

# Data Acquisition and Trigger Systems for Experimental Hall D

A Jefferson Lab Mini-WorkFest

November 8-9, 1999

## List of Attendees

David Abbott	Randy MacLeod
Paul Smith	Vardan Gyurjyan
Carl Timmer	Graham Heyes
Dieter Cords	Stephen Wood
Richard Sumner	Chris Cuevas
Shudi Gu	Mark Taylor
Jeff Wilson	Chip Watson
Stepan Stepanyan	Fabrizio Catarsi
Ed Jastrzemski	Fernando Barbosa
Elliott Wolin	Bernhard Mecking
Elton Smith	John Domingo
Andreas Rubin	Bob Michaels
Andi Klein	Franz Klein

## 1 Introduction - Specifications That Drive Design

A detailed investigation of QCD and confinement through the high-statistics photoproduction of mesons at Jefferson Lab has been proposed. This long term experimental program would require the construction of a new experimental hall and a hermetic  $4\pi$  detector capable of measuring energies and momenta of neutral and charged particles with very good resolution [1]. It is the goal of this workfest to identify the hardware and software requirements for the implementation of an efficient trigger and data acquisition system for this detector.

The Hall D detector is expected to run with tagged incident photon fluxes on target from  $10^7/s$  to  $10^8/s$ . With an incident electron beam energy of  $12 GeV$  ( $3 \mu\text{amps}$ ) and a  $10^{-4}$  radiator, The total  $\gamma p$  cross sections lead to a hadronic rate into the detector (at the high flux) of around  $360 kHz$ . By tagging the incident photon flux around the coherent bremsstrahlung peak (at  $9 GeV$ ) the rate of interesting events is reduced to about  $14 kHz$ .

The primary difficulty for triggering that must be addressed is that for the high intensity experiments a coincidence with the photon tagger becomes ineffective at discriminating the events of interest; Hence, the incident photon energy must be reconstructed from the event in the detector. Ultimately one would need a full reconstruction of the event to identify the photon of interest. We propose to implement a two stage trigger, one in hardware and one in software.

A Level 1 hardware trigger can be implemented by using all prompt information in the detector including track counts and energy sums. These can be gathered parasitically from the front-end data stream. If the data at the front-end can be pipelined one would be able to implement a deadtimeless readout scheme upon a positive Level 1 decision (within  $1-2 \mu\text{sec}$ ). In addition, by clocking changes in the detector on the order of its response ( $250 MHz$ ), one can effectively pipeline the Level 1 trigger itself. The goal of the Level 1 is to get a 3:1 reduction in the total hadronic rate by cutting on those events generated from very low energy photons (less than  $2 GeV$ ). This would result in a input rate into the Level 2 of around  $160 KHz$ .

Level 2 must be able to make a more accurate cut in incident photon energy - on the order of a 10:1 reduction. To make a serious cut, Level 2 must do a good job of reconstructing tracks through multiple detector elements. This stage is compute intensive and must come after "event building". To handle the  $100 kHz$  input rate, events will have to be grouped in blocks (e.g. 64-256). Built event blocks can then be passed to the Level 3 processor farm. Use of an online farm allows both flexibility and scalability for the Hall D DAQ system.

Running experiments at the higher luminosities has implications for the detector itself. First, for triggering and reconstruction without the tagger, the detector must be hermetic. Position and energy resolutions for tracking chambers and calorimeters are not too restrictive ( $200 \mu\text{m}$ , 5-10%), but the segmentation and response should be good to handle the high rates.

## 2 Data Acquisition at Jefferson Lab - Current Status

The Jefferson Lab Common Online Data Acquisition System (CODA) is in its second generation of release and is currently begin used in all three existing experimental halls. CODA was designed

as a modular, extensible software toolkit from which data acquisition systems of varied complexity could be built. However, because of its software oriented design, CODA's performance is largely defined by the hardware and operating systems upon which it runs.

The implementation of the CODA data acquisition system in each experimental hall will be discussed with an emphasis on current system limitations both in hardware and software.

## 2.1 Hall A

Hall A currently houses two high resolution spectrometers. One is used as an electron arm, the other a proton arm. A majority of the experiments run in Hall A are (e,e'p) coincidence. Three FASTBUS crates and 1 VME crate interface the spectrometer detectors and beam line. Typical running conditions do not exceed 2 *kHz* with 20% deadtimes. Readout latency for the complete system is about 300  $\mu$ s, but front-end module conversion times are only on the order of 50  $\mu$ s or less. However, because the front-end hardware in the VME crate is not buffered, its event by event readout drives the deadtime (which otherwise could be a factor of two lower).

