5 Front-end electronics, triggering, and data acquisition

In this section we examine the ramifications of luminosities ($\mathcal{L}$) between 10^{33} cm^{-2}s^{-1} and 10^{34} cm^{-2}s^{-1} on the CLEO front-end electronics, trigger, and data acquisition systems. We briefly review the present status of these systems and the implications of their unmodified use at $\mathcal{L} = 10^{33}$ in view of the expected backgrounds. Trigger and front-end architectures applicable to high luminosity operation are discussed, and a data-acquisition system having sufficient bandwidth and flexibility is presented. Since we are designing a new electronics architecture that will be in use beyond the initial CESR Phase 3 upgrade, it must be extensible to $\mathcal{L} = 10^{34}$, the upper limit that CESR can achieve.

5.1 CLEO II system

CLEO II data recording is controlled by a three level hardware trigger system and a software trigger level that follows data acquisition readout. The design is matched to the single-hit nature of the front-end electronics systems. The performance of the overall system as parameterized by dead-time is limited by this single-hit architecture and by the inability to perform simultaneous front-end acquisition and data readout.

5.1.1 Trigger

A level 0 trigger operates synchronously with the storage ring, the decision being made in the 360 ns between crossings of the seven beam bunches. In this way, failed level 0 triggers incur no dead-time. At 5.28 GeV and $\mathcal{L} = 2\times10^{32}$, the level 0 trigger reduces the 2.8 MHz beam crossing frequency by a factor of $\sim$300 to 10 kHz. Examples of existing level 0 triggers are:

- Any two separated barrel time of flight counters
- Opposite endcap time of flight counters
- One time of flight counter and a track segment in the vertex detector
- One high threshold barrel crystal region
- Opposite high threshold endcap crystal regions

A limitation of the present system is that mixed tracking and calorimeter triggers can not occur at this level because of their disparate signal development times. CsI energy trigger information is available after $\sim$2 μs. The self-resetting feature of the drift chamber electronics causes the wire hit information to be reset 800 ns after the discriminator fires if no trigger inhibits the reset.

Successful level 0 triggers incur experimental dead-time to freeze the contents of the front-end electronics and initiate the level 1 trigger hardware which has a decision time of 1.2 μs. Failed level 1 triggers entail another 1 μs dead-time to reset the front-end electronics. At 5.28 GeV and $\mathcal{L} = 2\times10^{32}$, the level 1 trigger reduces the 10 kHz level 0 trigger output by a factor of $\sim$300 to 30 Hz. This includes pre-scaling of QED processes in the endcap regions. Some examples of level 1 triggers in use today are:

- Back-to-back high threshold crystal regions in the barrel or in the endcap
- One time of flight counter, two track segments in the drift chamber, and one minimum ionizing crystal region
- One each high threshold barrel and endcap crystal regions

Successful level 1 triggers initiate, if appropriate, the level 2 trigger which consists of two hardware tracking processors. One uses the inner 10 layer drift chamber (VD) and the other uses the axial layers of the main drift chamber. One works with digital hit information formed from the OR of 4 drift cells. The other uses wire-level hit information to find tracks in predefined road-widths. The total processing time is 40 μs. Under 1994 $\Upsilon(4S)$ running conditions, this trigger rejected 50% of the level 1 triggers.

Accepted level 2 triggers initiate front-end readout by the data acquisition system where the data is buffered at the input to the VME-based event builder system. The trigger is then re-enabled and the data sparsification and event building proceed asynchronously. After event building, data are transferred to an on-line workstation where they are analyzed by the level 3 software trigger. Access to the full data allows this code to calculate such quantities as neutral energy balance and to use the inner tracking detector hits to eliminate beam-wall events. Level 3 reduces the trigger rate by about a factor of two.

The overall performance of the trigger chain can be simply determined by examining the fraction of events that are qualified as "good" by the off-line reconstruction code. Fifty percent of the events processed by the level 3 trigger are rejected. They are generally categorized as not containing charged tracks or not being interesting pure neutral events. The efficiency of the trigger system has been determined to be $> 91\%$ for $q\bar{q}$ events and nearly 100% for $\Upsilon(4S)$ hadronic events.

5.1.2 Gating and event flow control

The CLEO II gating system uses a subset of the timing modules developed to operate the storage ring. Different gates are generated for the various detector types and fanned-out to the crates. The gate widths and placement relative to beam crossing time are software controllable. An event trigger signal is also fanned-out to the crates to initiate the transfer of data to the raw data buffers. The data acquisition system (DAQ90) receives a copy of this signal so that it can check for the completion of the event read-in. On completion, DAQ90 sends a signal to the trigger system to re-enable it. The CsI electronics have an additional fast clear signal to reset the input integrators after a successful level 1 but a failed level 2.

5.1.3 Front-end electronics

There are three categories of front-end electronics in the present CLEO II detector: digital systems associated with the trigger systems, LeCroy FASTBUS ADC's for the CsI, and CLEO-specific time-to-voltage and charge-to-voltage sample/holds (S/H) for the various drift chamber, muon, and time-of-flight systems. The front-ends are susceptible to three sources of background: cosmic rays, stray beam particles, and synchrotron radiation. The first one can be dealt with at the trigger level and the rate is low enough that overlap with physics events is negligible. This is not necessarily true in the cases of stray beam particles and synchrotron radiation, however. For example, due to the CLEO II single hit electronics, which have a memory time of a couple of microseconds, $\sim$10% of the VD cells read for an event appear to be the result of synchrotron photons and stray beam particles.

5.1.4 Data acquisition

The data acquisition (DAQ) system was designed to work within the limiting features of the front-end electronics such that it introduces no additional dead-time. The analog-to-digital conversion time of the S/H's and the restriction that data transmission cannot occur while the front-end electronics are active determines the dead-time. This has led to a scheme where the data conversion occurs rapidly (1 μs) and the datum from each conversion is sent immediately to a remote buffer. Each front-end crate has a 16-bit wide, 12 MByte/s point-to-point link to a dedicated VSB dual-port memory. A typical crate has 2000 front-end channels and the total conversion and transfer time is 2 ms. This contributes 0.2% dead-time per Hz of level 2 trigger rate. The dual-port memories are serviced by a number of VME processors that do sparsification and event fragment assembly in a distributed fashion. The finished events are transferred via Ethernet to a workstation running the level 3 trigger code. The average CLEO hadronic event is presently about 8 kByte and the average readout event is 5 kByte.

5.1.5 System performance

The trigger/DAQ performance is summarized in Table 26. As can be seen, the level 1 latency time and the DAQ conversion/readout dominate the deadtime. They will contribute

Table 26: CLEO II trigger performance summary at $\mathcal{L} = 2\times10^{32}$ cm^{-2}s^{-1}.

