Chapter 6

Computing

6.1 Data Acquisition System and Trigger

The primary feature driving the redesign of the CLAS^{++} Data Acquisition (DAQ) and triggering system is the higher data rates associated with the approximate factor of 10 increase in instantaneous luminosity that will be delivered to the experiments at Hall B.

6.1.1 DAQ and trigger upgrade motivation

A fast and flexible DAQ is the cornerstone of a modern nuclear or high energy physics experiment. Similarly, DAQ performance highly depends on the quality of the triggering system. The JLAB 12GeV upgrade poses unique challenges to both the trigger and data acquisition systems for the CLAS^{++} detector.

Taking into account that the current DAQ system is operating at the limits of its capabilities, fulfilling the requirements of the CLAS^{++} DAQ without a major upgrade of the existing system will be a difficult job. During the years of the CLAS detector operation, the DAQ system was upgraded and optimized both in hardware and software to meet increasing requirements of the ongoing CLAS experiments. As a result, system performance exceeded the initial design criteria by almost factor of 3 and demonstrated that the present system performance is solely limited by the front-end digitization hardware.

Three main factors prevent further improvements of the existing system. First, the slow digitization sequence in the CLAS high precision TDCs (Lecroy 1872/75. 10000ns+2500ns/hit conversion time), second , the slow converting CLAS ADCs with 12000 ns per module conversion time and third, data rate limitations ($<$ 23Mbyte/s). The data rate problems can be solved by utilizing permanent storage units with fast access times, and by using different storage management software. Also, higher level software triggers (level3) can be implemented to reduce the data rate to disk. In order to reduce front-end dead time limitations, the high resolution TDCs and ADCs must be replaced.

6.1.2 Design criteria for the system

The CLAS^{++} DAQ must be able to handle level1 trigger rates up to 20kHz at 10kByte/event event size with less then 15% dead time. High data transfer rates over the network will be handled by using gigabyte Ethernet. The proposed DAQ system must assemble the event data from many front-end buffers to the Event Recorder (ER). The system must provide sufficient computing resources for executing physics algorithms which can substantially reduce the expected input rates. It must also provide continuous monitoring of the detector so that malfunctions may be readily identified and corrective measures taken. These functions will be performed by using a high performance readout network to connect the sub-detector readout units (ROC) via multiple event builders (EB) to the event filtering units (EFU, possibly implemented in computer farms). The flow of event data will be controlled by the triggering and timing system (TTS).

6.1.3 System upgrade

CLAS^{++} DAQ system design requirements are based on the rate estimates of the 12GeV leading physics programs. Current system will be upgraded to meet the mentioned design requirements and to minimize the efforts, which will be invested to satisfy increased requirements of 12GeV future physics programs. Following upgrade strategy will be used to achieve desired goal.

- Utilizing pipelined digitization hardware for the CLAS^{++} new detector components.

- Replace the CLAS existing detector components: ADCs and high resolution TDCs with the equivalent pipelined digitization hardware, ready to function in the free-running DAQ mode.

- Keep existing low resolution TDC (TDC1877/1877S) in the system.

Mentioned hardware upgrades, with conjunction of the fast and optimized readout system, will guarantee fulfillment of the CLAS^{++} DAQ design requirements. Using pipelined digitization hardware components in the system will provide most of the benefits, typical for the free-running DAQ systems. Gradual replacement of the remaining hardware (low resolution DC TDCs) with the pipelined equivalent will help us eventually implement full operational free-running DAQ system for future CLAS experiments.

6.1.4 Overview of the proposal

The system including free-running, pipelined DAQ front-end components is an ideal solution, which satisfies all CLAS^{++} data production requirements, and gives us flexibility for future improvements. Proposed system has the following benefits. First, pulse shapes are easily captured and stored for immediate or delayed analyses. Second, the system will have an enhanced ability to delay signals while accurately preserving time information between different events. A digital delay line, for example, is just a FIFO memory which can easily be 10s of microseconds long, with perfect signal fidelity. This property of the DAQ system will be extremely profitable for the CLAS^{++} new detector components, saving time, space, and money by eliminating signal delay cables. Finally, data processing is extremely flexible. This means that signal processing algorithms (level1, level2) in either firmware or software are easily modified so that different, triggering algorithms can be added without any hardware reconfiguration. However, before complete replacement of the low resolution TDC's, flexibility of the low level triggering system will be somewhat limited. The speed requirements of the low level triggering system will be defined by the TDC1877 multi-event buffer memory depth.

