8 Electronics, DAQ, Controls, and Online Computing

This chapter describes the LZ signal processing electronics, data acquisition, detector-control system, and online data processing.

8.1 Signal Processing

The processing of the signals generated by the TPC PMTs is schematically shown in Figure 8.1.1. The TPC PMTs operate at a negative HV supplied by the LZ HV system, described in more detail in Section 8.6. HV filters are installed at the HV flange on the breakout box. The PMT signals leave the breakout box via a different flange and are processed by the analog front-end electronics, described in more detail in Section 8.3. The amplified and shaped signals are connected to the data acquisition system (DAQ), described in more detail in Section 8.5. The digitized data are sent to Data Collectors and stored on local disks. The PMTs of the outer-detector system operate at positive HV. The processing of the signals from these PMTs, shown schematically in Figure 8.1.2. The same type of amplifier used for the Xe PMTs is used for the outer-detector PMTs. Gain and shaping parameters of these amplifier will be fixed once the operating conditions of the outer-detector PMTs have been finalized.

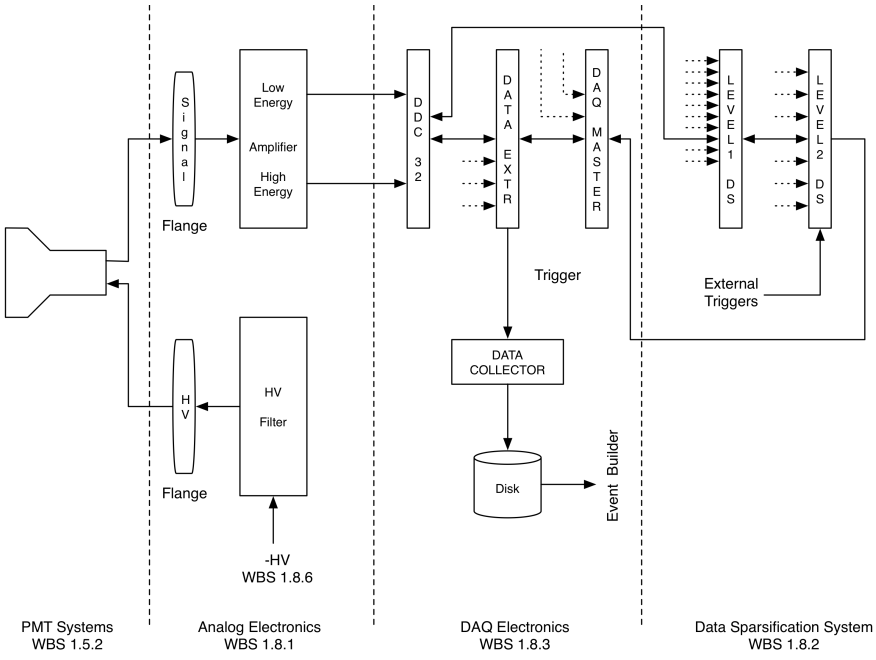

Figure 8.1.1: A schematic of the signal processing of the TPC PMTs. The TPC PMTs use dual-gain signal processing. The skin PMTs only utilize the low-energy section of the amplifiers.

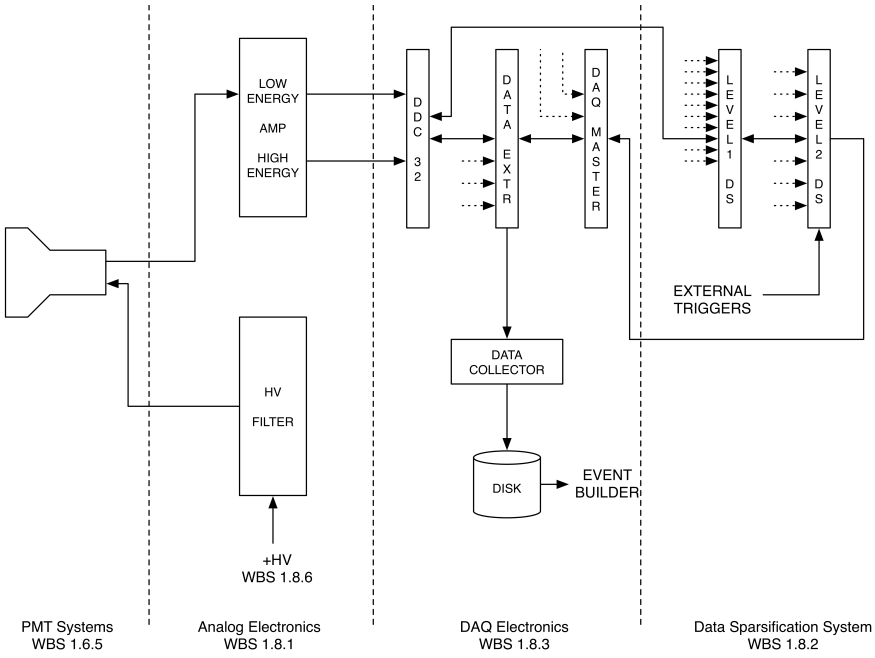

Figure 8.1.2: A schematic of the signal processing of the outer-detector PMTs. The outer detector PMTs use dual-gain signal processing.

The data flow is schematically shown in Figure 8.1.3. The event builder assembles the events by extracting the relevant information from the Data Collector disks, DAQ1 - 15. This step is discussed in more detail in Section 8.9. The event files are stored on local RAID arrays, RAID 1 and RAID 2, before being distributed to the data-processing centers for offline data processing and analysis.

8.2 Requirements

The parameters of the analog and digital electronics are defined based on the properties of the PMT signals, the required dynamic range of the different PMTs, and the expected calibration rates.

Three different PMTs are used in LZ. The TPC PMTs will see S1- and S2-type signals. The skin and outer-detector PMTs will only see S1-type signals. The relevant properties are listed in Table 8.2.1.

The cables that connect the xenon PMTs to the analog electronics are 40 to 50 ft long; the actual length depends on the final design of the conduits and the location of the area in which the analog electronics will be installed. The internal cables will be similar to the type used for LUX (Gore). Measurements with 45-ft-long LUX cables have shown an area reduction of S1 signals by approximately 13 to 23 %.

The design of the analog electronics is constrained by the required dynamic range of the LZ signals. Our design relies on the assumption that a single photoelectron (SPHE) detected in a TPC PMT generates a 11.5 mV ns pulse at the input of the amplifiers. The required dynamic range is defined by the sources used to calibrate LZ and the desire to detect high-energy events for background studies. Chapter 7 provides details on the LZ calibrations.

The DAQ system design is based on our experience with LUX and the required LZ calibration rates. Typical calibration rates are listed in Table 8.2.2. The TPC source calibration rates are limited by the maximum drift time of 700 µs in LZ. During source calibrations of the TPC, a 150 Hz calibration rate results in a 10 % probability of detecting a second calibration event within the drift time of the previous calibration event.

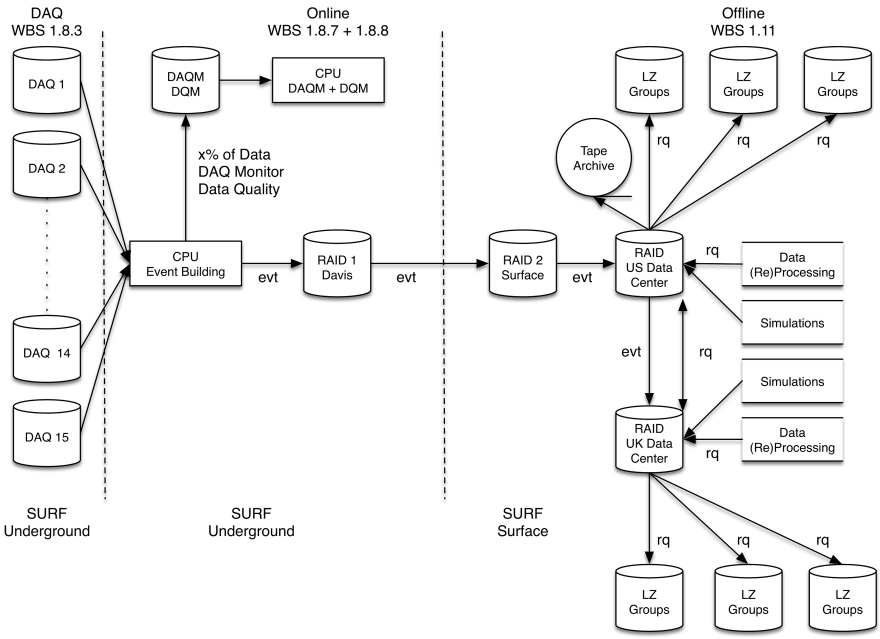

Figure 8.1.3: A schematic of the data flow.

Neutron calibrations are carried out to define the NR band in the TPC. These calibrations utilize external neutron sources and a neutron generator. External sources, inserted into the source tubes around the central cryostat, are used to calibrate the skin and the outer-detector PMTs. The count rates for these calibrations are not limited by the drift time in the TPC and they can be carried out at substantially higher rates. Weekly LED calibrations are done to examine the SPHE response of the PMTs and to monitor the PMT response. These calibrations can be carried out with rates as high as 4 kHz.

Table 8.2.1: Properties of the PMTs

System	Type PMT	Number	Gain	HV
TPC	R11410-20	494	$<5 \times 10^6$	<1,750 V
Skin	R8520	93	$<1 \times 10^6$	<900 V
Skin	R8778	38	$<5 \times 10^6$	<1,750 V
Outer Detector	R5912	120	$<1 \times 10^7$	<1,750 V

The data volume to be handled by the DAQ system can be estimated on the basis of our experience with LUX. In WIMP search mode, we will focus on events with energy depositions below 40 keV. During krypton calibrations, in which the total energy deposition is 41.6 keV, the average LUX event size was 203 kB (or 1.7 kB/channel). The size of each event was dominated by the width of the S2 signals. The event size of LZ can be estimated by scaling the LUX event size by the ratio of the number of PMT channels. Because LZ has four times as many TPC PMTs as LUX, and taking into account the dual-gains, we expect that the event size in LZ will be roughly 8 times as large as the LUX event size, or 1.6 MB. This estimate does not include the data volume associated with the outer-detector and the skin PMTs. If each outer-detector and skin PMT detects a single S1-like pulse, that increases the event size by about 45 kB. With compression, the typical

event size for LZ is estimated to be 0.53 MB. Monte Carlo simulations show that the total background rate in LZ will be about 40 Hz. The background rate in the WIMP search region (0 to 40 keV) will be about 0.4 Hz. At 40 Hz, the data rate is 21 MB/s. In 1,000 days, LZ will thus collect 1.9 PB of WIMP-search data. Other estimates, including energy depositions above 41.6 keV, result in an estimated total data volume of 2.8 PB. By optimizing the event selection, we expect to be able to reduce the total volume of WIMP-search data by a factor of 2 (see Section 8.5.2).

Table 8.2.2: Calibrations and expected count rates.

Calibration Type	System	Typical Count Rate	Frequency
Internal sources	TPC	<150 Hz	Twice a week (Kr)
External gamma sources	Skin and outer detector	TBD	TBD
Neutrons	TPC	<150 Hz	TBD
LEDs	TPC and outer detector	4 kHz	weekly

8.3 Analog Electronics

8.3.1 Design Criteria

The analog front-end design for the LZ experiment has benefited immensely from the experience with the LUX detector. In what follows, we have retained all the features of LUX electronics that performed well, while improving in some areas. The most important figure of merit of the LUX analog front end was its low noise characteristics, which allowed us to set our thresholds well below the SPHE level. Assuming the LZ noise characteristics will be the same as those for LUX, we expect to be able to run comfortably with a threshold of 0.25 PHE (photoelectron) for each PMT. Figure 8.3.1 shows a simulated distribution of total S1 pulse area for events in which one PHE each is detected in two PMTs (left) and in three PMTs (right). A threshold of >0.25 PHE is set for each PMT. The distributions are governed primarily by the ∼34 % rms width of the PMT signal (gain variation) and receives almost no contribution from electronic noise. Setting a threshold at 1 PHE for the total S1 signal yields an efficiency of ∼93 % for events that produce at least two PHEs in the PMTs. For the case of three PHEs, the corresponding efficiency is approaching ∼100 %. These electronics thresholds correspond to the lowest energy threshold achievable in such a detector, and hence drive the noise specifications for the front end. The choice between two- or three-fold coincidence will be governed by the dark currents in the PMTs, which have been specified to be in the 50 to 200 Hz range.

The second key design parameter for the LZ front end is the dynamic range, which is defined by the energy range of the sources used to calibrate LZ. Isotopes such as ^{83m}Kr (32.1 and 9.4 keV transitions), activated Xe (236 keV and 164 keV transitions), and tritium (endpoint at 18.2 keV) will be present or injected directly into the LXe volume and will be used to calibrate the detector periodically. LZ will retain the ability to detect high-energy events that saturate the PMTs of the top array by using only the light collected by the bottom array to determine the total S2 area.

The LZ electronics will provide excellent resolution for single liquid electrons, which are expected to yield at most 50 PHEs, depending on the height of the gas gap and the strength of the electric field in that region of the TPC. The typical duration of such pulses will be about 0.5 to 1.1 µs. At the same time, we will need to provide extremely clean measurements of SPHEs in order to have a sharp turn-on of the S1 efficiency. SPHE spectra also help with maintaining an in situ calibration of the PMT gains. To meet these requirements, the analog electronics provides one low-energy (high-gain) and one high-energy (low-gain) output for each PMT.

Figure 8.3.2 shows this concept. The high-energy channel has low gain and a 30-ns full width at tenth maximum (FWTM) shaping-time constant. Its dynamic range is defined by the 236-keV Xenon activation line. The low-energy channel has a 10 times higher area gain and wider shaping of 60 ns FWTM. It is optimized for an excellent SPHE response and dynamic range. The shaping times and gains are derived from one assumption: the DAQ will have a usable dynamic range of 1.8 V at the input and will sample the pulses at 100 MHz with 14-bit accuracy. A 0.2 V offset is applied to the digitizer channels in order to measure signal undershoots of up to 0.2 V. In summary, the relevant parameters are:

- Typical SPHE response of the PMTs: 11.5 mV ns pulse. The performance should also be verified if this value is as low as 6.5 mV ns for some PMTs
- The dark current for the PMTs at operating voltage has been specified to be in the 50 to 200 Hz range. The distribution within this range is not yet known.
- S1 light yield at 236 keV: 620 PHEs (for 650 V/cm).
- S1 light distribution: evenly distributed over all PMTs. The bottom array receives 80 % of the total S1 light.
- S2 light yield for ERs at 236 keV: 42 liquid electrons / keV and 70 PHEs / extracted liquid electron (conservative maximum).
- S2 light distribution: 22 % of the S2 light is in a single top PMT. The S2 light is distributed evenly over all bottom PMTs; the bottom array receives 45 % of the total S2 light.

The op-amps indicated in Figure 8.3.2 (left) are very similar to those used in LUX. The gain and shaping-time constants of the amplifiers were optimized using simulations. Figure 8.3.3 and 8.3.4 shows the results of simulations of S2 pulses associated with the 236-keV transition in activated Xe (top) and a 3-MeV energy deposition (bottom). Figure 8.3.3 shows that S2 saturation in the top PMT array for the 236-keV transition is not a problem if the amplitude of a single PHE is less than 1.5 mV (12 ADCC). If we allow one PMT to saturate, single PHEs of up to 3.0 mV can be accommodated. Figure 8.3.4 shows that the S2 associated with larger energy depositions will not start to saturate the PMTs of the bottom array if the amplitude of a single PHE is less than 15 mV (120 ADCC). The two outputs of the amplifiers have the following properties (assuming a 11.5 mV ns SPHE response of the PMT):

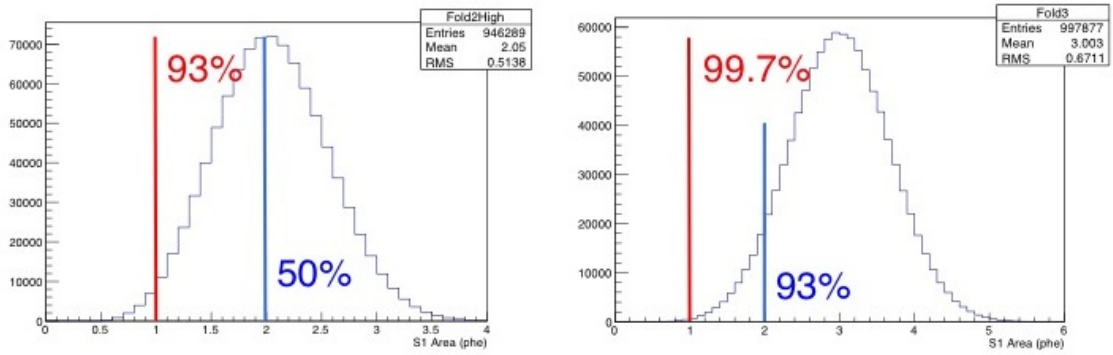

Figure 8.3.1: A simulated distribution of total pulse area for S1 pulses with two (left) and three (right) photoelectrons (left). Requiring a total area larger than 1 PHE provides nearly 100% efficiency for events with 3 or more S1 photoelectrons detected.