	Input rate	Output rate	Rejection ratio	Latency time	Dead-time %
Level 0	2.8 MHz	10 kHz	300:1	0	0.0
Level 1	10 kHz	30 Hz	300:1	2.2 μs	2.2
Level 2	30 Hz	15 Hz	2:1	40 μs	0.1
DAQ	15 Hz	15 Hz		2 ms	3.0
Level 3	15 Hz	7 Hz	2:1	variable	0.0
Total					5.6

to massive dead-times at $\mathcal{L} = 10^{33}$ if architectural changes are not made.

5.2 CLEO III

We present an overview for proposed general architecture of the CLEO III front-end electronics, trigger, and data acquisition systems. This has been driven by several conditions:

- Luminosity
- Backgrounds
- Live time
- Scalability
- Funding

The primary goal of the detector upgrade is to operate with CESR Phase 3 at luminosities up to $\mathcal{L} = 3\times10^{33}$ cm^{-2}s^{-1} and to easily accommodate $\mathcal{L} = 10^{34}$. The exact accelerator bunch structure that will achieve these luminosities is uncertain at this time, ranging from nine bunch trains with 5 sub-bunches to almost continuous bunches 14 ns apart. Clearly, any CLEO III design choices must be compatible with any machine bunch structure.

Particle and synchrotron radiation background rates are uncertain. Different CESR topologies make their prediction even more uncertain. Furthermore, extrapolations to luminosities fifty times greater than today cannot be viewed as reliable. From discussions below, we will see that background rates in detectors near the beam pipe can be as high as 100 kHz.

Some CLEO III electronics cannot easily be upgraded to to accommodate the full reach of CESR luminosity beyond Phase 3. Among these systems are the front-end electronics (preamps, data boards, and the trigger information derived from the front-end boards), the board and crate level event buffering, the crate back-plane technology, and the fast timing system including event flow control. The electronics architecture must accommodate operation beyond $\mathcal{L} = 2^{33}$ just as the CLEO I and II architecture will have accommodated 15 years of luminosity improvements.

Items that are required to operate CLEO and that can be enhanced or upgraded include the hardware trigger processors and decision logic, the event builder data collection and assembling, the level 2 trigger, the level 3 trigger implementation, and the data delivery mechanism to the event builder/level 3 trigger processor(s).

We briefly state the modifications required for the CLEO III detector. In the trigger system we eliminate the lowest level trigger which is driven by the beam crossing rate; track and energy information can be used in coincidence in the lowest level trigger; the deadtime is determined by the level 2 trigger latency, $\sim$10 μs. In the drift chamber we preserve sub-nanosecond timing precision with memory times up to 2.5 μs; both charge and time measurements are performed; and multi-hit capability is provided to deal with track congestion and pile-up at small radii. The silicon detector requires an on-detector front-end IC optimized for use in CLEO which digitizes and buffers data in 10 μs and may provide trigger information. The particle ID system requires development of a preamp and readout electronics. The CsI and muon systems have their readout electronics upgraded for 10 μs conversion and data buffering. A new DAQ architecture is used to readout the detector at up to 1 kHz with little or no dead-time impact by using several levels of buffering to allow overlapped readout and acquisition.

5.2.1 Trigger

We will see that if we want to contemplate higher background and physics rates, we must decrease the product of the rate times latency of the lowest level trigger as well as the digitization/readout time of the front-ends. We discuss the trigger solution first.

Backgrounds

A good understanding of the expected low-level trigger rate for $\mathcal{L} = 10^{33}$ cm^{-2}s^{-1} and greater is desirable when designing the architecture for the trigger and front-end electronics. Unfortunately, this is a difficult number to predict. To set the scale for the trigger rate, the cross sections of interesting processes at the $\Upsilon(4S)$ are given in Table 27. This does not include two photon induced processes. At the $\Upsilon(1S)$ the resonance contribution to the hadronic rate is five times larger than at the $\Upsilon(4S)$. At $\mathcal{L} = 10^{33}$, 1 nb of cross section gives 1 Hz of trigger. In round terms, the physics trigger rate will be $\sim$100 Hz at $\mathcal{L} = 10^{33}$. Assuming that any new trigger system will provide the same rejection and efficiency as the present one, we would expect a hardware trigger rate of about 200 Hz to 2 kHz without any pre-scaling of QED events.

While the interesting physics event-rate can be scaled from present CLEO experience,

Table 27: Cross sections at 5.28 GeV for an active detector region extending down to 330 mrad from the beam line.

	σ_{Total} (nb)	σ_{Barrel} (nb)	σ_{Endcap} (nb)
$e^+e^- \to e^+e^-$	72	19	53
$e^+e^- \to \gamma\gamma$	6.2	3.7	2.5
$e^+e^- \to \mu\mu$	0.72	0.60	0.12
$e^+e^- \to \tau\tau$	0.72	0.60	0.12
$e^+e^- \to$ hadrons	4		
total	84		

any fast (level 0) hardware trigger rate is largely driven by the IR background conditions. There are several ways to estimate this background rate for the $\mathcal{L} = 10^{33}$ configuration, although none is ideal. The first of these is to accept the rough estimates of the CESR machine designers, namely that changes to the IR and machine lattice will no more than double the background rates seen by CLEO II for each of the next two machine modifications (the completion of phase 2 and the introduction of phase 3). This means that at $\mathcal{L} = 10^{33}$, the level 0 trigger rate should be 40 kHz or less.

One can also use an extrapolation of the CLEO II experience operating over a factor of ten in luminosity. Figure 51 shows the level 0, 1, and 2 output rates over a factor of 3 in luminosity. All rates show a nearly linear dependence. Extrapolating to a factor of 3.3 increase in luminosity to $\mathcal{L} = 10^{33}$ gives a level 0 rate of 50 kHz.

Finally, lost beam particles and synchrotron radiation can be carefully simulated. The advantage of this technique is that the algorithm used can be tested by simulating the present machine/detector configuration and comparing with CLEO II observables, such as level 0 trigger rates and wire chamber currents. Comparisons of this sort have shown good agreement between measured and simulated quantities, giving us some confidence that the simulation will be correct for the 10^{33} and future CESR configurations as well. This method predicts that the level 0 trigger rate will remain below 25 kHz at $\mathcal{L} = 10^{33}$.

Since a definitive prediction of the level 0 trigger rate is not possible, the conservative approach we adopt is to take the largest of the above estimates (50 kHz). Furthermore, to allow for the initial period of machine commissioning when the backgrounds will inevitably be large, any new trigger/front-end designs should accommodate a level 0 rate of 100 kHz.

Summarizing, at a luminosity of 10^{33} cm^{-2}s^{-1} we should expect a level 0 (*i.e.*,prompt) trigger rate of up to 100 kHz, and a final hardware trigger event rate of about 200 Hz.