Proposals for the energy upgrade in Hall A include a possible third arm spectrometer with associated front-end crates. DAQ throughput and deadtime could be improved with the introduction of a pipelined ADC and high-resolution TDC.

## 2.2 Hall B

The CLAS detector in Experimental Hall B consists of almost 40000 instrumented channels interfaced to 19 FASTBUS and 5 VME crates. It is, by an order of magnitude, the largest DAQ system implemented using CODA at Jefferson Lab. It is of similar (or even greater) scale than the proposed Hall D detector. Typical "high" running rates in Hall B are 2.5 *kHz* and 10 MBytes/s. Readout latency for the front-end crates run between 60 to 200  $\mu$ s. Crates holding the TOF high resolution TDCs (Lecroy 1872) not only have the longest readout latency but also the highest conversion deadtime (around 50  $\mu$ s), hence they drive both the DAQ rate limitations as well as the deadtime for the system. In tests where these crates were dropped from the system, they were able to achieve rates up to 7 *kHz* to tape. The DAQ system for CLAS would certainly benefit from the development of a pipelined high resolution TDC.

## 2.3 Hall C

Hall C supports two spectrometers as standard equipment, the High Momentum Spectrometer (HMS) and the Short Orbit Spectrometer (SOS). Most experiments run in the Hall are coincidence between these or with a User supplied "third arm" detector. There are typically 4 FASTBUS crates and 2 VME crates that make up the front-end for the DAQ system. As with the other experimental halls the limitation to the DAQ is with the conversion and readout of the digitizing modules (primarily the Lecroy 1872 TDC).

Requirements in the energy upgrade proposal include the building of a new Super HMS spectrometer for 12 *GeV* electrons. This will add several thousand detector channels. At the higher electron energies and more forward angles, larger backgrounds may require the implementation of

a fast hardware Level 2 trigger. Hence, they would like to upgrade all the DAQ electronics to support pipelined readout (1-2  $\mu\text{s}$  after the event) and maximum event rates in excess of 10  $\text{kHz}$ . Going to a full pipelined readout scheme would also allow the elimination of a large amount of cable delay currently used for all detector signals.

### 3 Hall D DAQ Architecture Overview

The data acquisition goal for the Hall D detector is to accept on the order of a 160  $\text{kHz}$  Level 1 trigger rate while not incurring any DAQ system deadtime. These data will be passed through a Level 3 farm and reduced to a 14  $\text{kHz}$  rate to permanent storage. These requirements put restrictions on the design of the DAQ system at every level from front-end to tape.