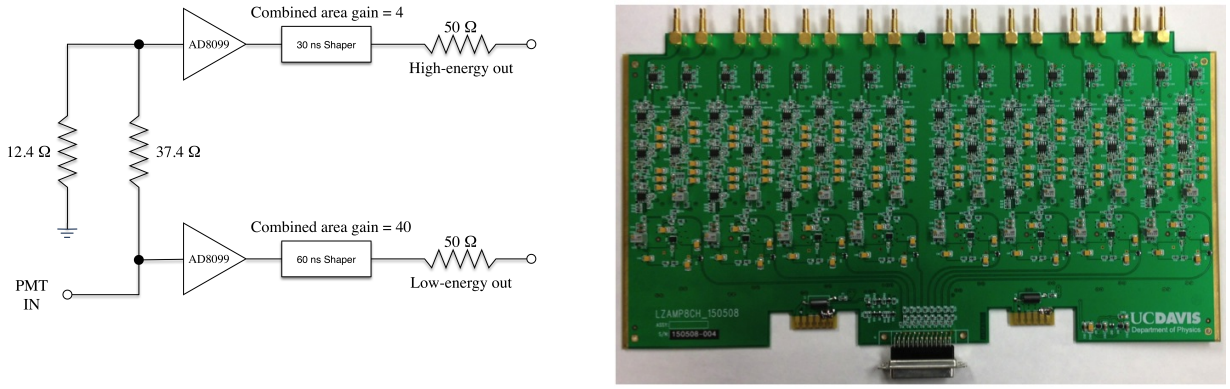

Figure 8.3.2: Left: A schematic diagram of a single channel of the LZ amplifier. Right: A photograph of the 8-channel prototype board.

- Low-energy output: Area gain = 40, 1 PHE = 105 ADCC (amplitude), S1 dynamic range: 140 PHEs (1,700 keV), S2 dynamic range 12 to 26 keV (top) and 1,300 to 2,800 keV (bottom) for 0.5 to 1.1 µs wide pulses (1σ) with no saturation.

- High-energy output: Area gain = 4, 1 PHE = 21 ADCC (amplitude), S1 dynamic range: 700 PHEs, S2 dynamic range 120 to 260 keV (top) and 13,000 to 27,000 keV (bottom) for 0.5 to 1.1 µs wide pulses (1σ) with no saturation.

The dynamic range for S2 signals is shown in Figure 8.3.5 [1]. The low-energy channel of the amplifier provides the dynamic range required for the tritium and krypton calibrations. The high-energy channel is required to provide the dynamic range required to measure the activated Xe lines. S2 signals due to $0\nu\beta\beta$ will saturate one or more channels of the top array, but the S2 pulse area can still be reconstructed using the low-energy channels of the bottom PMTs.

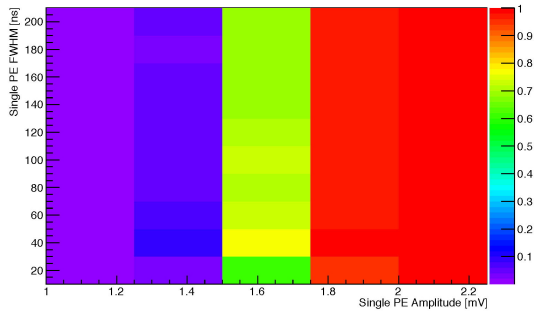

Figure 8.3.3: A simulation study of the S2 response for a 236-keV Xe transition in the center of the detector, as seen by the top PMTs. The gain and the shaping width of the SPHE response are varied. The color code shows the fraction of events in which the peak PMT saturates.

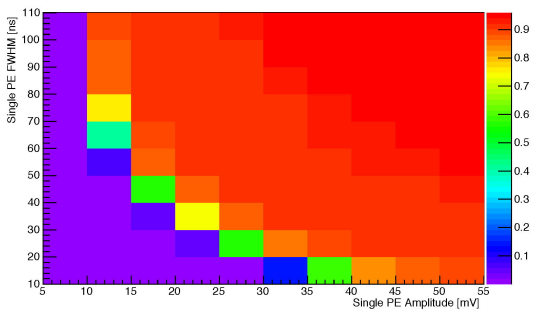

Figure 8.3.4: A simulation study of the S2 response for a 3-MeV energy deposition in the center of the detector, as seen by the bottom PMTs. The color code shows the fraction of PMTs of the bottom array that saturate.

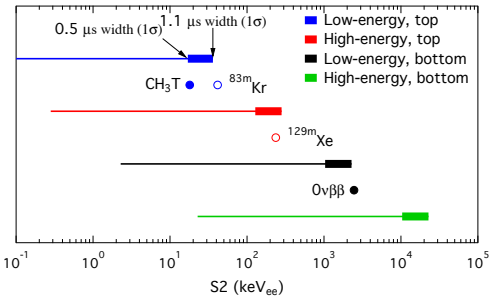

Figure 8.3.5: S2 dynamic range, expressed in terms of electron-recoil energy depositions for the low- and high-energy channels of the top and bottom PMT arrays [1]. The bars indicate variations in the upper level of the dynamic range due to the variations in the width of the S2. The lower and upper ends of these bars show the dynamic range for 0.5-μs-wide and 1.1-μs-wide pulses, respectively. Energy depositions for various calibration sources are indicated by the circles.

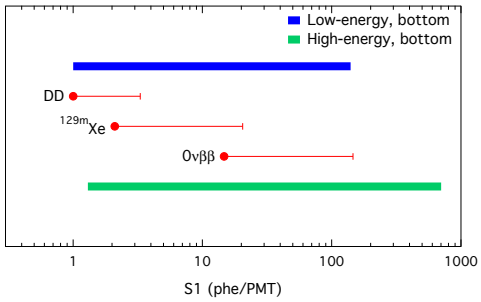

Figure 8.3.6: Dynamic range for S1 signals detected in the bottom PMTs [1]. The range required for DD neutrons, the ^{129m}Xe activation line, and $0\nu\beta\beta$ decay of ^{136}Xe are also shown. The left-end of each source line corresponds to an energy deposition in the center of the TPC; the right-end corresponds to an energy deposition 1 cm above the cathode.

The dynamic range for S1 signals is shown in Figure 8.3.6 [1]. The figure shows the dynamic range of a bottom PMT. Also shown are the number of PHEs associated with the full-energy deposition of the 236-keV Xe activation line and $0\nu\beta\beta$ decay of ^{136}Xe. The dynamic range provided by the dual-gain channels is sufficient for all LZ calibrations.

Minor changes to the gain and shaping parameters of the amplifiers can be easily accommodated via changes in resistor and capacitor values. This will be done if more detailed simulations, including the electronics response and the noise of all components of the electronics chain, indicate a need for modifications.

The same amplifier design will be used for the skin and outer-detector PMTs although the gain and shaping may be adjusted slightly, if needed. Dual gain amplifiers will be used for the outer-detector PMTs, in order to make measurements of cosmic muons. For the skin PMTs, only the low-energy channel will be instrumented. A summary of the number and type of analog signals is shown in Table 8.3.1.

Table 8.3.1: Summary of the number and type of the 1,359 analog signals.

	High-gain Signals			Low-gain Signals		
PMT Type	**#**	**Area Gain**	**Shaping (FWTM)**	**#**	**Area Gain**	**Shaping (FWTM)**
Top TPC	253	40	60 ns	253	4	30 ns
Bottom TPC	241	40	60 ns	241	4	30 ns
Skin (1")	93	100	60 ns	0	NA	NA
Skin (2")	38	40	60 ns	0	NA	NA
Outer Detector	120	40	60 ns	120	4	30 ns
Total	745			614		

8.3.2 LZ Amplifier Prototype

Figure 8.3.2 (right) shows a photograph of the first amplifier prototype with eight input channels. The amplifiers connect to the PMT signal lines using the DB-25 connector visible at the bottom of the figure. The DB-25 connector allows the signal lines to be interleaved with two ground lines for minimizing cross-talk. The 32-channel flange is shown in Figure 8.3.7 (middle). The amplifiers will be housed a 5-card mini crate that mounts on the signal flange, as shown in Figure 8.3.7 (right). It houses four amplifier cards and a fifth card for power distribution and monitoring.

Waveforms captured with an oscilloscope at various stages of the analog chain are shown in Figure 8.3.8. The response of a PMT, operating at 1,300 V, to an SPHE at the start and end of the cable in the vacuum space is shown. Also shown are the outputs of the LE and HE channels of the amplifier, at its output and at the end of external co-axial LMR-100 cables. The table at the bottom of the figure shows the losses and gains at each stage.

Figure 8.3.9 shows the measurements of the noise power in both LE and HE channels. The RMS ADC noise of the free-running DDC-32 digitizer channels was measured to be (1.19 ± 0.01) ADCC. In the spectra shown, the measurements were made using a DDC-32 and also using an oscilloscope. In both cases the contributions from the digitization have been subtracted in quadrature. The resulting input referred noise voltage in the "white" part of the spectrum were measured to be $11\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$ and $19\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$ respectively, for the LE and HE channels. These values are in excellent agreement with the circuit simulation model.

Figure 8.3.10 (left) shows measurements of the power dissipation. A fully loaded crate was powered up in a lab with an ambient temperature of 80 F, without any active cooling or forced air through the crate. Within 30 minutes the crate reached a temperature of 100 F and stabilized at that value. This is well within the operating range of the electronics.

Figure 8.3.10 (right) shows measurements of crosstalk between the individual amplifier channels. The central channel was pulsed such that it had an output of 240 mV. The neighboring channels shown in cyan and green had bipolar induced pulses with a pulse height of 0.6 mV, or 0.25 % crosstalk. The next-neighbor channel, shown in yellow, had a pulse height of 0.3 mV. These small values of crosstalk are completely acceptable. Moreover, bipolar pulses contribute very little area when integrated.

The linearity of the amplifier was studied using S2-like test pulses with a width of 1 µs (FWHM). The test pulses were propagated through 45 ft of gore cable before reaching the amplifier. The output signals saturate when their amplitude exceeds 2.6 V. Examples of the results of these linearity studies with S2-like pulses

Figure 8.3.7: Photographs of the PMT HV flange (left), the 32-channel PMT signal flange (middle) and a 32-channel crate that mounts on the signal flange.

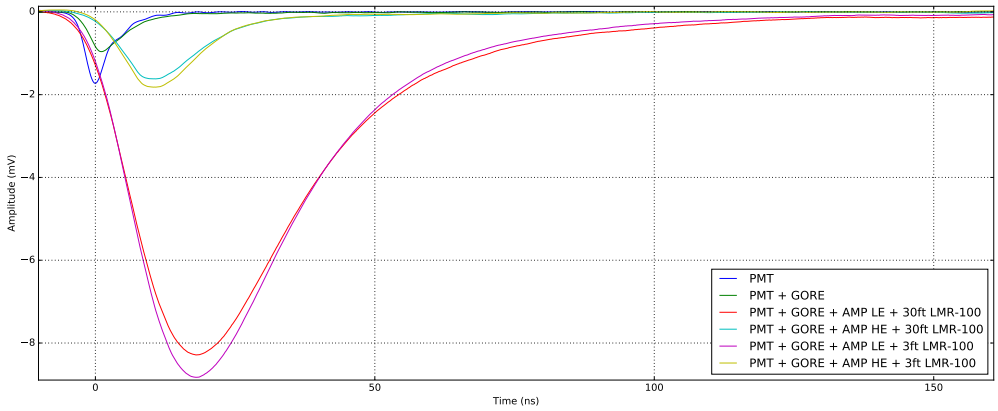

Configuration	Area from -7 to 150 ns (mVns)	Area from -7 to 20 ns (mVns)
PMT	-10.88	-10.65
PMT + GORE	-9.46	-8.22
PMT + GORE + AMP LE + 30ft LMR-100	-336.77	-117.62
PMT + GORE + AMP HE + 30ft LMR-100	-32.71	-22.82
PMT + GORE + AMP LE + 3ft LMR-100	-338.00	-122.77
PMT + GORE + AMP HE + 3ft LMR-100	-34.09	-25.03

Figure 8.3.8: Waveforms for an SPHE at various stages of the analog chain. The PMT was operated at −1,300 V. The effect of the internal Gore cable and the external LMR-100 cables is shown. The desired shaping and gain specifications have been achieved.

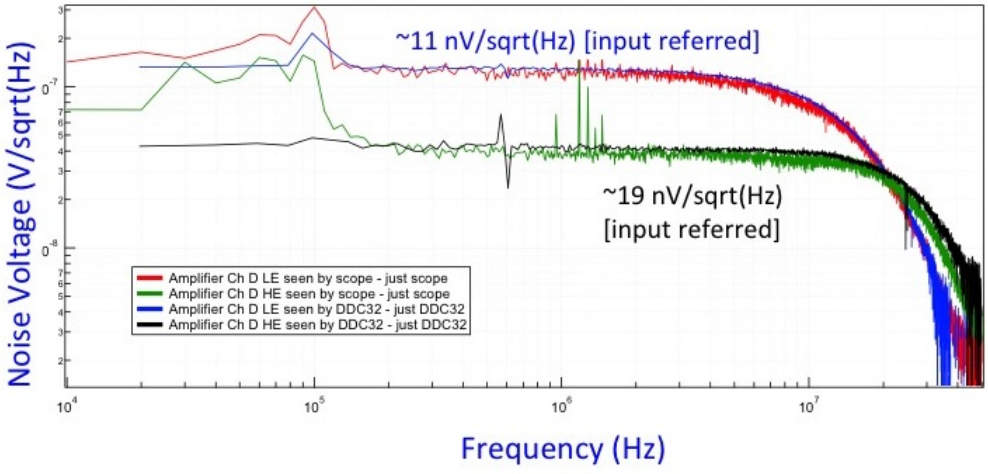

Figure 8.3.9: The measured spectrum of noise power in the amplifier channels.

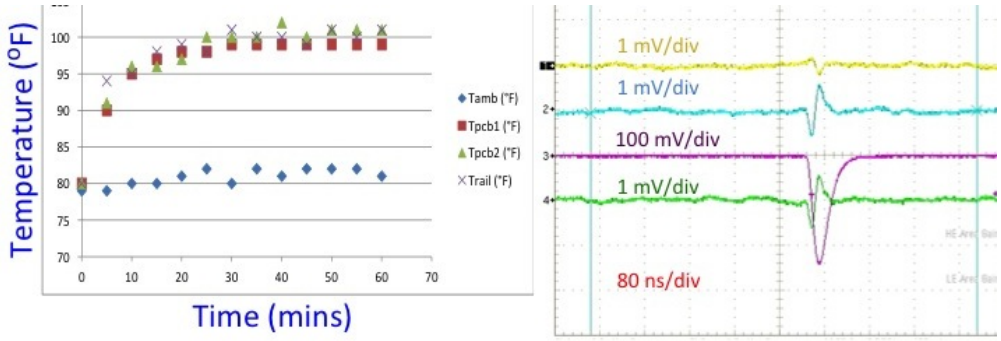

Figure 8.3.10: Left: The power dissipation and equilibrium operating temperature of a crate without any active cooling. Right: Measurements of the cross-talk for two neighboring and one next-neighbor channels.