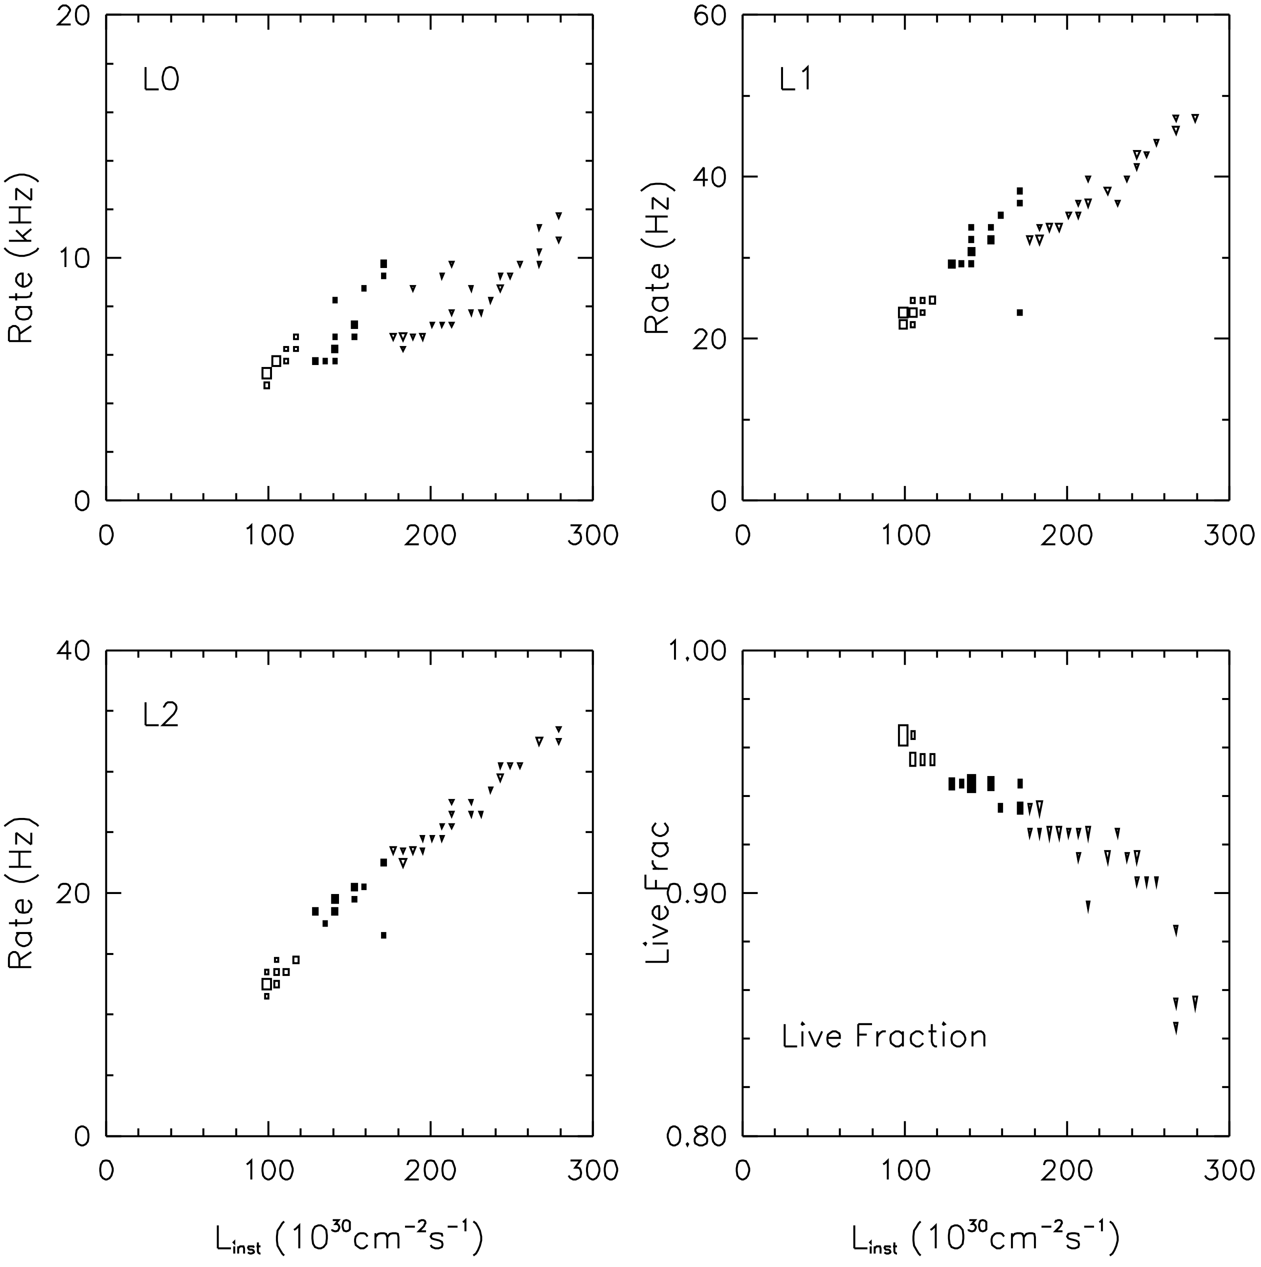

Figure 51: CLEO II measured trigger rates as a function of luminosity. Open squares are from December 1992, closed squares are from Summer 1993, and the triangles are from December 1993.

Level 1

A trigger architecture similar to that in CLEO II will lead to large dead-time introduced by the level 1 latency if it is driven at the level 0 rates described above. We propose to do without a trigger level that attempts to track the beam crossings. While technologically challenging, it is not impossible to build a trigger that operates at 72 MHz, but there is no need for this since drift chamber signals used in the trigger are not available for up to 500 ns after a beam crossing and since the interesting event rate is less than 1 kHz compared to the many megahertz of beam crossings. Event overlap is negligible but events can have about 1% occupancy due to background.

Level 1 will be the lowest level trigger combining the features of both the present level 0 and level 1 while maintaining the overall 10^5 rejection. The total processing time for this level can be as long as 2.5 μs, an upper limit set by the need to start gating the CsI ADC's. The rate at which the level 1 trigger decisions can be generated is naturally set by the longest drift times used in the tracker trigger processor. This will be between 200 and 1200 ns depending on the drift chamber gas and cell design choice.

Modest levels of pipelining will be required to implement this trigger. The pipeline also makes it possible to time-align the data from the CsI and trackers which generate trigger information on very different time scales relative to the event beam crossing. Since some detector components will have an uncertainty in the arrival time of their trigger information that is larger than the pipeline bucket width, their information must be stretched across several buckets if the logic interrogating the trigger pipelines is to obtain information from a unique set of pipeline positions. Equivalently, the trigger information could be stored only once and the interrogating hardware be made smart enough to examine several buckets in a single pipeline.

Note that we are not designing a dead-timeless system. Level 1 is dead-timeless until a successful trigger occurs. A small dead-time, 10–100 μs, is then incurred for data conversion and the level 2 decision time. This greatly simplifies the design. Furthermore, dead-time after a successful level 1 may be required for silicon readout systems that cannot perform simultaneous read and write.

