CompactDAQ: Compact Data Acquisition for Pulse-Shape Discrimination

Prince John
Roger Chamberlain
George Engel
Jeremy Koertzen
A. B. M. Rafi Sazzad
Jinke Xu
Jonathan Elson
Lee G. Sobotka

Prince John, Roger Chamberlain, George Engel, Jeremy Koertzen, A. B. M. Rafi Sazzad, Jinke Xu, Jonathan Elson, and Lee G. Sobotka, "CompactDAQ: Compact Data Acquisition for Pulse-Shape Discrimination," in *Proc. of Great Lakes Symposium on VLSI* (GLSVLSI), June 2026, pp. 862-869. DOI: 10.1145/3787109.3815236

Dept. of Computer Science and Engineering
Dept. of Chemistry
Washington University in St. Louis

Dept. of Electrical and Computer Engineering
Southern Illinois University Edwardsville

Lawrence Livermore National Laboratory

CompactDAQ: Compact Data Acquisition for Pulse-Shape Discrimination

Prince John
Computer Science and Engineering
Washington University in St. Louis
Saint Louis, MO, USA
johnp@wustl.edu

Roger Chamberlain
Computer Science and Engineering
Washington University in St. Louis
Saint Louis, MO, USA
roger@wustl.edu

George Engel
Electrical and Computer Engineering
Southern Illinois University Edwardsville
Edwardsville, IL, USA
gengel@siue.edu

Jeremy Koertzen
Lawrence Livermore National Laboratory
Livermore, CA, USA
koertzen1@llnl.gov

A. B. M. Rafi Sazzad
Lawrence Livermore National Laboratory
Livermore, CA, USA
sazzad1@llnl.gov

Jingke Xu
Lawrence Livermore National Laboratory
Livermore, CA, USA
jingkexu@llnl.gov

Jonathan Elson
Chemistry
Washington University in St. Louis
Saint Louis, MO, USA
jmelson@wustl.edu

Lee Sobotka
Chemistry
Washington University in St. Louis
Saint Louis, MO, USA
lgs@wustl.edu

Abstract

Discriminating between neutron and gamma interactions in scintillators is essential for many nuclear science experiments. Pulse-shape discrimination (PSD), based on comparing charge integrations over different regions of a scintillation pulse, is a well-established technique for this purpose. Fully digital PSD data acquisition systems offer high precision and flexibility but generate large data volumes that hinder scalability across thousands of detector channels. Conversely, traditional analog PSD approaches and existing ASIC-based implementations achieve low-latency, high-fidelity processing but rely on bulky, difficult-to-scale readout architectures.

We present CompactDAQ, a modular and highly integrated, end-to-end data acquisition system that unifies these approaches within a scalable and extensible architecture. Each 16-channel CompactDAQ board combines custom analog integration and triggering ASICs with on-board CMOS delay devices to perform real-time, pulse-shape processing and data sparsification. This design preserves the precision and timing performance of analog systems while dramatically reducing size, interconnect and readout complexity, enabling compact, scalable deployment of PSD-capable detector arrays with 1000s of channels.

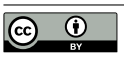
CCS Concepts

• **Hardware → Sensor devices and platforms**; **PCB design and layout**; • **Applied computing** → *Physics*; *Chemistry*.

Keywords

Nuclear science, pulse shape discrimination, gamma ray, neutron

ACM Reference Format:
Prince John, Roger Chamberlain, George Engel, Jeremy Koertzen, A. B. M. Rafi Sazzad, Jingke Xu, Jonathan Elson, and Lee Sobotka. 2026. CompactDAQ: Compact Data Acquisition for Pulse-Shape Discrimination. In *Great Lakes Symposium on VLSI 2026 (GLSVLSI '26), June 22–24, 2026, Canandaigua, NY, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3787109.3815236



This work is licensed under a Creative Commons Attribution 4.0 International License.
GLSVLSI '26, Canandaigua, NY, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2431-2/26/06
https://doi.org/10.1145/3787109.3815236