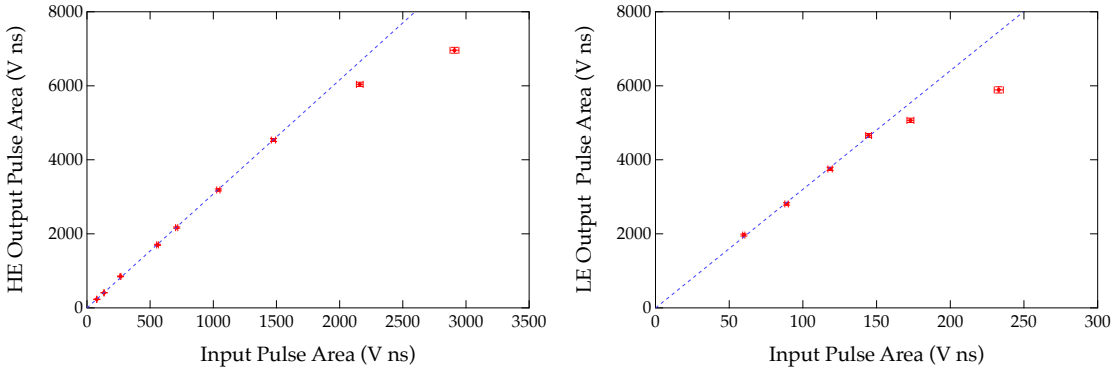

Figure 8.3.11: Results of linearity measurements. The area of the output pulse is plotted as function of the area on the input pulse. The results obtained for the high-energy (low-gain) channel are shown on the left while the data collected for the low-energy (high-gain) channel are shown on the right.

are shown in Figure 8.3.11 The low-energy and high-energy outputs becomes nonlinear when the area of the input signals exceed 150 V ns and 1,500 V ns, respectively. Although the response is nonlinear in this region, there is still a one-to-one correlation between the area of the output pulse and the area of the input pulse.

Finally, the response of the amplifiers was measured for the case of PMTs that fail to provide the nominal gain. For this test, the HV on the PMT was lowered to $-1{,}220$ V, corresponding to a gain of 1.9×10^6. Figure 8.3.12 shows a histogram of the S1 filter output for a three-sample wide filter for signals from the LE channel. The SPHE peak can be seen to have a peak value of 84.2 mV ns and an rms width of 28 %. The noise peak is also shown. Placing a cut at 50 mV ns results in an SPHE detection efficiency of 92 %, which meets specifications. The singles rate of 1 Hz due to noise at this cut is negligible compared to the 50 to 200 Hz dark noise expected from the PMTs.

8.4 Digital Electronics

The LZ digital electronics is based on a digital platform (a motherboard); a prototype of this platform is shown in Figure 8.4.1. The final LZ motherboard will be based on this design, but will operate with a more powerful Series-7 Kintex field-programmable gate array (FPGA) from Xilinx [2]. The final motherboard will provide gigabit Ethernet, RS-232, and low-voltage differential signaling (LVDS) interfaces, and four

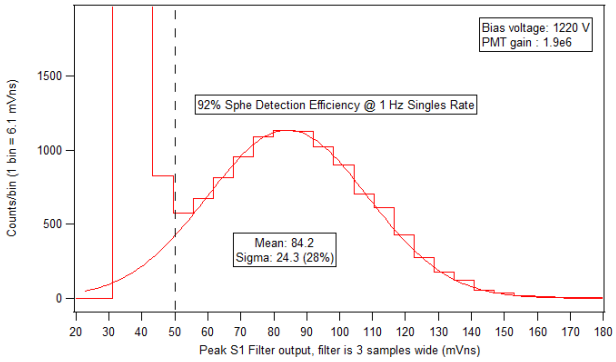

Figure 8.3.12: Measurements of the PMT response to SPHE at a gain much lower than the nominal value.

logic outputs, either TTL or NIM. Waveform memory (3,578 kB) will be provided by the FPGA. A large event-buffer memory of up to 128 MB will be provided by the dual-core processor. The onboard clock can be driven externally in order to synchronize multiple boards to the same clock source. Very high processing power, nominally 52 giga-operations per second, will be provided by the onboard FPGA. Two daughter card connectors can host two separate daughter cards, or one daughter card of twice the size. The I/O pins of these connectors are arranged as differential pairs, supporting either the differential or single-ended signals. A dual-core processor will be connected to the FPGA with the 32-bit memory bus, as well as two dedicated 16-bit-wide FIFOs. Readout of the FPGA data can be performed either via the memory bus or via the FIFOs, depending on the application. The board can be hosted in a 6U VME crate, or it can be powered with a tabletop power supply. Power consumption is minimized by using low-voltage chips.

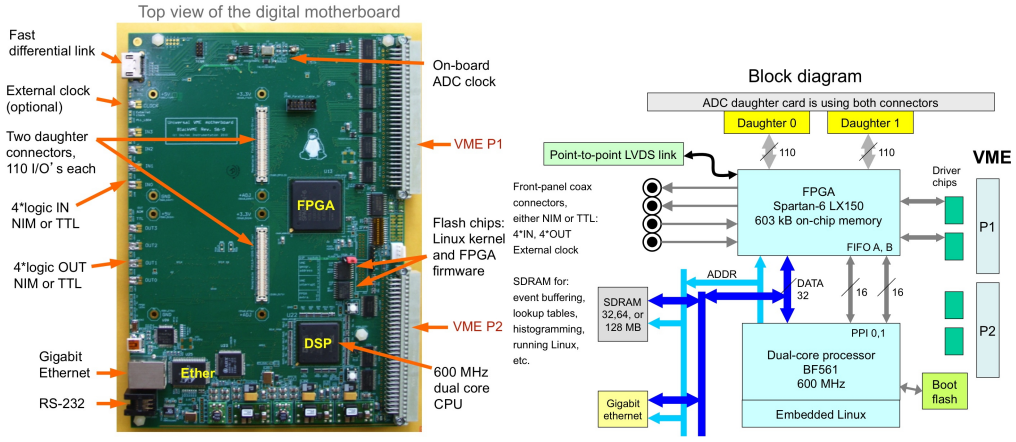

Figure 8.4.1: The digital motherboard used to develop the LZ digitizers. It provides gigabit Ethernet, RS-232, USB-2, VME, and LVDS interfaces. The FPGA and the dual-core processor are rated at 52 and 2.4 giga-operations per second, respectively.

The onboard processing power and multiple interfaces provide flexibility that can be applied to almost any project. The LVDS links enable custom communication architectures. The Ethernet provides support for distributed experiments and/or standalone remote applications. The processor is running Linux, which is popular, free, and fully customizable, allowing each board to perform on-the-fly data processing and online diagnostics.

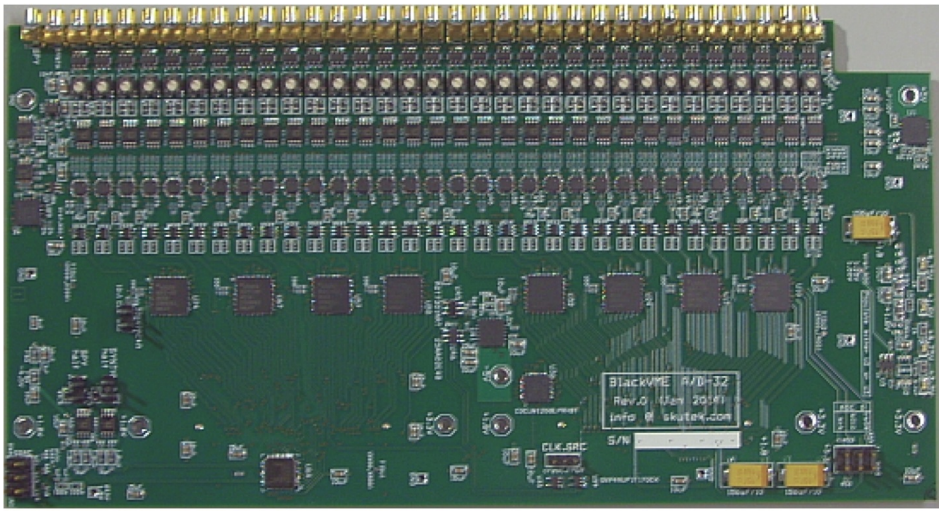

Figure 8.4.2: Prototype ADC daughter card with 32 channels.

The 32-channel ADC card shown in Figure 8.4.2 implements the digitizer front end. It provides 32 channels of digitization and two waveform reconstruction outputs for diagnostic. The ADC channels feature remote DC offset control. The card is connected to the two daughter connectors of the digital baseboard that provide the control signals and power. The printed circuit-board layout can accommodate quad A/D chips with sampling frequency up to 125 MHz. This card, installed on the digital motherboard, is referred to as the DDC-32 in the remainder of this chapter.

8.5 DAQ

The top-level architecture of the LZ DAQ system is shown schematically in Figure 8.5.1 [3]. The DDC-32s continuously digitize the incoming PMT signals and store them in circular buffers. When an interesting event is detected, the Data Extractor (DE) collects the information of interest from the DDC-32s. The DEs compress and stack the extracted data using their FPGAs and send the data to Data Collectors (DCs) for temporary storage. The Event Builder (EB) takes the data organized by channels and assembles the buffers into full event structures for online and offline analysis. The DAQ operation is controlled by the DAQ Master (DM) for high-speed operations such as system synchronization and waveform selection, and by the DAQ Expert Control/Monitoring (DECM) system for slow operations such as running setup/control and operator diagnostics. The entire system runs synchronously with one global clock.

8.5.1 Data Extraction

Figure 8.5.2 shows a more detailed overview of the different key elements of the DAQ system. The digitizers are sampling at 100 MHz with 14-bit resolution over a 2-V range. During normal operation, the boards will collect waveforms in a Pulse Only Digitization (POD) mode, which is expected to effectively reduce the raw waveform volume by a factor of 50 [4]. The amount of memory assigned to each channel is set so that no data truncation is expected even if the POD mode reduces the data volume by only a factor of 20. The POD waveforms are stored in dual-buffered memory that is divided into sections that hold the header information and the actual POD samples, as shown in Figure 8.5.3. Separate POD header and payload memories will improve the performance of extracting waveform data when the Data Sparsification Master (DSM) detects an event of interest [5].

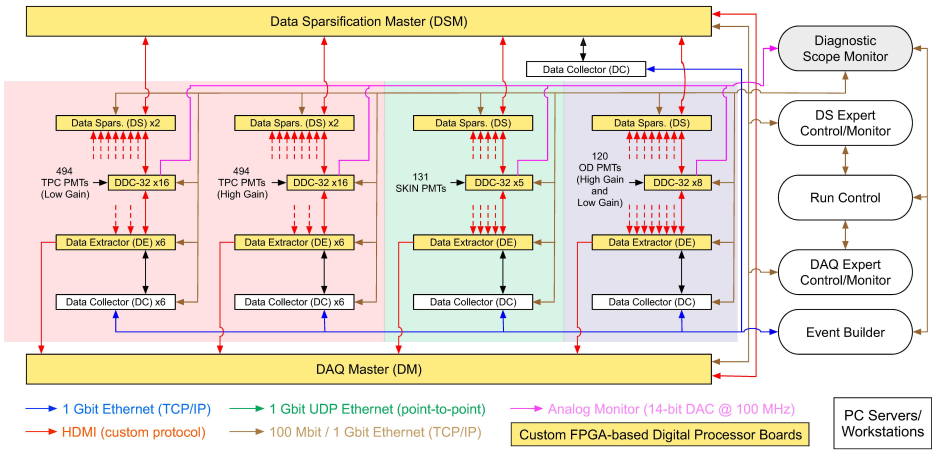

Figure 8.5.1: Diagram of the DAQ architecture [3]. Groups of digitizers (DDC-32) capture the amplified and shaped signals from the Xe, skin, and outer-detector PMTs. The waveforms of interest are extracted from the DDC-32s and compressed by the Data Extractors (DEs) before they are passed to Data Collectors (DCs) for temporary storage. One additional Data Collector will capture reduced sparsification quantities to be merged by the Event Builder with the waveform data in full event files. The DAQ Master Board (DM) coordinates the high-speed operation of the entire DAQ system when the Data Sparsification Master (DSM) signals the detection of waveforms to be preserved. The global clock will be distributed over the shown HDMI links or dedicated NIM clock inputs that are not shown in this diagram.

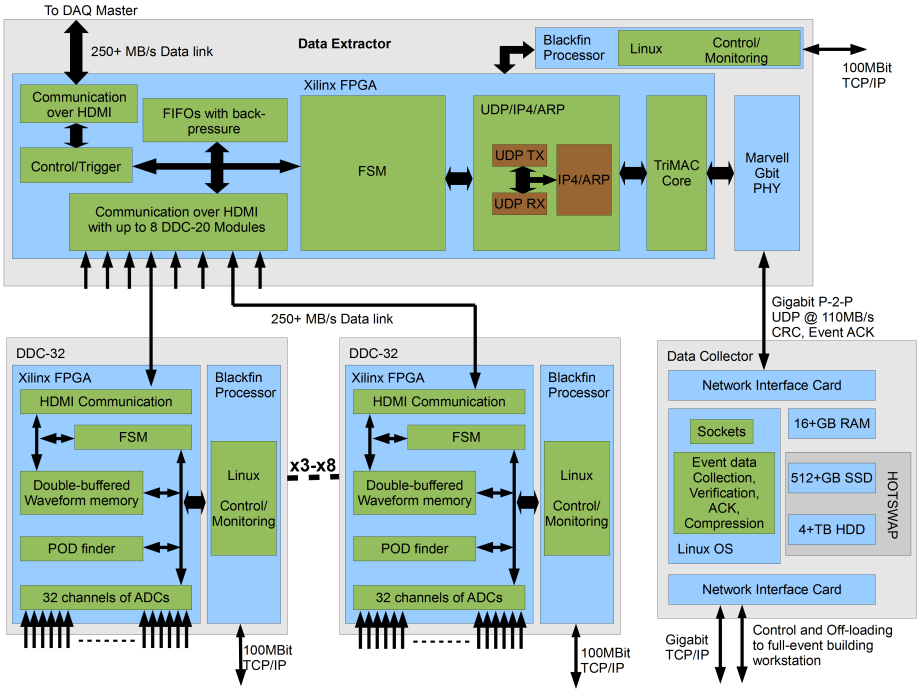

Figure 8.5.2: Detailed depiction of the inside of and interaction between the key elements of the proposed DAQ system.

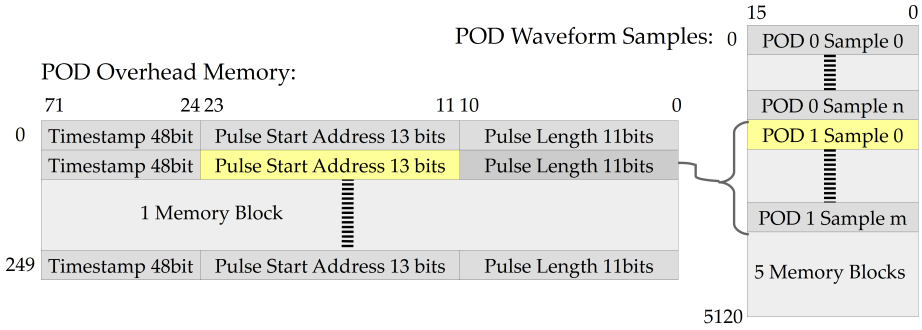

Figure 8.5.3: Depiction of the proposed memory organization of POD waveform storage in the FPGA for a single buffer (out of two) of a single channel. Separation of POD overhead and POD samples will improve the performance and ease of extracting information from memory.

The extreme flexibility that comes with using FPGAs and their internal memories allows us to assign the entire on-chip memory to just one specific channel when needed. This feature will be used for system diagnostics and noise measurements where capturing long, continuous (non-POD) waveforms is important. In such a mode, the DDC-32 will be able to capture 10-ms-long waveforms, suitable for power-spectral-density analysis.

The DEs and the DM use the same hardware but different firmware. They use the same motherboard as the digitizer boards, with different daughter cards that enable communication with multiple DDC-32 modules and the DCs. Each daughter card can serve up to 14 DDC-32s over HDMI links and one DC over a dedicated gigabit Ethernet connection.

The HDMI link has seven single-ended lanes used for communicating states of the finite state machines (FSMs) between boards and four LVDS lanes used for fast offloading of the waveforms from the digitizers. On the current motherboard, we have used input/output serial/deserializer (IOSERDES) elements, offered in the Spartan-6 FPGA series, and have confirmed the advertised 1-Gbit throughput per LVDS lane.

Table 8.5.1: Key parameters of the prototype Data Collector.

Processor:	Intel Xeon E3-1270V3 3.5GHz Quad-Core	**HDD 1:**	SAMSUNG 840 Pro Series 256GB SSD
Motherboard:	ASUS P9D-V ATX	**HDD 2:**	Western Digital RE 4TB 7200 RPM
Memory:	16GB Kingston DDR3 SDRAM ECC	**Case:**Case:	NORCO RPC-270 2U Server Case
NIC:	Intel Ethernet Server Adapter I350-T2	**Hot Swap:**	ICY DOCK 3.5" and 2.5" SATAIII 6Gps HDD Rack Tray

Table 8.5.2: Summary of the performances of the DAQ links and their expected utilization levels.

