SLAC – PUB – 4611
April 1988
(E/I)

TRIGGERING, FRONT–END ELECTRONICS, AND DATA ACQUISITION FOR HIGH–RATE BEAUTY EXPERIMENTS*

MARVIN JOHNSON
Fermi National Accelerator Laboratory, Batavia, Illinois 60510

A. J. LANKFORD
Stanford Linear Accelerator Center, Stanford University, Stanford, California 94309

Working Group Participants:

S. Amendolia
INFN/Pisa University

J. Appel
Fermi National Accelerator Laboratory

E. Barsotti
Fermi National Accelerator Laboratory

M. Bowden
Fermi National Accelerator Laboratory

S. Conetti
McGill University

R. Crittenden
Indiana University

T. Darries
Fermi National Accelerator Laboratory

A. Denisov
INP, Leningrad

A. Dzierba
Indiana University

S. Errede
University of Illinois, Urbana–Champaign

M. Fortner
Northern Illinois University

L. Gladney
University of Pennsylvania

H. Gonzales
Fermi National Accelerator Laboratory

E. Hartouni
University of Massachusetts, Amherst

S. Kartik
Indiana University

B. Knapp
Columbia University

S. Loken
Lawrence Berkeley Laboratory

R. Morrison
University of California, Santa Barbara

D. Potter
Carnegie–Mellon University

L. Rossi
INFN, Genoa

W. Selove
University of Pennsylvania

M. Sokoloff
University of Cincinnati

C. Swoboda
Fermi National Accelerator Laboratory

R. VanBerg
University of Pennsylvania

D. Wagoner
Prairie View A&M University

V. White
Fermi National Accelerator Laboratory

---

*Work supported in part by the Department of Energy, contracts DE–AC03–76SF00515 and DE–AC02–76CHO3000.

Summary Report of the Working Group on Triggering and Data Acquisition at the Workshop on High Sensitivity Beauty Physics at Fermilab, Batavia, Illinois, November 11–14, 1987

ABSTRACT

The working group explored the feasibility of building a trigger and an electronics and data acquisition system for both collider and fixed target experiments. There appears to be no fundamental technical limitation arising from either the rate or the amount of data for a collider experiment. The fixed target experiments will likely require a much higher rate because of the smaller cross section. Rates up to one event per RF bucket (50 MHz) appear to be feasible. Higher rates depend on the details of the particular experiment and trigger. Several ideas were presented on multiplicity jump and impact parameter triggers for fixed target experiments.

INTRODUCTION

The rare nature of beauty production and the rare nature of its interesting decays will pose challenges for the trigger and data acquisition electronics of a beauty experiment at either the Collider or in the fixed target program. Beauty production is only four parts in ten thousand of the inelastic cross section at the collider and only three parts in ten million of the inelastic cross section in fixed target experiments. The decay modes of interest have branching ratios of 10^{-4} or less. Taking into account the inefficiencies of tagging and reconstruction, a total of as many as 10^{14} to 10^{17} hadronic interactions is necessary to produce sufficient statistics to study rare beauty decays or CP violation in the B-meson system. High interaction rates are necessary to produce this many interactions. Table I summarizes some of the relevant numbers discussed for fixed target and collider experiments.

Table I. Summary of typical rates for high sensitivity beauty experiments.

Type of Experiment	Inelastic Interaction Rate	$\frac{\sigma(b\bar{b})}{\sigma(\text{inelastic})}$	$b\bar{b}$ Production Rate
Fixed Target Open Geometry	10^7 Hz	3×10^{-7}	3 Hz
Fixed Target Closed Geometry	10^9 Hz	3×10^{-7}	300 Hz
Fixed Target Remote Imaging	5×10^{10} Hz	3×10^{-7}	1.5×10^4 Hz
Collider $\mathcal{L} = 10^{30}$ cm^{-2} sec^{-1}	5×10^4 Hz	4×10^{-4}	20 Hz
Collider $\mathcal{L} = 5 \times 10^{31}$ cm^{-2} sec^{-1}	2.5×10^6 Hz	4×10^{-4}	10^3 Hz

The trigger for such an experiment must be extremely effective to reduce these rates to manageable levels. In addition, it must be efficient since the maximum available production rates are limited by either available luminosity or practical detector considerations. The trigger must also select events at high speed to keep pace with the high interaction rate. Since the ideal trigger is completely unbiased, the trigger is limited in the selection criteria which can be used. Interaction rates higher than 10^8 Hz will produce more than one interaction per bucket. Consequently, the trigger will possess a sophisticated multi-level architecture with large amounts of processing power. The data acquisition system will need to provide large amounts of buffering during the trigger process and also be capable of providing large amounts of data to the trigger processors at high rates. The data acquisition system may be required to write data to tape or other mass storage media at very high rates so final event selection can be done off line.

The general problems of the trigger and data acquisition system for these beauty experiments are similar in nature to the problems of triggering and acquiring data at the SSC. At the SSC, interaction rates are 10^8, and many of the processes of interest have rates of a few per second. At the SSC, beauty production is still only about four parts in a thousand. A workshop was held at Fermilab in November 1985 in order to study triggering and data acquisition at the SSC.[1] Many of the solutions and conclusions of that workshop are valid for high-rate beauty experiments. The problems of data acquisition are particularly similar in the two cases. For triggering, however, an important difference between the general-purpose SSC experiment discussed at that workshop and the beauty experiments considered here is that total transverse energy and missing transverse momentum triggers, useful at the SSC and easy to form, will not be powerful for beauty experiments. Consequently, beauty triggers will generally need to be more sophisticated. Indeed, the difficulty of extracting a beauty signal off line will be mirrored in the difficulty of triggering on line.