1 Introduction

Modern nuclear science experiments rely on high-channel-count detector arrays to measure position, energy, timing, and pulse-shape characteristics of radiation events with high precision. Historically, such measurements have been performed using chains of discrete analog modules—constant-fraction discriminators, charge-to-digital converters, and time-to-digital converters—all assembled and housed in large CAMAC, NIM or VME crates. While these systems have proven reliable and remain widely used in large-scale detector arrays, their bulk, wiring complexity, and synchronization requirements limit their practicality for desktop-scale setups. Over the last two decades the significant reduction in cost and increase in availability of high–sample-rate digitizers and flash ADCs have enabled direct waveform capture at the detector front end, but this approach produces massive data volumes and does not inherently address system scalability. CompactDAQ, the system developed in this work, takes a different approach: an ASIC-based data acquisition (DAQ) platform that unifies analog front-end processing, on-board timing, and digital communication within a single, compact, and scalable architecture. In short, it is a complete system with a small footprint. The accompanying firmware and software stack are designed from the outset to support both single-board, desktop-scale configurations and large multi-board arrays through a unified, extensible data architecture.

The particular challenge we address is the need to perform neutron/gamma (n/γ) pulse-shape discrimination (PSD). This refers to a particle identification technique that uses the shape of the signal generated from a particle detector to determine the type of particle that interacted with the detector [2]. For example, it can distinguish whether the incident particle was a photon or a neutron. Particle detectors used for PSD are generally comprised of a scintillator optically coupled to a light sensor, like a photomultiplier tube, that converts and amplifies the light output as an electrical signal. The light produced by scintillating materials capable of PSD (e.g., solid

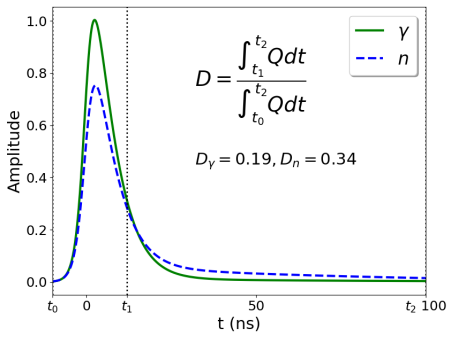


Figure 1: Plots of equivalent energy neutron and gamma waveforms with t_0, t_1, and t_2 markers separating the pulse peak and tail sections.

plastic EJ-276, organic liquid BC501A, etc.) has both prompt and delayed components of fluorescence, with the amount of delayed light produced being a function of the kind of particle that interacted with the scintillator [9, 10]. Therefore, particle interactions produce different characteristic pulse shapes with varying pulse peak to tail amplitude ratios.

Exploiting this fact, particle identification and discrimination is achieved using the well-established technique of comparing the ratio of the integrated charge on the trailing edge of the scintillation pulse to the overall integrated charge. Figure 1 illustrates the distinct tail decay differences between ideal neutron and gamma pulses. The PSD ratio D illustrated in the figure is the ratio of integrated charge over the tail, $t_1 \rightarrow t_2$, and total integrated charge, $t_0 \rightarrow t_2$. The waveforms are created based on the n/γ pulse shape fit for BC501A organic liquid scintillators by Friesen and Howell [6].

In this paper, we describe the design, implementation, and evaluation of CompactDAQ, a compact data acquisition platform that performs the pulse shape processing in real time using analog integrated circuits to independently integrate different sections of a pulse. Rather than multiple racks of equipment, small channel-count systems can be deployed on a desktop and the system design readily scales to support large channel-count systems as well.

2 Background and Related Work

The CompactDAQ platform relies on a pair of ASIC chips, one that measures the integrals illustrated in Figure 1, and one that performs essential timing functions. The first chip, the PSD ASIC [5], performs signal integration on 8 channels. There are 3 gated integrator sub-channels (A, B, and C) per input channel. Each of the sub-channels can be set independently to integrate the input pulse over a wide range of start delays and overall integration widths. Every channel also has a Time to Voltage Converter (TVC) that provides timing information.

The PSD chip has configuration settings that can independently set the sub-channel integration delays and widths in one of 4 sensitivity ranges. All eight channels of the PSD chip share the external configuration ranges, but each channel's actual integration is started independently with respect to the leading edge of its own trigger signal. Each of the sub-channels can also independently configure one of eight different gain settings to change the sensitivity of the integrators. Each of the TVCs is started by the trailing or falling edge of its trigger signal. The leading edge of an external `common_stop` signal stops all the triggered TVC channels.