Link	Expected Performance	Maximum Expected Usage	Usage
LVDS over HDMI	250+MB/s	8.6 MB/s	<3.5 %
DE → Gigabit UDP → DC	109 MB/s	34.4 MB/s	<33 %

Table 8.5.3: Summary of the expected storage and buffering capabilities on the Data Collectors during calibrations.

Source	Data Rate per Data Collector (uncompressed)	Solid-state Drive (512 GB)	Hard-disk Drive (1 TB)
^{83}Kr	~5.2 MB/s	~1.1 days (raw) ~4.4 days (7z)	~9 days (raw) ~35 days (7z)
LED	~34.4 MB/s	~4 hours (raw) ~16 hours (7z)	~1.4 days (raw) ~5.3 days (7z)

The gigabit Ethernet link between the DE and the DC utilizes the User Datagram Protocol (UDP). To mitigate the limitations of UDP, we will add for each event a cyclic redundancy check (CRC) and acknowledgment, and also use the dedicated Ethernet link in a point-to-point topology to eliminate congestion and packet loss. All UDP packets formed by the DEs FPGA will be sent over a single gigabit Ethernet cable to a dedicated network port on the DC.

The DCs will be implemented using server-grade, rack-mountable (2U) workstations. We have tested a prototype DC with the parameters shown in Table 8.5.1. We were able to confirm reliable data transfer from the DE to the solid-state drive (SSD) in the DC. With appropriate modifications to the network interface card (NIC) drivers, we were able to transfer data at a constant rate of 109.8 MB/s, with no data loss or corruption. Table 8.5.2 summarizes the tested link performances and their expected utilization.

The data stored on the DCs are processed by the Event Builder (EB). The EB builds events by sampling the different DC disks and collecting all information associated with a given event. The resulting event (EVT) files are compressed in order to achieve the smallest possible file size for storage and transmission. The EVT files are written to RAID array #1, located in the Davis Cavern, before being transferred to the RAID array #2, located on the surface at SURF. The RAID arrays have sufficient storage capacity for a full month of data taking.

The DAQ system is designed to be able to allow LED calibrations of the TPC PMTs in about 10 minutes. This requires an event rate of 4 kHz, resulting in a ~340 MB/s total waveform data rate. The fast (400+MB/s for SSDs) storage drives in the individual DCs allow for sustaining such collection rates, and the amount of space offered by each DC permits significant buffering in case of connection problems to the off-site permanent storage, as shown in Table 8.5.3.

We are planning to use 7z compression. Nominally, 7z uses the LZMA algorithm, but we have found that the PPMd algorithm [6] is better suited for the waveforms we are going to collect. Based on comparisons made using LUX data, the 7z PPMd compression offered an additional 33 % data reduction for waveforms and a 15 % data reduction for reduced quantities over gzip, which was used in LUX.

Because the POD mode relies on time stamping, the entire DAQ system will run with a global 100 MHz clock. We are considering two ways of distributing the clock, either by using a proven method of dedicated NIM logic clock FAN-IN/OUTs or using clock recovery from the serial links of the HDMI cables; the latter method is currently being evaluated. We are also looking at the possibility of using one of the single-ended channels of the HDMI cable for a dedicated clock signal.

The configuration, acquisition control, and monitoring will be done using the DECM workstation. For improved robustness, the DECM will communicate with all DAQ elements using an isolated 100/1000 LAN; the run-control (RC) system, described in Section 8.9, will not have direct access to the front-end digitizers. We have successfully tested and are planning to use ICE middleware as our communication framework [7].

8.5.2 Data Sparsification

The design of the LZ data sparsification system is based on our experience with event selection in the LUX trigger system. The top architecture of the LZ data sparsification system is shown schematically in Figure 8.5.1. The DAQ and data sparsification firmware operates in parallel on the FPGAs of the DDC-32 digitizers.

The DDC-32s continually process the incoming pulses and extract specific required quantities such as pulse area, pulse height, and time of occurrence. By using digital filters with various filter lengths, the system can distinguish S1 and S2 pulses. The parameters extracted from the waveforms by the DDC-32s are sent to the Data Sparsifiers (DSs) for further processing and generation of secondary quantities such as the multiplicity vectors of groups of PMTs. These secondary quantities are sent to the Data Sparsification Master (DSM), where the final waveform selection decision is made. The decision to preserve the current waveforms is sent to the DAQ Master. It is important to reiterate that much of the information based on which a given event selection was made is going to be off-loaded alongside event waveform data and stored in the full event files. This will allow for simpler cross-checking and verification of the system performance off-line.

A new feature under development is the creation of a total waveform digital sum that can be created at the DSM from all or subset of DDC-32 channels as shown in Figure 8.5.4. This allows for including total area cuts in the event selection decision process.

The LZ DAQ/Sparsification system is capable of handling event rates up to ~250 kHz. This is much higher than the highest event rates expected in the LZ detector, which will occur during LED calibrations (~4 kHz).

As in LUX, the PMT signals will be processed by parallel digital integrating filters, allowing for discrimination against area of the incoming pulses. The filter with an integration width of about 60 ns to 100 ns is tailored to S1-like pulses; the second filter with a few-microsecond width is optimal for wider S2-like pulses. The filters are designed to also perform automatic baseline subtraction.

As shown in Figure 8.5.1, the TPC, skin, and outer-detector PMTs will use separate DDC-32s and DSs because these PMTs have different purposes and the firmware will be tailored to their specific needs. Such separation also makes it easier to apply different scaling factors for the digitization of signals from different PMT groups. The three major waveform-selection modes for the central TPC PMTs are summarized in Table 8.5.4.

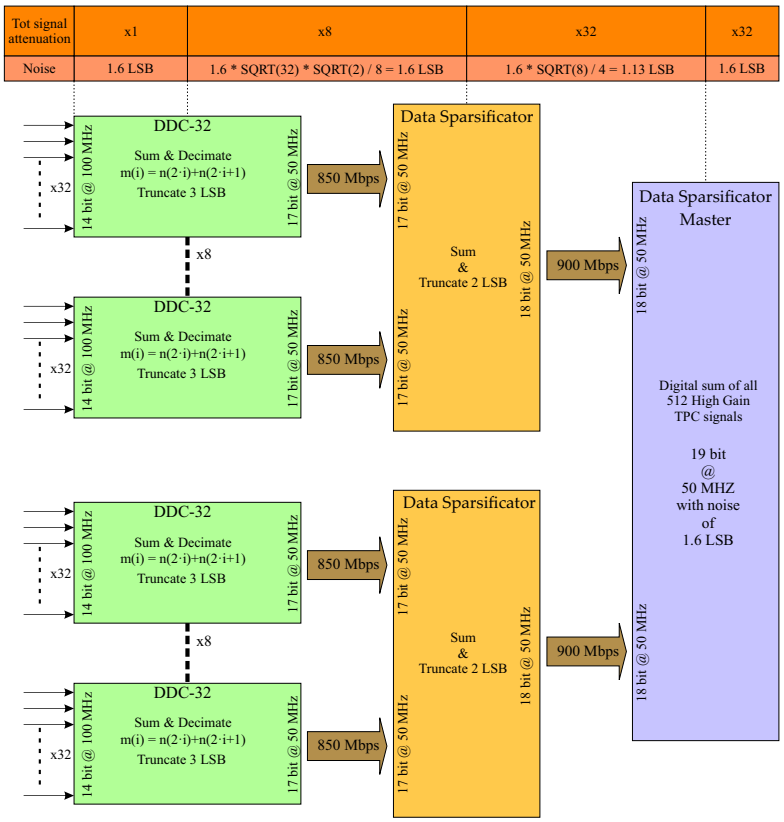

Figure 8.5.4: Schematic of digital waveform sum creation from all or subset of DDC-32 channels that allows the total pulse area to be a part of the sparsification decision making.

Table 8.5.4: Summary of three major waveform-selection modes for the central TPC PMTs.

Trigger Mode	Summary
S1 Mode	Detection of coincident S1-like signals across selected channels. No fiducialization.
S2 Mode	Detection of coincident S2-like signals across selected channels. Fiducialization in the x,y plane.
S1 and S2 Mode	Detection of S1-like signals followed by S2-like signals within a selected drift time range. Fiducialization in x, y and z planes.

If a waveform selection condition is met, the DSM sends a signal to the DM. At the same time, a packet containing the selection parameters used to make the decision is merged with the captured waveform data. This allows for offline evaluation of the selection decision for every individual event. This feature has proved extremely valuable in monitoring LUX data quality.

Similarly to the DAQ system, the DS system will be controlled and monitored by a dedicated DS expert control/monitoring (DSECM) workstation over an isolated 100/1000 LAN network. The DS system will operate on the same 100 MHz global clock as the DAQ.

The LZ DS system will adopt and expand on the monitoring capabilities of the LUX system. Capabilities such as performing continuous noise sweeps (monitoring S1 and S2 filter crossing rates as a function of threshold) or monitoring channel hit distributions of the selected events, all in parallel to regular uninterrupted data-sparsification operation, have been invaluable in LUX and surely will be in LZ.

Each of the DDC-32s has a dual 14-bit, 100 MHz analog reconstruction output (SPY) allowing for diagnostics and scope monitoring. The SPYs can be sourced with individual incoming channels or their digital sum. They can also be sourced with S1 and S2 filter outputs, individual or summed, or any other signal internal to the FPGA. This feature will be very useful, especially in development and deployment stages, as proven in the LUX project.

8.6 PMT HV

LZ will use the Wiener Mpod LX HV system to bias the xenon and outer-detector PMTs. This same system is currently being used for LUX. The system will provide negative HV to the xenon PMTs and positive HV to the outer-detector PMTs. The HV modules are part of the EDS 201-30x 504 series. Each module provides 32 HV channels and uses a common floating ground. The voltage ripple is less than 5 mV. The HV can be set with a resolution of 10 mV, and the current on individual channels can be measured with a resolution of 50 nA.

HV connections to the HV filters are made using Kerpen cable with Redel connectors on both ends. Each Kerpen cable carries HV for 32 channels. Important properties of the HV system are listed in Table 8.6.1.

Table 8.6.1: Details of the PMT HV system.

HV Module	Maximum HV HV	Maximum Current	PMTs	Channels	Number of Modules
EDS 201 30n	−3,000 V	500 µA/ch	TPC/Skin	625	21
EDS 201 30p	+3,000 V	500 µA/ch	Outer-detector	120	4

8.7 Cables

All network, signal, and HV cables are low smoke zero halogen (LSZH) cables. LZ uses the same cables that have been approved for LUX use in the Davis Laboratory. The exact lengths of the HV and the signal cables will be fixed once the location of the analog amplifiers is fixed and the route of the cable trays has been finalized.

The types of cable and their lengths are listed in Table 8.7.1. The networking cables will have various lengths; individual lengths will vary based on the location of the network switched and the location of the devices to be connected. About half of the network cables are used by the slow-control system; the other half will be used for the networks associated with the DAQ, trigger, and online systems.

We have not been able to identify a suitable HDMI cable that provides good performance and uses LSZH materials. The total length of these cables is small and the electronics racks are enclosed and vented directly into the exhaust system of the Davis Campus. This arrangement will be assessed by the SURF safety team. Note that rack enclosures and connections to the exhaust system are also used in LUX for this same reason.

Table 8.7.1: Information on LZ signal, logic, HV, power, and network cables.

Cable Type	Type	Length (ft)	Number of Cables	Notes
Network	Belden 7936A	10,000 (total)	750	Cat 6, LSZH Various lengths
Signal	LMR-100A-FR	56	143 bundles (8 cables/bundle)	LSZH
Logic	LMR-100A-FR	2	500	LSZH
HV	Kerpen	56	25	LSZH 32 channels per cable
Power cords	CordMaster	6	100	LSZH
HDMI	TBD	TBD	99	

8.8 Slow Control

8.8.1 Functions of the slow control

The slow control system performs supervisory control and monitoring of all the major subsystems of the experiment. The functions of slow control can be classified in several categories as described below:

- **Experiment parameter monitoring and logging.** This includes periodic readout of various sensors distributed throughout the xenon circulation, storage and recovery systems (temperature, pressure, gas flow etc.) and of advanced sensors installed in the detector (capacitive level and distance sensors, acoustic sensors and loop antennas). For the electronic racks the status of the uninterruptible power supplies (UPS) as well as temperature and current consumption of the front-end amplifiers will be monitored. Also, the environmental data, such as ambient temperature and pressure and radon concentration in the air, will be collected.

- **Control over experiment subsystems.** This includes controlling (by means of electrically-triggered pneumatic valves) pathways for performing operations of xenon circulation, storage and recovery, controlling the detector temperature via thermosiphons and heaters, and controlling distribution of liquid nitrogen required for thermosiphon operation. The slow control will also serve as a front-end to the calibration systems (krypton injection, radioactive source delivery, optical calibration) to guarantee that they are not interfering with normal detector operation. Finally, high voltage power supplies for the field-shaping grids and the PMTs in the detector will be programmed and monitored through slow control.

- **Safety.** One of the most important functions of the slow control system is to ensure safe operation of the detector and all the experiment subsystems in order to prevent equipment damage and xenon loss. This is achieved by implementing a set of interlocks that would protect the system from erroneous operator command. Additionally, the alarm system will be programmed to alert on-site crew and/or system experts when certain key parameters are out of the predefined range. The emergency xenon recovery, activated by pressure rise in the detector, will safely transfer xenon to the storage facility if the thermal control of the detector fails.

- **User interface.** The slow control system will provide the GUI to display all relevant telemetry information and controls organized in pages by subsystem. The plots of parameter history with be available as well. The controls available to the user will depend on privilege and expertise levels as determined by querying the centralized access control and privilege management facility provided by the experiment IT infrastructure. All control changes performed by users will be logged.
- **Automation and interfaces with run control and offline systems.** In addition to the set of scripts (procedures) directly accessible to slow control, a set of standard high-level operating procedures will be made compatible with remote procedure calls from run-control or to the offline systems. Programming these scripts on the slow-control (server side) will make these sequences faster, more reliable, and will eliminate possibility of an operator error. In some cases these scripts will include break points requiring explicit feedback from the run control client to advance to the next step, further ensuring system safety and accuracy.

8.8.2 Requirements

The main requirement to the slow control system is to provide the monitoring and control infrastructure that will guarantee safety of the experiment subsystems and xenon supply in all modes of operation. The Level 2 requirement R-180004 states: *Guarantee the safety of xenon supply and the xenon circulation system. Use Programmable Logic Controllers (PLCs) to control and monitor the xenon purification, circulation, and storage systems.* The corresponding Level 3 requirements specify the necessary parameters of the monitoring and control infrastructure:

- **R-180501:** Monitor state of detector. Monitor up to 2,500 channels.
- **R-180502:** Provide access to detector controls. Control up to 1,400 channels.
- **R-180503:** Provide an alarm system. Allow for alarm priorities, alarm profiles, and alarm pipelines.
- **R-180504:** Record system operations and system state. Provide a history of all monitored parameters and changes in the status of controls as well as a log of user access and actions.

8.8.3 Overview

The functional diagram of the slow control system and its interaction with the experiment subsystems and infrastructure is shown in Figure 8.8.1. The core of the slow control system is composed of two components: (1) the integrated supervisory control and data acquisition (SCADA) software platform *Ignition* from Inductive Automation [8], and (2) the Siemens SIMATIC S7-410H programmable logic controller (PLC) with associated I/O modules [9]. The Ignition server works as the main hub of the slow control system and is responsible for collecting and archiving the sensor data and providing the authorized users with access to the controls. Additionally, the Ignition server runs the alarm system, provides the scripting engine for experiment automation and provides the ICE server for interfacing with run control. It also provides a GUI for accessing historical data in the form of plots by local and remote clients. The configuration data (properties of sensors and controls, alarms, user preferences etc.) are stored in the local configuration database.