Although the problems of triggering on beauty events, particularly in an unbiased fashion, and of acquiring data from these events are formidable in either fixed target or collider experiments, sophisticated tools exist for confronting the task. The Working Group on Triggers and Data Acquisition concentrated on considering existing and emerging trigger techniques and electronics as tools in forming the trigger selection. Detailed trigger criteria, and recipes for combining criteria, were not discussed in depth. Some focus was placed on techniques suitable for vertexing and tracking in the trigger since these will be important selection criteria in any experiment. Lepton identification, a similarly important selection criterion, was not considered extensively since it was studied by the particle identification working group.[2] Architectures for trigger and data acquisition systems were also discussed.

This summary report will review the discussions of the working group. The presentation attempts to be generic, not discussing differences between experiments, although important differences between the fixed target and collider cases are noted. Some of the trigger and data acquisition tools are discussed, and examples of applications are presented. A number of excellent individual contributions to these proceedings, which are based on presentations made to the working group, are referenced by this summary rather than duplicate their content here.

OVERVIEW OF TRIGGER SYSTEMS FOR B QUARKS

Typical Triggers

A number of possible triggers can be used for B physics. Table II contains a sample listing. For some of the triggers in the table, crude estimates of the rejections they provide and of their efficiency in tagging beauty events are also given. These triggers select event characteristics which enhance the proportion of beauty events. All but multiplicity jump can be used in both collider and fixed target experiments. The first five triggers each select a different basic event topology or characteristic and would normally be used as separate, parallel triggers rather than in combination. These triggers use event characteristics that can be measured without the silicon tracking system. The last three triggers are directly related to the detection of B decays some distance from the primary vertex. Since vertex aspects of the event are independent of the first five, it is possible to use them in combination with the other triggers. A common feature of all eight triggers is that they require at least some tracking information. The need for tracking in the trigger is in contrast to triggers in some collider experiments which require only calorimetric information. In fact, for beauty physics only very loose overall transverse energy or missing transverse momentum criteria can be applied without seriously decreasing the efficiency for beauty detection. This is because beauty production, particularly at the collider, does not arise from the highest energy interactions of the constituents.

Table II. Sample b tags and rough estimates of rejection and efficiency.

Trigger	Rejection	b Efficiency
A single lepton at high p_t	1000	25%
A number of particles at high p_t		
Di-leptons at high mass	10^5 @ 10^7 Hz 10^4 @ 10^8 Hz	10^{-3}
Di-hadrons at high mass		
Multi-hadrons at high mass		
A track with large impact parameter	25	50%
Multiple vertices		
A multiplicity jump downstream of the target	10	33%

As discussed in a later section, these different triggers require different time scales. For instance, a multiplicity jump trigger can be performed faster than reconstructing multiple vertices, and a single lepton at high p_t can in most cases be selected faster than di-leptons at high mass. These different triggers are also affected differently by multiple interactions per bucket. For instance, multiplicity jump triggers may become confused; whereas, triggers which find rare high p_t leptons may not be as affected. In choosing to use some of these triggers, the bias of the criterion, as well as the efficiency, rejection and speed, must be considered. This consideration was beyond the scope of our working group.

Detector Systems

For the sake of discussion of trigger and data acquisition systems, the typical detector for either a collider experiment or a fixed target experiment is illustrated in Fig. 1. The detector components are similar in the two cases, although the geometries are different. Starting at the beampipe or target there is a very high resolution tracking system, usually with silicon strips. The next layer is conventional tracking with wire chambers. Outside the tracking, or perhaps integrated with it, is the particle identification system. This system uses techniques such as transition radiation detection or Cerenkov ring imaging. Electromagnetic and hadronic calorimetry and muon detectors complete the standard detector. Not every experiment would look like this detector; for example, the wire chambers might be combined with the particle identification system. Although the component subsystems of the detector are similar in the fixed target and collider cases, the channel counts and response times of the components are different. Order-of-magnitude estimates of the number of channels and response times for these two types of high-rate experiments are given in Table III.

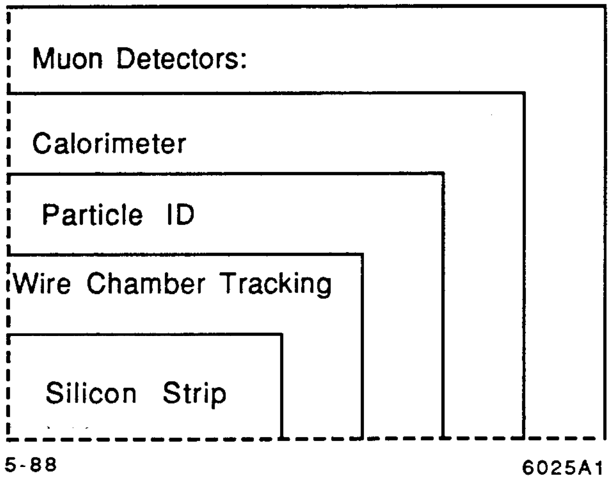

Fig. 1. Generic detector used for data acquisition design.

Table III. Parameters for typical beauty detectors.