Precise triggering of the PSD chips is handled by the second chip, the Constant Fraction Discriminator (CFD) ASIC [12]. The 16-channel CFD ASIC generates timing pulses that indicate the time of arrival of detector signals while minimizing the dependence on pulse amplitude variations that would affect simple threshold-based timing methods. Instead of triggering at a fixed voltage level, the circuit triggers at a constant fraction of each pulse's peak amplitude. This produces timing signals that are largely independent of the energy of the detected event. On each CompactDAQ circuit board, one CFD chip delivers trigger signals to a pair of PSD chips, thereby supporting 16 input channels.

A representative example of a large-scale, ASIC-based data acquisition system is the TexNeut neutron detector array [16, 17], developed at Texas A&M University. TexNeut was designed to detect coincident neutrons emitted in reactions studied with the active target and is composed of modular organic scintillator detectors coupled to photomultiplier tubes (PMTs). Each detector segment uses p-terphenyl crystals arranged in optically coupled pseudo-bar configurations, with readout from both ends providing position and timing information. The data acquisition system used for TexNeut illustrates the complex, multi-layered architecture often required for these types of systems. Data are acquired by 16-channel cards that employ the same PSD ASICs; however, they rely on discrete amplification and signal-splitting units (ASIs) for analog pre-conditioning. These units divide and route signals through several hundred feet of coaxial cable to introduce timing delays into the PSD boards, while one branch of each signal is sent to external, non-ASIC-based CFD trigger modules. The acquired data are then routed through hierarchical XLM and VME readout modules before final collection via VM-USB interfaces.

Although powerful, this setup underscores the significant physical and architectural overhead of scaling DAQ systems to large detector arrays, where multiple layers of logic, synchronization, and interconnects are required simply to extract digitized charge data from the front end and makes smaller desktop-scale setups using the same acquisition stack impractical. Figure 2 shows a picture of the electronics and data acquisition rack used by the current electronics setup for the TexNeut array.

To address the bulk and interconnect complexity inherent to discrete and ASIC-based analog PSD systems, recent detector and DAQ architectures have increasingly adopted fully digital approaches. These systems digitize the entire waveform as early as possible with high sample rate digitizers(> 200 MS/s) and use digital signal processing (DSP) algorithms for PSD. The fully digital techniques have demonstrated performance exceeding that of conventional analog systems [11], with more recent machine learning–based methods achieving even greater PSD performance. The ML-based techniques are especially better for PSD with pulse pile up conditions—a situation where a secondary detection happens on the same channel before the first pulse has completely decayed—with high pulse count rates (10^5 counts/s) [7].

Proof of concept work has been done to combine the flexibility of fully digital systems with the high density and low cost of analog ASIC-based systems [13]. Mixed-signal systems try to address the demands of DSP techniques requiring high speed digitizers [1]

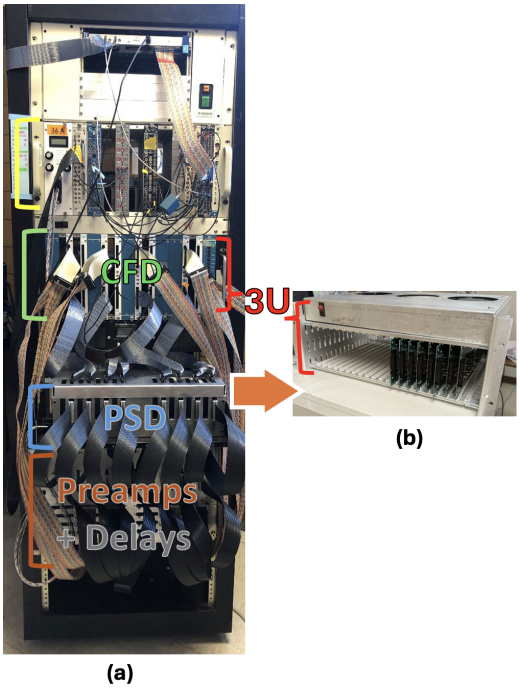


Figure 2: (a) The current PSD DAQ system for the TexNeut neutron detection array at Texas A&M University [15]. (b) Equivalent multi-card CompactDAQ system in a single 3U tall Eurocard rack.

for good performance while keeping the systems scalable. While the availability of high speed digitizers is no longer a substantial limitation as was the case for Aryaeinejad et al. [1], completely digital systems still present challenges when scaled to detector arrays with 1000s of channels. As investigated by Garnett and Byun [8], the reduction of sampling rate from the max 3.2 GS/s for the inputs to their ML-based models had overall predictable reductions in PSD performance.

