The CDMS II Data Acquisition System

D.A. Bauer,[b] S. Burke,[e] J. Cooley,[c] M. Crisler,[b] P. Cushman,[f] [1]
F. DeJongh,[b] L. Duong,[f] R. Ferril,[e] S.R. Golwala,[a] J. Hall,[b]
D. Holmgren,[b] R. Mahapatra,[d] H. Nelson,[e] A. Reisetter,[f]
J. Sander,[d] C. Savage[e]

[a]*Physics Dept., California Institute of Technology, Pasadena, CA 91125, USA*

[b]*Fermi National Accelerator Laboratory, Batavia, IL 60510, USA*

[c]*Department of Physics, Southern Methodist University, Dallas, TX 75275, USA*

[d]*Dept. of Physics & Astronomy, Texas A & M University, College Station, TX 77843, USA*

[e]*Department of Physics, University of California, Santa Barbara, CA 93106, USA*

[f]*School of Physics and Astronomy, University of Minnesota, 116 Church St SE, Minneapolis, MN 55455, USA*

Abstract

The Data Acquisition System for the CDMS II dark matter experiment was designed and built when the experiment moved to its new underground installation at the Soudan Lab. The combination of remote operation and increased data load necessitated a completely new design. Elements of the original LabView system remained as stand-alone diagnostic programs, but the main data processing moved to a VME-based system with custom electronics for signal conditioning, trigger formation and buffering. The data rate was increased 100-fold and the automated cryogenic system was linked to the data acquisition. A modular server framework with associated user interfaces was implemented in Java to allow control and monitoring of the entire experiment remotely.

Key words: DAQ, data acquisition, CDMS, dark matter, WIMP, Soudan
PACS: 07.05.Hd, 95.35.+d

[1] Corresponding author: Priscilla Cushman, School of Physics and Astronomy, University of Minnesota, 116 Church St SE, Minneapolis, MN 55455 USA. Email: prisca@physics.umn.edu

1   Introduction

The Cryogenic Dark Matter Search (CDMS) experiment is designed to detect Weakly Interacting Massive Particles (WIMPs) via their interactions with germanium and silicon nuclei inside solid state detectors. WIMPs are a well-motivated candidate for the dark matter which makes up most of the mass in the universe. As the Earth moves through the dark matter cloud which surrounds our galaxy, the WIMP particles pass through our detectors and can occasionally interact. Direct detection of this WIMP 'wind' would confirm indirect evidence of dark matter from astrophysical observations.

The CDMS II [1–3] configuration of 30 germanium and silicon detectors ran at the Soudan Underground Lab in northern Minnesota until 2009, now replaced with the SuperCDMS detectors. Particle interactions in each 7.6 cm diameter by 1 cm thick crystal deposit energy via both ionization and phonon production. The resulting charge drifts under a 3V bias and is collected on one detector face by an inner electrode disk and an outer fiducial ring. The other face is divided into four phonon quadrants read out by transition edge sensors. This results in two charge and four phonon signals. The detectors are maintained below 50 mK in a cryostat cooled by a helium dilution refrigerator. The first stage of amplification occurs within the cold volume via SQUIDs for the phonon signals and FETs for the charge signals. Copper-kapton cables carry signal and control lines from the 4K stage of the cryostat to a room-temperature vacuum bulkhead. On the warm side of the bulkhead, 9U VME front-end boards provide control and additional amplification of the signals, which are read out by waveform digitizers. Global trigger signals are formed by a separate set of trigger boards.

The cryostat is surrounded by a layered shield of copper, polyethylene and lead. Just outside of this shield is an active veto against cosmic rays, consisting of 5.2 cm thick plastic scintillator panels read out by photomultiplier tubes and waveform digitizers. A time history of veto hits is stored with every triggered event to allow identification (and subsequent rejection) of any event that is within a time window of 200 mus around a cosmic ray or above-threshold activity in a panel.

The detectors are calibrated using radioactive gamma (^{133}Ba) and neutron (^{252}Cf) sources that can be inserted as needed. The normal background trigger rate is 0.2 Hz, dominated by environmental gamma rays, but the calibration rate depends on the strength of the radioactive source and can be hundreds of Hz. Each background event contains approximately 1 MB of data, while each calibration event has about 50 KB of data due to selective readout of the detectors. The volume of data per day averages about 50 GB. The data is stored on SuperDLT 350 GB tapes and simultaneously transmitted over the

network to a surface computer farm and then to Fermi National Accelerator Laboratory for further processing. The data acquisition system described in this paper covers the system from the front-end electronics through the chain to the raw data output, including both hardware and software.

2 Design Requirements

CDMS II is a rare event search experiment with an expected background rate of $\sim$0.2 Hz from gammas and betas. High DAQ trigger rates only occur when taking calibration data with in−situ radioactive sources. In order to accumulate a large amount of calibration data in a short amount of time, the DAQ must be able to operate at a rate of $>$ 10 Hz. This ensures that our livetime in the WIMP search mode remains at $>$ 80%.

