipeline and buffer lengths. We will consider how the needs of the trigger to access information for detector components affects the design of the data acquisition system and conversely how the needs and constraints on the data acquisition system design will limit the possible trigger algorithms which can be used.

We will generate events using well established SSC Monte Carlo programs such as ISAJET and use this data to drive a simulation of the data acquisition subsystem (see the next section). This simulation will tell us the response of proposed DACQ designs to a realistic environment and will give us confidence that we are able to understand how well a detector can accept interesting events and reject background.

3.3. Simulation of the Data Acquisition Subsystem

We propose to develop a computer model of an SSC data acquisition (DACQ) subsystem which will incorporate the data acquisition, event trigger, data transmission, data collection, and event building components that make up the data acquisition subsystem. This approach extends the simulation and modelling work done at the University of Illinois[2,3] as part of the SSC generic detector R&D program. The generic R&D project has had two goals:

- To demonstrate the feasibility of realistic DACQ modelling and simulation, and
- To develop the tools necessary for the application of modern CAE design methods to high energy physics systems.

The first goal was met by modelling and simulating the CDF muon trigger[4] (see figure 2). The second goal will be met by applying the method to several current projects, one of which is the SSC data acquisition subsystem described in this proposal. By participating in collaborative efforts, we will not only apply the newly developed methods to useful projects, but will also disseminate these tools throughout the particle physics community. Widespread use of CAE design methods will be required for the successful design of the large and complex detectors that will be built for the Supercollider.

The design of large data electronics and computer systems incorporates two major aspects. The first is the ability to simulate the behavior of proposed designs so that flaws can be uncovered and adequate diagnostics built in. The second is to coordinate the work of many people at many places to reduce the problems of poor communication which often leads to incompatible system components and degraded performance.

Realistic simulation of the DACQ subsystem can help to answer questions at all levels of detail. Issues of the overall architecture as well as issues of detailed gate level implementation can be addressed by a good, hierarchical modeling tool,

such as the one that we have at Illinois.[5] With hierarchical modelling, one can describe a system at the highest functional level (*i.e.*, as a set of connected black boxes) or open up the boxes to consider the lower levels of implementation (at the lowest level are the individual electronic gates). A hierarchical model of an SSC data acquisition system is a feasible project. The highest level model of the CDF muon trigger took a masters level graduate student nine months to implement. The structure of the modelling language allows naturally for several projects to be ongoing in parallel - each designer works inside one of the black boxes.

We thus think that the simulation of the DACQ subsystem should consist of at least six parallel projects:

- "Black box" description of the global architecture,
- The event trigger and readout control,
- The data collection front end,
- The parallel event builder,
- The processor farm, and
- The data links.

We do not propose to work on the details of parallel processor farms or data links, but it is necessary to include a high level model of these components in order to simulate subsystem performance. We expect to establish contacts with groups which are working in these areas, in order to achieve a design of a complete subsystem.

The first project will be a high level functional model of the entire subsystem. In addition to testing the feasibility of proposed DACQ architectures, it will provide a design specification that will aid in the coordination of a distributed design effort. As the component designs proceed from a high (functional) level to a low (gate or IC) level, continual design simulation will ensure conformance with the overall design. (The development of tools to aid the coordination of mutilevel, multidesigner work is a part of the Illinois generic R&D project.)

We anticipate that the middle three projects, the modelling and simulation of subsystem components, will eventually take place at the institutions where the detailed design work is being done. However, in order to get to that point, we will need to accomplish two tasks, first to construct at Illinois the high level model, and second to transfer Illinois expertise with the design tools to the appropriate people at each institution. Before actual SSC detector construction begins, a third task must be completed - the creation of an efficient design management network.

We make the following estimate of the effort required (also see the milestones and personnel chapters):

First year:

- Construction of the black box model. (Twelve months, 1 student and 0.5 engineer.) We leave an entire year for this project to allow for the study of several alternative architectures. Funding for a second graduate student would speed it up.
- Transfer of the design methodology to our collaborators (Twelve months, 0.5 engineer, plus travel for collaborators) Travel is an important part of the project overhead, since efficient use of the design tools will require a learning period.