The CLEO II trigger has been generally successful in providing the physics of interest to the collaboration. For $\mathcal{L} = 10^{33}$ we can maintain the general trigger categories used today, including low multiplicity final states of τ and $\gamma\gamma$ physics. Improvement is possible in tracking by using finer granularity hit information and in energy triggering by providing physics-oriented quantities such as total energy and energy flow.

Tracking triggers will continue to rely on linking track segments found in adjacent clusters of axial layers. The loosest pure track trigger would comprise one track spanning the entire tracking radius. If this is too loose, an additional track that exits the main drift chamber at its half radius can be demanded. One possible improvement to the present track segment approach where four or more wires are OR-ed to form a hit bit is to use each axial wire in the segment finding. This is a simple brute force extension of the present track segment processor design. It will provide tighter x-y vertexing by reducing the segmentation at small radii. Also, segments can be momenta binned and this additional information used in the linking of segments.

The CLEO II energy trigger is essentially topological in nature, examining with very coarse granularity whether regions of the calorimeter are above either or both of two thresholds (100 MeV and 500 MeV). The primitives used are two level-discriminator bits available for each 4×4 crystal array. The present trigger OR's these into 16 regions in the barrel and 16 regions in the endcap. This is adequate for triggering on a few large showers or on constrained geometry of lower energy showers. An upgraded trigger could use these bits at their full granularity. In addition, there is another set of unused discriminator bits where the leading edge is determined by a level-threshold and the trailing edge is determined by the zero-crossing of a fast shaped analog sum of the 4×4 crystals. These may be important to "bunch-stamp" neutral energy triggers or to correlate with track triggers.

While the digital approach should be maintained, it should be augmented with an analog energy trigger. There are unused analog current sums from each 4×4 crystal array which could easily be made available at the front-end crate level. This trigger can be done in analog form with current sums or digitally with sums from FADC's. In an analog implementation, one would conditionally combine the currents by digital thresholding to reduce noise contributions from quiet crystal arrays, and then threshold the combined analog sum. By appropriate weighting of the currents, total, transverse, and longitudinal energies can be formed. A level 1 trigger of this type could serve as the seed for a more sophisticated global energy trigger at level 2. These changes are attractive but their efficacy should be studied by Monte Carlo first.

It is conceivable that the particle ID system would be fast enough to contribute to the first level trigger. The major benefit could be to identify the event's beam crossing number which would give a start time for the drift chamber and CsI gate.

A desirable feature of the level 1 trigger is that it be in continual operation even when the DAQ system is incapable of handling more events. The delivery of the level 1 triggers to the crates and level 2 can be vetoed by the buffer overflow signals from the front-end boards and crates and from the event builder. In this way, the level 1 trigger can be used to measure actual experimental dead-time since we will not be measuring the gated or ungated luminosity of each beam individual beam bunch crossing.

As for any trigger, the input data and hardware algorithms must be easily modeled by the detector simulation codes.

Level 2

It should be noted that a level 2 trigger may not be required at all if level 1 can reduce the event rate to less than 1 kHz. In this scenario, the on-line software trigger—level 3—

will be required to play a bigger role on the overall rejection factor of the system. Any decision to proceed along this route would have to take into account the increased bandwidth requirements within a front-end crate and from the crate to the level 3 node(s). Advantages would include a simpler hardware system with great flexibility that is easier to understand and simulate.

If it is needed, the level 2 trigger is initiated by a level 1 trigger. Since front-end data conversion and storage are also initiated by level 1, level 2 can take as long as the slowest converting detector device without increasing the dead-time. If the front-end systems support being interrupted during conversion, a failed level 2 trigger can abort the conversion and re-enable data taking. The present level 2 track processors require 50 μs and a new design should be able to cycle in $\sim$10 μs.

The level 2 tracking trigger would make use of the same wire hit information used at level 1. It would perform a wire-by-wire track linking algorithm and report a histogram of found momenta. The track finding would be in $r - \phi$ only since z track information is encoded in the stereo layer drift times and in the silicon. The former is hard to use in hardware and the latter may not be available to the trigger. After digitizing the $\sim$500 crystal energy currents, the data can be processed to provide refined total and projective energies. It is also possible to provide a list of cluster energies and photon pair masses.

Figure 52 shows a block diagram of the trigger.

5.2.2 Gating and event flow control

The gating logic interfaces to the CESR timing systems using components from the accelerator master timing system. The various gates and clocks used by the front-end and trigger systems and their synchronization with the beam bunch crossings are implemented with these CESR components. The timing system will be packaged in VME and be remotely software programmable.

The event flow control system and its network coordinate the DAQ, trigger, and gating systems. Gates, level 1 and level 2 trigger accepts, level 2 trigger aborts, and event numbers are distributed to the crates. Control of programmable calibration pulsers is also done here. Data collected via this network are the event buffer availability in front-end boards, front-end crates, and event builder memories. A grand OR of the status of these systems is used to inhibit—throttle—further event triggers to the front-ends.

It may be impossible to measure the individual luminosity of each bunch crossing in CESR. This is the motivation for having the level 1 trigger in continuous operation even when the DAQ system cannot handle more events. The throttle signals from DAQ convert successful level 1 triggers into failures. The gating and event flow control system provides scalers to record this activity allowing a reliable dead-time correction to be made to the off-line calculated luminosity.