The CDMS II DAQ system was designed to digitize analogue signals from both detectors and muon veto panels at this relatively high rate and provide flexibility for upgrades, different trigger conditions, and the addition of more detectors. Since it is operated in a mine with limited access, the control and monitoring must be remotely accessed, with appropriate safety protocols and interlocks. It must be capable of autonomous operation, even when control and monitoring functions are interrupted. This is achieved by use of a server/client framework in which DAQ control servers run locally on computers physically located in the underground lab, while the client interfaces may be run from any other computer connected via ethernet to the mine network. While the need for fast readout necessitated a LINUX−based system at Soudan, many of the CDMS test facilities use MAC OS or Windows software. Thus, the DAQ was designed to be portable across different computer operating systems.

3 Systems Implementation

3.1 Overview

These design requirements lead to the general architecture for the CDMS−II DAQ displayed in Figure 1, where the data flow is upward from the detectors (labeled ZIP for Z-sensitive Ionization and Phonon) to the EVENT BUILDER, and the user control is downward from the user via RUN CONTROL.
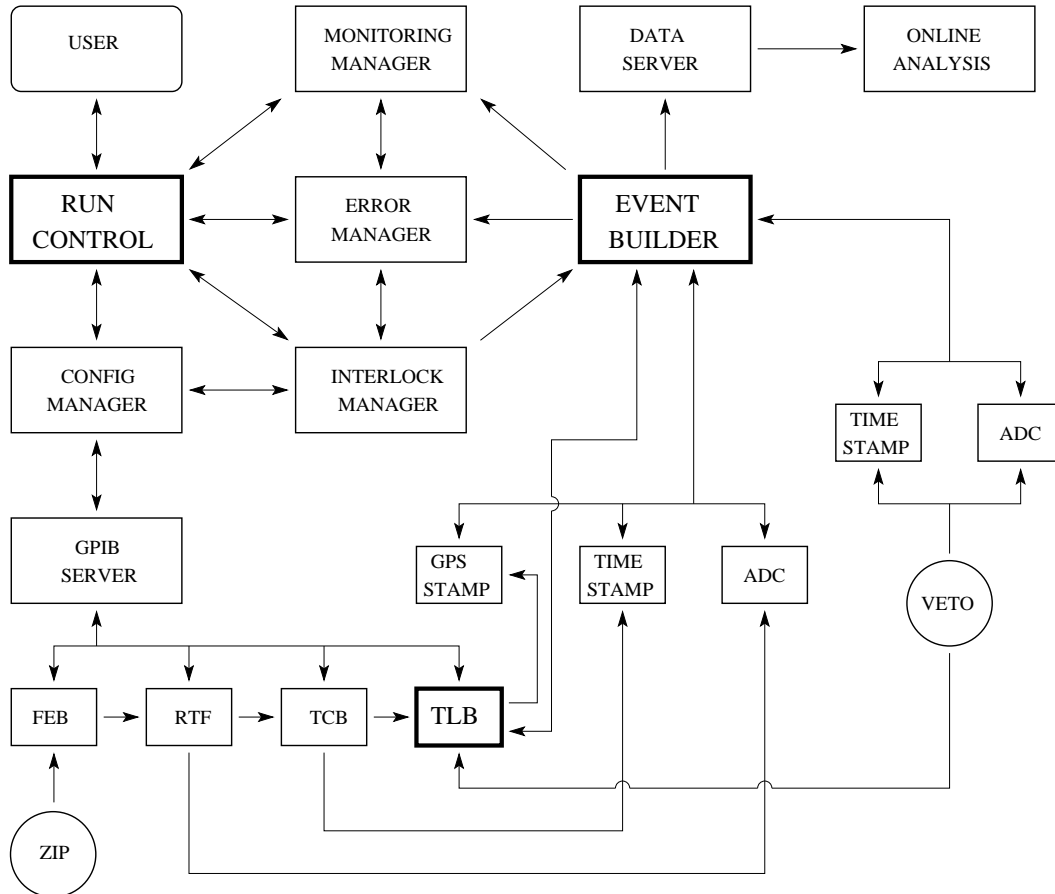
Fig. 1. Overall schematic of the CDMS II DAQ system showing data flow and user control sequencing. The FEB, RTF, TCB, TLB are custom-built electronics boards which condition the signals and form the trigger. The user interacts with the main functional block, RUN CONTROL, and fast data processing takes place in the EVENT BUILDER block.

Data Flow: Each of the 30 detectors delivers six electrical signals (four phonon and two ionization) which exit the cold vacuum cryostat via the electronics box and enter the Front End Boards (FEBs). The signals are then carried by twisted pair cables to the Receiver/Trigger/Filter board (RTF) where they are filtered and then read out by waveform digitizers (ADC). The RTF also sums ionization and phonon signals to form logic signals for the final trigger.

The trigger logic board (TLB) then uses the logic signals from all the RTF boards to create an experimental global trigger. The resulting logic signal is used to set the experimental time reference and initiate data transfer from the waveform digitizers. Structured raw data events are formed and written to a file, after which the TLB is re-enabled for the next trigger.

Control Blocks: The user interacts with the DAQ system through RUN CONTROL, which has the following functions:

(1) Interaction with the custom electronics boards. This is done through CONFIG MANAGER via the GPIB SERVER, since the boards use a General Purpose Interface Bus (GPIB) for communication across crates and slots.

(2) Management of the EVENT BUILDER. Notice in Figure 1 that the semaphoring between RUN CONTROL and EVENT BUILDER is through the INTERLOCK MANAGER and then back via the ERROR MANAGER. In this way, RUN CONTROL can configure the electronics and detectors for data-taking and then semaphore the EVENT BUILDER for a data run.