Proposed system will open areas of new capabilities for CLAS^{++}, for example:

- Pulse specific corrections and shape analyses to provide particle identification.
- Inner tracking chamber transient signal analyses. Analyses of the induced charged signals can provide interaction location information.
- Hierarchies of processing complexity, which readily support simple-fast (level1) versus complex-slow (level3) decision making on an event by event basis.
- Complex coincidence criteria, which will support capturing and/or processing data either immediately or after delay.
- Digital communication, which will enhance operating convenience by allowing remote operation, digital hardware calibration, and restoration of previous setup data from files or databases.

Figure 6.1 shows possible data flow diagram of the CLAS^{++} DAQ system. Here is a possible scenario of implementing higher level triggering algorithms, which demonstrates flexibility of the proposed system. In order to optimize the data flow, the event filtering computer farm performs event selection in two stages. First, a level2 filtering decision is made on a subset of the data from a programmable set of the detector components. This will help us to avoid system bandwidth saturation by reading out large volumes of tracking data at high level1 rates. The remainder of the full event data are only transferred to the filtering farm on a level2 accept, and a final level3 algorithm is then applied to the complete event. This activity is controlled by the Filtering Supervisor (FS) system.

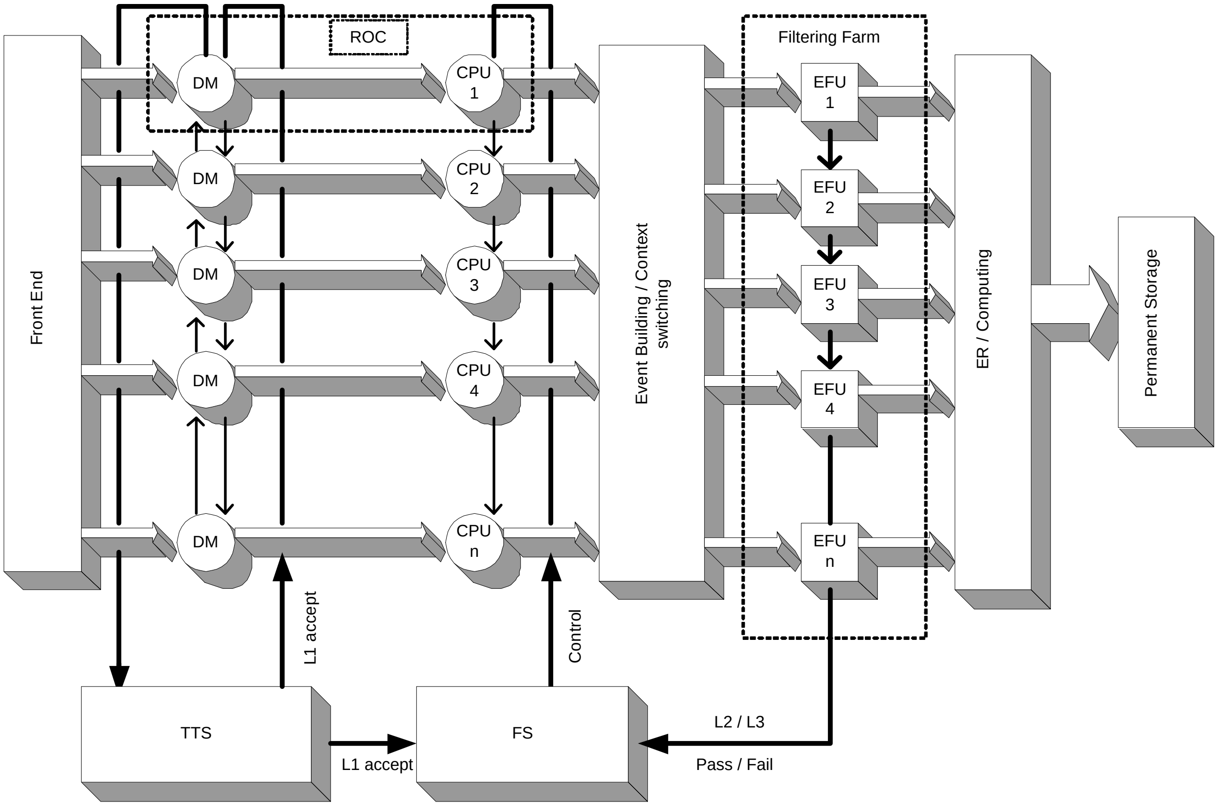

Figure 6.1: CLAS^{++} DAQ bloc diagram. DM - Digitization Module, TTS - Triggering and Timing System, FS - Filtering Supervisor, EFU - Event Filtering Unit, ER - Event Recorder.