System	Collider		Fixed Target	
	Channels	Response	Channels	Response
Silicon	10^5	200 ns	10^4	18 ns
Wire Chambers	10^5	200 ns	10^4	18 ns
Particle ID	10^4	200 ns	10^4	18 ns
Calorimeter	10^5	200 ns	10^5	18 ns
Muon Detectors	10^4	>200 ns	10^3	>18 ns

This typical detector is not novel; however, trigger considerations demand that component response times be reasonably prompt in the collider case and quite fast in the fixed target case. Such response times may challenge detector and front-end electronics designs with respect to collecting charge from the detector in a short time. Particularly in the case of silicon tracking detectors, noise performance of the electronics will limit achievable response times. Prompt response times, along with general segmentation considerations for beauty reconstruction, also lead to large channel counts which will demand large bandwidth in the data acquisition system. These bandwidth considerations are true for the data path to the trigger processors as well as the data path to tape; however, the bandwidth problems are generally no more severe than those considered for the SSC.[3]

Life History of a B Trigger

Data from each detector component becomes available to the trigger at a different time. Likewise, individual trigger criteria built from the detector data are ready at various times, and some of these criteria are dependent on others as increasingly complex criteria are constructed. The time history of a typical trigger decision is illustrated in Fig. 2. The time history stretches from the one nanosecond scale to as long as one second between writing out events. The Fermilab fixed target program has a natural cycle of 18.8 nanoseconds corresponding to the time between RF buckets. The collider has a minimum time between crossings of 200 nanoseconds.

The multiplicity jump trigger can be implemented using differences in pulse height between layers of silicon and can be available very quickly, possibly in less than ten nanoseconds with very fast electronics. Digital information on hit cells in silicon vertex detectors and small-cell drift chambers will be available in tens of nanoseconds. Energy clusters in the calorimeter should be available in less than about 100 nanoseconds for any technology that will work at the high event rates of the B physics program. These trigger criteria can be used by fast Level 1 triggers.

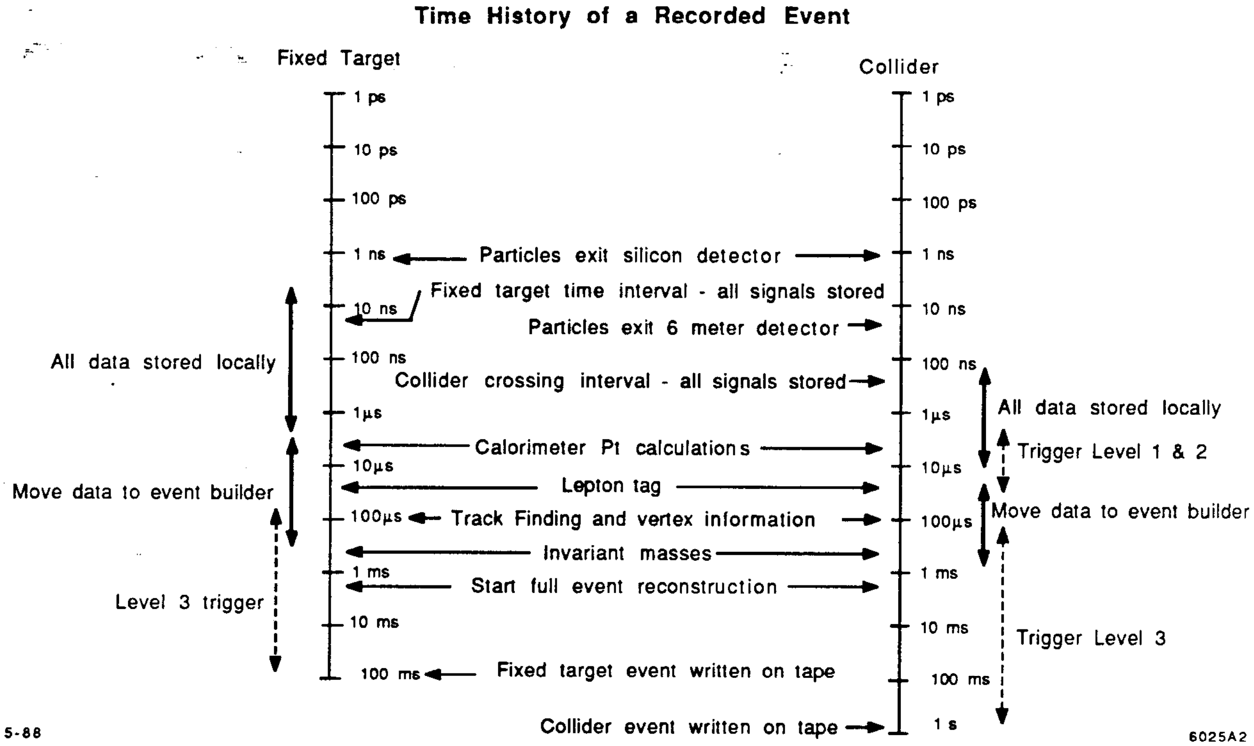

Fig. 2. Schematic diagram of the time structure of an accepted event for a fixed target or collider experiment.

Chamber hits and local track segments and slopes will be available in approximately two microseconds. This task may be performed by electronics on the chambers. The track segments can be ordered in decreasing p_t to speed track finding. Fast track processing should give full tracks within 10 to 20 microseconds. This information can be combined with information from the particle identification or the silicon tracking at that time. The tracks can also be combined to form invariant masses at about 200 microseconds. These tasks will be performed by the Level 2 trigger.

By approximately one millisecond it is possible to start full event reconstruction. The Level 3 trigger can process these events to further reduce the trigger rate to acceptable output rates. For a typical fixed target experiment, events must be written to tape or other mass storage medium at about 100 milliseconds. For the collider, events will be written at about one-half to one per second.

These considerations define the needs for local data storage and event buffering. All data must be stored locally for about five microseconds. The data will then be moved to and held in buffers for the full event reconstruction. We expect that data from all detector components will be in full event buffers by approximately one millisecond.