(3) Interfacing to the computer that controls and monitors the cryogenics system, enabling fully automatic data-taking, coordinated with independent management of the cryogenics system.
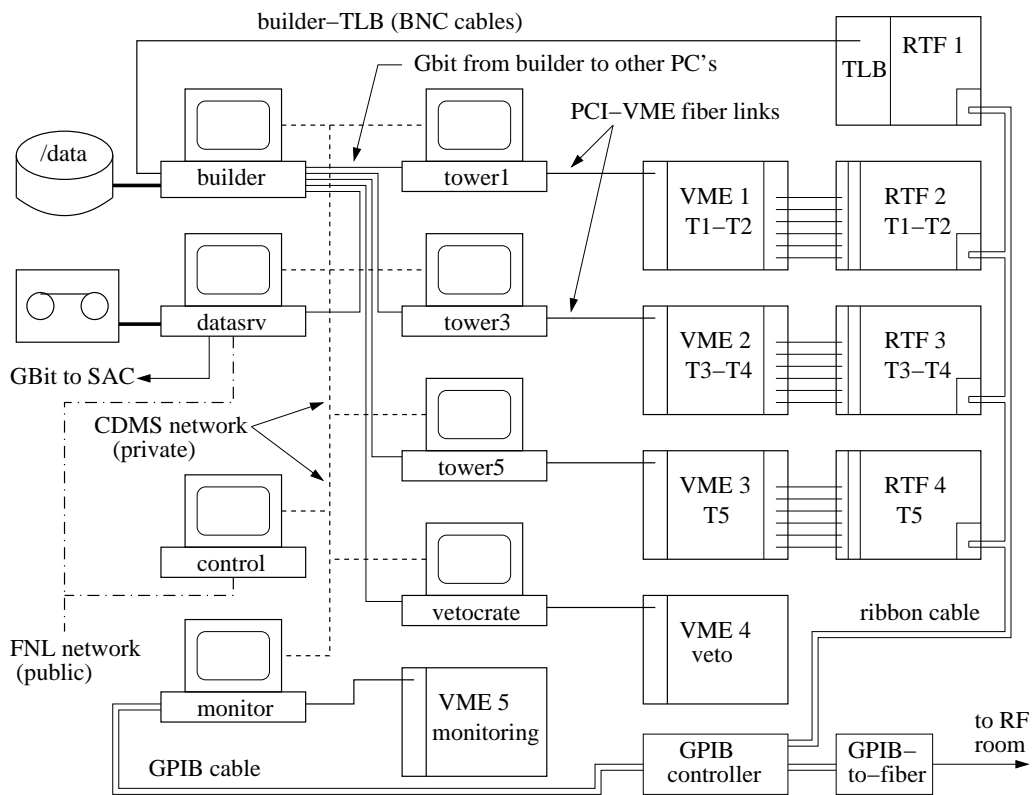

Fig. 2. Layout of the CDMS II control room at the Soudan Underground Lab, showing the information flow between the physical computers and the crates which contain the custom electronics.

Figure 2 shows the layout of the physical computers and electronics inside the control room at the Soudan Underground Laboratory. Generally speaking, the functional blocks from Figure 1 match specific computing or electronics components of Figure 2 in an intuitive way. For example, both the

CONFIG MANAGER and GPIB SERVER operate on the physical computer named `monitor` since it has a GPIB hardware interface to the GPIB controller box. The EVENT BUILDER functional block is actually a distributed LINUX−VME system, which is comprised of the various `tower` computers, which access commercial ADCs via a fiber link to their VME controller.

3.2 Detector Electronics FEB/RTF

Each detector is associated with its own Front End Board (FEB) in a custom Wiener [5] UEV-6021 9U crate located in the same RF-shielded clean room. Copper-kapton flex transmission lines transfer the primary signals from the combined SQUID and FET card within the shielded cold volume, and a meter of twisted pair cable delivers the signals to the FEB at room temperature after passing though a vacuum feed-through panel. The FEB amplifies the four phonon and two ionization signals, and then sends the amplified signals to the RTF in the control room via a shielded differential twisted-pair cable through two common-ground patch panels in the walls of the clean room and control room respectively.

The FEB also controls a variety of important detector functions, namely:

- Bias voltage across the germanium detector
- SQUID bias and lockpoints for the phonon signal
- FET biases for the ionization signal
- FET heating controls at startup
- Parameters related to the automated infrared flashing process used to remove charge trapping centers.

The Receiver-Trigger board (RTF) conditions the signal by sending it through a 0.5 MHz low-pass filter and providing a -2 V, offset in order to utilize the full dynamic range of the waveform digitizer to which the filtered signal now passes. A custom flat coaxial cable is associated with each set of six stacked detectors, called a tower. This cable plugs into the front of the RTF and its six lemo connectors on the other end plug into the Struck SIS3301 FADC waveform digitizer.

The RTF board also has logic and monitoring capabilities. Each board determines if the signal is above threshold by forming an analog sum of the phonon and ionization pulses independently and comparing them to user-specified high and low thresholds. The summed phonon pulse is compared to an ultra-low threshold as well. Signals above threshold generate a trigger pulse, which is passed to the TCB across the backplane of the shared crate. The TCB stretches and converts the RTF logic signals and sends them to a scalar. The thresholds and offset settings are recorded by a slow ADC. Connections to

scalars and ADCs are via ribbon cable from the back of the RTF crate.