The FADC data can be used to form the level 1 trigger by utilizing a tree of digital adders to sum information from various channels in the detector system (Figure 6.2).

If we assume a reasonable time window of 100ns for FADCs with 8bit sampling depth and 250MHz sampling frequency, we would expect 25bytes/FADC channel to get energy and timing information. Taking into account approximately 5300 FADC channels, and assuming 2% occupancy, the data volume from all FADCs will be 2.6kBytes (no headers included). However, using the FADC built-in computational resources (digital signal processors or DSP) one can perform FADC data reduction in real-time. Currently we are studying a FADC prototype developed by the Hall-D collaboration to determine time, energy and spatial resolutions possible to get using that particular module. For the pulses, corresponding to the very small energy depositions in the calorimeter, FADC sampling depth and sampling frequency will be simulated to determine the optimum FADC design. We hope that at the time of the CLAS^{++} DAQ commissioning there will be FADC with higher sampling depth and frequencies in the market within current estimated price range. Otherwise conventional fast ADCs will be used (currently available 6000nsec conversion time in the same price range), with the conventional triggering system. This backup plan will satisfy CLAS^{++} DAQ design requirements, however it will complicate further

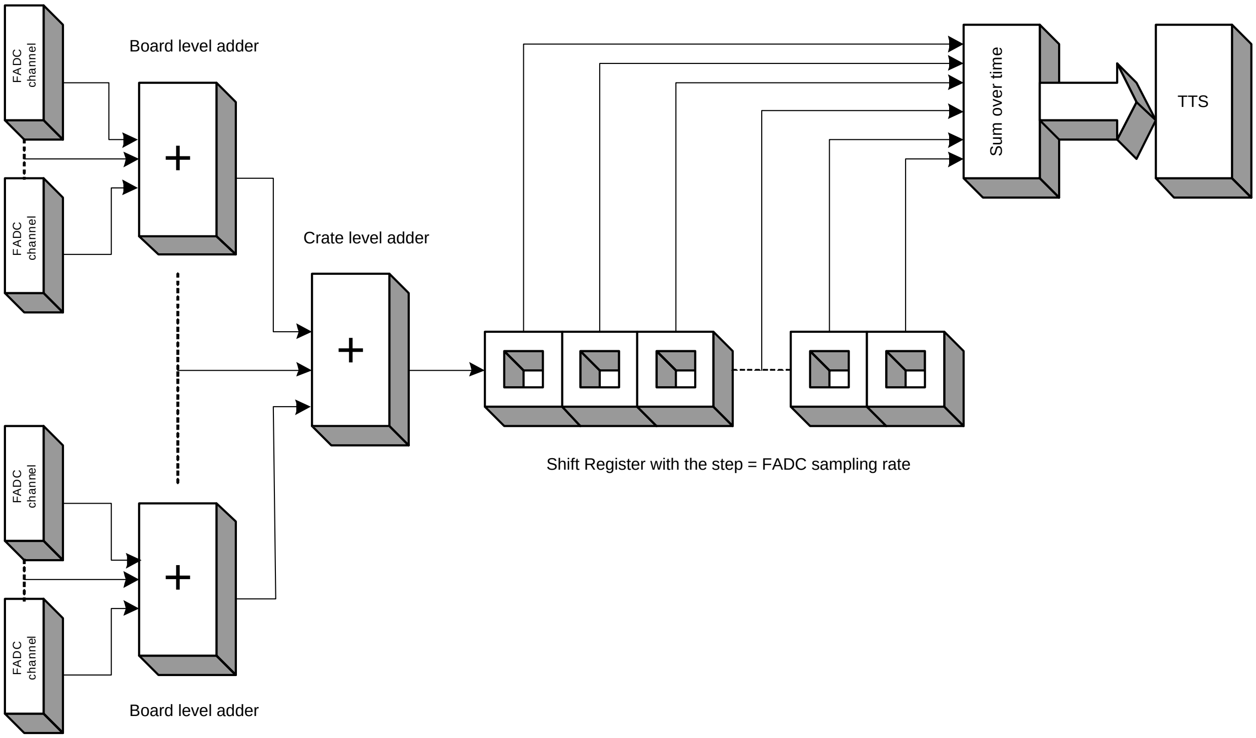

Figure 6.2: Level1 triggering system based on flash ADCs.

improvements of the entire system.