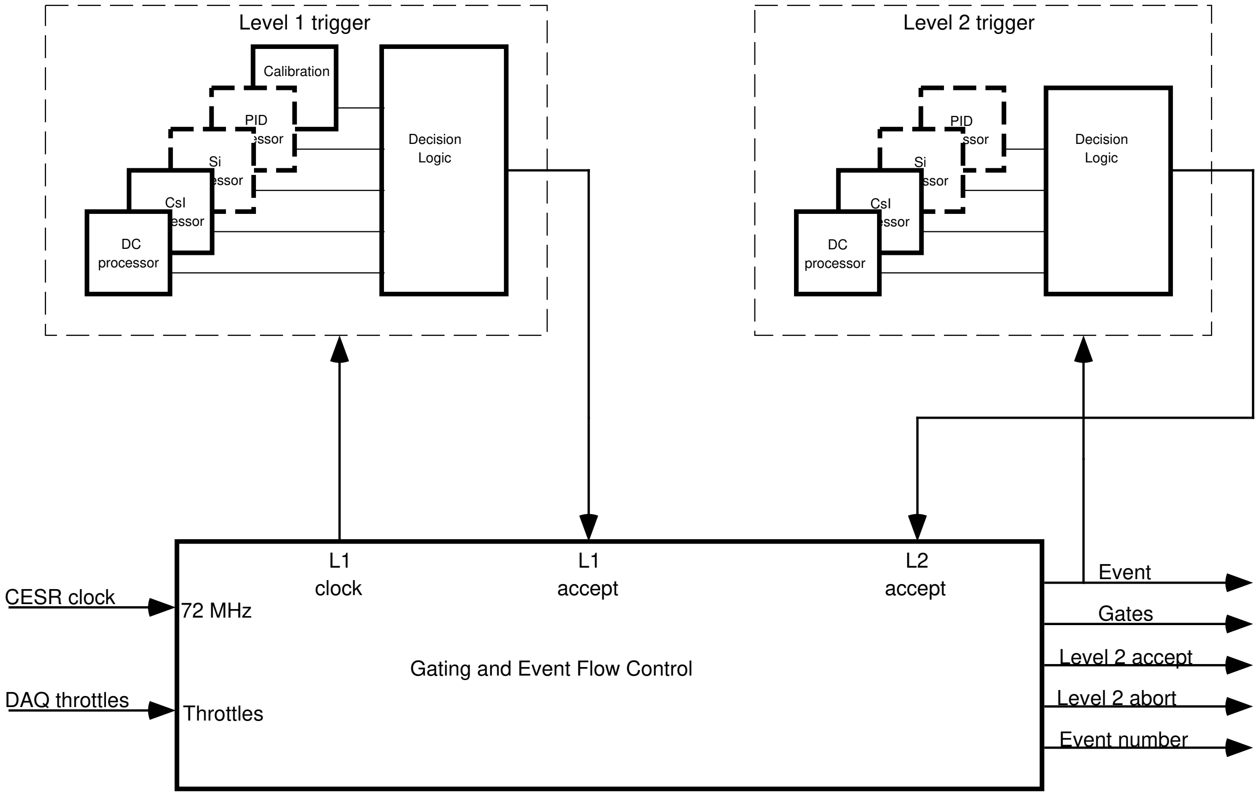

Figure 52: Block diagram of CLEO III trigger system showing level 1 and 2 triggers and their interaction with the gating and event flow logic. The silicon and particle ID systems are shown as contributing to both levels of the trigger but this is yet to be determined.

Table 28: Estimated performance of present CLEO II trigger at $\mathcal{L} = 10^{33}$ cm^{-2}s^{-1} with single-hit, ADC per channel front-end electronics. The level 2 decision time is assumed to be 10 μs.

	Input rate	Output rate	Rejection ratio	Latency time	Dead-time %
Level 1	18 MHz	200 Hz	10^5:1	0 μs	0
Level 2	200 Hz	100 Hz	2:1	10 μs	0.2
DAQ	100 Hz	100 Hz		0	0
Level 3	100 Hz	50 Hz	2:1	variable	0.0
Total					0.2

An additional feature required of the event flow control network is a mechanism to perform a back-plane reset in the VME and FASTBUS crates. This is usual to recover from processor "hangs" that require system resets.

5.2.3 Front-end electronics

New front-end electronics will be required for several reasons. First and foremost, we must significantly reduce the readout dead-time. Although the present hardware supports a rather cumbersome way of halving this time, the improvement of a factor of twenty that we seek can not be realized without a complete redesign of the front-end digitizing and buffering hardware. Secondly, waiting $\sim$2.5 μs after a good event's crossing for the trigger to assert dead-time means that single-hit electronics will effectively integrate over this time. This may not be acceptable in view of synchrotron-radiation background levels, especially in the inner regions of the tracking volume. The solution is to implement multi-hit front-end electronics, particularly in the regions of high occupancy, but preferably throughout the entire tracking system in order to maintain uniformity of architecture.

The design goal of the front-ends will be to digitize and buffer the entire detector in less than 100 μs, preferably less than 10 μs. This will give a readout induced dead-time of 1–10% at a 1000 Hz event-rate limit. This is acceptable for any foreseeable luminosity upgrade. Dead-time from a system having the trigger and front-end characteristics discussed in the above sections is shown in Table 28. Block diagrams of the front-end system components are shown in Fig. 53.

Drift chamber

A preamp will be mounted at the drift chamber endplates. This will provide some shaping and line driving ability. The signal is not discriminated here because we are interested in the wire charge for dE/dx measurements. To help reduce the mass at the endplates, monolithic preamps will be investigated. Smaller connectors and thinner cables to the crates boards