3.3 Triggering Electronics TCB/TLB

The custom trigger conditioner boards (TCBs) and trigger logic board (TLB) provide flexible triggering options using field programmable gate arrays (FPGAs). Each TCB handles the logic pulses from the six RTF boards associated with a single tower. TCB and TLB logic is loaded into FPGAs from EPROMs on power up. New triggering modes are introduced by simply swapping an EPROM.

The TCBs and TLB work together to define a global trigger for the experiment and record the time of the trigger in a history buffer. Figure 3 shows a schematic view of how these two cards work. The TCB creates an 800 ns-wide differential TTL copy of the trigger pulse received from the RTF board. The resulting logic levels are then sent to the history buffer and the TLB via SCSI-2 cables. The TLB forms a global trigger from the OR (rising edge) of up to seven TCBs as well as three external lines: random triggers used to monitor the noise baseline, the muon veto multiplicity trigger, and a veto pulser trigger used to monitor the veto system.

Four different triggering modes are defined separately for each phonon and ionization signal: **off**, **high**, **gated** (low, not high), and **low**. During normal CDMS running, the TLB was configured in **low** trigger mode for all the phonon channels and in **off** mode for all the ionization channels. The mask provides a record of which detectors triggered. After a global trigger occurs, the TLB is disabled until the EVENT BUILDER toggles the trigger enable line.

The history buffer also records the time that any of the following logic levels reach it from the TLB:

- Trigger enable
- Global trigger
- Veto pulser trigger
- Veto multiplicity trigger
- Random trigger

The TLB also provides EVENT BUILDER with a mask of logic values for each event. The mask record contains trigger information for each detector and tower, as well as the random and veto trigger bits and an 8-bit event number. Each detector's trigger mask bit may be set in one of 4 modes: on, off, triggered, and exceeds ultra-low threshold. The mask record was updated to allow for readout of neighboring detectors when obtaining calibration data.
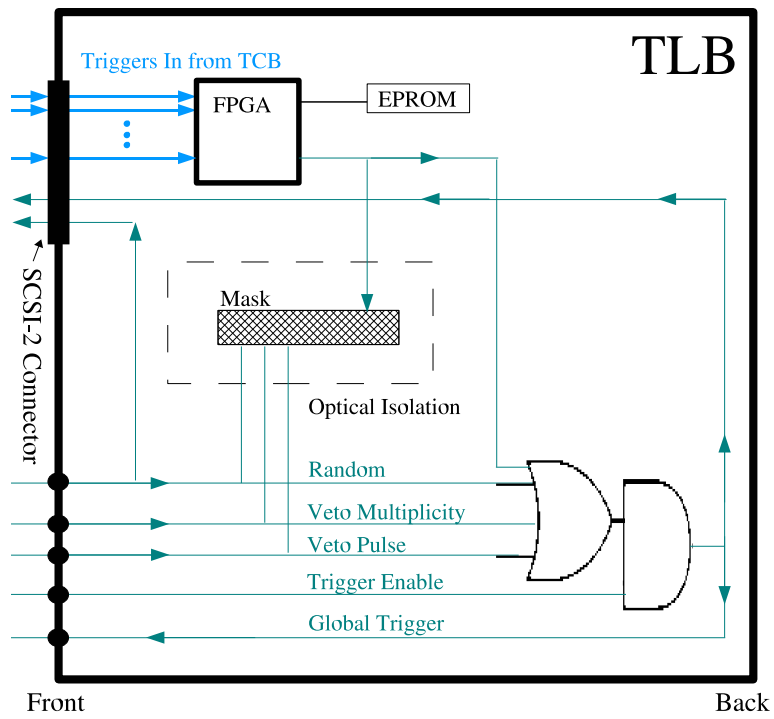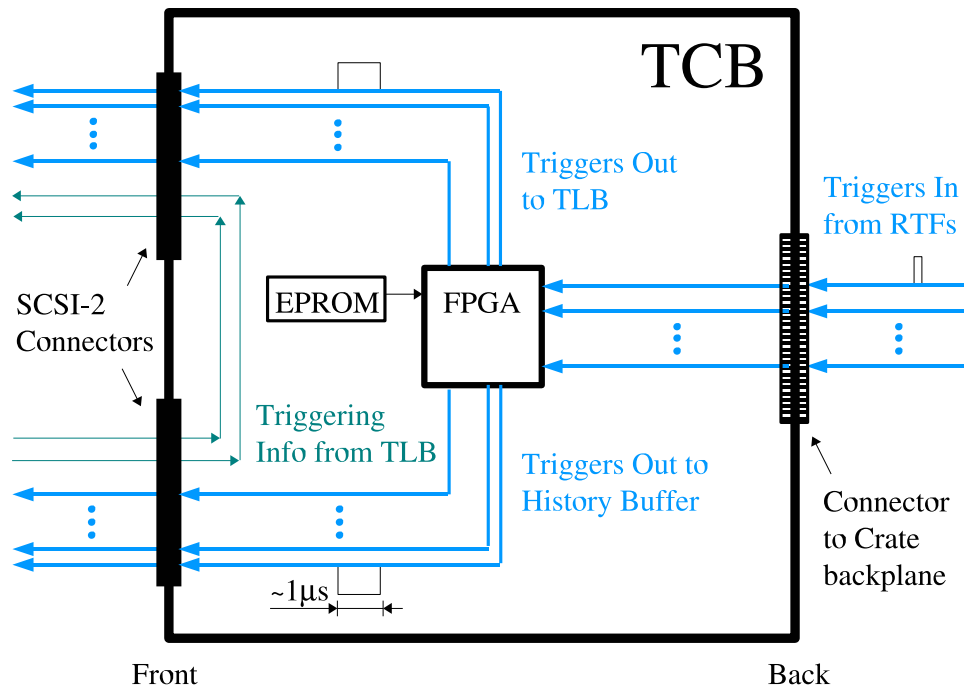
7

Fig. 3. A schematic view of a TCB and TLB.