Single Lepton at High p_t Trigger

As an example of the manner in which trigger criteria are combined in the trigger selection process, consider the case of a trigger on a single electron at high p_t in a collider experiment. The first criterion to be applied is that an energy cluster in the calorimeter exceeds some threshold. The trigger then requires that some track segment in the outer layer of tracking be associated with the energy cluster. These criteria can be applied within the first few microseconds and could provide a rejection factor of about a thousand. The track segment will then be linked with the inner layers of tracking into a track, and its momentum will be determined. This momentum must then match the energy of the cluster. This requirement could provide an additional rejection factor of about five. Finally, the electron identification can be confirmed by associating information from the particle identification system with the track. The particle identification system should provide an additional rejection factor of about ten.

Overall, this trigger provides a rejection factor of 5×10^4 and takes between 50 and 500 microseconds. The first step identifies charged electromagnetic showers while rejecting most hadron showers. The second step, which matches momentum to energy, eliminates overlapping charged and neutral pions. The third step uses the particle identification to reject hadron showers that fake an electromagnetic shower. More about the criteria utilized to identify a high p_t electron are discussed in the report of the Particle Identification Working Group.[2] The trigger can be further enhanced be looking at impact parameter or vertex multiplicity.

TRIGGER LEVELS

The trigger for a beauty experiment would consist of three levels, as is now typical of many experiments. The first level, or Level 1, would use prompt data from the detector. It would employ fast electronics in a manner designed to minimize buffering. Techniques used would include analog summation and thresholds as well as extensive use of memory look-up. The second level, Level 2, would have more processing time available. It could use recursive and serial processing. Less parallel hardware would be needed. It most likely would employ a specialized processor to do rudimentary tracking and vertexing. The third level, Level 3, would have still more processing time available. It would make use of an array of general-purpose microprocessors for convenience and flexibility. It might in some cases do complete event reconstruction.

The division of the trigger system into three levels is arbitrary, and is done here for the sake of discussion; fewer or more levels may be used. Moreover, the allocation of tasks among levels is rather arbitrary. For example, the division of work between a Level 3 and a Level 2 processor is very dependent on the technology used for Level 2. Some experiments might not have a Level 3 trigger and record all the Level 2 triggers using the new, high-speed data recorders. Even in this case, however, the equivalent of Level 3 will exist off line.

First Level Triggers

For fixed target experiments there was considerable discussion of multiplicity jump and impact parameter triggers. The increased number of particles from b decays relative to c decays should allow good multiplicity jump triggers. The basic idea of a multiplicity jump trigger is to use the increase in the number of tracks in an event from a B decay as a trigger. This is done by placing a silicon plane very close to the target (or using the silicon itself as a target), measuring the multiplicity before the B decay and then downstream of the decay. This measurement can be done either by comparing pulse heights from the various planes and triggering on a large increase in pulse height or by counting the number of strips hit in the various planes and looking for an increase. The pulse height trigger can probably be done in 10 nsec or so. W. Selove estimated that a typical B event will have a pair of B mesons which will increase the average charged particle multiplicity by about ten particles in 1 or 2 inches downstream of the primary vertex.

D. Potter described experiments done using pulse height measurements from 14×14 mm^2 silicon wafers without strips. He and N. Reay estimated that a rejection factor as high as 50 with a 30% efficiency might be achieved. In his contribution to the Workshop, W. Selove[4] lists several backgrounds for this method. Among these are 1) the Landau tail on the pulse height spectrum, 2) nuclear interactions in the downstream silicon planes, 3) large angle tracks which miss the downstream planes and 4) amplifier noise. He describes a method of combining the analog method with digital track counting to reduce these backgrounds.

Impact parameter triggers are implemented by measuring tracks and determining if some of the tracks miss the primary vertex by a certain minimum amount (the impact parameter). A closely related trigger looks for more than one vertex. Impact parameter triggers are quite powerful. WA82 found that the minimum bias background was reduced by a factor of 25 while the acceptance for charm remained at 50%. M. Sokoloff,[5] in a paper in these proceedings, describes a study using minimum bias events from E691 and a Monte Carlo which shows that an impact parameter has about a 20 to 1 rejection factor and a 75% acceptance.

The difficulty in an impact parameter trigger is doing the calculations fast enough. During the workshop, S. Amendolia described a method of building a content addressable memory. This technique is described in more detail in the paper by M. Dell'Orso and L. Ristori[6] in these proceedings. He gave an example of a four-plane, 1024 strip per plane silicon detector. The content addressable memory is composed of a matrix of 10-bit memory cells each connected to a vertical bus and a horizontal bus. The vertical height of the matrix is at least equal to the number of strips in the detector (1024 is the minimum but it is likely to be much more), and the horizontal width is equal to the number of planes of detectors (four in this case). This arrangement is shown in Fig. 7 of Ref. 6. A hit from each plane is placed on the four vertical busses. If a match is found with one of the memory cells, it raises a line indicating a match. If all four cells on a horizontal bus raise their lines, then a track has been found. S. Amendolia estimated that the four-plane system described above would require 10^7 patterns and would fit into five FASTBUS crates. This system can be extended slightly to give the vertex position in a thin target. The presence of two different vertex positions would then indicate a decay and an impact parameter can be extracted. This system is, in principle, capable of making a decision in one clock cycle. Because of cost, it may be necessary to subdivide the system in a way that reduces the chip count but takes several clock cycles. With clock times comparable to the 18.8 nsec bucket spacing, this system could not function alone at a 53 MHz rate.

An interesting possibility is using an analog multiplicity jump trigger to reduce the trigger rate from one every 18.8 nsec to one every 200 nsec or so. This would give enough time to apply a content addressable memory method to estimate the impact parameter. This combination might reduce the event rate by a factor of 50 to 100 and still have good acceptance.