The PLC component is designed to guarantee safety of the critical subsystems (xenon circulation, storage and recovery, detector temperature control and liquid nitrogen distribution) by implementing necessary interlocks to safeguard their correct operation. It is also programmed to ensure safety of the xenon supply even in the case of a major infrastructure failure, for example a prolonged power or network outage with no access underground. The data flow between the Ignition server and the PLC passes through a dedicated TCP/IP network. The Ignition server directly communicates with the non-critical experiment subsystems

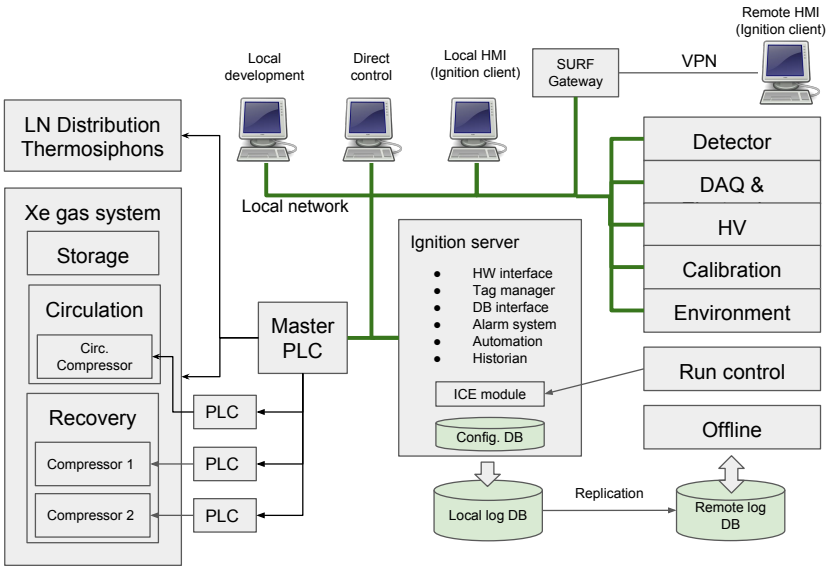

Figure 8.8.1: Slow control functional diagram.

through the same network. The computers used for local programming and control of both PLC and Ignition server are connected to this network as well. Communication between the Ignition server, the PLC, other non-critical experiment subsystems, operator consoles, and PLC software development workstations uses Ethernet. Multiple TCP/IP subnets in conjunction with router access control lists will be used to separate the slow control network into security zones with appropriate access restrictions.

The historical data log is stored in a local MySQL slow control database mirrored to a replica database from which these data can be accessed by clients on the LZ experiment network including the offline system.

8.8.4 PLC System

The PLC system is the core component of the slow control which is expected to maintain the detector and xenon system in a safe state even in the case of failure or shutdown of the rest of the slow control infrastructure. It is designed to provide continuous real time monitoring and control of the critical subsystems, including:

- Xenon purification, circulation and storage systems,
- recovery compressors,
- detector temperature control and liquid nitrogen distribution systems,
- vacuum pumps.

The master control of the the critical subsystems is performed by a Siemens S7-410H Redundant Hot Backup PLC with S7-300 associated I/O modules. This PLC allows bumpless transfer from one active CPU to a backup CPU, with synchronization and self-testing of both CPUs. Each CPU is connected to a separate interface module (IM), allowing fault tolerant process control and bumpless programming and hardware

changes in run mode. Members of the LZ run control team have extensive positive experience with the Siemens S7-410. This PLC system will be programmed using the Siemens PCS7 engineering programming software (meeting IEC 61131-3 standard).

Four compressor units (primary and secondary recovery compressors and two circulation compressors) are controlled by local Beckhoff PLCs, sharing data with the master PLC via Profibus. The Master PLC provides system data and interlocks to these subordinate PLC systems. Compressor-related instruments and valves are connected directly to Beckhoff modules. In the unlikely event of the failure of the master PLC, local PLCs can operate independently. Control logic resides in both master and locals PLCs. Local control of PLCs is also provided via touch screen so the autonomous units can be controlled directly by the on-site operator in case of an emergency. All PLCs are connected to a single dedicated private network with TCP/IP connectivity to 3 development machines (2 underground and one at the surface). The development machines will be remotely accessible through a secured connection (e.g. VPN).

More information on the PLC system and its support of xenon recovery operations is given in section 6.4.6.

8.8.5 Ignition

The choice of Ignition as the software platform for LZ slow control is based on the wide range of highly customizable, easy to use core functions and tools for development of efficient and robust SCADA systems. This set of tools includes: the tag management system, scripting and automation facilities, a versatile alarm system, a user management interface, a historian function, a database interface, a versatile toolbox for GUI desgin as well as a number of other functions and utilities. Ignition comes with a comprehensive library of device drivers supporting most of the hardware used in LZ.

Ignition follows a web-based client-server model where the server is deployed on site and takes care of communication with all the hardware while GUI clients can run both locally and remotely. It has a modular, easily expandable architecture: an extensive SDK allows developers to build their own custom modules and device drivers in Java or Python. At a higher level, Sequential Function Charts (SFC) or Python scripts can be used for development of automation procedures.

A high level of system availability is maintained even in the case of a system fault through the automatic swap to a mirror server to which configuration is constantly backed up (only a few minutes are lost in such an event). If even higher system availability is determined to be necessary, virtually zero downtime can be achieved by running two redundant servers in parallel with automatic hot swap (for the cost of the second server license). Security is provided by support of secure communication with clients, user authentication (site-wide authentication is possible through an Active Directory), user role management and action logging.

Ignition is also expected to be more cost-efficient for scientific automation than most commercial SCADA solutions as it is licensed by server with no limitation imposed on number of devices, tags or clients. Also, due to its multi-platform nature, it does not limit the developers or users to a single operating system.

8.8.6 Tags and device interfaces

A *tag* is the term used for an elementary unit of information in industrial automation. A tag can be bound to a device (made readable for monitoring or writable for control), a record in a database or a variable in control software. Ignition tags follow an object oriented philosophy. *User Defined Tags* (UDT), equivalent to classes used in programming, allow the user to define a set of generic tags each associated with a certain device type from which all the tags for individual devices are later derived. A well-structured UDT scheme speeds up the setup, configuration and modification of large groups of sensors. UDTs can be associated with GUI templates with positive impact on the GUI usability as well as a reduction in the development time and effort. A UDT library developed for the System Test slow control system will be used as the base for LZ slow control UDT.

The information transfer between the tags and the associated devices can be handled in several different ways. The preferred one is by using the OPC-UA protocol supported by many PLCs either directly or through middleware solutions (for example the Kepware server). In addition, Ignition has its own OPC-UA server with drivers for many well-known PLC brands, including Siemens PLCs. For the rest of the devices the preferred protocol will be MODBUS which is also supported by the Ignition OPC-UA server. The devices not supporting MODBUS (e.g. RGA units) will be interfaced via custom made Java drivers written in the Ignition SDK. All these interfacing schemes were successfully implemented and field-tested for the System Test slow control system.

8.8.7 Historian, alarms and automation

Ignition offers highly customizable historian features for logging the values of the selected tags in a database. There are options for both periodic logs and those on an evaluate-on-change basis. It is also possible to have different timing parameters for different tag classes or temporarily increase the readout rate for a particular tag. The GUI history explorer module developed for the System Test provides tools for rapid historical data visualization and analysis.

Ignition also provides a powerful alarm framework whereby each tag can have multiple alarm levels associated with it. Alarms can be configured with different profiles that include priority, escalation scheme and notification pipeline. Notifications can be both local (client GUI, speaker, siren) and remote (email, VoIP or text message). The configuration of these alarm profiles is done by the developers in the Designer. Individual alarm activation/deactivation, setpoint configuration and acknowledgments can be done in the client GUI by a user with adequate privilege level. This allows developing a unified interface for setting alarms and storing pre-defined alarm configurations on a per-user basis; for example, a gas system expert may have a particular set of alarms that they want to be active during a troubleshooting exercise.

For further customization, Ignition provides a Python scripting engine. Python scripts can run both in the client or gateway and can be attached to event handlers (e.g., mouse clicked, key pressed, a tag changing state events), a timer or run inside an SFC. Since Python scripts running on the client block the client execution, these are only suited for GUI customization or very simple automation routines. On the other hand, scripts running in the SFC which are executed in the gateway, are intended to control complex automation routines and procedures. SFCs use the IEC61131-3 programming specification (also used for PLCs) and can be developed in the designer using a drag and drop programming tool. SFCs can be started and monitored both from gateway or Clients GUI but they always run on the gateway. Several SFCs can run in parallel and can be programmed to run for extended periods (even surviving a gateway restart).

8.8.8 Graphical User Interface

Ignition offers an efficient and intuitive environment for high performance GUI development. The GUI design concepts were extensively tested and refined during the development of slow control for the System Test until the GUI configuration met with the approval of both project scientists and engineers. The LZ slow control GUI will be based on this configuration and will offer different levels of supervision and control of the system.

At a higher level, several panels will show the real-time status of the system and subsystems from which an operator will be able to tell the overall health of the system. These panels will allow the authorized user to verify and control key system parameters and assess high level information, for example alarm status, current operation mode, status of automation scripts, etc. These panels will be designed so that even a novice operator could recognize an abnormal situation. Furthermore, access to the solutions for the common problems and tasks will be provided.

At a lower GUI level, detailed information about every sensor and actuator in the system will be provided for debugging and development purposes. The panels will be organized in form of complete but simple P&IDs, as shown in the example in Figure 8.8.2. Upon a click on a P&ID element, a new window will open and the system experts will be able to see and change the sensor configuration or status of the control, allowing complete access and control for the run control system experts.

At a very high level, a single summary plot of the entire system can be prepared by slow control and sent to run control for display in a single status pane on the run control GUI. This high level display can ensure that key information is available to any shift personnel regardless of run-control expertise, and regardless of the specific arrangement of windows, dialog boxes or other information on the slow-control system.

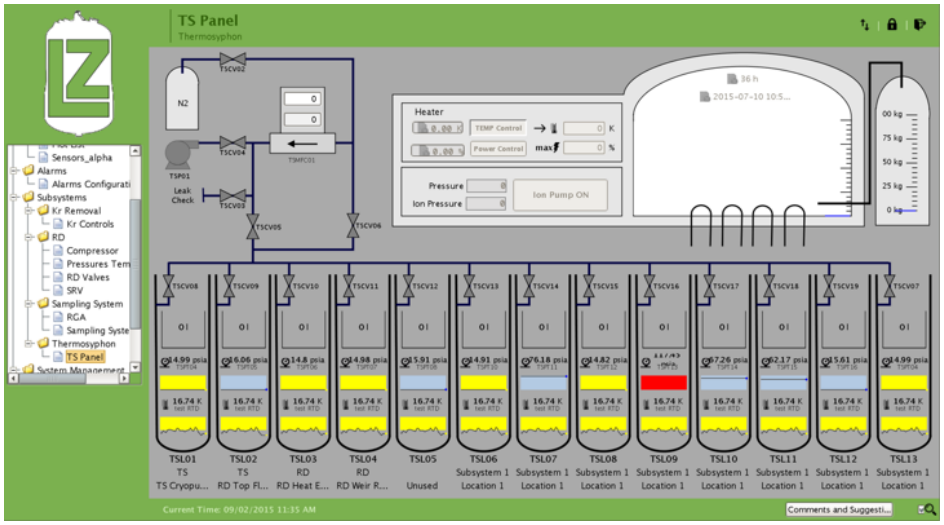

Figure 8.8.2: An example of a low-level panel in the System Test slow control GUI.

In a addition, the GUI will provide panels for historical data plotting (individual, multi- and scatter-plots) and analysis. The lists of sensors, controls and alarms similar to those used in the LUX slow control system will be provided as well. For increased usability, all the viewing panels will be fully customizable so that each user can save personalized views between usages. Finally, administrative panels will give access to advanced system configuration such as tag metadata, user management, etc.

8.8.9 Integration into experiment online IT

The slow control system will seamlessly integrate into the LZ experiment online IT and take advantage of the network security architecture and the authentication/authorization infrastructure provided for LZ online systems (see 8.12). The main servers for slow control (Ignition and the MySQL database) will be run as virtual machines, allowing for an additional layer of redundancy, fault tolerance and easy backup/restore.

With the possible exception of operator consoles where physical access is restricted, all users will authenticate to the slow control system with their own credentials (no shared accounts). All roles and permissions for slow control are maintained in the central directory and made available to Ignition through LDAP.

8.9 Online System

The core of the online system is the run-control system (RC), schematically shown in Figure 8.9.1. It comprises the software and hardware required to allow the operator to define and initiate data collection runs

of different types, control and monitor subsystems (DAQ, slow control, the event builder (EB), and offline software), and log key information to the database. The physical system is hosted on a single rack-mount computer, identical to the hardware used for the event builder. The use of identical hardware with an on-site spare minimizes any downtime of this critical system. RC software is committed to the LZ GIT repository and is maintained on a mirror system throughout the project and operation periods to allow off-line debugging and code development.

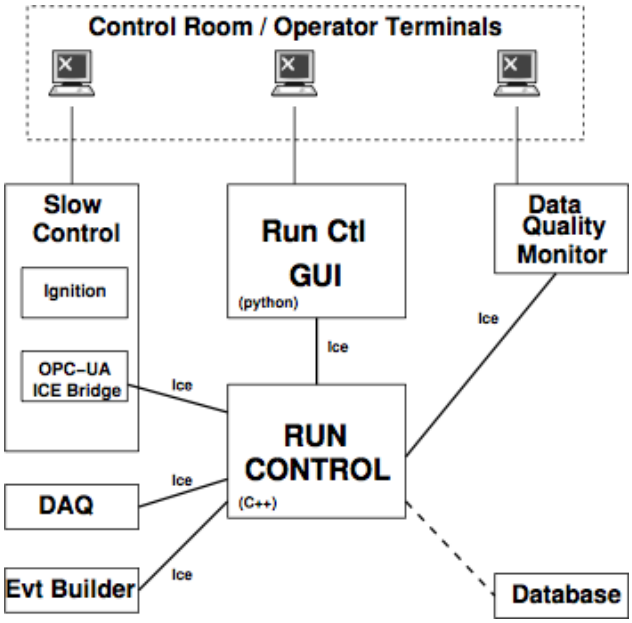

Figure 8.9.1: Block diagram showing the primary interfaces between the RC system and other LZ subsystems. The users use a GUI to interact with the RC system.

The work on the run-control software development started with a definition of the architecture (adopting various software standards), the development of detailed specifications for the software and the identification of key interfaces with the various subsystems. The interplay of run control with slow control is of central importance: Where slow control provides detailed control and monitoring of the detector hardware, run control both initiates high-level slow control scripts (to configure the instrument for various run types) and is the sole entity that configures and controls the data acquisition subsystems (DAQ, EB and DQM). OPC/UA (and similar) standards are used by much of the hardware for slow control, but the middleware (communication protocol) needed for coordinating other systems has different requirements. Thus one critical decision in determining the RC architecture was the selection of the middleware to be used for communications between RC and the subsystems. We have past experience with a number of middleware packages, including the industry standard CORBA (Common Object Request Broker Architecture) system. Given the limitations of CORBA (e.g., scalability, throughput, reliability, documentation and thread-safety) a new solution was considered with similar functionality but improved performance. A recent study [10] to identify middleware to run CERN accelerators ranked ICE (Internet Communications Engine) and ZeroMQ as the top two evaluated systems. The criteria for this evaluation included reliability and speed, and the ability of the system to handle a large number of messages per second (both small and large messages) and publish to a large number of clients in minimal time. Based on our comparisons of different Middleware options, ICE has been identified as the best choice for the LZ run control system. ICE is a remote-procedure-call (RPC) system with open source C++ and Python implementations. In ICE there is no clear distinction of client or server, but it is convenient to think of the subsystem ICE nodes as acting primarily as servers, while the

RC process is the main ICE Client. The RC process also acts as a server for the RC GUI. The low-level RC program is implemented in C++, while the RC GUI is implemented in Python. An API is being developed to allow LZ subsystems to implement ICE communications either through Python or C++. The RC Gui is being developed in PyQt [11] with Qwt for real-time displays.

To provide a seamless interface to the DAQ system, each data collection mode (e.g., calibration runs, data runs) is associated with a corresponding set of command sequences to the various subsystems. The list of data collection modes includes normal data-taking modes (S1 only, S2 only, S1 and S2; see Table 8.5.4 for more details) and a number of calibration modes, including the LED, krypton injection, tritium, and neutron calibration (see Chapter 7 for more details). The RC group interacts closely with the groups responsible for the various subsystems to determine the required interfaces, including defining what status information is shown, what commands are provided, and how status information should be displayed. Based on this information, the RC group has been developing skeleton (server/daemon) program for each subsystem, including the appropriate handles for communication with RC, covering commands and sharing of status information. These skeleton server programs provide the starting point for further development by each group for their specific subsystem, working in close consultation with the RC programmer. This model has already been used to develop the DAQ interface, where the Washington University RC engineer provided a simple DAQ ICE server skeleton program to the DAQ system developer, who subsequently used this code to quickly develop the DAQ server side communications. The initial RC DAQ system is now complete, and was used for instantiating runs in the LZ electronics chain test. The run control/slow control interface is currently under development. As before, the RC engineer is providing code to the SC system programmer to be used in the development of an ICE/OPC-UA bridge to be used for the RC-SC command and communication link.