To communicate between the DAQ computers and the VME hardware modules holding the digitizers, digital input modules, and ADCs, we use commercially available DAQ interface boards (SIS1100 PCI), VME controller (SIS3100), fiber optic link, and device drivers from Struck. Physical locations on the VME backplane are addressable as 32-bit memory regions by DAQ processes running on the computer.

Waveform digitizers record the signals from the solid state detectors (Struck SIS3301 FADC [6]) and from the scintillator veto panels (Joerger VTR812 [7]). Both flavors of digitizer are labeled *ADC* in Figure 1. The Struck waveform digitizers have 14-bit digitization over the range of $\pm$ 2.5 volts with a maximum digitization rate of 80 MHz. The Joerger units have 12-bit digitization over the range of $0 - 5$ volts with a maximum digitization rate of 40 MHz. To achieve the best resolution of the analog phonon and ionization pulses, the Struck modules are set to digitize at 1.25 MHz for a length of 2048 samples ($\sim$1.6 ms), where 512 samples are located before the trigger. The Joerger modules are operated at a digitization rate of 5 MHz for 1024 samples where 924 are located before the trigger. The length of the Joerger samples are shorter because of the shorter integrated length of the veto pulses.

The TCB sends copies of the logic levels to CDMS-customized Struck SIS3400 TDC/time stamp modules, which use the rising edge to create a digital record of the trigger times, which we call the history buffer. The modules are set to operate at 10 MHz. With 128k time stamps stored in its memory buffer (16 bytes per time record), the SIS3400 module allows for a full recording of all triggers between successive global triggers of the TLB. The buffer only fills up occasionally during in situ calibration due to the much elevated trigger rates, but "buffer full" events are flagged.

Each VME crate also contains a digital input module (VMIVME$-$1101) which records the trigger masks of all detectors serviced by the crate. This 32$-$bit mask provides redundant trigger information. It is also used to implement a selective data readout scheme, which reduces dead time by limiting detector trace readout to trigger mask hits only. Before implementation, event rates were limited to about 30 Hz. Selective readout reduced the average DAQ deadtime to $\sim$6 ms, allowing us to take calibration data at $\sim$150 Hz. Since surface events are the dominant background for CDMS, large quantities of [133]Ba gamma calibration data are required for full characterization of the data and determination of the depth profile. Selective readout allowed us to accumulate the required statistics without unduly impacting the WIMP search livetime.

After the DAQ was designed, we added a GPS time stamp module (Symmetricom bc635VME) to the readout chain. The neutrino beam from Fermilab is aimed at the MINOS neutrino oscillation experiment at Soudan. This flux of high energy neutrinos can interact with the rock and produce neutrons in the CDMS detectors. The estimated rate of these neutrons is extremely small over the entire life of the CDMS experiment, however, we have still chosen to veto events which are in time with the MINOS beam spill periods. These correspond to millisecond wide windows every few seconds when the Fermilab accelerator is running. Since MINOS needs the GPS time to synchronize their near and far detector spills, this GPS IRIG B signal is available to CDMS as well. The global trigger from the TLB is used to associate a GPS time with each recorded event, which can be used to veto events offline by comparison with MINOS spill times. The subsequent loss in livetime is negligible.

The DAQ is also responsible for monitoring the proper functioning of the detectors and experiment via slow ADCs (VMIVME-3128) and Joerger VSC64 scaler modules, both of which are connected to the RTF board. The slow ADCs are used to monitor the trigger thresholds for the detectors as well as the phonon channel bias points with 14-bit resolution over a maximum range of ± 10 volts.

4 CDMS II Software

A new CDMS II Software package was designed for Soudan to replace the LabView-based system used by CDMS I at the shallow Stanford site. The CDMS II software refers to a system of servers, GUIs (graphical user interfaces), configuration files, and an EVENT BUILDER. The EVENT BUILDER incorporates a fast data readout scheme described in Section 4.1. The rest of the system, often referred to as Slow Control, is covered in the remaining sections.