Second year:

- Modelling of the subsystem components. (1 student and 0.5 engineer for each component.) These will obviously proceed in parallel, and their schedules will be closely linked to the understanding of the issues that will confront us.

Second year and later:

- Coordination of the design, and maintenance of consistency. (0.5 engineer) This will be an important ongoing effort at Illinois

It is important to note that the utility of design tools depends on their timely availability, that is while design is taking place. Thus a prompt schedule is important.

3.4. Parallel Event Builder

We propose to develop an event builder based on a dual-port memory architecture. This approach extends the data acquisition architecture developed[6-8] at Brown University for the DØ experiment to the much higher bandwidths needed for the SSC (10 GB/sec versus 320 MB/sec for DØ). A very different event builder architecture[9] based on a high-speed switching network is being supported by the generic detector R&D program. The principal advantage of the switching network appears to be its extremely high bandwidth, possibly as high as 100 GB/sec. While this extremely high bandwidth is needed for a B physics detector, a detector emphasizing high-p_T physics probably does not require 100 GB/sec bandwidth and may choose the dual-port memory approach because of its greater flexibility and fault tolerance. We believe both event builder architectures should be pursued at this time to demonstrate the feasibilty, advantages, and limitations of each approach. We describe below a proposed architecture for a dual-port memory based event builder.

The parallel event builder is responsible for assembling complete event records and transmitting these event records to a selected node in the processor farm. Events coming from the input data stream must be routed onto one of the data paths leading to a group of processing nodes. Figure 3 illustrates this approach, showing data coming from the detector on three "data buses" and being routed onto four "farm buses," each of which leads to three processing nodes. An incoming event will be selected for routing to one of the available farm buses where it will be picked up by one of the available processing nodes. Future events will be routed to other farm buses, thereby dividing the high bandwidth of the incoming data stream into multiple lower bandwidth output data streams. The key element in this architecture is the router/buffer modules shown in Figure 3 wherever the data and farm buses cross-connect. These modules are dual-port memories that pick-off data for a selected event from the data buses and then retransmit the data at the bandwidth of the farm bus. A subsequent event will have a different

farm bus and row of router/buffer modules selected, thereby dividing the incoming data bandwidth among the various buses. This design is easily scalable to higher bandwidth by increasing the number of buses and router/buffer nodes.

One possible scenario for the SSC would incorporate 128 high speed (100 MB/sec) fiber optic data buses connected to 128 fiber optic farm buses, giving a maximum bandwidth of 12.8 GB/sec. Such a system would have 16K router/buffer nodes. By grouping several farm bus outputs to a single data bus receiver, control chip, and memory module, we expect to achieve 64 router/buffer nodes on a router/buffer board. Assuming 16 router/buffer boards per crate, a complete 12.8 GB/sec event builder for the SSC would occupy 16 crates. By the mid-90's we anticipate fiber optic links of 200 MB/sec or more will be available, allowing even higher system bandwidth and/or fewer router/buffer crates.

This architecture has many advantages:

- The architecture is extremely flexible and robust. Any particular event can be routed to any available node. Thus, calibration data can be routed to dedicated processors, and one event can be sent to two or more processors for data verification and integrity checks. New versions of the software filter program can be tested and compared with the current version without interfering with data taking. Processors can easily be brought online or taken offline without disrupting the data flow.

- The architecture is fault tolerant. Any component of the event builder can be taken out of the system with only a small reduction in the available bandwidth. Transmission errors, overrunning buffers, and unavailable processing nodes are examples of errors that can be easily handled with minimal loss of data.

- The architecture is scalable. Higher data bandwidths can be achieved by adding additional data paths, processing nodes, and router/buffer modules. Additional computing capacity can be obtained by increasing the number of processors on a farm bus. Testing of detector components can proceed with

a minimal set of components.

- No data ordering is required for this architecture. This may prove to be extremely useful as different detector elements are trying to send data from different events. A simple tag field on the data can be used to direct data to the correct processor, at the block or word level.
- No synchronization of the data stream is required. Each detector element can send data at its maximum speed, subject to data bus arbitration, with no need to synchronize its output with other detector elements.