Throughout the course of the development, the run-control and slow-control groups have worked together closely and try to adhere to common standards for coding, code-management, and common libraries for communication and interfaces. Most commands and status data communication will be handled by the ICE/OPC-UA bridge, providing a bridge between the RC ICE RPC mechanism (optimized for data acquisition), and a tag-based interface primarily using OPC-UA (the standard for slow hardware interface) for Ignition. As with other subsystems, data sharing between run-control and various subsystems can occur through the database, but commands and critical data communication will occur only through the ICE connection between slow-control server computers and the run-control system.

The final RC system provides the user with a simple GUI to the DAQ, data quality monitoring, and the slow-control systems. The GUI starts and stops all subsystems, and displays key status information and alerts if communication links are lost. The system will guide the user through setting up and starting/stopping various data-collection modes and starting runs, making sure all settings, data, and operator comments are logged into the database. The GUI includes key alerts for out-of-bound values, as determined by the slow-control system or other subsystems, and provides a small number of user-selected plots that allow the user to quickly monitor the health of the LZ system.

A skeleton version of the RC software is already complete and was used for the electronics chain test. See below for a screen shot of the RC GUI as of mid-January 2016. During this initial development period, the options for the basic system framework were determined. In the next R& D stage, system hardware will be purchased. The RC group will work with software developers for the other subsystems to develop interface control documents (ICDs) specifying command sequences, settings, and critical data to be exchanged directly through an ICE link or indirectly through the database. Lessons learned from the electronic system test and the feedback from early users will provide input for the final design phase, during which ICDs will be finalized and the full set of interfaces between RC and the subsystem implemented.

The other major component of the online system is the Event Builder (EB). The Event Builder reads the binary files from the 15 individual data collectors (14 for the TPC, skin and outer detector channels,

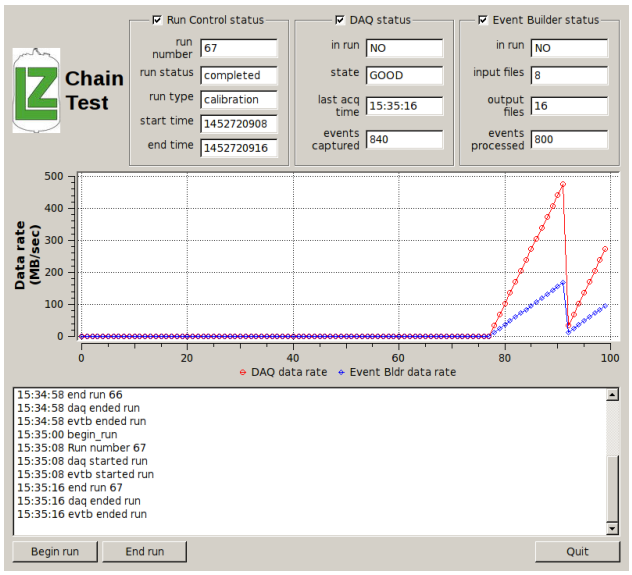

Figure 8.9.2: Run Control GUI for for electronics chain test.

and one for the sparsification/trigger system), assembles the events, and writes out the raw data files on the underground RAID array.

The format and naming convention for both binary input and raw output files is documented in an ICD between WBS 1.8 and 1.11, i.e., between the online and offline systems. Based on the experience from the Daya Bay experiment, it has been decided that the size of the raw output files should be kept at about 1 GB/file. The choice of the raw data format was based on three factors, namely a self-defining data structure, machine independence, and the availability of a generic event viewer (which allows a quick inspection of the data without relying on an experiment-specific framework). This narrowed down the choices to HDF5 and `root`. The raw output format was chosen to be based on `root`. An HDF5-based EB was also developed, but its performance (both in speed and internal compression) was shown to be inferior to that of the `root`-based EB. The EB needs to be able to keep up with the highest data rate expected in the LZ detector, e.g., about 140 Hz during ^{83}Kr calibrations. Therefore, assuming a safety factor of 2 to 3, the EB must be able to operate at 280 Hz to 420 Hz. The internal compression is defined to be acceptable if it is comparable to the level achieved by the standard gzip compression algorithm on the binary files.

The EB is controlled by Run Control, which initiates the start of the EB for a particular run. The EB reports back to Run Control the current event and file number with a frequency that can be adjusted through a parameter. The end of run (of predetermined duration or number of events) is propagated to the EB directly through a flag included in the binary files, but can also be initiated directly by Run Control if the DAQ is non-responsive. Once a run has been successfully completed by the EB, Run Control is notified by the EB, which in turn releases the binary files on the data collectors to be deleted. All communications between Run Control and the EB occur through ICE.

Because the RC code will be developed and maintained by an experienced professional programmer, it is expected that the RC system (and other subsystem communication links) will be well maintained throughout the LZ project and operation periods. The RC software engineer will ensure that other online systems adhere to the ICDs, and use the software development standards decided by the WBS 1.8 group. The RC software engineer will assist with other online programming tasks as needed, lending support and responding to user requests for changes in the online system as needed.

8.10 Electronics Chain Test Facility

8.10.1 Description and goals of the Electronics chain test facility

To test all the components of the signal processing chain, an electronics chain test facility (ECT) has been set up at the University of Rochester [12]. The main goals of the ECT are:

1. **Evaluate Analog Electronics:**
   - **Noise and cross-talk.** To measure the noise and cross-talk of the amplifiers, digitizers, and PMT bases.
   - **Signal propagation.** To measure and optimize shaping, gain, linearity, and saturation.
2. **Evaluate Digital Electronics:**
   - **Firmware and software development.** The ECT provides a direct hardware testing platform minimizing the need for time-consuming simulations.
   - **Clock distribution.** Determine if we can distribute the clock over the HDMI links or if we have to use dedicated clock distribution resources.
   - **System performance.** Verify that the system is capable of handling the expected data rates:
     - SerDes links utilization.
     - P-2-P Gigabit UDP link utilization.
     - Data Collector CPU/disk speed utilization (storing on local disks and providing the data to the event builder).
3. **Develop Online Systems:**
   - **Run Control.** Optimize interactions between run control and DAQ and DS systems.
   - **Event Building.** Optimize event building with data collected with the entire electronics chain.
   - **Data Quality Monitor.** Develop software tools to monitor the quality of the data collected.

In the first phase of the ECT, we have a single chain setup shown in Figure 8.10.1. This stage allows us to concentrate on the quality of signal propagation from PMT to disk. The first results from phase 1 are described in more detail in the remainder of this Section.

Figure 8.10.2 shows the chain test setup in its second phase. The second phase is geared towards developing and scaling the DAQ system.

8.10.2 Single photoelectron pulse propagation

The initial measurements carried out at the ECT focused on SPHE pulses from a single R11410-20 PMT. These studies allowed use characterize signal-to-noise ratios and timing resolution of the electronics chain.

Figure 8.10.3a shows an intensity map and an average of 1000 SPHE pulses captured at the signal flange (before amplifiers), using a Tektronix MSO4104 oscilloscope with 250 MHz set bandwidth and 2.5 GHz sampling. Figure 8.10.3b shows a corresponding SPHE pulse captured at the output of the LE amplifier channel.

The R11410-20 are expected to operate at 3.5×10^6 gain, but since the manufacturer guarantees them to 2×10^6, we measured the SPHE response for both scenarios. Figure 8.10.4 shows two SPHE spectra obtained at PMT gains of 3.25×10^6 and 1.9×10^6. The signals were propagated through the entire analog chain composed of 45 feet of gore cable, the amplifiers, and 30 feet of LMR-100 cable, before being digitized

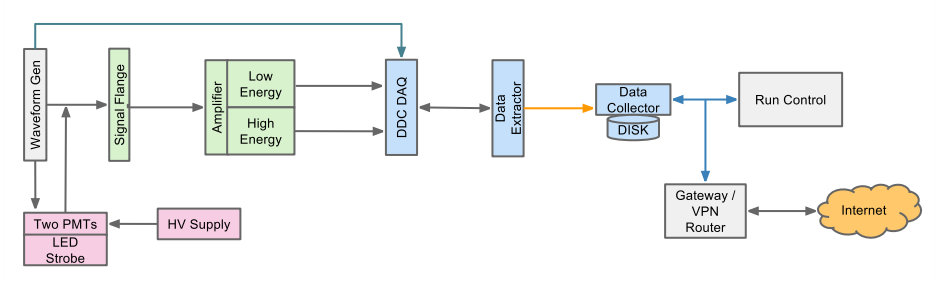

Figure 8.10.1: Diagram of the first stage of the Chain Test setup. The setup includes: two LZ PMTs, LZ internal signal and HV cable, LZ signal flange, LZ amplifier cage with one LZ amplifier board, one digitization chain, Run Control, and OpenVPN remote access.

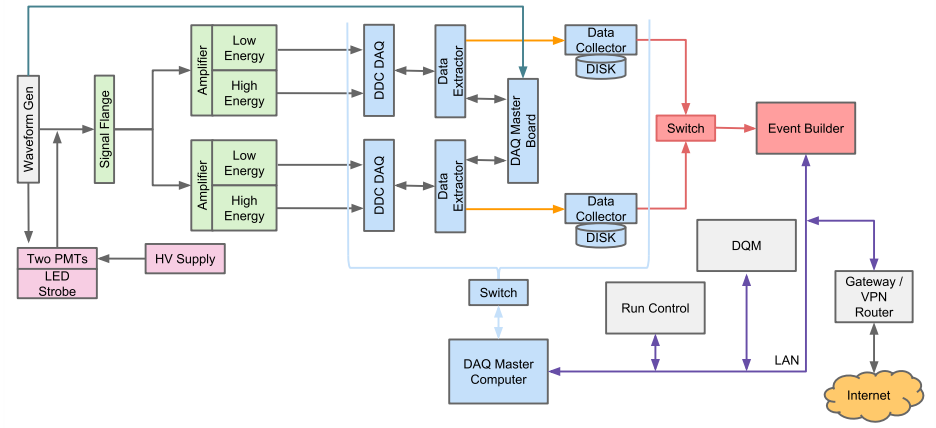

Figure 8.10.2: Diagram of the second phase of the Electronics Chain Test Facility [12]. The setup includes: two LZ PMTs, LZ internal signal and HV cable, LZ signal flange, LZ amplifier cage with two LZ amplifier boards, two parallel digitization chains, Event Builder, Run Control, DQM, DAQ Master Computer and OpenVPN remote access.

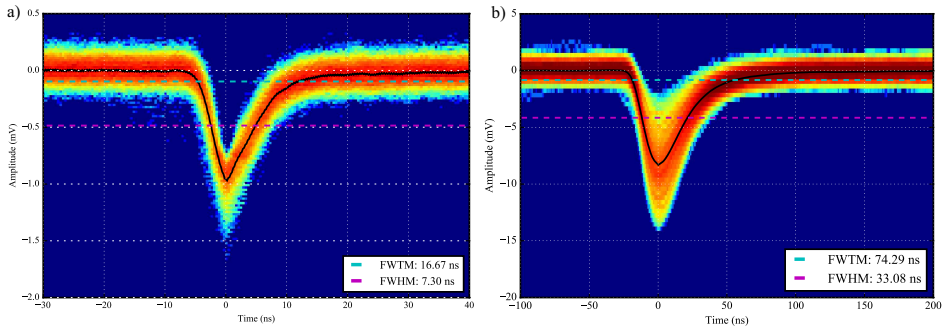

Figure 8.10.3: Single photoelectron intensity plots. (a) LZ PMT @ 1,300 V → 45 ft GORE cable → Signal Flange → 6 ft LMR-100 → Scope, (b) LZ PMT @ 1,300 V → 45 ft GORE cable → Signal Flange → LZ Amplifier LE output → 30 ft LMR-100 cable → Scope

on the DDC-32. The waveforms were passed through an S1 filter and the filters' peak output (proportional to area) was histogrammed. The filters' width was set to 3 samples to optimize the peak to baseline noise ratio. The measured detection efficiency of 98 % (at 3.25×10^6) and 92 % (at 1.9×10^6) are obtained with only 1 Hz of singles rate (where the S1 filter captures a fluctuation in the baseline noise).

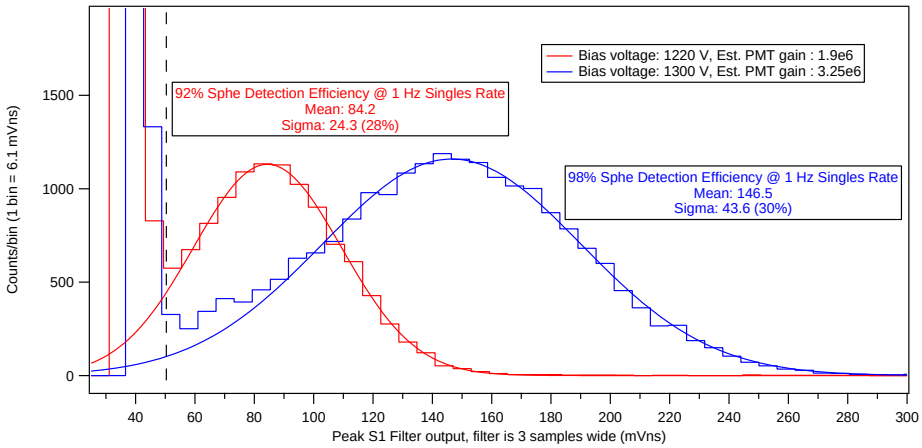

Figure 8.10.4: Single photoelectron spectrum at estimated 1.9×10^6 and 3.25×10^6 gain. Detection efficiency at 1 Hz singles rate is 92 % and 98 %, respectively.

In order to determine how well we can reconstruct the relative timing between SPHE signals with a 60 ns shaping time (FWTM), digitized at 100 MHz, we carried our a series of measurements with short light pulses from an LED. As light source we used the CAEN SP5601 LED pulser, which has a measured resolution of 0.44 ns (1σ) [13]. To limit the impact of the PMT's transit time spread (TTS) which has been measured to be 2.55 ns (1σ) [14] for photons distributed across the entire photocathode, an aperture (1 cm diameter) was placed in front of the PMT. The PMT was operated with a gain of 3.25×10^6. Figure 8.10.5a shows that measured reconstructed time difference between the LE SPHE signals and the LED pulser start signal has a resolution of 1.65 ns (1σ). It should be noted that the resolution obtained between two LED pulser start signals, also digitized with the DDC-32, is 54 ps (1σ) (Fig. 8.10.5b).

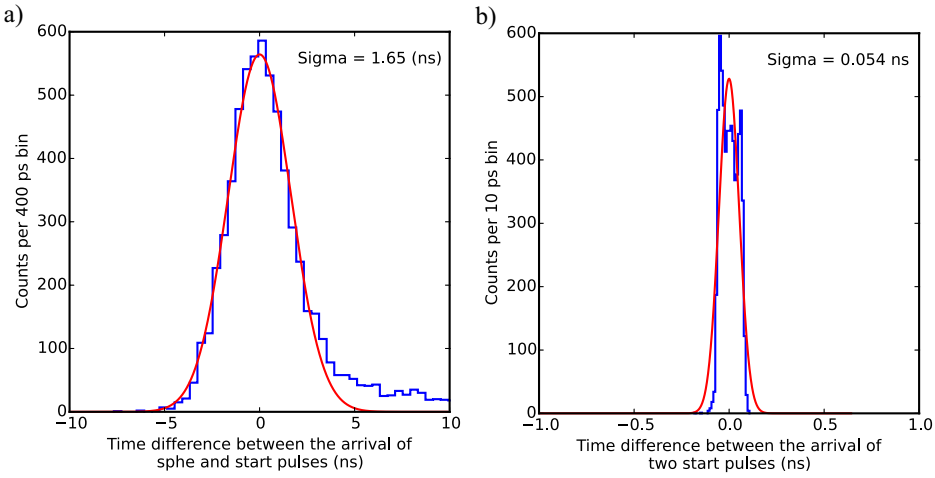

Figure 8.10.5: Time difference distributions for: a) LE SPHE signals obtained with a gain of 3.25×10^6, b) two LED start pulses. The time range in b) is ten times smaller than the time range in a).

8.10.3 PMT Base Saturation

The PMT base was evaluated for saturation with S2-like pulses. Wide pulses were sent to a 425 nm LED and varied in amplitude to produce S2-like responses in the PMT with a width of 0.5 µs (1σ). To determine the amount of light emitted by the LED, we collected data in two modes: 1) with a 1 cm diameter aperture in front of the PMT to limit the amount of incident light in order to calibrate the LED, and 2) without the aperture where we expect PMT saturation for large LED signals. Signals from the PMT are propagated through 45 feet of Gore cable to the DB25 connector on the signal flange and digitized on the Tektronix MSO4104 oscilloscope along with the signals sent to the LED. Figure 8.10.6 shows the averaged PMT pulses collected without the aperture for different LED pulses.