Because of the different geometry of collider experiments, information from silicon tracking is more difficult to use in first level triggers at the Collider. On the other hand, lepton triggers, as discussed in the section *Single Lepton at High p_t Trigger,* are possible in the Collider. In addition, some loose cuts on the sum of p_t or identification of a single track with p_t above about 1 GeV may provide sufficient rate reduction at Level 1 for a more sophisticated Level 2 trigger.[7]

Most people agreed that the Level 1 processor is some type of special-purpose electronics that would use the raw data directly from the detector. Since this trigger is likely to take longer than the time between buckets or beam crossings, it must be pipelined and the data stored until the trigger decision is complete. The trigger decision must be made as soon as possible to minimize the amount of front-end buffer storage. A fixed target experiment sees interactions at about ten times the collider rate; however, the compactness and the favorable geometry of a fixed target experiment, as well as the possibility of a multiplicity jump trigger, allow a faster Level 1 decision. For instance, the multiplicity jump trigger might reduce the fixed target trigger rate to about the same as the Collider crossing rate. Consequently, the amount of front-end storage per channel may be similar for the two detectors.

Second Level Triggers

Second level triggers will play a critical role in beauty experiments due to dependence on tracking and on vertexing for beauty triggers. The rejection available at the first level without tracking is insufficient, and software processing at the third level cannot handle the rates. Typically, a second level trigger uses specialized hardware processors to perform, in a fast and usually approximate fashion, tasks that are normally performed as part of event reconstruction and selection off line. For example, second level triggers frequently reconstruct charged tracks and compute invariant masses. R. Crittendon and A. Dzierba described at the workshop such a trigger used by E672. Second level triggers can also associate particle identification information with tracks. They often have the flexibility to combine information from detector components with different geometries. They may also refine decisions performed by the first level trigger, for instance, by improving the definition of track segments and of energy clusters. Since the input rate to the second level processors has been reduced by the first level trigger, the second level trigger processors can utilize more sophisticated hardware tools than first level processors. The input rate, however, still does not give second level processors sufficient time to use fully programmable microprocessors.

An example of a second level trigger is the CDF fast track-finding processor, described at the workshop by L. Gladney. This processor is typical in that it uses a combination of serial and parallel processing to achieve adequate processing times without excessive amounts of electronics. The segment finding step is performed by sequentially comparing mask patterns to patterns of chamber hits within a cell for all cells in parallel. This technique finds track segments in fixed time, independent of the complexity of the event. The time required depends on the number of masks, which also determines the minimum momentum of the search. In CDF the search is limited to momentum above 3 GeV, and takes about five microseconds. The segment finding is of a nature that could be integrated into the front-end electronics of a cell and placed directly on the chamber. Linking of segments into tracks involves bringing together the segments from all cells and requires from 300 nsec to 100 microseconds depending on the complexity of the event. This processor is discussed in more detail in a separate contribution to these proceedings.[8] That discussion also describes some of the concerns and limitations in constructing a track-finding processor.

Another example of a second level trigger processor discussed by the working group is the Data Driven Pipelined Processor used by E690 at Fermilab. This technique, pioneered by Knapp, Sippach and collaborators, was presented by E. Hartouni at this workshop.[9] This processor aggressively uses familiar techniques to achieve a very high trigger throughput. The E690 processor fully reconstructs tracks in four bending views at a rate of 10^5 per second. It was recently used to reconstruct 10^9 events off line in a six-week period.

The pipelined structure of the E690 processor increases throughput by allocating sequential elements in the processing decision to separate processors. Each processor can then be highly specialized for its particular element in the process. The pipeline utilizes parallel processing in order to avoid bottlenecks. If an element in the process requires more time and cannot be decomposed into pipelined subelements, parallel processors are used to shorten the time required. The structure is data-driven to efficiently make use of the processors in the pipeline. This means that processors are followed and preceded in the pipeline by buffers. The result of one step in the pipeline is put into a buffer where the next processor in the pipeline can take it out as soon as the processor is ready for more data. This intermediate buffering smooths irregularities in data rates and helps keep processors fully employed. The extremely high performance of the E690 processor results from the thoughtful use of these techniques to eliminate bottlenecks and efficiently use processors. Additional throughput is available using further parallelism.

The E690 processor characterizes an approach to the problem of constructing a hardware machine for a second level trigger. The structure is flexible and to a large extent programmable; however, the large throughput is a product of understanding the tracking problem completely before the running of the experiment so the processor can be configured in the most efficient way.

Technologies have improved since the development of the processors described above. For example, larger and faster memory chips and the content addressable memories already discussed will provide second level processors with more power. In addition, field programmable logic and semi-custom gate arrays will become standard tools in these processors. Custom chips using predefined arithmetic cells are now used in high-energy physics.[10] These can be used as building blocks in the next generation of specialized processors of the E690 variety. Digital signal processors (DSP's) will also be incorporated into processing elements, as they are currently for DELPHI.[11] These very fast processors can also be built into building blocks in an E690-style processor. Transputers will also be interesting tools. Basically, high speed computational tools are available in a size and a price which allows them to be used in a massively parallel approach to the processing problem.

In fact, new high-performance digital signal processor chips expected within a year may offer the possibility of a new approach to second level triggers. This possibility was discussed by E. Barsotti[12] at the workshop. He described an array of DSP-based processors capable of performing general-purpose computing. These processors could make trigger decisions on separate events in parallel, as Level 3 processor farms do. For example, 250 processors in 63 FASTBUS modules might handle rates of 5×10^5 into Level 2 by performing 5000 to 10,000 floating point operations each in 500 microseconds per event. Data transfer and buffering aspects of such a system are formidable; however, they are manageable, as discussed in Ref. 12.