We propose to demonstrate the feasibility of this architecture through extensive simulation of the architecture and development of a prototype event builder. Simulation studies will help us make design choices and determine design parameters to optimize the bandwidth and functionality of the architecture. Construction of a working prototype will help us evaluate performance, study packaging options (high density of router/buffer nodes is critical), and greatly increase our understanding of how to implement the architecture.

We plan to purchase prototype router/buffer boards that incorporate a 2×2 array of router/buffer nodes operating on 100 MB/sec data and farm buses. We anticipate working with ZRL, Inc. in the specification, design, and construction of these boards; ZRL, Inc. has extensive experience in the design of dual-port memories and is responsible for the multi-port memory boards used in the DØ data acquisition system. These boards will be used to assemble a small scale demonstration of the SSC data acquisition subsystem, including the necessary 100 MB/sec data links, to test the feasibility and performance of the dual-port memory architecture. If these activities are successful, we anticipate proceeding to the design of high density router/buffer boards using custom application specific integrated circuits to implement the necessary control logic. The eventual feasibility of building a large scale system depends critically on our ability to package large numbers of router/buffer nodes onto a finite number of circuit boards.

3.5. On-line Processor Farm

On-line processing by a "farm" of general purpose computers will form an important high level component of the data acquisition. The I/O capabilities of the farm, the speed and flexibility of the software, and the reliability of the calibration that will be used for event reconstruction are all issues that will interact with the design of the Data acquisition subsystem. We plan to collaborate with groups from Fermilab and Pennsylvania who are developing processor farm technology.

3.6. Data Collection Front End

The data collection system must gather relatively sparse data from a large number of front-end IC's and output this data at the bandwidth of the high-speed data links leading to the event builder. We plan to include the data collection system in our studies of architecture design and simulation, as described above. We plan to coordinate our efforts with the groups at Fermilab and LBL which are working on front end electronics development.

3.7. Data Links

The KEK group is collaborating with Fermilab on a project[10] to develop a high speed fiber optics data link. At the SSC, such links must be capable of data rates of at least 100 MB/sec with low error rate and high reliability. We anticipate fiber optics will be used to transport data from the detector to the event builder due to its excellent noise immunity, low mass, and high bandwidth. The shorter data links that transport data from the event builder to the processing farm and then to the host computing facility may be either copper or fiber.

Rapid advances are being made in commercial communications systems to build these links, including

- High speed GaAs Multiplexer and Demultiplexer chips that take a low clock speed parallel data path and multiplex/demultiplex it onto a very high speed (1 Gb/sec) serial data path,

- GaAs laser diodes to power fiber optics,
- Low-loss single mode fiber optics,
- High speed PIN photodiodes to detect light pulses, and
- High speed transimpedence amplifiers.

The event builder demonstration system will initially incorporate high speed data links using copper conductors. We then plan to upgrade these links to fiber optics to investigate the feasibility, error rate, and reliability of this technology.

4. Milestones

We do not include milestones which depend on the creation of new collaborations with other groups.

FY90:

Complete a "black box" model

Define functional requirements of DACQ components

Specify data and control flow requirements

Transfer Simulation technology to collaborators

Specification of prototype router/buffer boards

Identify R&D needs and expand participation in the project

FY91:

Detailed models of subsystem components

Studies of performance as a function of component parameters

Refinement of DACQ architecture

Construction of prototype router/buffer boards

Integration of fiber optic data links

Later:

Coordinate distributed design efforts

Design of control ASIC for router/buffer boards

5. Personnel

We only list those people for whom we are requesting support.

FY90:

Brown University	0.5 engineer
University of Illinois	1 engineer
	1 graduate student
Pennsylvania	0.5 engineer
Rutherford	0.5 engineer

FY91:

Brown University	1 engineer
Colorado	0.5 engineer
University of Illinois	0.5 engineer
	1 graduate student
Pennsylvania	0.5 engineer
Rutherford	0.5 engineer

6. Budget

Project Budget (Subprojects are detailed on the following pages)

FY90:

Simulation	\$335k
Event builder	\$185k
Trigger	\$20k
General	\$20k
Total:	\$550k

FY91:

Simulation	\$395k
Event builder	\$335k
Trigger	\$20k
General	\$20k
Total:	\$760k

All amounts include benefits and overhead where appropriate.