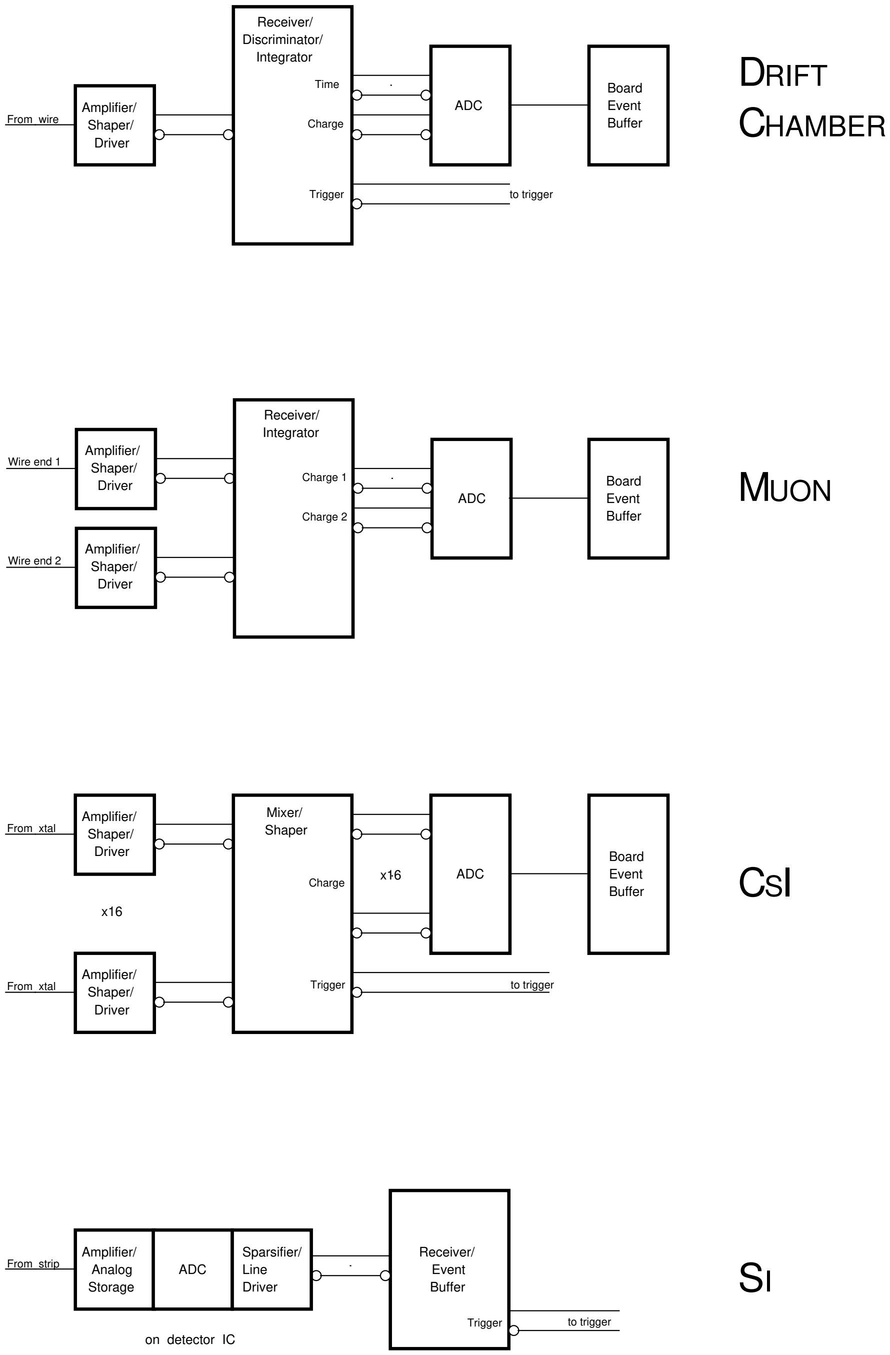

Figure 53: Block diagrams of CLEO front-end electronic systems.

will help to free up valuable space.

There are possible commercial solutions available for implementing a multi-hit front-end system. LeCroy has introduced an interesting timing circuit, the MTD133, that records the value of a 16-bit counter for up to 16 input transitions. A common stop signal records a final clock value and prepares the circuit for readout at which time the final clock value is subtracted from each of the up to 16 time recordings. The time resolution is 0.75 ns and the maximum depth is 40 μs.

Set up to record times on both rising and falling edges, the MTD133 can also be used as the counter in a Wilkinson run-down ADC. An experiment at CEBAF will multiplex timing and charge information into the same MTD133 channel using pulse width as markers. In future versions of the MTD133, LeCroy promises several additional features, including 0.5 ns resolution over a 32 μs range. Since Lecroy has not yet produced an MTD133 with the required performance, we must carefully evaluate other commercial solutions and the system used in other experiments.

The use of commercial TDC/ADC's will require a separate development of a system to split the preamp signals into timing and charge chains. The timing chain requires a discriminator whose output is compatible with the TDC. The charge chain may need to integrate the signal depending on the charge measuring scheme—TDC or ADC. The timing chain may need to provide the wire hit information to the trigger system. For some LeCroy FASTBUS modules this is available on the auxiliary connector.

Calorimeter

The present ADC system provides 12-bit resolution on two scales to give 15-bit dynamic range. It requires 750 μs to convert. It is not self-sparsifying, does not have multi-event storage, nor does it allow simultaneous read and write. Desirable features of any new ADC system are increased resolution, 14 bits with a 17-bit dynamic range, and lower coherent noise. The present ADC resolution is comparable to the noise and it would be nice to better resolve the pedestal to study coherent versus incoherent noise contributions. The preamps and mixer-shaper system will be retained with some re-work of the trigger section. To use the analog trigger sums at the mixer-shaper card level requires a new upper back-plane.

Multi-hit electronics for the calorimeter is not a requirement. The design of the level 1 trigger has been constrained to make a decision before 2.5 μs specifically for the purpose of gating the calorimeter conversion hardware, and we will use this to greatly simplify these front-ends. A single-hit system with sufficient dynamic range can be constructed by viewing each crystal with a pair of 14-bit ADC's in a dual-range configuration. The ADC outputs are stored in moderate depth FIFO's that are read out over a digital back-plane asynchronous to the trigger. The requirement that this system convert and buffer its data at the 10 μs time scale (for a reasonable cost) is not particularly onerous given recent developments in commercial ADC technology. While two ADC's per channel may appear excessive, it avoids the expense and bulk of fast, precision analog multiplexors. In addition, massively parallel digitizing and buffering at the front-end opens up the possibility of doing intelligent sparsification in the front-end crates, removing any bandwidth bottleneck presently associated with the "slightly parallel" DAQ90 sparsification scheme.

The use of the MTD133 for precision calorimetry is unclear at this time. With a dual range system where each charge is measured to 12-bit precision, a given crystal signal can occupy up to 4096 of the 65535 counts of the MTD133 time span. This essentially introduces an additional 3 μs crystal dead-time for each hit that starts the integrate/ramp-down process. This is probably not a problem since the occupancy of the crystals is small. There are other fast converting ADC systems commercially available.

Muons

A single range version of the gated integrator ADC's used for the CsI will work in the muon system. A preamplifier-shaper specific to the muon system will need to be developed.

Silicon

On-detector silicon readout electronics is always multiplexed due to space, density, and power dissipation limits. It can easily define the performance of the DAQ system. We require an analog-storage silicon front-end (for best spatial resolution) that converts and buffers its data in 10–50 μs and allows overlapped readout and acquisition. On-chip sparsification can help to meet this goal. Serial-bus interfaces to the IC's can reduce the bulk of the cable plant.

It is probably beyond the means of the CLEO collaboration to develop an entirely new silicon readout integrated circuit. Happily, this is an area of much activity in the HEP community. D0 and CDF are building new detectors based on the SVX-II circuit that provides analog pipelining, on-board digitization, and sub-10 μs readout. Two LHC R&D projects, RD2 and RD20, are developing analog pipelined, simultaneous write/read silicon front-ends. The time scale of these developments and others may be compatible with CESR phase 3 and it behooves us to leverage off these efforts. Modification of another experiment's electronics to CLEO specifications is definitely within the abilities of the collaboration. One area of customization is to provide strip OR's for triggering purposes. The selection of an existing front-end IC will depend on its compatability with the detector strip pitch, noise performance, and tolerance to an integrated radiation dose of $\sim$1 Mrad.

The crate-based data boards that will control and receive data from the detectors will probably be specific to CLEO. These entail sequencers, FIFO's, sparsification logic, and multi-event buffers.

Particle ID

Since a particle ID scheme has not been chosen, the design of the relevant front-ends is not possible at this time. We will simply decree that the parameters of such a system will have to fit within the constraints described above, namely that it must be able to tolerate a

2.5 μs delay between the actual event and the time it is "strobed" and that converting and buffering the data should take $\sim$10 μs.

Trigger systems

The trigger system has another role as a source of event data. Like the other front-end systems, it must buffer multiple events at the individual board level and at the crate level.

5.2.4 Data acquisition

The DAQ architecture must accommodate operation up to 1 kHz readout rates. Some aspects of DAQ will support this initially since this choice affects the design of the front-end systems and the data collection network. Other aspects of DAQ are scalable and upgradable. This is especially true of the event builder and level 3 trigger nodes.

Front-end issues

To minimize dead-time from the DAQ system, the front-end boards will need to provide a small amount of event buffering. The readout of these events must be able to occur in parallel with either recording data between triggers or with digitizing data from a later trigger. Only the conversion time of $\sim$10 μs drives the dead-time. Data transfer to the crate controller can take 1 ms on average and occurs asynchronously.