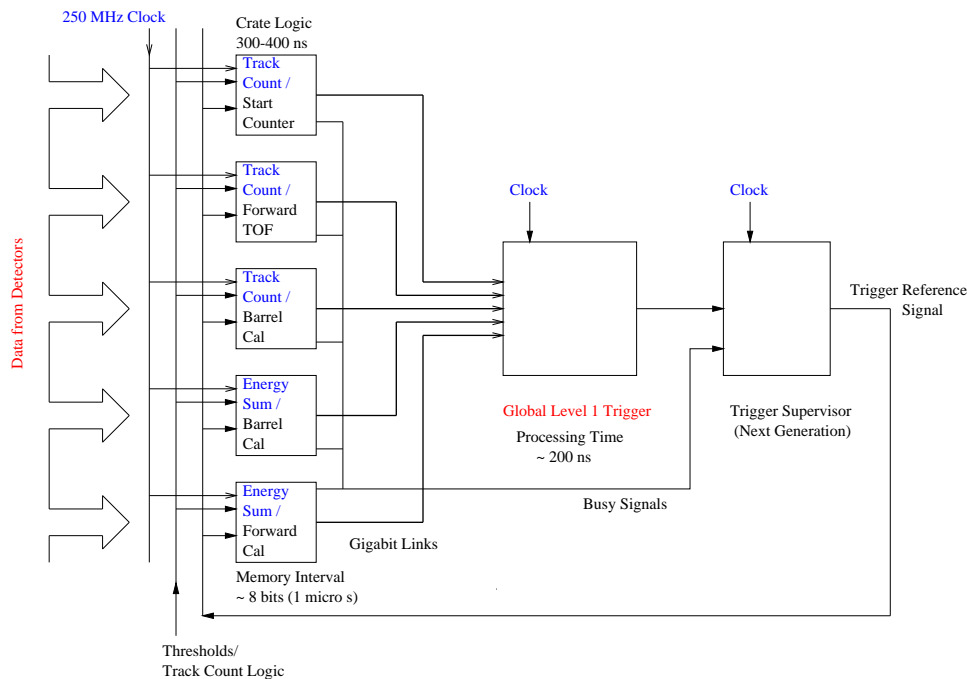


Figure 1: Global Level 1 Trigger

#### 3.1 Global Trigger

The Global Level 1 Trigger is responsible for establishing the first “rough cut” on incident photon energy using only prompt information from individual detector subsystems. Ideally it should run “deadtimeless” by continually updating state changes in the detector. It must run as fast as possible as this time determines the pipeline length for the front-ends. Output to the front-ends is basically a time index into their data buffers.

With the current detector design the primary candidates for input to the trigger include: energy in the Forward Calorimeter, energy in the Barrel Calorimeter (preferably segmented in theta), track counts in the Forward Time of Flight, Barrel Calorimeter, and “start counter”. Trigger definitions would include: Min/Max/Exact numbers of tracks, Min/Max total energy deposited, time windows for mtching detector subsystems, and an output delay from a trigger “match”.

The Global Trigger Processor (GTP) is shown as a schematic diagram in Figure 1. It consists of five separate subsystems located at or near the front-end. Each subsystem computes continuously at the pipeline rate of the digitization (e.g. 250 *MHz* ). Eash subsystem “timestamps” its output based on a synched distributed master clock (10 bits at 250 *MHz* would provide a 4  $\mu$ s window before rollover). The output for each subsystem would be in the format of a Subsystem Event Report (SER - might look like <Hdr:6><Time:10><Func1:8><Func2:8> ) and would be sent via a fast (greater than GB/s) optical link to the main GTP. Output need only be sent to the GTP if there were changes to report from a particular subsystem.

The GTP would run at the same synched distributed clock speed as the subsystems. It would receive and buffer the SERs from each subsystem and use multiple internal processors to find a match for one of many triggers (8-16 types). Output of a valid Level 1 trigger should be programmably delayed to the front-end electronics. This simplifies the front-end readout in that it becomes a fixed lookback into the pipeline.

Preliminary estimates for Level 1 trigger timing are listed in the table below:

Flight/Detector Time	30 ns
PMT Latency	30 ns
Cables to FEE	30 ns
FEE to trigger out	64 ns
Subsystem Processing	200 ns
Transfer SER to GTP	100 ns
GTP Processing	400 ns
Level 1 output to FEE	50 ns

The total time is less than 1  $\mu$ s. If the front-end electronics are designed for a pipeline of 1.536  $\mu$ s (384 stages at 250 *MHz* ) this would probably be sufficient.

### 3.2 Front End Processing and Data Flow

A possible pipelined Trigger/DAQ architecture is shown in Figure 2. Detectors which measure charge (calorimeters, dE/dX) are digitized by 250 *MHz* flash ADCs (FADC). The data pass through a shift register thus making available a time window. Successive samples within this time window are added together. This is equivalent to the gate width in a conventional charge sensitive ADC. Multiple channels are added together to form an energy sum. These adders are pipelined at the same 250 *MHz* rate, giving an energy measurement every 4 ns for use in the Level 1 trigger.

After passing through the shift register, the data are stored in a RAM. 16K locations would store 64 uS worth of time information. If this RAM has 2 ports, the data can be read out while the FADCs continue to run, making the front end electronics completely deadtimeless.