4.1 Fast Data Readout

The CDMS II data readout system is composed of two principal components: (1) an instance implementation of the general R2DM (Run-2 Data Manager) software suite developed at FNAL (Fermi National Lab) for prototyping event builders for the CDF and D0 experiments, and (2) a suite of custom C programs written specifically for interface with the commercial VME modules.
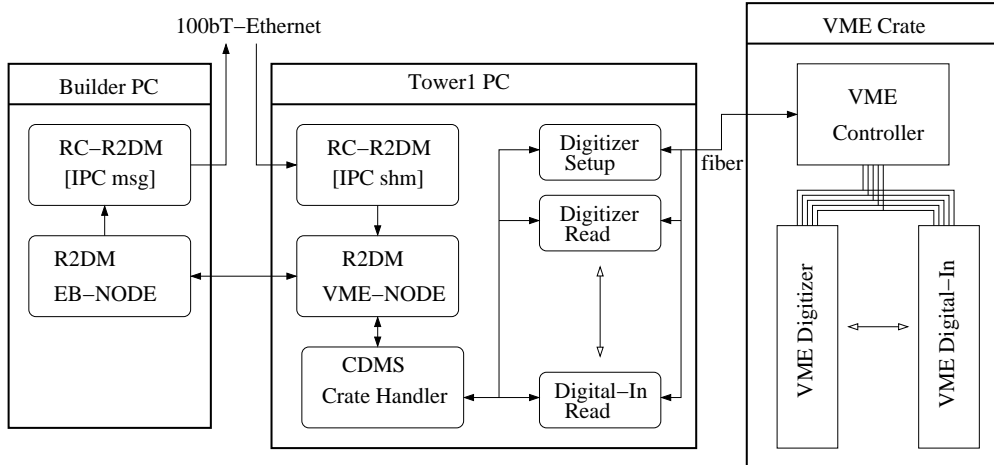
Fig. 4. Block diagram showing the software functional blocks of the CDMS II DAQ system and how these blocks are embedded inside the physical components of the system. This diagram shows one chain in a structure of many parallel EB−NODE to VME−NODE links. CDMS has a total of 4 such R2DM links.

R2DM is a general purpose C++ package based upon the various lower level UNIX utilities of TCP/IP sockets and System V IPCs (Inter-Process Communications tools). R2DM itself is built using three software packages: ACE, ITC, and thread_util. ACE [4] is a well-known TCP/IP communications package built upon the UNIX socket and was developed by researchers at Washington University. Both ITC and thread_util were developed at FNAL. ITC is an intertask communications package which uses ACE to move data and messages between processes and machines, and thread_util is a general purpose utility used to manipulate and manage threads. Based upon these three suites of utilities, R2DM provides the end-user with a logical framework into which to implement event building across a distributed UNIX system that has a number of subevent collector elements and a central event processor. R2DM also provides the user with the further flexibility of data queues and timeout intervals so that fast synchronous data-taking may be implemented along with delayed processing and decision making prior to data logging. This is the desired framework for high rate experiments in high energy physics such as CDF and D0. However, for rare event search experiments such as CDMS, the integrity of the full event is more important so that in the R2DM implementation for CDMS, only a single buffer in the data queue is used.

Figure 4 shows how the R2DM components of CDMS fit within the CDMS data readout functional blocks as well as how these blocks fit within the physical components of the system. The CDMS instances of the R2DM components are the EB−NODE and VME−NODE representing the central event processor and subevent collecting component, respectively. Also shown is the IPC shm (shared memory segment) into which the INTERLOCK MANAGER writes to communicate RUN CONTROL messages to the EVENT BUILDER. The

11

EB answers RC via the IPC Message queue.

The link between the `Tower1 PC` and the VME crate is provided by a commercial DAQ card (Struck SIS1100 PCI). Struck also provided LINUX device drivers for the SIS1100 so that from a user standpoint, simple data accesses to the VME bus can be coded as simple read/write commands that directly map onto the VME bus address space.

The final software component of the CDMS DAQ system is an addendum C++ class called Crate Handler which encapsulates all the public calls from the R2DM event collector class to interface with the VME hardware modules. Crate Handler also implements much of the required data formatting. This has the advantage of parallel processing since CDMS has a total of 4 such R2DM VME−NODE elements.

4.2   User Interface and Slow Control Servers

The DAQ slow control component, consisting of servers and their associated GUIs, allows the user to configure the detector and veto hardware. To help facilitate control and monitoring, CDMS II utilizes a GUI-server relationship, implemented in pure Java, through Remote Method Invocation (RMI) which provides communication channels between the GUI and servers. GUIs used to control the experiment are only accessible from the run control GUI. Each GUI-server combination was designed for a single, specific task. A command issued on a GUI is sent to the relevant server via RUN CONTROL. Since the software is modular, the servers are able to operate on different computers dedicated to specific tasks.

Each server obtains experimental run settings from configuration files and issues commands to the hardware. Hardware configuration settings are provided to the servers by means of configuration files, either statically present on server computers in the Soudan Underground Laboratory or user-provided configuration files through GUIs.

Most of the servers themselves are written in Java, but some hardware-level servers, such as the EVENT BUILDER and the high voltage server are written in C and C++. For inter-process communication among servers written in different languages, we use the Common Object Request Broker Architecture (CORBA) technology. A wrapper program allows us to integrate CORBA channels of communication between servers.

The _RUN CONTROL_ server organizes the actions of the other servers and routes user commands. It reads in the top level configuration file to obtain the run settings for the current data-taking mode. The RUN CONTROL server

also enforces security by verifying the user names and IP addresses of operators accessing the run control GUI. The run control GUI communicates directly with the RUN CONTROL server and is an interface allowing a user to control or monitor a run by enabling a user to configure, start, and stop data runs, as well as view recent error messages and access the GUIs that are connected to the other servers.