DACQ Simulation Budget

FY90:

1 Engineer	$100k
1 Graduate Student	$25k
3 Workstations	$60k
Software	$100k
Travel	$50k
Total:	$335k

FY91:

2.5 engineers	$250k
1 Graduate student	$25k
Network File Server	$50k
Software	$50k
Travel	$20k
Total:	$395k

All amounts include benefits and overhead where appropriate.

Parallel Event Builder Budget

FY90:

0.5 Engineer	\$50k
Prototype Router/Buffer Boards	\$100k
Associated Test Equipment	\$25k
Travel	\$10k
Total:	\$185k

FY91:

1 engineers	\$100k
Development of ASIC Control Logic	\$150k
Workstation and Software	\$50k
Fiber optic hardware	\$25k
Travel	\$10k
Total:	\$335k

All amounts include benefits and overhead where appropriate.

Trigger Studies Budget

Much of this effort will be funded by UK HEP programs. We include only related US travel needs here.

FY90:

Travel	\$20k
Total:	\$20k

FY91:

Travel	\$20k
Total:	\$20k

All amounts include benefits and overhead where appropriate.

REFERENCES

1. A.J. Lankford, Edward Barsotti, and Irwin Gaines, Proceedings of the 4th Pisa Meeting on Advanced Detectors: Frontier Detectors for Frontier Physics, La Biodola, Isola d'Elba, Italy (May 21-26, 1989), SLAC-PUB-5042, to be published in Nuclear Instruments and Methods.

2. J.J. Thaler, International Conference on The Impact of Digital Microelectronics and Microprocessors on Particle Physics, p. 32, Trieste (March 1988).

3. J.J. Thaler, "Data Acquisition Modeling and Simulation", International Conference on Computing in High Energy Physics, Oxford (April 1989), to be published.

4. E. Hughes, "A Verilog model of the CDF Muon Trigger', University of Illinois MS Thesis (unpublished).

5. Verilog, by Gateway Design, Lowell, MA.

6. D. Cutts *et al.*, International Conference on The Impact of Digital Microelectronics and Microprocessors on Particle Physics, p. 81, Trieste (March 1988).

7. D. Cutts *et al.*, IEEE Trans. Nucl. Sci. **36**, p. 738 (1989).

8. D. Cutts *et al.*, "Data Acquisition at DØ", International Conference on Computing in High Energy Physics, Oxford (April 1989), to be published.