For precise time measurements we will use multi-hit TDCs. Currently, a high resolution multi-hit TDC based on the F1 ASIC, which was developed at Jefferson Lab for Hall-D collaboration, is at the test stage. This TDC design can provide up to 60ps of resolution and will sustain high rates.

The manufacturer of the digitization hardware has not been finally determined. However, we prefer digitization hardware developed in-house for Hall-D over the commercially available alternatives due to channel and maintenance cost considerations.

6.1.5 Level1 triggering and timing system

The TTS is responsible for selecting physics events and for suppressing background as efficiently as possible. It will adopt a real-time, parallel, pipelined architecture for the trigger electronics. High speed triggering algorithms will be designed and implemented in Field Programmable Gate Arrays (FPGA). This approach will help us to minimize the trigger latency and shrink DPRAM depth of the digitization modules. After getting a trigger decision, the TTS will initiate detector readout by sending an L1Accept signal to the ROCs. The TTS will also be responsible for injecting calibration pulses as well as SYNC pulses used to resynchronize the entire DAQ system. It is obvious that we will need a series of trigger simulations to help define trigger algorithms and understand the real requirements of level1.

6.1.6 Higher level triggering systems

Level2 and Level3 filtering processes are performed by the computer farm processors. Currently, one full event analysis using RECSIS, with momentum resolution of less then 1%, requires 40miliseconds on a 500MHz Pentium III processor (21SPECint). It is clear that at the trigger level we do not need high resolution momentum reconstruction. Based on a recent development of the fast reconstruction program, we are quite confident that we will be able to analyze the events at level3 with less then 5% momentum resolution, spending 3-4 milliseconds on a 21SPECint processor. Therefore, filtering 50KHz events at level3 will need 4200 SPECints processing power. This can be achieved with a filtering farm, composed of 42 computers each with 200 SPECints of processing power (50% cpu utilization is estimated, due to I/O overhead, brakes, etc.).

The FS will be designed to maintain coherency between the two filtering stages. It will be important for error detection too. Several recovery protocols will be defined to restore and maintain system synchronization and data flow. The Level3 pass will finally send events for recording, and monitoring.

6.1.7 Control and monitoring

The control and monitoring system must be able to deal with network distributed hardware components and their heterogeneous software environments. This system must have the ability to efficiently retrieve, organize, and manage information from widely dispersed sources within as well as from outside the CLAS^{++} control environment. Feedback systems between the DAQ, trigger and slow control systems will be of paramount importance.

Statistical information about data taking, as well as detector performance will be accumulated and visually presented. On-line databases will store run conditions, configuration parameters and calibration constants for later processing and evaluation. Data integrity will be checked continuously. Physics performance histograms will be accumulated and presented. The system will detect, record, and analyze error conditions. Serious conditions will generate informative alarms for operators.

In order to achieve scalability and robustness, the control and monitoring system will be designed under a single framework. This will allow interconnection and interoperation with multiple legacy systems (EPICS, Smart Sockets, etc.). The dynamic and distributed nature of both data and applications for the CLAS^{++} DAQ will require that software not merely respond to requests for information but intelligently anticipate, adapt, cooperate, and actively seek ways to optimized the performance of the entire system. Thus the control system will be designed using software agents. A very basic definition of an agent is atomic entities that communicate to implement the functionality of the control system with the following properties: autonomy, reactivity, social ability, proactivity, temporal continuity, goal orientedness, mobility, collaboration and adaptivity.

To achieve these goals using a traditional client-server approach will be difficult. The agent paradigm is different then the client-server approach, since agents can interact on a peer-to-peer level, collaborating, and cooperating to achieve their goals.

Agents engineering aspects will be addressed by adopting the domain independent software standard formulated by FIPA (Foundation for Intelligent Physical Agents). Currently, a control software framework based on agent technology is under development by the JLAB Data Acquisition group.