A spectrum of solutions to the problems of second level triggering are possible, ranging from specialized circuits to general-purpose processors. The best solution will probably be either processor chips built into a system with a specialized architecture, or specialized circuits added as "hardware subroutines" to general-purpose microprocessors.

Third Level Triggers

Third level triggers are typically characterized by the use of general-purpose processors and by the fact that the data from the entire detector is available. The use of an array (or "farm") of fully parallel general-purpose microprocessors provides a flexible environment in which trigger cuts are easy to implement, including sophisticated event topology cuts. This environment is capable of evolving as experience accumulates and as the experiment accepts increasing data rates. In addition, information generated by the trigger process, such as track parameters, could be readily used to reduce the quantity of data per event.

Discussions of third level triggers for future experiments are usually vague, as were such discussions at this workshop. Nonetheless, some general observations can be made. Such discussions are vague because third level triggers are, in principle, powerful and fully flexible in that they are composed of processors programmable in high-level languages. Harnessing this power will be crucial to high-rate beauty experiments since the overall trigger rejection and efficiency demand very sophisticated event selection. On the other hand, the total CPU power available on a "farm" of third level processors is limited by practical numbers of processors and by data throughput available into the farm. To make efficient use of third level processors in available time, trigger strategies will have to be carefully refined. In many cases, nearly the full off line reconstruction and event selection will need to be performed on line. Moreover, in experiments with high rates into the third level, the third-level decision process will also need to be tiered, to reject events as quickly as possible. In practice, rejections obtained by third-level processing to date have not been high. For instance, J. Appel reported that in one Fermilab fixed target experiment a trigger algorithm which executed in 30 to 40 milliseconds on an ACP processor[13] while performing vertexing using three layers of silicon strips provided only a factor two enrichment in the charm sample. On the other hand, if the physics can be pulled from the data off line in finite time, then the off-line processing power can, in principle, be put on line to make the event selection there. The challenge is to craft a set of efficient and effective selection criteria in advance of data-taking, so that they can be used on line before data is written to tape.

In constructing a third level trigger, some of the strategies used in developing first and second level triggers can be employed. For instance, pipelining, buffering and parallel processing are applicable in a third level trigger which itself is multi-tiered. Moreover, specialized fast hardware co-processors can be used for certain CPU-intensive, specific tasks. In fact, the data-driven pipelined processor approach characterized by E690 is not restricted to the second level trigger. Software and general-purpose processors can be a logical part of this approach. On the other hand, as the power of commercial microprocessors evolves, efficient use of processing power may become less of an issue.

DATA ACQUISITION

Overview

A general model of data flow through the data acquisition system is shown in Fig. 3. This model describes the data acquisition system of either a fixed target or a collider experiment. The data is buffered in levels as the multi-level trigger decision is made. After each level of the trigger decision, the data is transferred from one buffer level to the next. The data rate is reduced at each level because only accepted event candidates are transferred. The data can be further reduced by processing, such as zero suppression and local cluster definition, or even local track segment finding. Such processing occurs in parallel with the trigger decisions and may include some of the processing used in the trigger.

As the data is reduced, it is channeled onto fewer and fewer data paths. Ultimately, the data converges onto less than about ten high-speed digital links which carry the data from portions of the detector into Level 3 buffers where the full event is assembled. The Level 3 trigger processors, as they perform event analysis for sophisticated trigger decisions, can also perform further data reduction. Finally, the data is written to a mass storage medium. Potential bottlenecks arise within the data acquisition process, such as into Level 3, where the practical number of parallel busses is limited, and such as onto mass storage. These points may ultimately determine the degree by which the trigger must reduce the interaction rate; however, the demands on the trigger are eased if adequate data reduction can be performed prior to the bottlenecks.

The data rates posed by beauty experiments either at the Collider or in the fixed target program are less than, but similar to, the rates expected by a large SSC detector. A fixed target experiment will have about one-half as much data at the front end as a Collider experiment, but will have a comparable amount of data in reduced form. The rate of event writing to mass storage will probably be higher in a fixed target experiment than at the Collider.