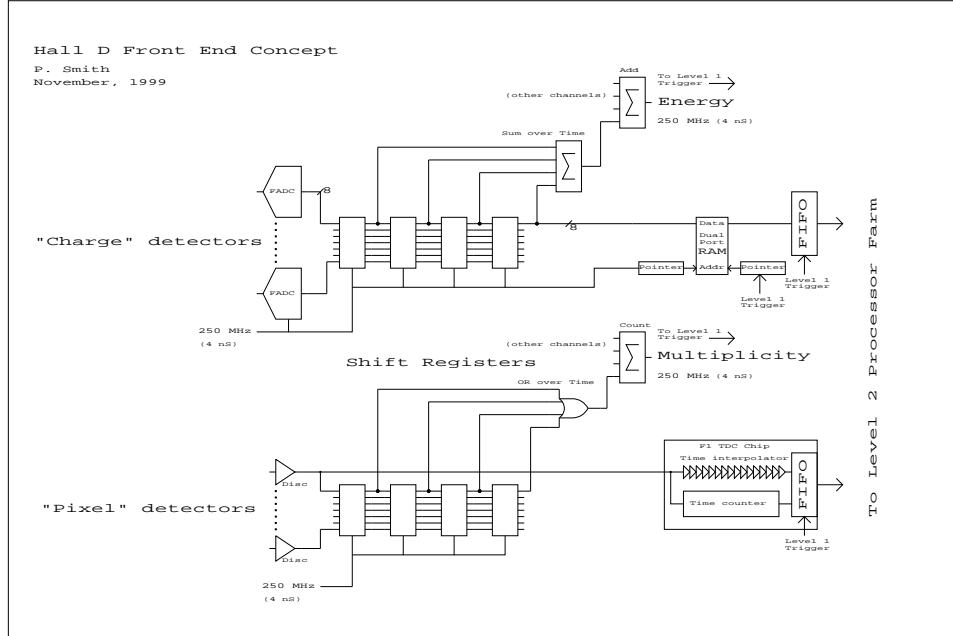


Figure 2: Front End Trigger/DAQ

Upon receipt of a Level 1 trigger, a pointer address is calculated and data are read from the FADC dual port RAMs and passed to Level 2 hardware/software. To achieve a practical data transfer rate, the zeros must be suppressed in hardware at this point.

The shift registers, adders, and dual port RAM would be implemented as an application specific IC (ASIC) or possibly in a programmable gate array (PGA). Adders and memory lead to extremely regular structures and are efficiently implemented in an IC.

"Pixel" detectors such as drift chambers would have a similar shift register for the Level 1 trigger. Stages of the shift register would be ORed to provide a time window and counted by adders pipelined at 250 MHz, thus giving a multiplicity value every 4 ns for use in the Level 1 trigger. A momentum cut could be implemented by looking at the outer layers of detectors in the solenoid since low momentum tracks "curl up" in the magnetic field and don't reach the outer layers. Forming tracks or at least track segments is also a possibility.

Before discussing CODA, one area that has to be studied further for Hall D is the processing required to reduce the tremendous amount of data generated by the FADC channels. This data must be analyzed and reduced to a few words per event. Many questions remain to be answered. What are the algorithms involved? How much processing does this require? What type of processor is best suited to this task? How do we handled the flow control? For the following discussion, we make the conservative assumption that the output of the FADC reduction results in about 12 bytes/channel. If sufficient reduction can't be achieved, then the data volume could become a

problem for both the VME backplane and the communication network from ROCs to EBs.

### 3.3 CODA Back End

CODA is the DAQ software system currently used by all experiments at Jefferson Lab and with some modifications could be used in Hall D. CODA is a flexible high performance collection of software components which allow dynamic construction of DAQ systems. There are four main components:

- ROC - The Read-Out Controller runs in a CPU housed in the same crate as the digitizing electronics. It is dynamically programmed to read-out events which are later sent to the EB.
- EB - The Event Builder collects event fragments from all ROCs outputs these events to the event transport system.
- ET - The Event Transport system moves events efficiently from one Unix process to another either on the same CPU or across the network. It frees CODA components from knowing the details of where the event stream has to go.
- ER - The Event Recorder gathers events from ET and writes them to permanent storage.

CODA comes with a lot of other tools for configuring and running a high energy/nuclear physics experiment (<http://www.jlab.org/coda>).

The data flow architecture is shown in Figure 3. This design is similar to that of Hall B at JLab except there are several EBs and ERs and a large farm. DAQ system of such a scale have been constructed at many High Energy labs around the world. Currently in CODA, the ROC reads out the electronics for each and every event. In Hall D, the event rate is too high for this scheme to work with a general purpose real-time operating system and CPU. There needs to be some other method that is either built into the electronics or into a single device in the crate to allow the ROC to read out events in blocks. The approximate data rate per crate is in the neighborhood of 20 MB/sec. This can be handled easily by a VME64x system and with sufficient computing power (250 SpecInt) by the ROC CPU equipped with gigabit ethernet. The total data throughput that the main DAQ switch needs is around 600MB/sec. This capacity is currently available in gigabit ethernet switches. The data filter program running on the farm nodes will connect to an EB event transfer system to read events and write them to an ER event transfer system. Control and monitoring software for these nodes and processes will need to be written and interface with the experiment control software. One important issue is how often calibration runs need to be done to keep the filtering efficiency high. The expected data rate out of the farm is about 20 *kHz* and < 100 MB/sec. Five or fewer Event recorder processes should be able to handle this. Note that the events will not be written to mass storage in any particular order. Dual ported fiber channel RAID disk will be used as in CLAS to allow the DAQ and the tape system to access the data in turn.