High-rate digitization in detector arrays presents a fundamental data-volume challenge. Scalable systems such as the Pixie-16 digitizer from XIA can sample 16 channels with up to 16-bit ADCs at 500 MS/s, producing up to 50 MB/s of data per 16-channel module when operating in full-waveform capture mode at high trigger rates [14]. Although modern digital pulse-processing firmware can reduce this load by performing real-time charge integration or PSD parameter extraction using on-board FPGAs, these approaches still require continuous high-bandwidth sampling and substantial backend infrastructure for buffering, configuration, and data movement. At its core, this approach still depends on deploying high-speed ADCs for every channel, and scaling this requirement quickly compounds overall system complexity: a single Pixie-16 module incorporates four dedicated signal-processing FPGAs along with an additional system-control FPGA for coordinating acquisition, memory, and host communication. In comparison, each 16-channel CompatDAQ card uses a single mid-tier FPGA to handle all configuration and readout functions.

These challenges are amplified in TexNeut-style neutron detector arrays, where the expected event rate per detector channel is

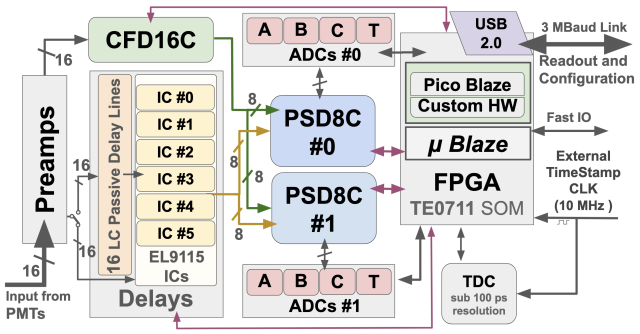


Figure 3: CompactDAQ block-level architectural overview.

significantly lower than [7] at around 1000 counts/s. With sparse detection rates, a substantial portion of the data stream contains no information. Systems employing the custom PSD ASICs become ideal in these scenarios. However, all current implementations, as previously discussed, have significant infrastructural overheads that make desktop scale setups impractical.

The PSD ASICs have been previously deployed in compact systems like the handheld PSD capable system using elpasolite CLYC scintillators created by Budden et al. [3], but they provide no ability to scale the handheld detectors to a larger detector array and do not require or implement any precision event timing measurements.

3 Design Description

3.1 16-channel CompactDAQ Circuit Board

A block diagram of the CompactDAQ circuit board is shown in Figure 3. We first describe the operation of the analog front-end (the left portion of the diagram), followed by the operation of the digital back-end (the right portion of the diagram), and finish with the description of the firmware and software.

3.1.1 Analog Front-end. The analog front-end signal path handles the high-frequency, short rise time pulse signals that are generated from nuclear radiation detectors. Each CompactDAQ board is capable of accepting inputs directly from 16 detectors or PMTs. The first stage in the analog signal path is the on-board pre-amplifier. The *preamp* stage uses extremely high bandwidth current feedback op-amps to amplify and buffer the detector inputs. It has jumpers to toggle between 1× and 3× gain and between positive and negative polarity input signals. Analog signals after the *preamp* stage are all positively ended. The analog signals exit the *preamp* stage and are split into two paths.

The first path enters the CFD16C ASIC, which can generate 16 independent constant fraction trigger pulses. The CFD ASIC produces a simultaneous fast `OR` pulse—all sixteen CFD triggers OR'd together—that is routed off the CompactDAQ board for experiment-specific trigger qualifying logic. The trigger pulses enter a pair of PSD8C ASICs, with PSD#0 receiving CFD triggers for channels 0-7, and PSD#1 channels 8-15. The PSD ASICs integrator and TVC output signals are digitized using Analog Devices AD7687 16-bit differential ADCs, which are controlled by and provide their signals to the FPGA. Since the integration of the input pulses has taken place in the PSD chips, the analog-to-digital conversion only needs to take place once per triggered channel.