With the aperture, the PMT pulse area increases linearly with the LED pulse area, demonstrating that as the amplitude of the pulse to the LED increases the amount of light being generated increases. This allows us to calibrate the amount of light incident onto the PMT. In the measurements without the aperture, we know that the LED pulse area is proportional to the light emitted by the LED. Figure 8.10.7 shows a plot of the measured PMT pulse area when the aperture is removed versus the PMT pulse area we would have obtained if the aperture would have been in front of the PMT. This figure shows that for the two largest LED amplitudes, the PMT base saturates.

In the region of saturation, the shape of the response of the PMT changes, as can be see in Figure 8.10.6. The saturation manifests itself in the form of narrowing the shape of the pulse for larger light input to the PMT and no increase in pulse area. Saturation occurs when the amplitude of the signal from the PMT at the input of the oscilloscope exceeds 0.6 V.

8.10.4 DAQ Code and Performance

The main Data Collector acquisition code is named DCCore and has been rewritten in C++ using Qt based multithreading framework. It is responsible for receiving UDP datagrams from the Data Extractors, performing consistency checks on it and storing it to a local SSD drive. The DCCore successfully communicates with the Run Control over ICE by responding to run start/stop commands and reporting the acquisition status and progress.

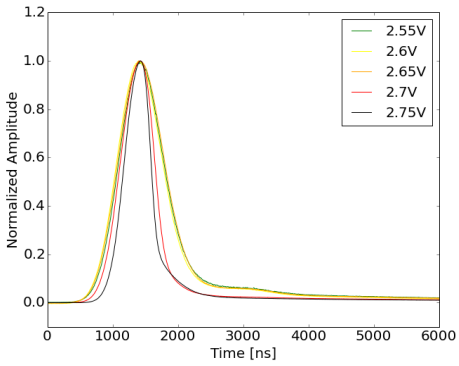

Figure 8.10.6: Average peak aligned and normalized pulses showing the S2-like response of the PMT. The legend corresponds to the peak amplitude of the pulse used to drive the LED. The first signs of saturation begin to manifest itself for an LED amplitude of 2.65 V.

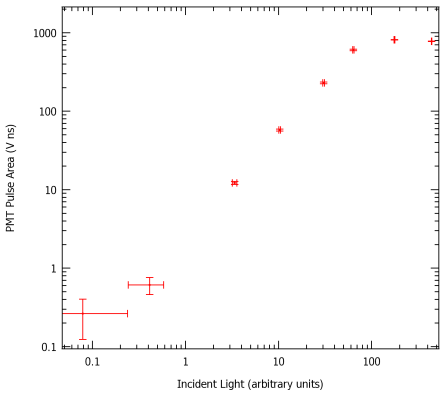

Figure 8.10.7: Result of measured pulse area detected by the PMT without a hole versus the expected pulse area of the PMT with a hole. As the expected amount of light increase the PMT's measured response increases till it begins to saturate for the last two points.

During active data collection, at the rate of 70 MB/s, the DCCore has been shown to utilize a single CPU core at only ∼15 % level. Currently the measurement was done with the peak rate of 70 MB/s, because the FPGA waveform capture firmware is implemented in a single-buffering mode and thus not push the Gigabit UDP link to its maximum of 109 MB/s. The single-buffered rate is understood and agrees perfectly with what we estimate in this scenario.

Additionally a Python based DAQScope code has been developed. It communicates with the DCCore using ICE. It is a GUI tool that allows for real-time preview of the waveforms incoming into the Data Collector. It also allows for continuous calculation of the waveforms' power spectral density (PSD) plots. The screenshot of the tool in operation is shown in Figure 8.10.8

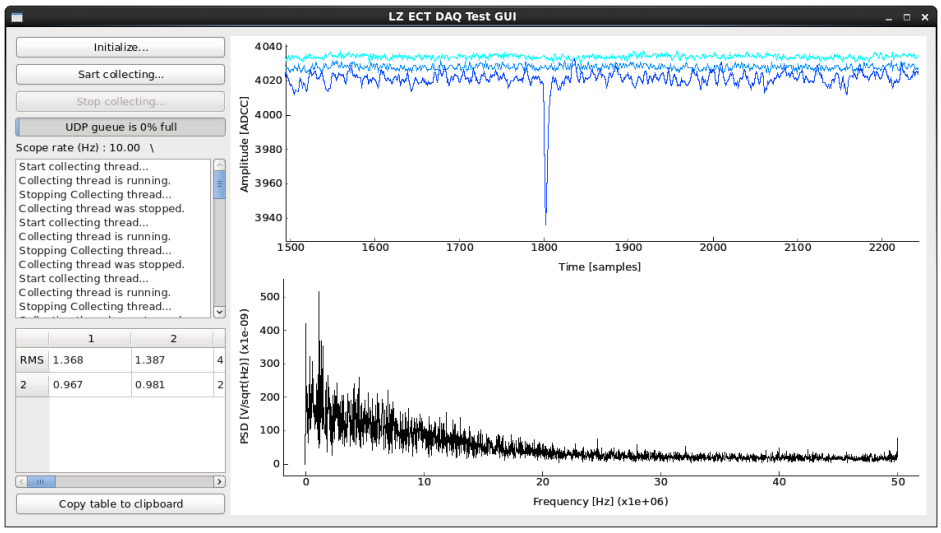

Figure 8.10.8: A screenshot of the DAQScope tool displaying a subsample of the waveforms incoming into a given Data Collector and showing a continuously calculated PSD spectrum for one of the channels.

8.10.5 Real Signals from LUX

At the end of LUX, we had the opportunity to replace the existing signal chain with the LZ signal chain from amplifier to disk. This allowed us to benchmark the performance of our electronics against those used in an active dark matter detector and to look at real signals. For these tests, two LZ amplifiers were mounted on a LUX signal flange that had signals originating from both the top and bottom PMT arrays. Dual gain signals from the amplifier were routed through 45 feet of LMR-100 cable to be digitized on a prototype DDC32. Data was extracted from the DDC32 via and HDMI link to a DE board and then pushed to a DC where it was stored in a partial event file format. The LUX grounding scheme was used for the electronics chain.

Noise measurements, SPHE calibrations, and real event signals from LUX were collected during these tests. The gain corrected noise measurements are shown in Figure 8.10.9. These measurements show that, even though no special consideration was given to optimize the grounding scheme, the LZ system performs as well as the LUX systems. From SPHE calibrations we measure a SPHE detection efficiency greater than 90%. An example of the S1 filter output for SPHEs is shown in Figure 8.10.10.

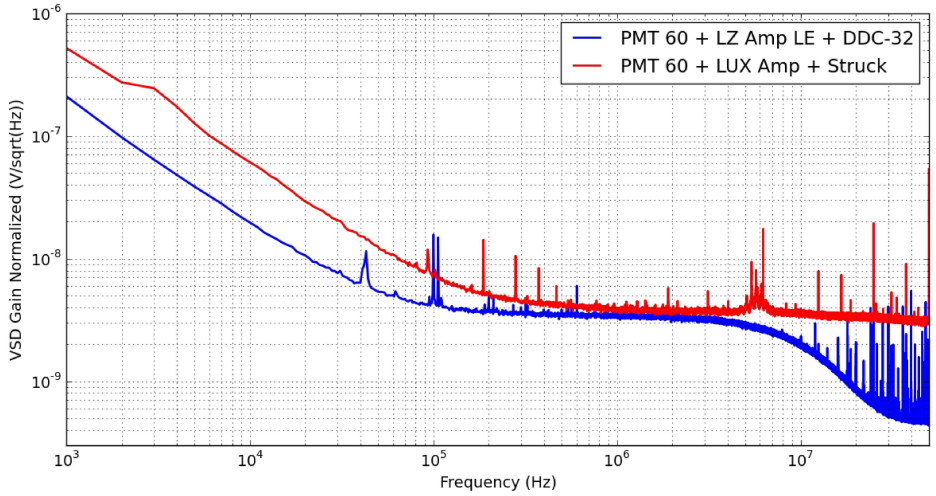

Figure 8.10.9: Gain normalized noise spectra compared between the LUX and LZ DAQ chains.

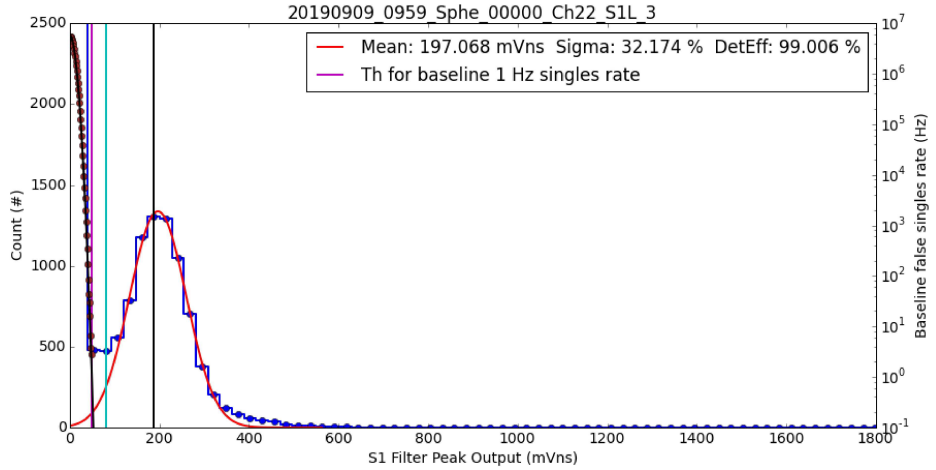

Figure 8.10.10: Maximum S1 filter output for SPHEs obtained with the LZ DAQ electronics installed on the LUX detector.

8.11 Grounding

The LZ experiment comprises many electrical and electronic sub-systems that may be adversely affected by electrical noise. Each sub-system in the LZ experiment will be guarded from noise induced by other sub-systems. Guarding will also reduce noise introduced into other sub-systems. The PMT signal collection chain is one of the most sensitive and most critical sub-systems in LZ. PMT signals are in the sub-mV range before amplification so even a small amount of noise has a potentially large impact. PMT signal quality directly affects the quality of the physics data therefore signal integrity is of great concern.

Unwanted signals or noise can be coupled from an "aggressor" circuit into a "victim" circuit through capacitive coupling, inductive coupling and conduction. Proper grounding can be critical to reducing interference through these mechanisms. Grounding plans for experiments of the size of LZ are complex and require significant engineering effort.

The LUX grounding system serves as the starting point for the LZ grounding scheme. The main lesson learned in implementing the LUX grounding system was the importance of a single-point or "star" ground configuration (Figure 8.11.1 a). All of the equipment in the electronics racks was grounded to the rack grounding bars using flat braided cable (Figure 8.11.1 b) and all of the grounding rods have dedicated grounding paths to the central star point using 2 AWG copper welding cable.

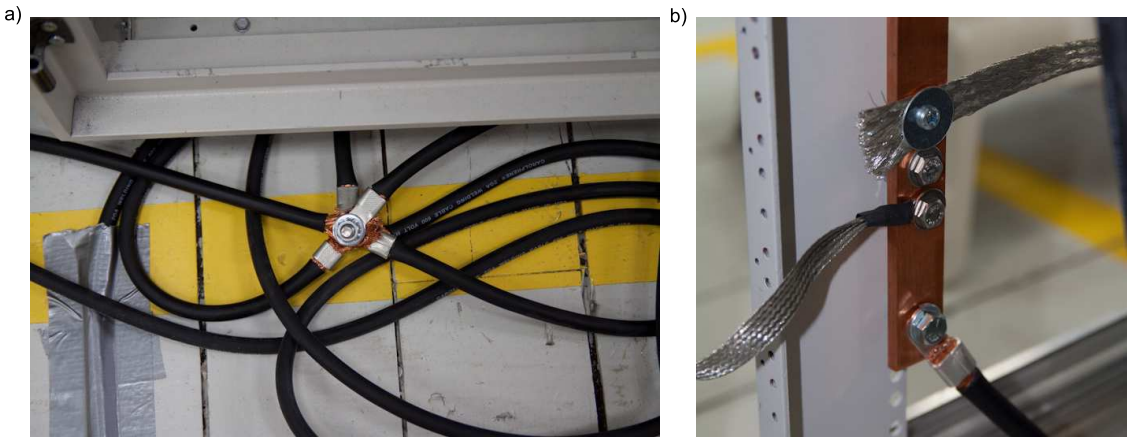

Figure 8.11.1: a) Star ground central point to all the rack grounds. b) Sample grounding connections within one of the racks, where individual equipment and the rack door are grounded with steel braid to the rack grounding rod. The grounding rod is connected to the center star grounding point using a high gauge copper welding cable.

For many of the circuits in LZ, the signal cable shield also acts as the signal return path. In order for the shield to be effective, it must not conduct significant amounts of current. If excess current does flow in the shield, it will show up as noise in the signal. This can occur when the "ground" potential is different at each end of the cable. In situations such as this, galvanic isolation can break the flow of current and eliminate the induced noise. We are currently testing the feasibility of using isolation transformers on the signal lines between the amplifier outputs (at the breakout) and the digitizer inputs (30 ft of cable away in the electronic racks above). We are testing two candidates from Mini-Circuits: FTB-1-6+ and FTB-1-1+. As these are typically meant for RF applications, we need to quantify how these transformers impact the pulse time characteristics of typical amplified LZ PMT signals. If the transformers do not significantly degrade the signal quality they should allow for minimizing potential ground loops between the breakout area and the upper electronics racks.

Significant noise can also be generated by electrical equipment such as pumps, chillers, and blowers that draw significant current or have switched loads that might cause electrical transients. Interference from these sources will be eliminated or greatly reduced by isolating them on separate electrical circuits and by physically locating them as far from the sensitive PMT signal cables as practical. Capacitive and inductive coupling to PMT signal cables can occur if cables from other sub-systems run in close proximity to the PMT cables for significant distances. Again, by placing the PMT signal cables as far as is practical from other sub-system cables, we will greatly reduce the effects of capacitive and inductive coupling.

Ideally we will be able to identify all sources of potential noise and eliminate them during the design stages. But experience shows that it is unlikely that every noise source will be identified prior to the installation and commissioning of the system. During deployment/construction of the LZ experiment we will, early on, have a subset of the DAQ system installed such that we can constantly monitor PMT signals for any noise pickup as equipment is installed. If we see any noise being picked up from other systems during the installation of equipment for the project, the data we gather from the DAQ system will greatly simplify pinpointing the source and allow us to put in place fixes to ensure excellent signal integrity.

8.12 Network Infrastructure, Security, Remote Access

The LZ network and computing infrastructure will be designed according to best practices to meet the LZ requirements for performance, security and availability / operational reliability and uptime. A prototype has been developed and successfully demonstrated at the SLAC LZ System Test Facility, discussed in Section 3.10.

8.12.1 Local Area Network

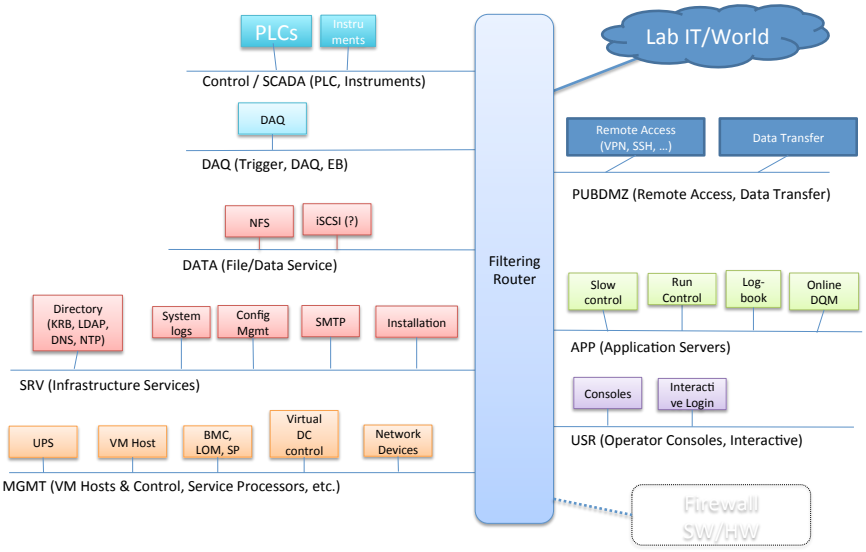

Figure 8.12.1: The LZ Online Network

The LZ experiment LAN in the Davis Campus will be implemented with a single central router/switch that acts as communication hub for the experiment, provides separate network zones using Virtual LANs (VLANs) and IP subnets and enforces inter-zone network restrictions using simple access control lists (ACLs). Network zoning will be used to implement a standard network security architecture with "Demilitarized Zones" (DMZ) and restricted networks. If necessary, enhanced network security controls can be implemented with host-based firewalls or a dedicated firewall appliance. Figure 8.12.1 shows a logical diagram of the LZ experiment LAN.