In summary, the CODA DAQ system can be adapted to work in Hall D. Parallel event building will be added next year. The run-control system is in the process of be rewritten. Hall D's needs will be incorporated into the design requirements. The collaboration should decide on a data format

Prelim. Hall-D DAQ Schematic  
 Oct 1999 - R.W. MacLeod

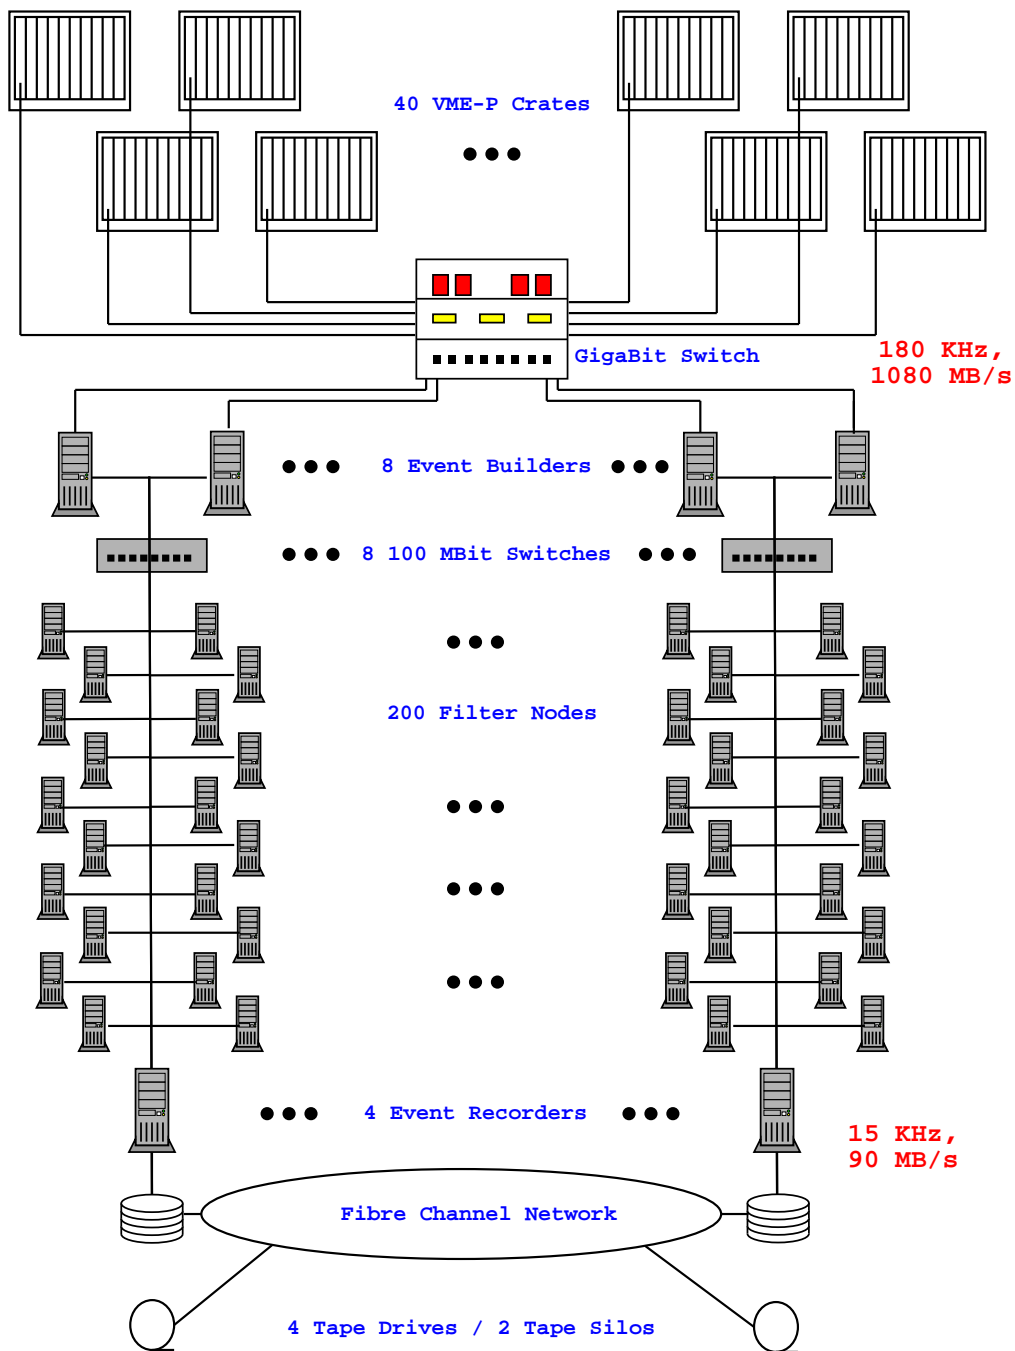


Figure 3: CODA Back-End

and ensure that it will work efficiently within CODA. The JLAB DAQ Group will ensure that CODA is ready to acquire data at 160 *kHz* when the Hall D collaboration is ready to make it so.

## 4 Hardware Technology - Current Status

Current commercially available hardware designs and standards vary in their applicability and suitability to the data acquisition requirements of Hall D. In this section we present a review of some of these commercial products and the directions the industries are moving in the next several years.

### 4.1 Digitization and Data Bus Designs

Commercial Analog to Digital electronics (ADCs and TDCs) for physics applications have historically been developed for use in both CAMAC and FASTBUS crate systems. These systems provide a flexible, modular design that supports high power and low noise for the analog side, and a fast parallel backplane for efficient readout of the digitized data. However, in recent years the number of companies that are providing electronics in these form factors is declining. More importantly, the development of new hardware and technology designs for these form-factors is dwindling.

In the past 10 years the physics community has slowly embraced a different modular electronics bus standard (VME) for many applications. The VME standard has many advantages that include a faster bus speed (up to 160 MBytes/s), but perhaps more importantly it has broad support in many other commercial industries. This should make hardware and electronics using the VME standard viable and affordable in the marketplace for many years to come.