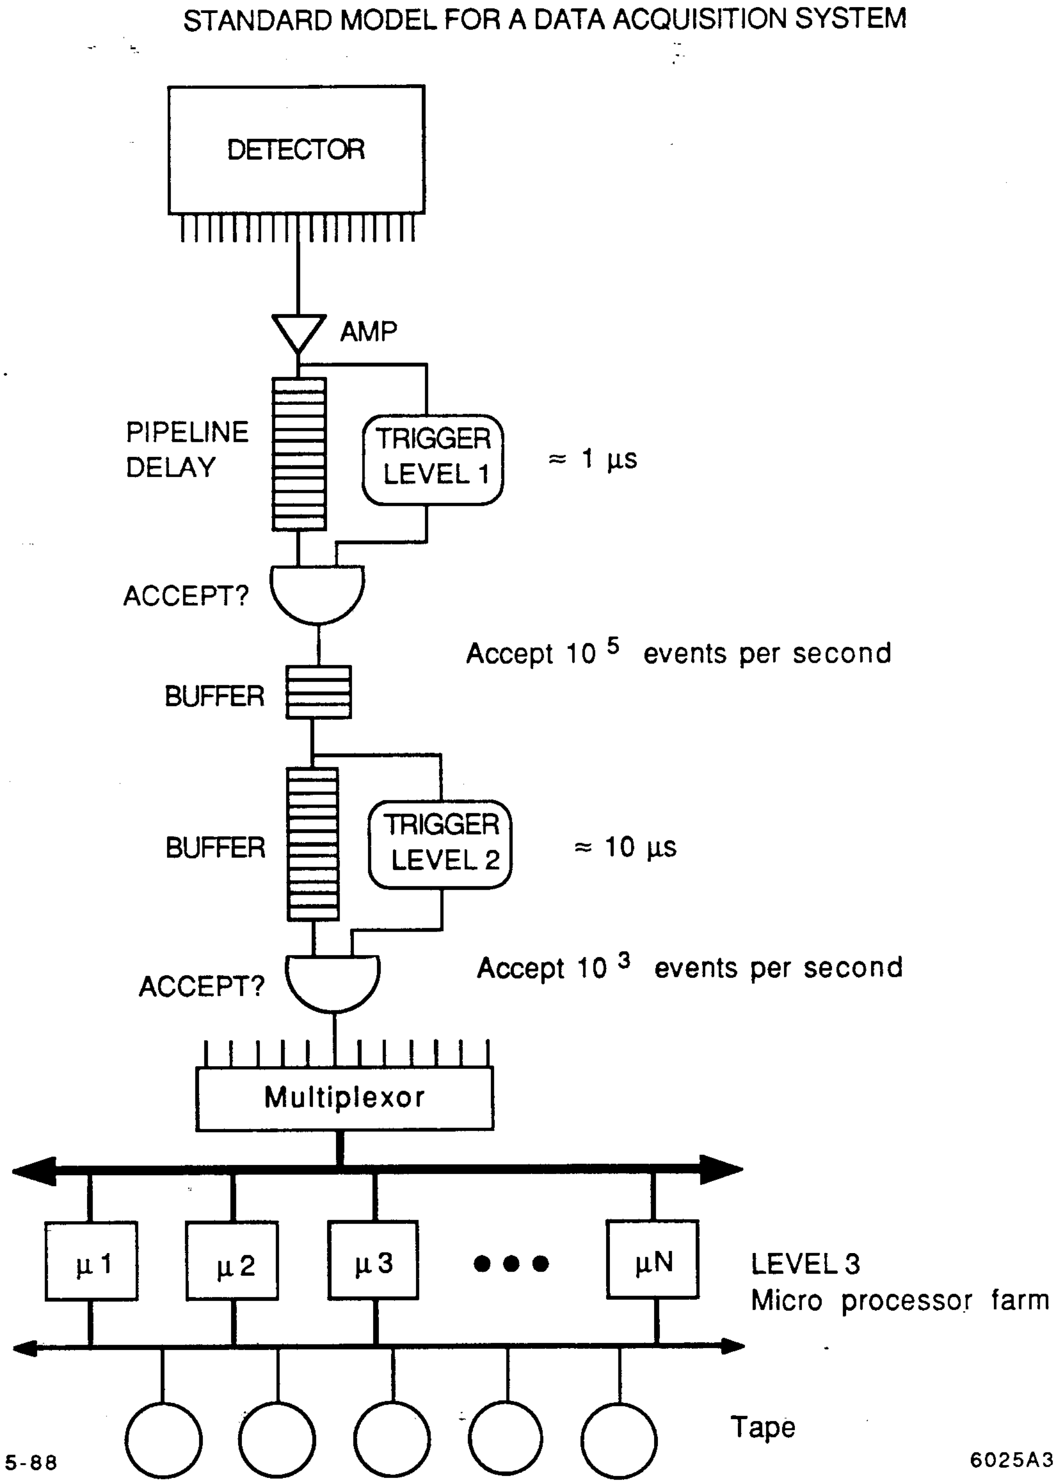

Fig. 3. Schematic diagram of a trigger and data acquisition system for either a collider or a fixed target experiment.

Data acquisition systems can be constructed to handle these rates. An example of such a system was described at the workshop by Barsotti, Bowden, Gonzalez and Swoboda.[12] This system is capable of very high performance. It includes front-end electronics composed of two types of custom VLSI circuits. The first is "microplex-like" and performs scanning, digitization and buffering for 64 channels. The second performs data compaction and is used to form a pipeline of data compaction which channels the data into progressively fewer data paths. High-rate fiber-optic data transmission transfers the data from the front end to Level 2 and Level 3 arrays of processors. The system provides architectures for moving data from a number of parallel but different data paths into another number of parallel processors. See Ref. 12 for a more thorough discussion of this system.

Front-End Electronics

The open geometry of fixed target experiments makes the front-end electronics system much easier to design. There is typically lots of space transverse to the beam direction for mounting amplifiers, cables, and so on. Also, the small beam size and the strong forward opening angles make small planar detectors very practical; there is no need to try to resemble a sphere. On the other hand, the demand for very high rates is a severe constraint. A maximum event rate for fixed target experiments is very dependent on the type of trigger. Once one gets more than one event per RF bucket, however, the problems of triggering increase substantially because of accidentals.