The *CONFIG MANAGER* is responsible for writing all configuration files. When a user updates the experimental configuration, the CONFIG MAN-AGER translates those updated settings into the appropriate GPIB format. The updated settings are then written to the appropriate hardware device's configuration file. The run configuration GUI allows the user to specify which FEB and RTF settings files are to be used for a given hardware configuration. It also allows the user to specify each detector's phonon and ionization trigger modes. After selecting the run settings, the user can issue a command to generate new configuration files. The specified run settings are then sent to the CONFIG MANAGER to create the new configuration files.

The *GPIB SERVER* delivers user commands from GUIs or configuration files to the GPIB devices (FEBs, RTF boards, TCBs, and the TLB). At the beginning of a run, the GPIB SERVER is responsible for preparing the GPIB devices for data-taking. It reads in the top level configuration file and those of the individual FEB, RTF, TCB, and TLB. The GPIB SERVER then sends all commands in the hardware configuration files to the appropriate GPIB devices. When the experimental run state is changed, the GPIB SERVER's responsibilities include updating detector configurations, setting trigger thresholds, and setting detector trigger modes.

The *FEB* and *RTF* GUIs allow the user to change settings on any FEB or RTF board. The FEB GUI allows the user to modify detector settings and the RTF GUI provides control of trigger levels and offsets. A user can save the current FEB or RTF settings into FEB and RTF settings files. These settings files can be reloaded into the GUIs to restore a useful detector state or be loaded into the run configuration GUI to specify experimental settings.

The *CAEN SERVER* monitors voltage and current from 40 active channels in the CAEN SY2527 [8] high voltage power supply on a continuous basis, since deviations of a few volts can significantly affect the PMT gain and resulting muon veto efficiency. Any deviations beyond 2% in voltage or 10% in current generate an alarm message that is transmitted through the error channel to the run control GUI. A high voltage GUI enables the user to reset individual channels. The high voltage server also keeps track of internal power supply error conditions and communicates them to the error channel.
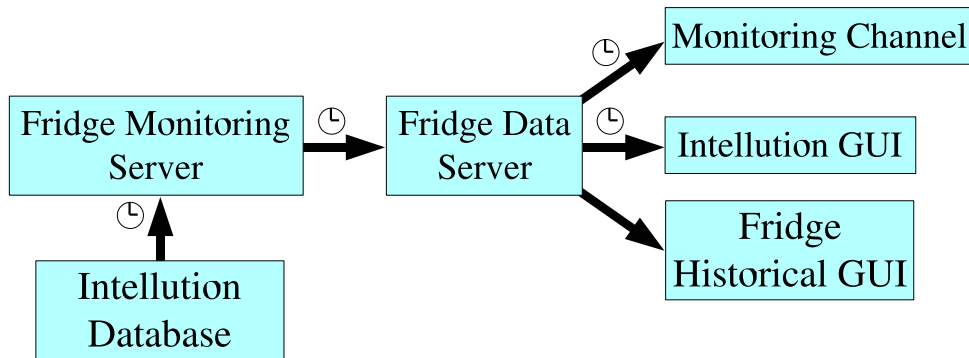
Fridge Monitoring

Fig. 5. A schematic view of the fridge monitoring software. The clock symbol above an arrow indicates that the subsequent program retrieves data from the previous program every minute.

4.3 Cryogenics Monitoring

The cryogenics system, including the helium dilution refrigerator, uses an Intellution [9] industrial control system and proprietary software. The monitoring system associated with it was designed to provide performance information to many simultaneous users without the purchase of additional Intellution licenses. To accomplish this goal, it had to satisfy two criteria: (1) data needed by the experiment should be available from a stable and shared DAQ computer and (2) users should be able to view the data by either following user selected fridge variables as they change over time, or by viewing a snapshot of all the current fridge variables.

Figure 5 illustrates the flow of information from the cryogenics system. The fridge monitoring server must run on the Intellution control system computer in order to have access to the fridge via the proprietary Intellution software. It obtains the value of all fridge variables every minute from the control system's database by means of an executable file provided by Intellution. The server is designed to allow the fridge variables obtained to change; currently 153 fridge variables are retrieved from the Intellution control system. The monitoring server then bundles these values together and makes them available upon request of the fridge data server.

The fridge data server runs on the DAQ computer in the CDMS II control room (labeled `datasrv` in Figure 2). It retrieves the fridge data every minute from the monitoring server and saves it to file. The fridge data server then provides this information to the monitoring channel, the Intellution GUI, and

the fridge historical GUI. The "historical GUI" was designed to allow users to see trends in the fridge data. The "Intellution GUI" was designed to allow users to easily view a snapshot of the most recent fridge data.