The VME standard does suffer from some deficiencies with respect to Physics specific applications. These include a lack of a standard for high-power, low noise voltage supplies to support analog electronics. For readout, there has been no standard for slave-terminated or multi-module block transfers. For several years some vendors (CAEN) have supported a CERN specified extension to VME called V430. It provides for additional voltages and more power to VME modules via a third (P0) connector. More recently, the VME International Trade Association (VITA) has adopted a new standard for VME in Physics Applications (VIPA) which addresses all of the issues. Early adopters of the new VIPA standard include several experiments at Fermilab.

### 4.2 Future directions

Currently, there is only one vendor (Wiener) manufacturing a complete VIPA system (i.e. backplane, power supplies, fans, remote control and monitoring), and the price is somewhat expensive. CAEN has indicated a willingness to provide some of its new front-end electronics using the VIPA standard (e.g. a new high-resolution pipeline TDC). The extent to which applications for Hall D need the VIPA form-factor should be studied and weighed with respect to the increased cost. It is possible that some or all of the commercial and custom front-end designs can be implemented within the current VME standard (VME64x).

## 5 Hall D Hardware requirements

Two basic types of digitizing electronics will be used as an interface with the Hall D detector. In order to support a deadtimeless readout scheme a high-speed sampling (250 MHz) flash ADC as well as a high resolution (50 ps) multi-hit pipelined TDC are required. Although there are no commercially available products that are suitable the technology does exist to custom design boards.

In addition to the digitizing electronics, the Hall D trigger system requires distribution of a high speed clock for purposes of passing trigger information to the front ends as well as trigger synchronization for event building.

### 5.1 ADC Design

A majority of the detector channels will require energy and/or low resolution (1 ns) time measurements. By instrumenting these channels with high speed (250 MHz) flash ADCs, one can pipeline the energy information and provide enough sampling for coarse timing measurements. The primary drawback to such a design is that a large amount of data can be generated (e.g. 1  $\mu$ s at 250 MHz and 8 bits/sample gives 250 bytes/channel). This data must be “processed down” at the front-end to give more specific energy and time information.