9. M. Bowden *et al.*, "A High-Throughput Data Acquisition Architecture Based on Serial Interconnects", Fermilab.

10. C. Swoboda, *et al.*, IEEE Trans. Nucl. Sci. **36**, p. 774(1989).

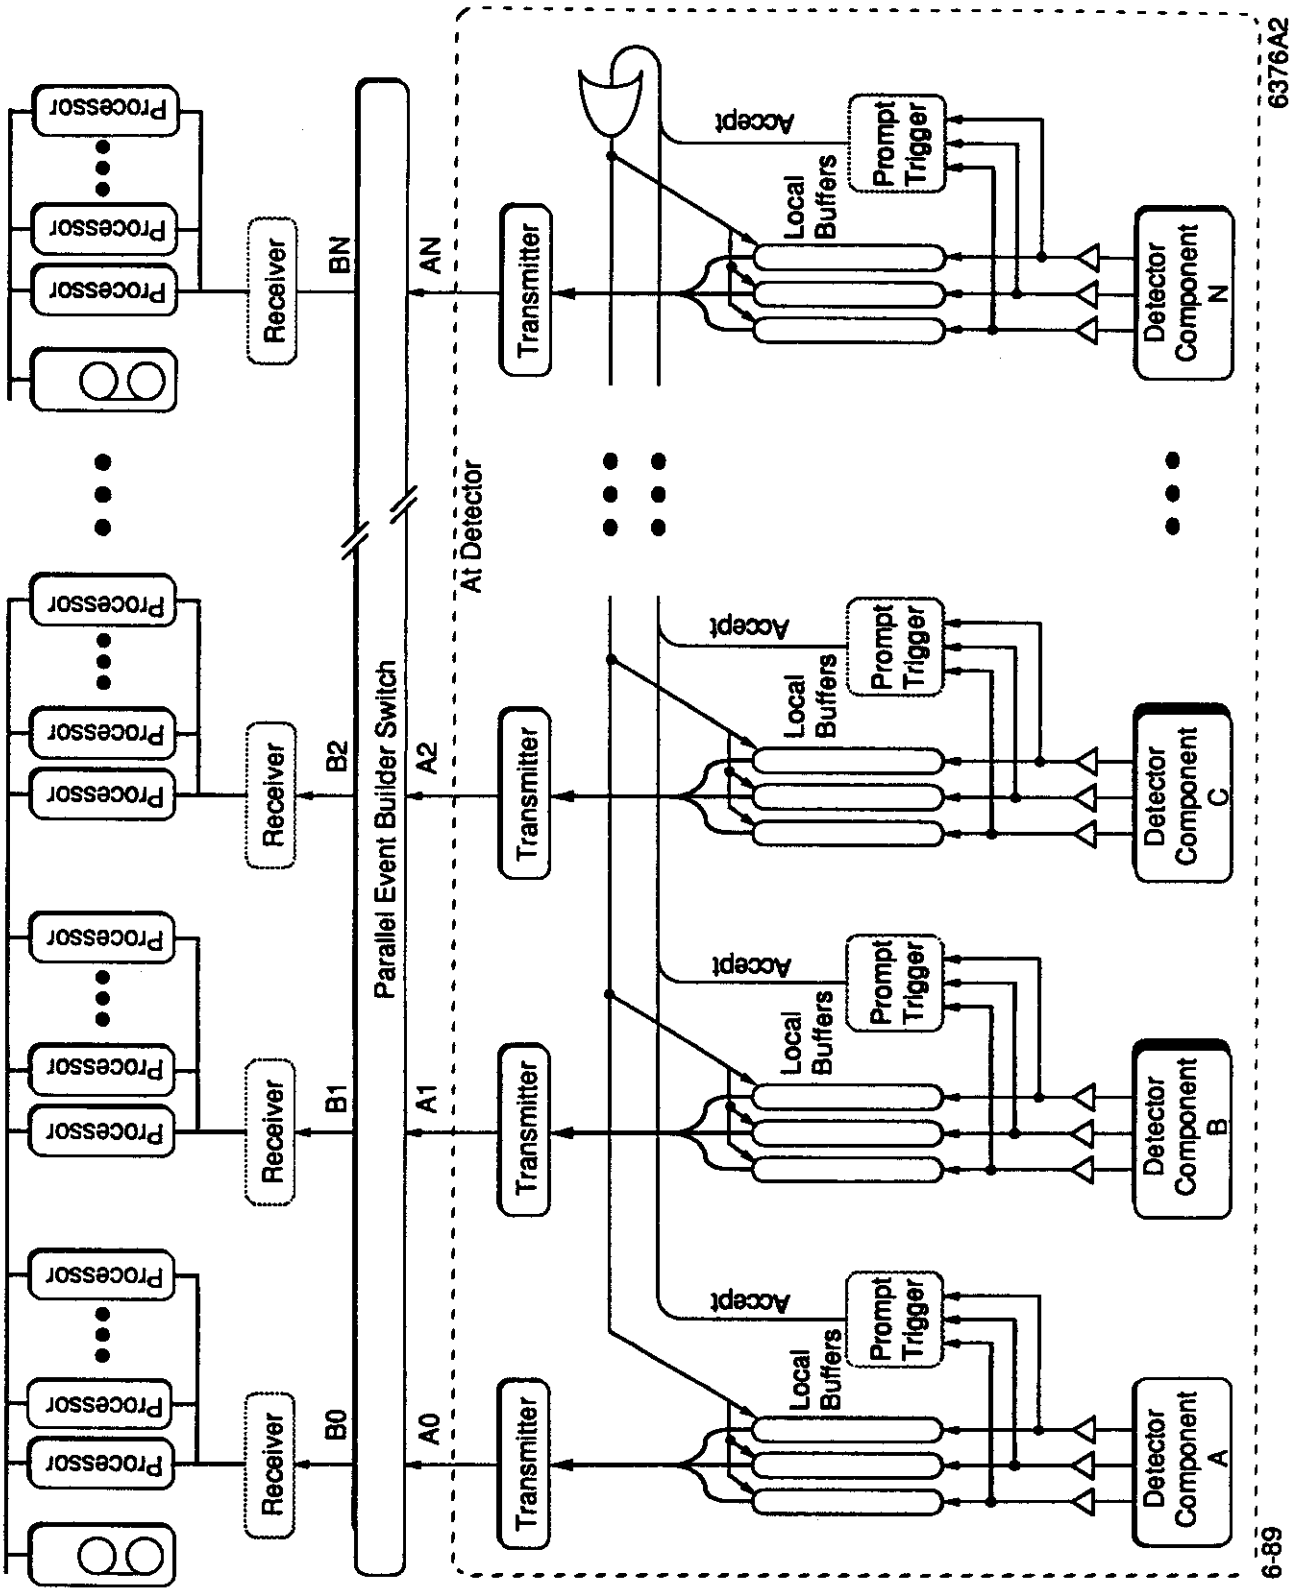