6.2 Event Reconstruction and Offline Computing

The expected on-line data rate of 100 to 200 MByte/s cannot efficiently be handled when using the same procedures adopted by CLAS in the past: an efficient online event reduction is required to keep the data transfer rates to the Jefferson Lab tape silo within the bandwidth limits. Furthermore, the on-line event reduction and event processing will allow for a comparatively fast access to processed data for further analysis. We consider establishing a “grid-based” cluster of analysis and simulation centers to optimize the data processing.

6.2.1 Data Reduction and On-line Event Reconstruction

Major parts of data reduction will have to be performed on-line. The first stage of data reduction (noise reduction) will be performed at the crate level, the second stage in form of fast event filtering and tagging of events (Level 3) which will remove data which are not of interest for further analysis.

A fast on-line reconstruction will provide sufficient information on data quality and first-pass analyses of basic reactions. Depending on the available CPU power it is possible to perform a full event reconstruction of all events: using an on-line farm in the Hall-B Counting House and/or part of the JLab CUE farms which requires either a second output stream or large pre-silo stage disks. Taking into consideration the current achievements (7-10 ms/event) and the projected increase of CPU speed over the next 6 years, a full event reconstruction can be performed within 3-4 ms despite the more complex detector setup of CLAS^{++}. The output will be written to disk in form of data summary files which contain all information required to perform a first-pass event analysis.

The (quasi-) on-line event reconstruction requires zero order calibration constants which have to come from analyses of previous run periods (and commissioning data) as well as the analysis of data taken during the first days in the specific run period. Additionally, small subsets of data will be selected for further analysis to perform a quasi-online calibration of all detector components. This continuous calibration process will provide a quick response to any changes in crucial detector parameters. In parallel, these subsets will be transferred to university-based analysis centers which will be responsible for a refinement of the on-line calibration.

The data summary files as well as the data subsets to be used for calibration purposes will be stored for several days on large RAID arrays to allow for fast access. We expect that the Jlab Computer Center will provide the necessary infrastructure of fast network connections as well as short- and long-term storage media.

6.2.2 Offline Data Analysis

Since the first-pass reconstruction is performed (quasi-) on-line, it is convenient to distribute the data summary files to university-based analysis centers for further analysis: refinement of detector calibrations and quality controls, preliminary physics analyses. The results will be used as input for a final (second-pass) reconstruction process which can most efficiently be performed at Jefferson Lab. This strategy enables collaborators in their home institutions to contribute effectively to the analysis process. It requires to create appropriate software that minimizes the efforts to complete these – mostly labor intensive – tasks. A grid-based computing environment including IT services customized for scientific computing will be developed in close collaboration with the Hall-D efforts and groups in the high-energy physics community.

6.2.3 Event Simulation

For most experiments to be performed in Hall B, the quality of the results will be limited by systematic uncertainties rather than statistical errors. Therefore it is crucial to understand the acceptance and efficiency of the detector very well. Extensive Monte Carlo simulations play a key role to this end. They will be distributed over several university-based simulation centers extending the use of these institutions in the current CLAS projects. An effective infrastructure has to be developed to manage and coordinate the simulation centers (remote submission, job tracking, communicating of results, simulation database).

Storage and accessibility of simulated data will be an important issue since we expect that at least four times more simulated than real data will be required to ensure high precision of the results. With respect to the pressure on the Jlab tape silo it may be more convenient to add appropriate storage media to the simulation centers.

The design and construction of the new detector parts will be accompanied by the development of corresponding simulation software which will have to be continuously adjusted to an optimal description of the real detector in later years.

6.2.4 Software Development

The development of simulation, reconstruction, and analysis software requires a group of experienced programmers. To optimize the process, parts of the CLAS software as well as techniques and experiences of high-energy physics and other nuclear physics communities will be adapted to the needs of CLAS^{++}. It will be essential to integrate the on-line and offline software efforts as the calibration and event reconstruction will run quasi-online which requires speed-optimized, multi-threaded code.

University-based groups will be involved in the development and maintenance of software, especially those groups who take over the responsibilities to create an analysis or simulation center. The software efforts will be coordinated by a small group of physicists/programmers who are familiar with modern programming techniques and tools as well as the physics program of CLAS^{++}.