We assumed that the high event rates from either a fixed target or a collider experiment will require a pipelined trigger which implies that all the data must be stored in the front end for several events. Because of the large number of channels and, at least in the collider case, small space, some form of custom chip will be required. In a contribution to this workshop, R. VanBerg[14] discusses the motivation for custom front-end chips and some considerations in the design of the chips. He also sketches block diagrams of chips for readout of silicon strips in fixed target and in collider experiments as illustrative examples of front-end electronics for any detector component. There was a general consensus that it is now quite feasible to design custom analog chips that will meet the bandwidth and density requirements for either fixed target or collider experiments. These front end chips are expected to have some form of capacitor storage (microplex chip). This could either be pipelined storage with one capacitor per clock step or some sort of data driven system where a signal above a threshold is tagged with the appropriate time stamp. There was also some difference of opinion on when to digitize the data. One view is to keep data in analog form as long as possible and to make as many trigger decisions as possible before digitizing the data. The argument for this view is that the circuits are simpler and thus lower cost both in money and power. This point of view is discussed in R. VanBerg's paper.[14] The other view is to convert to digital signals as early as feasible and do most of the buffering and processing in the digital world. The basic arguments for this is that the driving force in the industrial world is digital processing and so there will be a large choice of commercially available, high-speed parts for data transmission and signal processing. E. Barsotti et al.[12] describe a large digital system for a colliding beam detector which leans in this direction. They suggest a microplex-like chip with a built in flash ADC for the front-end data storage which would be followed by a data compaction (zero suppression) chip. These two chips would form a pipeline system capable of processing data at a rate of a few hundred nanoseconds per pipelined stage. At the end of the pipeline, the data would be locally stored in memory. After the Level 2 trigger, the data would be transmitted over wide-band fiber optic cables to a microprocessor farm. The use of fiber optics substantially reduces the space needed for cables and also gives good noise immunity.

Conclusions

The general conclusion of the Working Group was that a system for triggering and data acquisition is feasible either for a fixed target experiment at 50 MHz or at a collider with a few hundred nanosecond crossing interval. In either case, such a system could be based upon existing approaches and is similar to models developed for the SSC. The system would of necessity be aggressive in its application of existing and emerging techniques. For instance, some number of custom VLSI circuits may need to be built, but they do not seem to be beyond current technology. New electronics technologies can be used to increase the power of familiar triggering techniques. No estimate of the quantity of either development or circuitry was made; however, these most likely match reasonably with the scope of an ambitious beauty experiment.

A fixed target experiment or a collider experiment each offers its own advantages. A fixed target experiment offers a geometry which makes tracking, and particularly vertexing, easier at the trigger level and which also makes a multiplicity jump trigger feasible. The geometry also offers practical advantages for packaging of front-end electronics. More than one interaction per RF bucket in a fixed target experiment excludes certain triggers and compounds rejection problems, and more than a few interactions per bucket may be feasible for only a few specific triggers. In the Collider, interaction rates are at least one, and perhaps three, orders of magnitude lower, and rejection requirements are about three orders of magnitude less. In some sense the trigger seems easier in the fixed target case; however, probably not enough easier to produce the rejection needed by a fixed target experiment. In either case, future studies must identify the physics criteria on which to select events in order that specific triggers on these criteria, and complete trigger systems, can be studied. The ability to distinguish beauty events from background in sufficient numbers is of course a prerequisite to the study of beauty physics. It was the sentiment or spirit of the Working Group that, if this distinction can be made, hardware can be developed to make the distinction at the trigger level.

Acknowledgements

The authors would like to acknowledge the contributions of all the participants of our working group to the material presented in this summary report. We particularly thank those who gave presentations to the working group and those who also made written contributions to the workshop proceedings. These contributions greatly facilitated the writing of this summary report. We are also indebted to Stewart Loken for preparing an early draft of a substantial portion of this report. Finally, the authors and the working group participants would like to thank the organizers for a fine and productive workshop.

References

1. *Proceedings of the Workshop on Triggering, Data Acquisition and Offline Computing for High-Energy/High Luminosity Hadron–Hadron Colliders,* November 11–14, 1985 (Batavia, Illinois, 1985). See also, T. J. Devlin, A. Lankford and H. H. Williams, Electronics, Triggering, and Data Acquisition for the SSC, *Proceedings of the 1986 Summer Study on the Physics of the Superconducting Supercollider,* June 23–July 11, 1986, edited by Rene Donaldson and Jay Marx, Snowmass, Colorado (1986), p. 439.
2. See the summary report of the Particle Identification Working Group at this Workshop.
3. A. J. Lankford and G. P. Dubois, *Overview of Data Filtering Acquisition for a 4π Detector at the SSC,* Trigger–85, p. 185.
4. W. Selove, *A High Rate Trigger System for E771,* contributed to this Workshop.
5. Michael D. Sokoloff, *A Vertex Trigger for Fixed Target Beauty,* contributed to this Workshop.
6. Mauro Dell'Orso and Luciano Ristori, "The Pattern Matching Machine", contributed to this Workshop.
7. N. S. Lockyer et al., *Proposal for a Bottom Collider Detector,* Fermilab proposal.
8. L. D. Gladney, "Considerations for Track Processors in High-Rate Collider Experiments," contributed to this Workshop.
9. Edward P. Hartouni, *Data Acquisition and Triggers, The FNAL E690 Approach,* contributed to this Workshop.
10. S. MacKenzie et al., The Digital Correction Unit: A Data Correction/Compactification Chip, IEEE Trans. Nucl. Sci. **NS–34** (1987), p. 250
11. D. Crosetto et al., *Parallel Arrays of Digital Signal Processors as Central Decision Elements for Upper Level Triggers in High Energy Physics Experiments,* presented at the IEEE Nuclear Science Symposium, San Francisco, CA., October 21–23, 1987.
12. E. Barsotti et al., *B Physics...Digital Triggers and Data Acquisition Using New Microplex and Data Compaction IC's, Many Parallel Math Processors, and Fiber Optics to Fast Digital Trigger & Higher Level Farms,* contributed to this Workshop.
13. T. Nash et al., *The ACP Multiprocessor System at Fermilab,* Proceedings of the XXIII International Conference on High Energy Physics, July 16–23, 1986, edited by Stewart C. Loken, Berkeley, CA (1987), p. 1459.
14. Richard Van Berg, *Front End Architectures for New Beauty Detectors,* contributed to this Workshop.