Appropriate commercial designs are not currently available; hence, we propose to build a flash ADC prototype board in a VME form-factor. A possible solution would be to use the Analog Devices 9054 (200 MHz 8 bit) flash ADC chip interfaced to two Cypress 100 MHz dual-ported RAMs. Prompt energy information would be needed for the Level 1 trigger, so prior to storage in the DP-ram the FADC data would be passed through a series of shift registers and sampled and summed over a specified time window. Upon receipt of a Level 1 trigger a section of data from the DP-Ram would be read and processed, and the reduced data would be stored in a FIFO available for readout from the VME side. A possible solution for the processor and FIFO interface would be the Xilinx Vertex FPGA (100 MHz).

### 5.2 TDC Design

The F1-TDC chip is a new, high resolution Time-to-Digital Converter that has been specifically developed for the COMPASS [2] experiment at CERN and implemented in 0.6 micron CMOS technology. It has been designed to address time measurements on a variety of detectors using a pipelined architecture and as such can be operated deadtimeless in high resolution (4 channels with 60 ps/LSB each), low resolution (8 channels with 120 ps/LSB each) and latch (32 channels with 5.7 ns/LSB each) modes. It is being marketed by acam [3] for under \$60 in quantities of 2000 units.

The core of the F1-TDC is composed of a 19-tap, asymmetric ring oscillator, Phase Locked Loop (PLL) capable of operating at an input reference clock frequency of up to 40 MHz and achieving a 16-bit dynamic range. Time measurements can be referenced to a Synch-Reset input, to a Common Start input or to a Trigger input. Each channel has dual port, 16 hit buffers, a trigger matching buffer and 8 readout buffers. The trigger matching feature of the F1-TDC, with programmable trigger latency and programmable hit selection window, is of critical importance to pipelined DAQ

architectures. Valid data for all channels is stored in a 16-register, 24-bit interface FIFO from which data can be retrieved at a databus readout rate of up to 50 *MHz* or in burst mode.

We expect to start the design of a multi-channel TDC prototype at Jefferson Lab in the very near future. This prototype will incorporate F1-TDC chips, readout and setup data storage, a processor and VME interface.

### 5.3 Clock Distribution/Trigger Synchronization

A synchronous pipelined data acquisition system operating at 250 *MHz* has been proposed for Hall D. Here we explore the possible solutions for distributing and synchronizing this clock. We assume in this study a board + crate architecture for the front-end electronics.

The goals are twofold: all channel pipelines of the system should be clocked at the same time, and all the pipelines of the system should all be enabled for clocking on the same clock interval. Accomplishing these allows the 4 ns detector data samples to be correlated across the entire system. Achieving a channel clocking skew of less than 1 ns across the entire system is our goal.

There are two signals we consider: clock and reset. Each front-end board has a counter that is incremented by clock. The count is held at zero while reset is asserted; counting is enabled when reset is deasserted. The count value serves to time stamp the data, and is used as the address in memory where the data sample is stored. The counter on every board in the system must be enabled to begin counting on the same clock edge so that all data fragments are properly aligned.

A centrally located master clock module generates a low frequency (25-50 *MHz* ) clock signal. The asynchronous reset signal is synchronized to the master clock in this module. Both the low-frequency clock and the synchronous reset are fanned out and are sent to each crate in the system. Each crate contains a local clock distribution board. A programmable delay element is inserted at the clock and reset inputs of the crate distribution card. This allows for a fine adjustment of timing at the crate level. The local 250 *MHz* clock signal is generated from the low-frequency input clock using a Phase-Locked Loop (PLL) circuit. The PLL guarantees phase alignment of the 250 *MHz* clock with the input clock. The 250 *MHz* local clock and the reset are fanned out and are sent to each module in the crate.

Fan out of the clock and reset signals on the master clock module, the crate distribution board, and the front-end boards is done using low-skew (50-100 ps) fan out buffers and employing careful layout techniques (e.g. equal length traces). Distribution of the signals to modules within a crate may be done using short (<1 ft), matched, high-quality coaxial cables.

Since the master clock module and front-end crates may be hundreds of feet apart, we are investigating the use of fiber-optic links in this level of distribution. Care must be taken not to introduce significant jitter into the signals here, as the PLL circuit may have problems locking onto the signal.