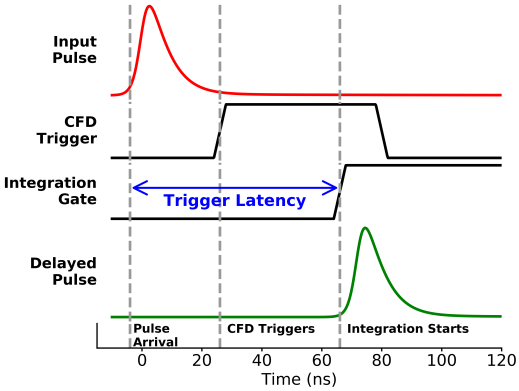


Figure 4: Timing diagram of the trigger signals.

3.1.2 On-board Analog Delays. Since the PSD ASIC integrates pulse energy in real time, it is essential that the integration gates begin and end at consistent points in the pulse's lifecycle. This timing alignment is established by the custom CFD ASIC, which generates triggers at a constant fraction of the pulse's peak amplitude. However, a finite latency exists between the arrival of the pulse at the CFD input and the corresponding trigger output that initiates integration in the PSD. Furthermore, the PSD ASIC itself imposes a minimum delay before charge integration can start. To account for these combined delays, the analog pulse signals entering the PSD ASICs must be time-delayed accordingly to ensure that the integration starts before the pulses arrive.

The timing of an individual channel is illustrated in Figure 4. The input signal from the pre-amps (red waveform) is delivered to both the CFD chip and configurable delay elements, this time is marked as the pulse arrival. If a trigger is initiated by the CFD chip, that trigger and the delayed input signal (green waveform) is delivered to a PSD chip for integration.

Two elements that distinguish the analog front end of the CompactDAQ design from that of the TexNeut system (and other similar existing systems) are the constant fraction functionality (implemented here by the CFD chip vs. discrete elements in the TexNeut system) and the implementation of the signal delay going into the PSD chips. Traditional systems employ fixed-length cabling, relying on the signal propagation in the cables, an approach that imposes severe constraints on the ability to build a compact system. The CompactDAQ design utilizes a combination of passive RL circuits (Rhombus SL7T-505) for long delays of slower signals and a set of programmable analog video delay line chips (Renesas EL9115) for short delays of faster signals.

3.1.3 Timestamps. Another feature of CompactDAQ is its sub-nanosecond timestamping capability, enabling precise correlation of detector signals across distributed acquisition systems. The most critical timing reference is the rising edge of the global `common_stop` signal, which halts the 16 individual TVCs inside each PSD ASIC. To maintain synchronization with other detector systems, CompactDAQ also receives a 10 MHz external reference clock. However, the 100 ns period of this clock limits the timing resolution.

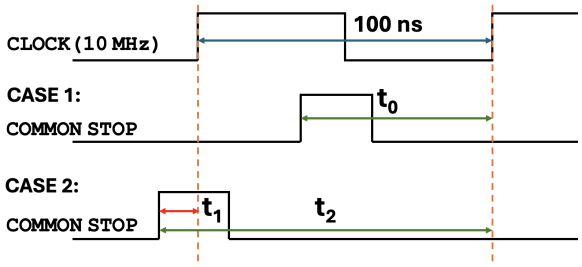


Figure 5: Timing diagram of the external 10 MHz reference clock and `common_stop` signal, indicating the time of arrival of a nuclear event within the 100 ns clock window.

To achieve sub-nanosecond precision, CompactDAQ integrates a high-resolution time-to-digital converter (TI TDC7200) that measures the fine offset of the `common_stop` edge within each 100 ns clock window. The TDC operates in its high-precision mode, which requires a minimum 12 ns interval between the start and stop pulses. Because the `common_stop` signal is asynchronous, it can occasionally arrive within 12 ns of the next clock edge, violating this requirement. Figure 5 illustrates the two possible cases within a 100 ns clock period. In Case 1, the arrival time satisfies the minimum interval ($t_0 > 12$ ns). In Case 2, the `common_stop` signal arrives too close to the next clock edge ($t_1 < 12$ ns). To handle these minimum-interval violations, our design exploits the TDC7200's multiple-stop capability. The first stop corresponds to the immediate rising edge of the clock, and the second occurs at the following edge, exactly 100 ns later. By comparing the two measurements, CompactDAQ reconstructs the true timing of the `common_stop` event even when the first interval falls below specification. The TDC's built-in real-time calibration compensates for temperature and systematic drift, enabling precise reconstruction of the event phase relative to the reference clock. The resulting fine-time measurement is combined with the coarse 10 MHz counter to produce a 56-bit timestamp recorded for every event.