Local buffers can overflow if they are not emptied fast enough by the crate controller. A trigger throttling signal needs to be generated from each crate to indicate that the board-level buffers are approaching overflow. This feeds into the event flow control network.

If data conversion time is longer than a level 2 trigger decision time, the front-end board conversion must be capable of being aborted. Similarly, a mechanism for flushing level 2 rejected events from local buffers must be available. It is desirable that data conversion and event buffering all occur in less than 10 μs to give a dead-time contribution of 1% at 1 kHz. Conversion times as long as 50 μs would not be disastrous.

Crate level issues

The packaging of the front-end electronics will probably not be homogeneous. The trigger system and other CLEO-specific electronics will probably be implemented in VME or VXI. At the moment, the most cost effective commercial electronics for time and charge measurements are packaged in FASTBUS.

In either case, a processor will be resident in each crate. Note that most commercial front-end boards are self-sparsifying and that the crate processor will be relieved of this chore. Primarily, it will read an event's worth of data from the front-end boards, assemble them into an event fragment tagged with an event number, buffer it, and transfer it to the event builder/level 3 complex. Data is pulled from the front-end boards by the controller, which works from a list of outstanding events to process. Data is pushed from the crate controller into the data collection network as long as there is buffer space on the receiving end. Our present thinking is that the current overhead for pulling the data from the crates at the request of the event building is too high. This view may change.

As with the front-end boards, there is a finite amount of event buffer space available in the crate controller. A throttling signal needs to be generated from each crate to indicate that the buffer is approaching overflow. This feeds into the event flow control network.

For VME, there is an abundance of commercial processors and network interfaces. For FASTBUS, the choices are few. The most modern and flexible FASTBUS based processor is the MIPS-based FNAL FRC. In both cases, a uni-directional, high-speed data transfer port is required for the event data and a local area network (LAN) connection is required for configuration and control. Whether event flow control can be incorporated into the LAN is up for discussion. Additionally, VxWorks will become the real-time operating system of choice since it is required to support the FRC and more processors and I/O devices are supported under it than under our present pSOS^{+m}.

We observe that VME does not provide a mechanism for sparse scan readout, *i.e.,*slave boards cannot signal that all data has been transferred. Of course, this is not a problem in FASTBUS. The impact on the effective back-plane bandwidth has to be carefully understood.

Bandwidth

The present CLEO II readout event size is about 5 kByte. The size of events for an upgraded detector will depend on the kind of multi-hit electronics used, the front-end solutions adopted for the silicon and particle-ID components, and the overall detector occupancy. A conservative estimate is 25 kByte per event, spread over tens of front-end crates. The aggregate bandwidth at 1 kHz readout rate is about 25 MByte/s. The multiplicity of crates means that the required bandwidths for the crate back-plane and data link from the crate to the event builder/level 3 complex is modest. The impact of the full bandwidth is on the design of the event builder. Note that this bandwidth is estimated to be needed at luminosities approaching $\mathcal{L} = 10^{34}$ and need not be supported in an event builder at $\mathcal{L} = 10^{33}$.

Data collection network

A variety of data collection schemes have been considered: extensions to the CLEO II system, packet and channel oriented switches, and central shared memory. For reasons of cost, timeliness, and implementability, the shared memory approach is preferred. It is believed that the CLEO II system can not scale much beyond 100 Hz. ATM and Fibre Channel switch-based systems are not yet well enough established for us to commit to them. Ring-oriented SCI is also early in its life cycle. The useful bandwidth of FDDI may not be adequate. They, along with HIPPI, are also relatively expensive, especially when implemented in each crate. A further restriction is that the use of FASTBUS greatly limits the commercial interfaces available. In-house engineering of ATM or other complex network interfaces is beyond the collaboration's capabilities.

An alternative is to implement a point-to-point link from each crate to a shared memory system. Each point-to-point link terminates in a individual dual-port memory as in CLEO II. The link technology could be a version of the present CLEO II data link, LeCroy ECL-line, or a TAXI-like serial link on copper or fiber at a few hundred megabits per second. The FNAL FRC allows personalization of its data output port. The existence of what is already available for this port or what can be easily implemented will determine the link structure of CLEO III.

The centralized memory system is the terminus of that part of the DAQ system which must support 1 kHz operation with 25 kByte events at the time of commissioning. The rest of the system can be scaled with luminosity and event rate.

Unlike the problem of off-line computing, DAQ on-line computing is primarily an I/O problem requiring modest computing performance—1 Mips/event—but capable of absorbing data at 25 MByte/s and generating 6 to 12 MByte/s of output. Note that these are effective rates which will require at least a factor of two greater peak rates. A distributed level 3 system comprised of workstations requires the development of a full bandwidth event builder, an event distribution scheme, and interface to the workstation. The merging of the output data streams over conventional LAN technology is problematic even with CLEO II's modest output data rates. A single box solution providing good I/O and CPU performance is appropriate here. Multi-processor computers with 1000 Mips and easy-to-interface I/O are available today. The merging of output data is easy since there is no data movement from the box over a network.

We observe that powerful computers such as the SGI Power Challenge series and the DEC AXP processors with FUTUREBUS+ will allow the attachment of multiple VME crates. A single node terminus for the links can serve as event builder and as level 3 processor. Essentially, we merge the present DAQ90 hardware with the level 3 workstation. Initially, processor CPU cycles can move data from the dual-port memories. As data bandwidths exceed 10 MByte/s, it may become necessary to add VME-resident data movement engines to DMA the data into the host memory. A possible configuration of a memory-based system is shown in Fig. 54 along with a switch-based system.

The cost of the above processor is not included as part of the CLEO III upgrade. The interfaces to the system are included.

This approach limits the custom hardware requirements to the link hardware, possibly the dual-port memories, and the host software. We leverage the computer vendors' ability to interface to VME and to provide the CPU performance in an easily maintainable and programmable environment.