Figure 1. Model SSC Data Acquisition Architecture

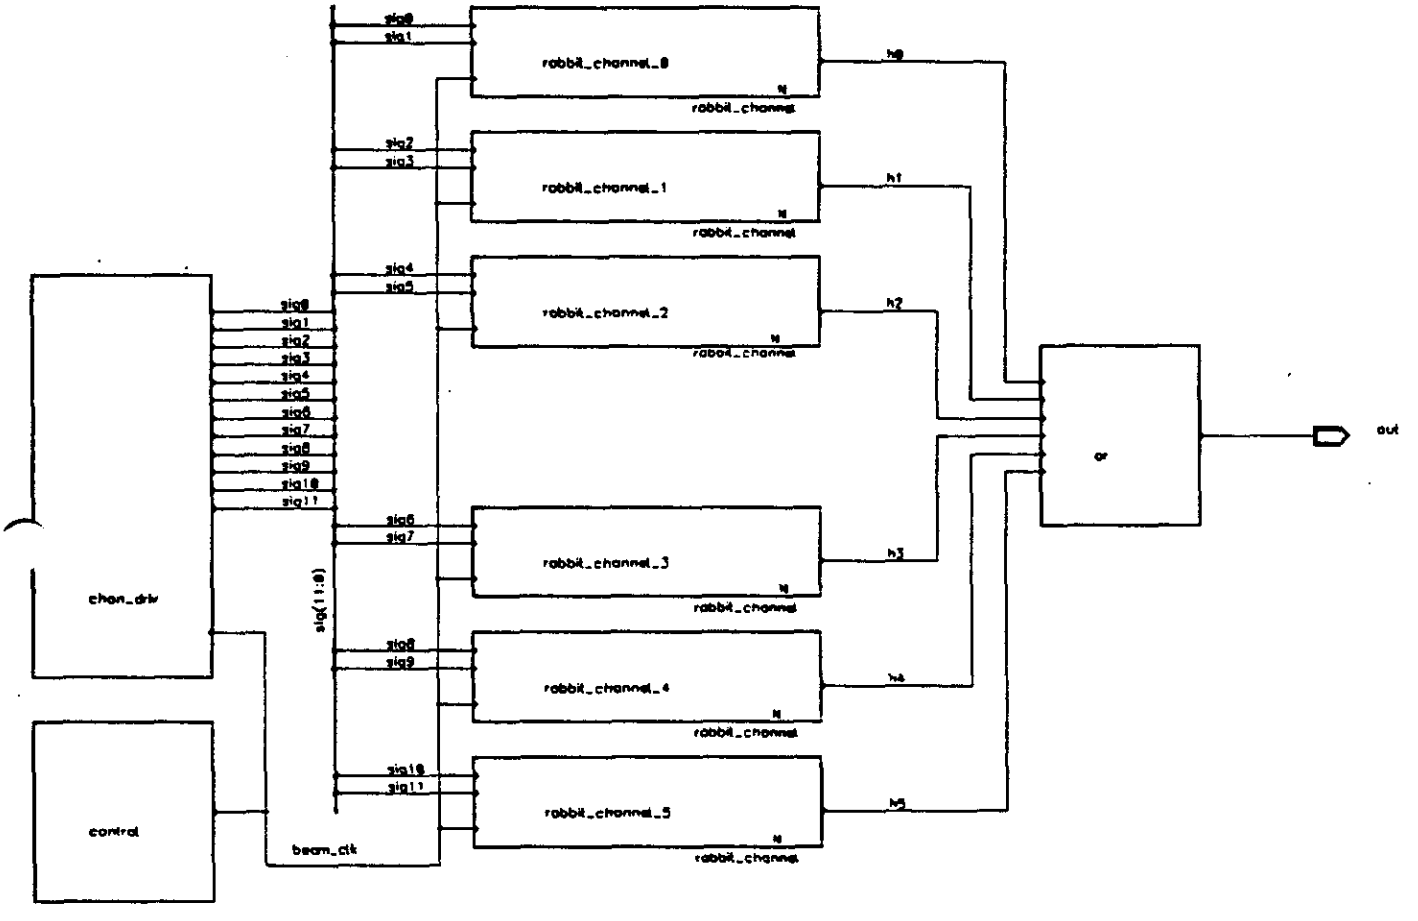

Figure 2. Block Diagram of the CDF Level-1 Muon Trigger