The TVCs integrated in the PSD ASIC achieve a timing resolution better than 300 ps FWHM, corresponding to approximately 130 ps (1σ) for a Gaussian response, while the TDC7200 contributes a sub-100 ps (1σ) fine-time measurement. Combined, this yields true sub-nanosecond timestamping accuracy for every recorded event, implemented entirely within the CompactDAQ card.

3.1.4 Digital Back-end. The digital back-end functions are encapsulated in an AMD Artix 7 FPGA deployed on a Trenz TE0711 daughterboard. Functionally, the FPGA must fulfill two distinct roles: The first is a **control role**, in which the FPGA configures ASICs and peripheral ICs, manages diagnostic signal routing, and communicates with the host computer over a USB link. The second is a **data-acquisition role**, in which the FPGA performs timing-critical operations such as triggering, timestamping, and reading from the TDC and PSD ASICs. Meeting both requirements necessitates an architecture that combines flexible software control with deterministic hardware execution.

To satisfy this duality, the FPGA design is divided into complementary subsystems. A MicroBlaze softcore processor, instantiated through Xilinx's IP integrator, provides a 32-bit microcontroller

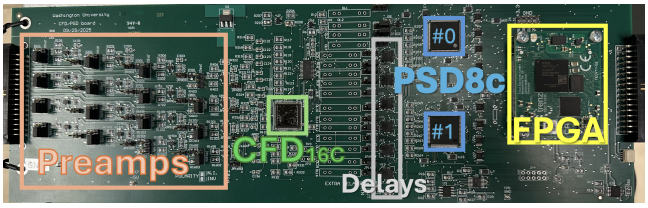


Figure 6: Picture of the fully assembled 16-channel CompactDAQ board Rev4.

environment for control and communication tasks. For fast, cycle-accurate orchestration of acquisition sequences, an additional 8-bit soft processor, the PicoBlaze, is embedded as part of the custom hardware. In addition to the PicoBlaze controller, the custom hardware section includes dedicated logic blocks for ADC interfacing and timestamp generation. Multiplexers and routing logic, collectively referred to as glue logic, enable dynamic switching of pin control between the MicroBlaze, PicoBlaze, ADC, and timestamp interfaces. This arrangement allows each subsystem to operate independently while sharing the FPGA's physical resources.

3.1.5 Auxiliary Onboard Components. The PSD chips require external control voltages to configure the integration gate delays and widths. This functionality is provided by an onboard octal Digital-to-Analog Converter (DAC), the LTC1660, which is controlled over SPI by the FPGA. The DAC generates the seven control voltages required by the PSD chip—six for the delay and width settings of the three integration gates, and one to set the auto-reset time. In addition, it supplies an offset voltage used to zero out the multiplicity output from the CFD chip at the backplane interface connector.

To monitor the input detector signals, the board includes an analog multiplexer (ADG1206). The multiplexer routes the 16 pre-amplified and delayed signal channels to the backplane connector through a unity-gain buffer amplifier and an analog switch (MAX4516), since all multiplexer outputs share a common line.

3.1.6 Physical Layout. The CompactDAQ circuit board is designed in a rectangular 100 mm × 315 mm card form factor, with detector inputs on one short end and all other interface signals on the opposite short end. This form factor was chosen to allow 16 of these cards to fit into a standard 3U Eurocard 19" rack. A picture of a fully assembled Rev4 PCB is shown in Figure 6.

The circuit board is a 6-layer FR4 PCB with four signal layers and two internal power planes. To meet thermal requirements, a split power plane is used, with a substantial −5 V plane section placed directly beneath the delay ICs for thermal dissipation.