5.2.5 Slow controls

Less challenging than the event data collection system, but critical for detector operation,

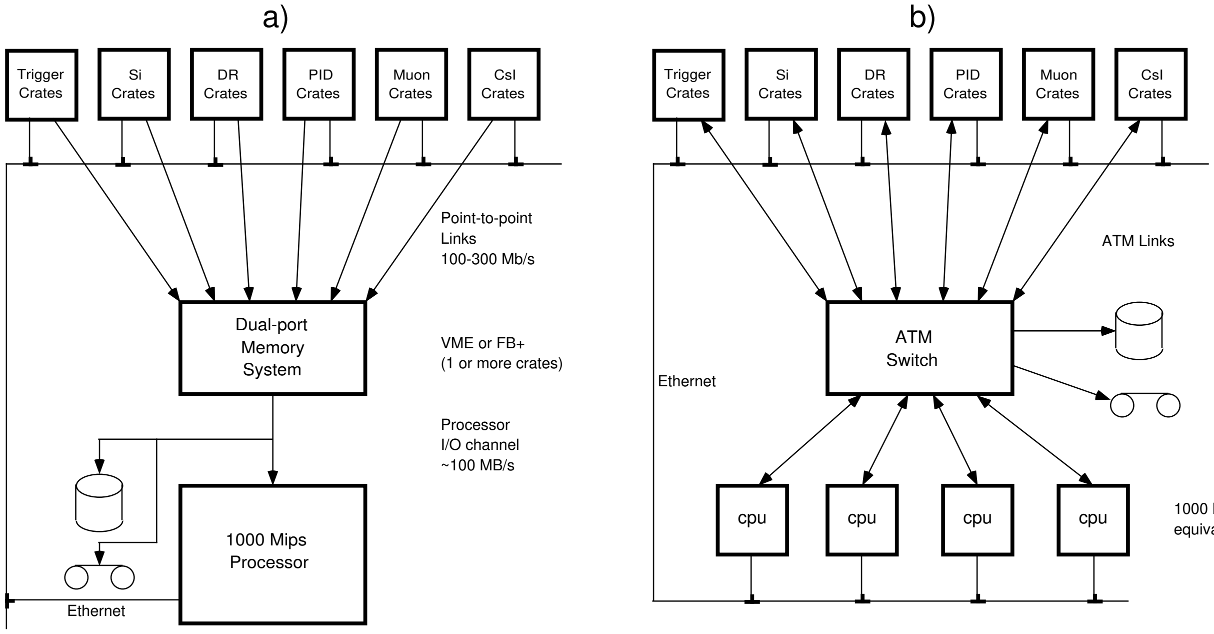

Figure 54: Block diagram of possible CLEO III DAQ systems: a) memory based and b) switch based.

is the slow control system for controlling and monitoring the myriad of support systems necessary for data taking. These include high and low voltage systems, gas mixing and monitoring, environmental conditions in and around the experiment, the cryogenics plant, and CESR conditions. A detailed inventory is too large to document here. We lump these together under the heading of slow controls. An important component of this is the shift operators interface. CLEO III will be moving away from VAX-based systems and their Q-bus. Many of the interfaces to the above slow control items will need to be re-invented. This presents an opportunity to take a new approach using distributed intelligence. For example, the present gas mixing system and the high voltage program share the same computer as the operator interface. These systems and others could be operated as autonomous systems using a distributed compute environment. This could be based on VME processors running a real-time operating system networked with Ethernet. Another approach would use PC's running Lab Windows or PC Spuds where the hardware is inexpensive, the development environment is friendly, and the graphical tools make complex, informative displays easy to code.

5.2.6 R&D

Silicon

- Front-end. Design and develop the on-detector integrated circuits which include the functions of preamp, analog memory unit, ADC, sparsifier, control/data links, calibration, and possibly trigger primitives.

- Data boards. This off-detector system is responsible for sequencing the on-detector electronics, receiving the sparsified data, buffering multi-events, and possibly delivering data to the level 2 trigger.
- Slow controls. Design the system for monitoring cooling, power, other environment quantities, and position.

Drift chamber

- Front-end. Design and develop a preamp/shaper/line-driver for mounting on the drift chamber endplate. Design and develop the calibration pulsing and distribution system. Design the cable plant between the on-detector electronics and the crate-based electronics.
- Data boards. Design and develop a line-receiver/discriminator/integrator and a trigger-primitive generator. Design or procure the TDC/ADC system.
- Slow controls. Design the low and high voltage distribution and monitoring systems, the gas mixing and monitoring system, and the environment monitoring system.

Particle identification

- Front-end. Determine if an analog or digital system is required and then develop the appropriate preamp, shaper, discriminator, and line-driver. If trigger information is desired, supply the requisite signal drivers.
- Data boards. Develop a line-receiver/discriminator and trigger-primitive signals if required. Procure an ADC/TDC system.
- Slow controls. Design the low and high voltage distribution and monitoring systems, the gas mixing and monitoring system, and the environment monitoring system.

CsI calorimeter

- Data boards. Procure the ADC system. Modify the digital trigger bits if required and implement physical trigger quantities based on available analog trigger information.
- Slow controls. Interface the mixer/shaper crates to the control system and provide monitoring for low voltages, crate and preamp cooling systems, and other environment monitors.

Muon

- Front-end. Design and develop a preamp/shaper/line-driver— this may be identical to that of the drift chamber but mounted near the data boards. Design and develop the calibration pulsing and distribution system.

- Data boards. Procure the ADC system.
- Slow controls. Design the low and high voltage distribution and monitoring systems, the gas mixing and monitoring system, and the environment monitoring system.

Trigger

- Level 1 Processors. This level of the trigger will require development of processors for tracking and calorimetry and the decision logic. Simulation of the trigger logic, development of codes for inclusion in the physics Monte Carlo, and testing on simulated physics and background events will be done.
- Level 2 Processors. This level of the trigger will require development of processors for tracking, calorimetry, possibly silicon, and possibly particle identification, and the decision logic. Simulation of the trigger logic, development of codes for inclusion in the physics Monte Carlo, and testing on simulated physics and background events will be done.
- Level 3. This software trigger requires the development and simulation of algorithms.

Flow control

- Clock/gating system. Develop a uniform clock and gating distribution system and interface to the CESR timing system.
- Event and ID distribution. Develop a uniform mechanism for distributing event triggers, event tags, and event aborts to the data crates.
- DAQ throttles. Develop a uniform mechanism for collecting the status of the various event buffers and inhibiting further triggering as indicated.
- Dead-time monitors. Provide a mechanism for measuring the deadtime induced by the trigger and data acquisition systems.
- Configuration software. Develop code to configure and manipulate the gating and throttling system.
- Simulation. Simulate the system to establish its viability and integrity.

Data acquisition

- Crates. Establish a small set of crate formats for the crate-based electronics. Provide an environment monitoring port to the slow control system. Establish the location of the crate based electronics and any special housing requirements.

- CPU and OS. It is assumed that each data crate will have a local processor. Select a small number of processors and a common operating system, develop coding standards and code. Select a LAN technology for non-data communications with the processors.
- Data links. Develop a common data link protocol and hardware for moving event data from the crates to the event builder memories. Develop the event builder buffer memories.
- Online computer. Establish coding standards and generate event building and level 3 trigger codes. Allow partitioned readout of the experiment for simultaneous maintenance of the sub-detector systems.
- Trigger throttles. Generate signals for the event flow control throttling network.

Slow controls

- Development environment. Provide a uniform coding environment for all slow control code development.
- Standard alarms/limits package. Provide uniform and rugged alarms and limits utilities.
- Cryo interface. Interface to the He cryogenics plant.
- CESR interface. Interface to the CESR control system to share information on quantities of mutual interest.
- Operator interface. Provide the interface for the shift operators including run control, alarms, and displays of time trends of detector systems.