4.4   Data Monitoring

To monitor data quality and detector performance, the DAQ system includes a noise analyzer. At the beginning of any given data run, the DAQ system generates a configurable number of software trigger events at a configurable trigger rate. During CDMS II operation, the typical number of software trigger events was 500 and the typical trigger rate was 10 Hz. These events represent the state of the electrical and thermal noise as recorded in situ. Within minutes of completion of this initial stage of data-taking, the DAQ plots averaged noise power spectral densities (see Figure 6) of all phonon and ionization channels using these trigger events. The DAQ system automatically compares them to standard spectra. Significant deviations result in error messages sent to the Control GUI. Personnel on shift must then investigate further and make a decision as to whether a reset is required.
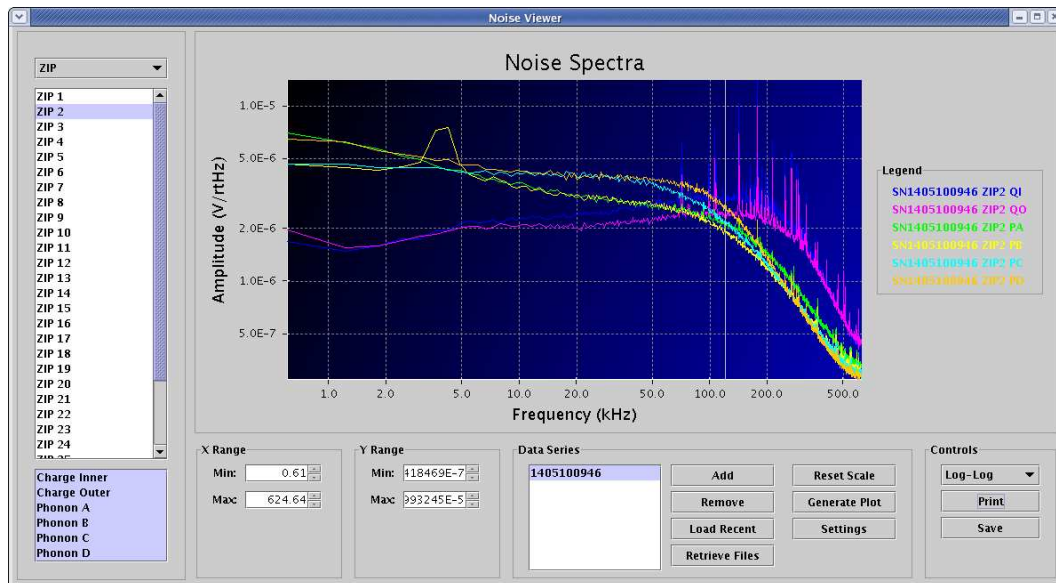

Fig. 6. Example noise power spectral density for detector 2 in tower stack 1.

Global indicators of detector functionality and the ambient noise environment are obtained from trigger rates during standard data-taking. The trigger rates, as recorded by the scalars, are averaged over a period of 1 minute (a user configurable parameter). Trending plots of trigger rates as a function of the last 24 hours of data-taking, are monitored for irregular spikes due to electrical noise.

15

Any significant deviation is automatically detected and warning messages are sent to the Control GUI to alert the shift personnel.

DC offsets for each phonon channel are monitored because changes indicate flux jumps in the SQUID arrays, which can affect data quality. If significant jumps are detected (> 0.2 mV in magnitude), the DAQ system implements an automatic flush and re-bias (re-lock) algorithm to reset the SQUID arrays. This ensures high data quality throughout the duration of a given run.

More sophisticated information is obtained from the fast online analyzer, which samples the data from standard triggers every 2 seconds (a user configurable value). During in situ calibration when the event rate is high, this 0.5 Hz siphon samples only a small fraction of the data. During normal WIMP search data-taking, essentially all background events are processed by the fast online analyzer. The fast online analyzer uses the ionization and phonon peak pulse heights as its energy variable. The event times are reconstructed from the raw ADC samples without any filtering or smoothing. This approach is fast, with accuracies a factor of 2-3 worse than more detailed pulse-fitting algorithms, but sufficient to create diagnostic data. This information is then converted to secondary variables for monitoring the detector response. Most important is the ratio of ionization to phonon energy and two measures of the event position derived from the relative timing and energy sharing between phonon quadrants respectively.

5 Conclusion

The CDMS II Data Acquisition system has functioned efficiently over the last decade, providing the timing and detailed pulse shape information crucial to our background rejection process. Complete waveforms were digitized for the phonon and ionization signals from germanium and silicon detectors, as well as pulses from the acrylic scintillator veto panels. This information, paired with trigger time and configuration via a separate history buffer, has made this experiment possible. This data was collected at 150 Hz during calibration mode, thus enabling us to acquire WIMP-search data at 80% livetime. Remote operation was implemented with a suite of user-friendly control GUIs and online data monitoring tasks. Automated tasks such as charge trap neutralization and cryogenic transfer schedules were interlocked with data-taking in order to minimize down time. Custom electronics and trigger boards equipped with FPGAs allowed flexibility in the definition of trigger conditions.

References

[1] Z. Ahmed et al.. (CDMS Collab) *Dark Matter Search Results from the CDMS II Experiment* Science 26 (March 2010) Vol. 327. no. 5973, pp. 1619 1621.

[2] Z. Ahmed et al.. (CDMS Collab) *A Search for WIMPs with the First Five-Tower Data from CDMS at the Soudan Underground Laboratory* Phys. Rev. Lett. 011301 (2009). [arXiv:0802.3530v2 astro-ph]

[3] D.S. Akerib et al.. (CDMS Collab) *Limits on Spin-Independent Interactions of Weakly Interacting Massive Particles with Nucleons from the Two-Tower Run of the Cryogenic Dark Matter Search* Phys. Rev. Lett. 96, 011302 (2006).

[4] http://www.cs.wustl.edu/ACE.html

[5] http://www.wiener-d.com

[6] http://www.struck.de/custom.htm

[7] http://joergerinc.com

[8] http://www.caen.it/csite/HomePage.jsp

[9] http://www.ge-ip.com/products/family/industrialsoftware