8.12.2 Reliability and Availability

Connection to the surface will use dedicated and redundant fiber connections on physically separate paths (one fiber path through Ross shaft, one fiber path through Yates shaft). The central LZ router/switch located in the Davis Campus will be implemented as redundant stack of two switches with full power, uplink and control redundancy. All critical devices such as main PLCs and critical computing infrastructure will have redundant Ethernet connections to both stack elements. A separate, dedicated network connection to the surface will allow emergency out-of-band access to the PLCs and the network and IT infrastructure underground.

8.12.3 Information Security

The LZ computing architecture will implement a defense-in-depth security design based on a zoned network, central authentication and authorization (Kerberos and LDAP) with appropriate account-lockout policies, a configuration management system for provisioning, managing and patching hosts, central logging and monitoring, and frequent backups for recovery from a wide class of problems. Advanced firewalls, network- or host-based intrusion detection, or other security controls can be easily added if deemed necessary.

Network zoning will be used to create a layered configuration where devices with weak security controls or network stacks (such as embedded systems, sensors, PLCs) are only allowed to communicate with trusted and managed systems.

Central management of authentication and authorization in a "Directory" based on Kerberos and LDAP will be used to manage user accounts and all privileges in the LZ system, including host access privileges and slow control roles. This will allow for easy enforcement of credential requirements (such as minimum password complexity or password expiration), secure auditing and logging of user privileges and lockout of user accounts if necessary. It will also allow to securely implement user self-service (e.g. for password reset or contact information change) and delegation of privilege management to subsystem managers (e.g. the slow control manager can assign roles to users without having to ask a system administrator to make the changes).

A configuration management system (such as e.g. SaltStack [15]) will be used to provision hosts with a secure baseline configuration and to maintain and update the host configuration as necessary. Using such a system minimizes the amount of effort required for system management and provides a way to detect unauthorized changes relative to the baseline configuration.

Central logging and monitoring provides another avenue to detect unauthorized activities, especially when combined with simple keyword-based automated log analysis, but is also a valuable tool for general non-security related health monitoring, diagnosis and debugging of the system.

Where possible, frequent downtime-less file system snapshots and backups of all important system configuration and user/application data will be performed. Access to backup files will be restricted so that even a system compromise will not allow the attacker to modify historical backup data.

Virtualization technologies may be used to reduce the effort required for system management and to provide virtual machines that can be easily controlled remotely by authorized users and that can be easily backed up.

8.12.4 Remote Access

The implementation of secure remote access will depend on the applications that need to be accessed and the client devices (desktop, laptop, tablet/phone) accessing them. All remote access protocols must use secure and encrypted communication and strong authentication mechanisms. While the details are still under study, we foresee some combination of web, command line, remote desktop and application client access over encrypted communication channels provided by a VPN and/or SSH.

8.13 Installation and Commissioning

The electronics will be installed in two locations. The analog electronics for the Xe PMTs will be installed on the mezzanine level, as shown schematically in Figure 8.13.1. The amplifiers will be arranged in 23 racks, with 4 amplifier cards installed in each rack. The racks will be located horizontally on the wall of the water tank in order to improve cooling and provide better access to each amplifier crate. Exhaust ducts may must be installed at this location to facilitate cooling of the electronics. Because it will be difficult to enclose the electronics crates at this location, it is important the air flowing over the electronics is directed into the exhaust system. This will prevent any smoke generation as a result of an electrical problem, e.g., overheating of components, from spreading into the Davis Campus. The outputs of the amplifiers are routed in cable trays to the DAQ and trigger systems installed in the electronics racks installed at deck level (see Figures 8.13.1 and 8.13.2).

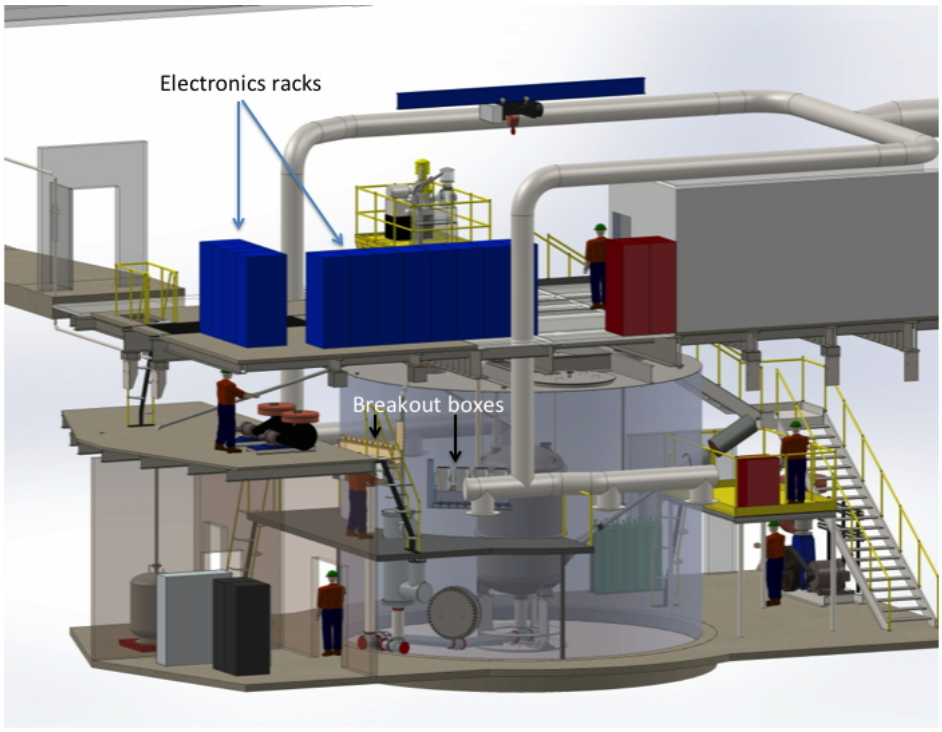

Figure 8.13.1: Installation of the analog electronics for the xenon PMTs at the mezzanine level and the electronics racks at the deck level.

A schematic of one element of the breakout box is shown in Figure 8.13.3. The internal signal and HV cables enter this section of the breakout box through the ports in he center. The signal cables are routed to

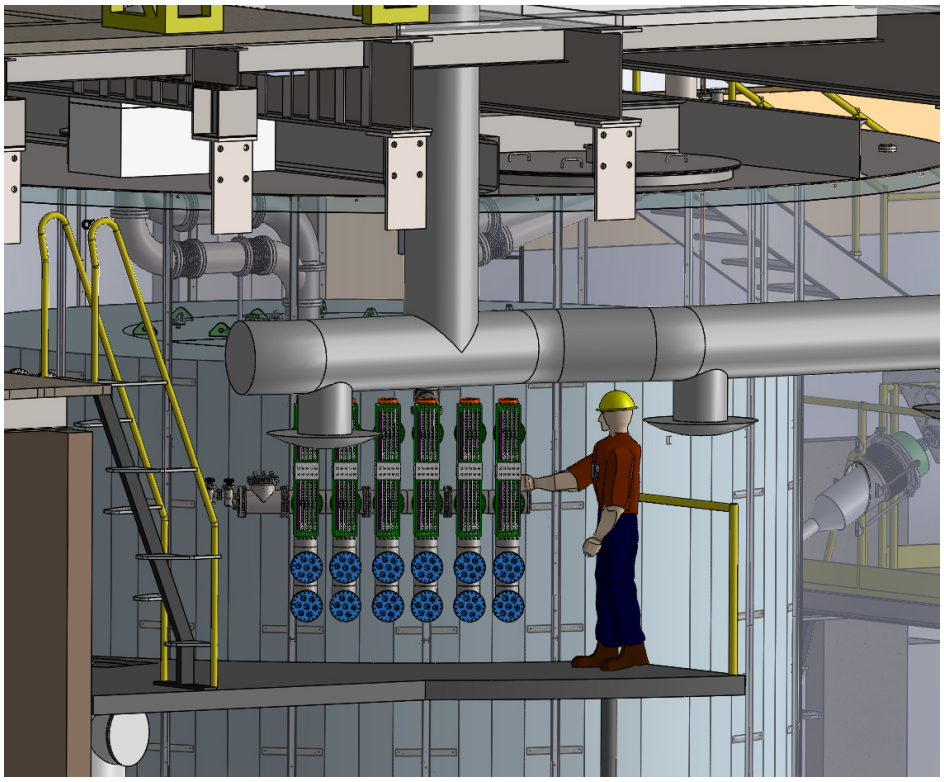

Figure 8.13.2: Close-up of the mezzanine level showing the amplifier racks installed on the top breakout box. The HV filter boxes are not shown.

the two flanges, shown in yellow, with four DB-25 connectors on which the amplifier crates mount. The HV cables are routed to the HV flange, shown in blue, visible at the bottom of this breakout section.

The amplifiers and the HV pickoffs for the outer-detector system will be installed in the electronics racks shown in Figure 8.13.1. The HV cables from the outer-detector PMTs are routed via one of the flanges on top of the water tank to this location. The design of the flange is such that no connectors are required to ensure a light-tight and radon-tight connection. Not using connectors on the flange improves signal quality.

A total of 11 electronics racks are installed at the deck level, as shown in Figure 8.13.1, to provide space for all electronics systems, except the xenon amplifiers. The racks are fully enclosed and the forced airflow across the electronics is directed into the exhaust system of the Davis Laboratory. Figure 8.13.4 shows the layout of the components installed in these 11 racks.

The power requirements of the electronics system in the Davis Laboratory have been estimated based on the measured power consumption of the various components of the system. The power required by the amplifiers, the DAQ, the computer systems, and the PMT HV systems is 26 kW. The power required for the various calibration systems for the outer detector and the xenon PMTs remains to be determined. The power required by the electronics is provided by uninterruptible power supplies, installed at the bottom of each electronics rack. The total maximum load of the electronics is 94 kW. This is the load of the equipment at startup.

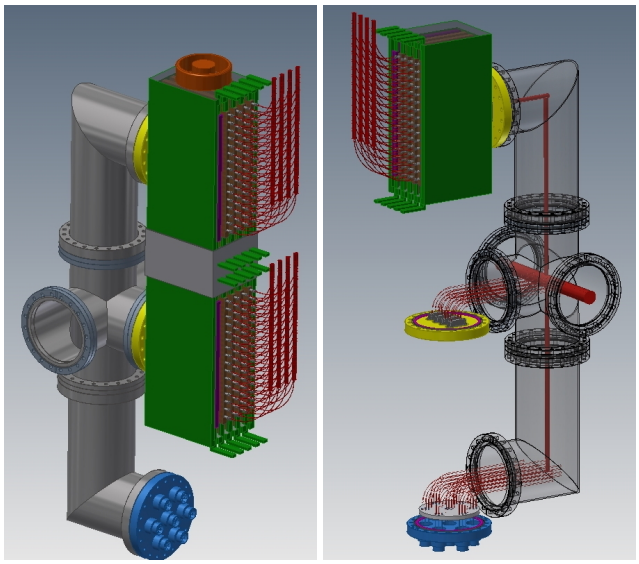

Figure 8.13.3: Close-up of one element of the breakout box. The figure on the left shows two amplifier cages installed on the top two signal flanges (yellow) of this section. The HV flange is visible at the bottom (blue). The figure on the right shows the routing of the signal and HV cables in this breakout section.

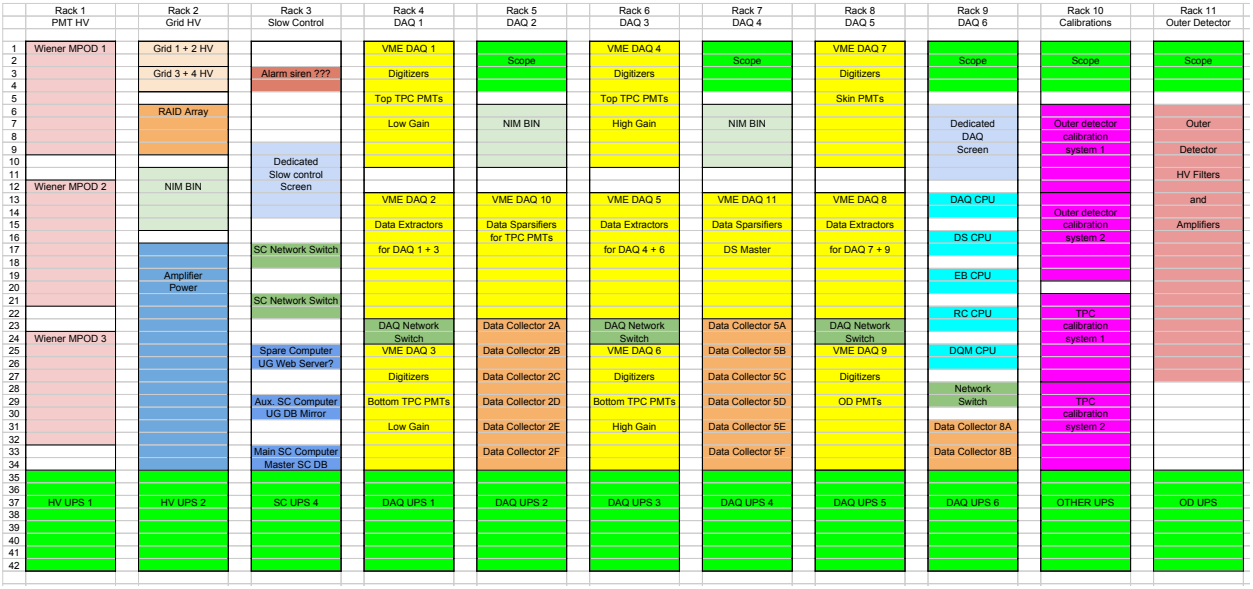

Figure 8.13.4: Detailed layout of the LZ electronics racks, installed at the deck level.

8.14 Bibliography

[1] J. Yin (LZ), *Proceedings, Light Detection in Noble Elements (LIDINE 2015): Albany, NY, USA, August 28-30, 2015*, JINST **11**, C02054 (2016), arXiv:1511.00260 [physics.ins-det].

[2] *All Programmable FPGAs and 3D ICs*, Xilinx Inc., 2100 Logic Drive, San Jose, CA 95124-3400, USA (2016).

[3] E. Druszkiewicz (LZ), *Proceedings, Light Detection in Noble Elements (LIDINE 2015): Albany, NY, USA, August 28-30, 2015*, JINST **11**, C02072 (2016), arXiv:1511.08385 [physics.ins-det].

[4] D. S. Akerib *et al.* (LUX), Nucl. Instrum. Meth. **A668**, 1 (2012), arXiv:1108.1836 [astro-ph].

[5] E. Druszkiewicz, *Digital Advances in Triggering and Data Acquisition Systems for Large Scale Dark Matter Search Experiments*, Ph.D. thesis, University of Rochester (2017), [manuscript in prepration].

[6] D. Salomon and G. Motta, *Handbook of Data Compression* (Springer London, 2010) ISBN 978-1-84882-902-2.

[7] M. Henning, IEEE Internet Computing **8**, 66 (2004).

[8] *Ignition SQL Bridge*, Inductive Automation, 340 Palladio Parkway, Suite 540 (4th Floor) Folsom, CA 95630-8833, USA (2015).

[9] *Siemens Simatic PCS 7 CPU 410-5H Contoller* (2016).

[10] A. Dworak, P. Charrue, F. Ehm, W. Sliwinski, and M. Sobczak, in *13th International Conference on Accelerator and Large Experimental Physics Control Systems, 10-14th October 2011, WTC Grenoble, France* (2011) pp. 1334–1337, *Middleware Trends And Market Leaders 2011*.

[11] *Introduction to PyQt*, Riverbank Computing Limited, Redcotts House, 1 Redcotts Lane, Wimborne, Dorset BH21 1JX, UK (2016).

[12] D. Khaitan (LZ), *Proceedings, Light Detection in Noble Elements (LIDINE 2015): Albany, NY, USA, August 28-30, 2015*, JINST **11**, C03029 (2016), arXiv:1511.07752 [physics.ins-det].

[13] *SP5601 Led Driver - Data Sheet* (2014), CAEN S.p.A., Via Vetraia, 11 55049 Viareggio (LU), Italy.

[14] K. Lung *et al.*, Nucl. Instrum. Meth. **A696**, 32 (2012), arXiv:1202.2628 [physics.ins-det].

[15] *SaltStack system management software* (2016).