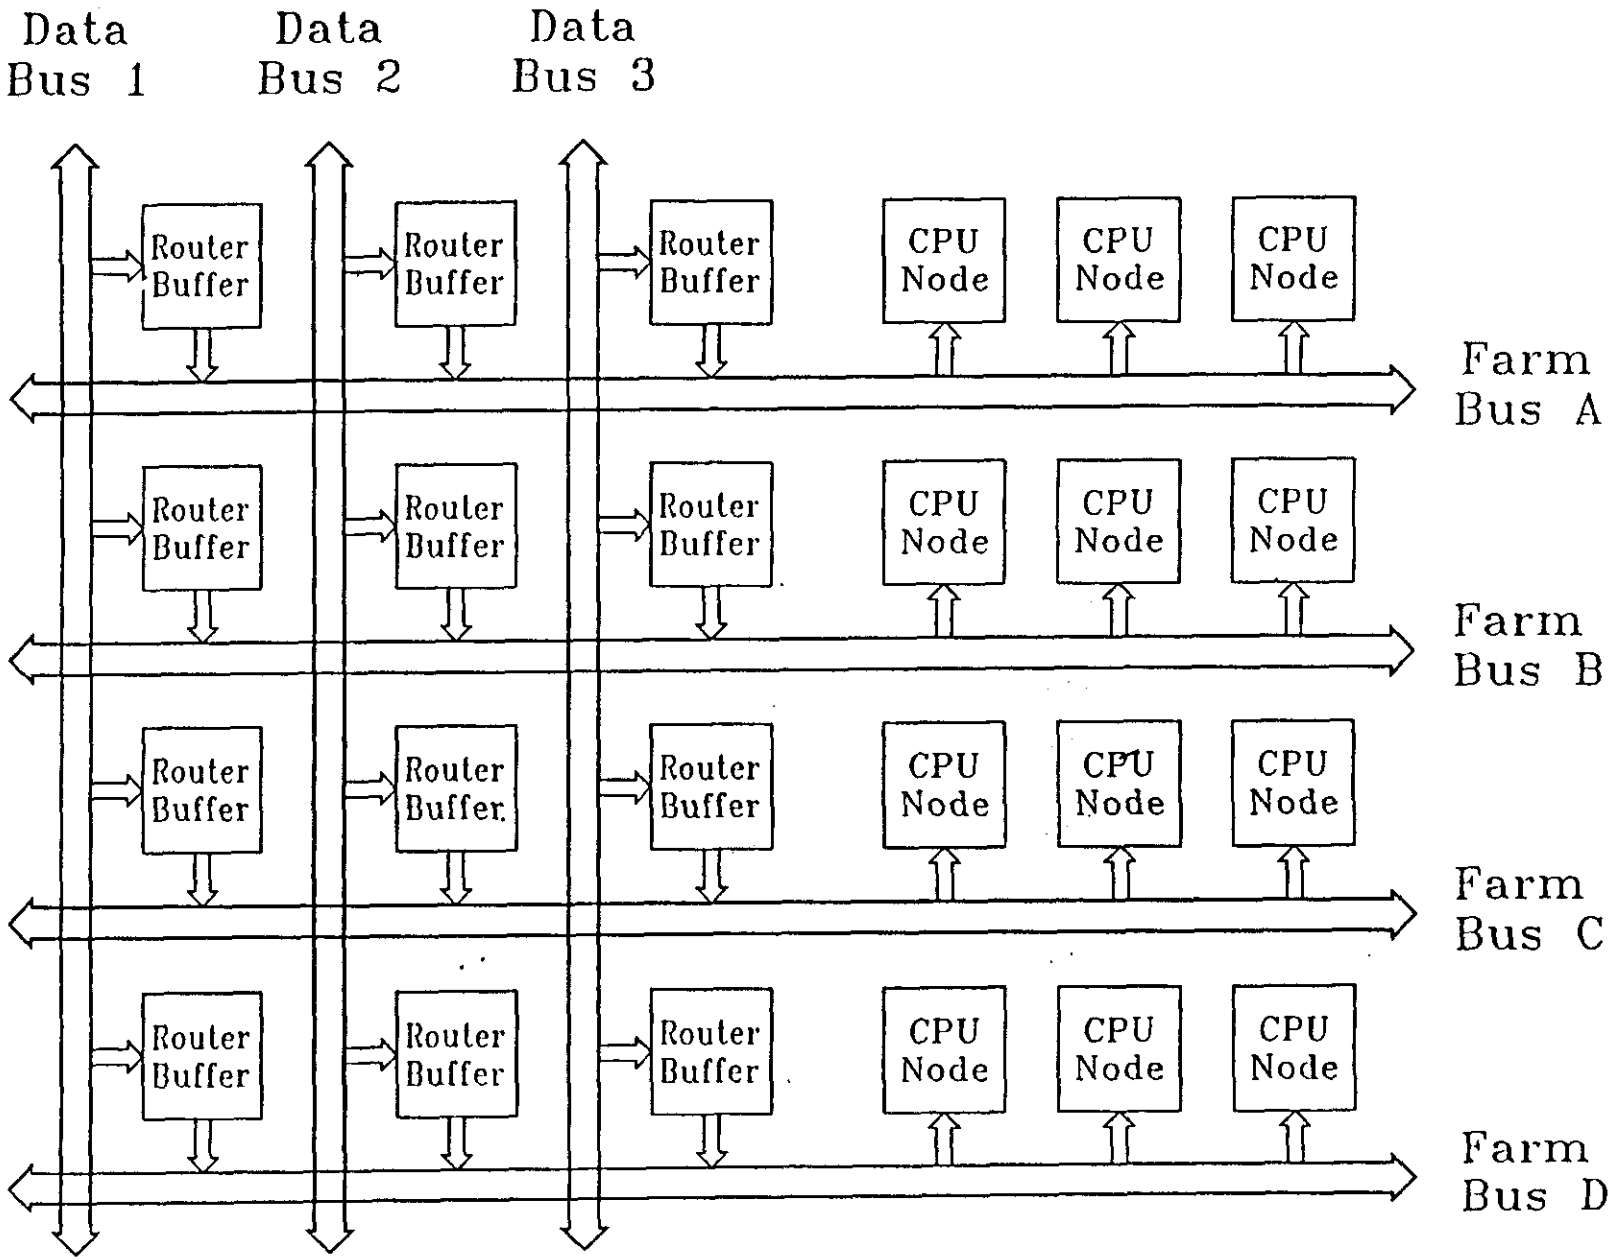

Fig. 3. Event Builder Architecture Based on Dual-Port Memories.