A scheme for distributing the low-frequency clock and reset signals to the crates uses G-link components. G-link is a virtual ribbon cable interface for the transmission of data. Parallel data (a frame) loaded into the transmit chip (Tx) is delivered to the receiver chip (Rx) over a serial channel (fiber) and is reconstructed into its original parallel form. The serial link is synchronous. Frames may be clocked into the Tx chip at a rate of 7.5 *MHz* to 75 *MHz* . The frame clock is recovered from the serial data stream by the Rx chip. Individual G-links connect the master clock

module to each crate distribution board. The master clock serves as the frame clock for each Tx chip, and the Rx chip in each crate recovers this clock. Reset is specified as a bit asserted in the parallel data word transmitted, and is synchronized to the clock at each Rx. Many other “timing events” could also be broadcast as unique data patterns to the crates.

## 6 Online Management

Data acquisition monitoring and control include tasks such as online event processing, hardware configuration and control, experiment bookkeeping, alarm handling, and messaging. Integration of this system with the high speed data flow channel will require careful planning to avoid common pitfalls such as incompatible online and offline analysis efforts, duplication of software development, and inefficient and/or inadequate data quality monitoring. Although we have 2 years before hard decisions have to be made concerning online strategies, software, etc, now is the time to start thinking about them.

Six broad, overlapping areas in the online were discussed (there are others not covered): slow controls, information management, information transfer, information storage and retrieval, information display, and configuration control.

The main question in slow controls is whether to use EPICS or not. The argument against EPICS is that it is a large, powerful, but cumbersome and expensive system. The cost per control point is high in EPICS for sparse systems like Hall D. It is also important to note that many modern commercial control systems for hardware (e.g. HV controllers) provide all the functionality that EPICS could. We should think in terms of “internet appliances” rather than in terms of dumb devices needing complicated computer systems to control them.

Two information management schemes commonly used in online systems are “push” (subscription based) and “pull” (client/server based) systems; CLAS uses both. Push systems are very powerful and convenient, but the two can coexist easily.

Information transfer via “middleware” is common in today’s computing environments. The leading solutions today are object-based (e.g. CORBA) and message-based (e.g. SmartSockets) systems. Use of a single middleware system for the entire online (offline?) is probably desirable, although multiple systems can easily coexist.

Databases are powerful and widely used, and Hall D thinking should be “database-centric” from the start. The leading contenders are relational and object-oriented databases. The former is widely used and very mature, the latter is more powerful but less mature and less widely used.

There many choices for information display: ROOT, hbook/paw, JAS/Java, HistoScope, web, etc. We probably will use more than one.

Detector configuration control is best done via a grand state machine, perhaps integrated with the CODA run control state machine. The alternatives are pre-run checklists, or operator skill and memory. The former is reliable but tedious, the latter unreliable but widely used (e.g. CLAS).

## 7 Summary and Conclusions

Much of what is required to implement a triggering and data acquisition system for Hall D is technically feasible today. More ambitious projects in the field of High Energy Physics have been or will be implemented within the time frame for the proposed construction. Most details dealing with bandwidth and computing power are ones of scale (i.e. How many Level 3 farm nodes and tape drives will be necessary).

There are some issues that will require R&D on the collaboration's part. Firstly, prototyping of the FADC and TDC front-end electronics should be done, probably sooner than later. There are applications in the existing experimental halls that could benefit now from their design. Secondly, triggering without a tagger at high luminosities needs to be better understood. Information on the input, programming and fault tolerance of the GTP is needed. More studies on detector response, rates, and occupancies under realistic physics conditions would be worthwhile. This would also provide information concerning the processing and bandwidth requirements for the front-end as well as for the Level 3 farm.

For the DAQ backend, including the Level 3 farm and Online/DAQ management careful consideration should be used in implementing various strategies. Increased complexity can directly translate into increased costs (e.g. more CPUs or more manpower for software development). For the Level 3 farm one can look to more efficient network solutions such as parallel cluster processing. This is something that the SSI (Strategic Super-computing Initiative) is working on, and Hall D could take advantage of the resources. EPICS is probably overkill for slow controls and monitoring in Hall D. Perhaps one could use some EPICS tools at the network level (like cdev) to integrate various independent control systems. As an online strategy all Hall D control systems should be "database-centric" (including the DAQ). Of all the computing based technologies that are emerging today, Java should be embraced at some level in the development of the online system.

## References

- [1] Hall D Collaboration, Photoproduction of Unusual Mesons at Jefferson Lab, Version 2 (August 1999)
- [2] H. Fischer et al., T TDC Chip and Readout Driver Developments for Compass and LHC Experiments, Fourth Workshop on Electronics for LHC-Experiments, Rome, September 21-25 1998
- [3] acam-messelectronic gmbh, Haid-und-Neu Strasse 7, D-76131 Karlsruhe