3.1.7 Firmware and Software. Software and firmware are central to CompactDAQ's goal of creating a reproducible and extensible acquisition platform. The system is designed to be maintainable and fully reconstructible without dependence on proprietary GUI workflows or legacy tool chains.

At the firmware level, the FPGA build environment is implemented entirely through scripted flows using TCL. All IP instantiations, address maps, and pin assignments are defined in text-based source files, allowing the complete firmware to be regenerated from version-controlled repositories. Custom generators automatically produce synchronized GPIO definitions, constraint files, and Verilog headers from a single specification, ensuring that every interface derives from a unified source of truth. This approach prevents mismatches between hardware revisions and firmware mappings, significantly reducing maintenance effort. A single top-level `run` script encapsulates all build operations: project creation, constraint generation, PicoBlaze assembly, and hardware packaging. This makes the CompactDAQ firmware environment portable and self-documenting.

Configuration and control occur through an ASCII-based command protocol executed on the embedded MicroBlaze processor. Inspired by the SCPI (Standard Commands for Programmable Instruments) format, the protocol uses colon-delimited, fixed-length keywords to identify hierarchical subsystems (e.g., `MUX:AMP:01` selects the preamp channel routed to the diagnostic output). A lightweight framing layer built from standard ASCII control characters (`STX/ETX/ACK/NAK`) provides reliable session control and acknowledgments.

The command protocol is intentionally decoupled from the graphical interface. A standalone Python API implements the full command set and serves as the primary mechanism for generating and transmitting configuration sequences. This enables automated testing, calibration routines, or external integration without reliance on the GUI or any specific host software.

For ease of use and to support collaborators, we provide a PySide6-based configuration GUI that layers real-time visualization (via PyQtGraph), subsystem controls, and device discovery on top of the same middleware. The GUI is optional: it adds convenience, but all system functionality remains accessible through the underlying API and serial command interface.

3.2 Readout

CompactDAQ employs a standard USB 2.0 interface for event readout, intentionally departing from earlier ASIC-based systems that relied on custom daisy-chained buses. Using USB allows each CompactDAQ card to operate as a self-contained unit directly connected to a host computer, eliminating the complex transport infrastructure typical of traditional analog DAQ systems. This significantly reduces setup time and hardware overhead for desktop scale setups.

The PSD ASICs generate compact, event-encapsulated data—9 bytes per channel, or a worst-case 149 bytes when all 16 channels trigger simultaneously. Because this size is well below the 254-byte payload limit of the Consistent Overhead Byte Stuffing (COBS) [4] framing scheme, each event can be transmitted in a single frame without fragmentation, greatly simplifying host-side parsing.

Instead of implementing a full USB 2.0 stack on the FPGA, CompactDAQ transmits data via a 3 MBaud UART stream bridged by an FTDI FT232Q USB–UART converter. USB serves purely as a transport wrapper, while COBS handles event framing. This architecture removes the need for custom USB stacks or host drivers, relying instead on the widely supported FTDI virtual COM interface available natively on most operating systems.

3.2.1 Scaling to Larger Channel Counts. USB-based readout scales efficiently for small and medium setups—additional CompactDAQ cards can be connected directly via USB. However, as the number of cards increases, scaling becomes constrained by the underlying USB topology and resource limits. Although USB 2.0 permits up to

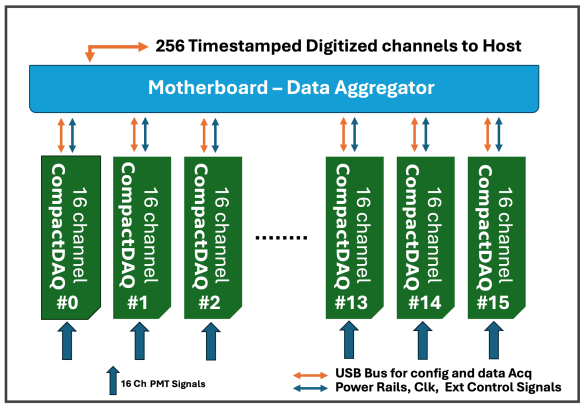


Figure 7: Layout of boards slotted in one 3U Eurocard rack mount frame, holding up to 16 boards for 256 channels.

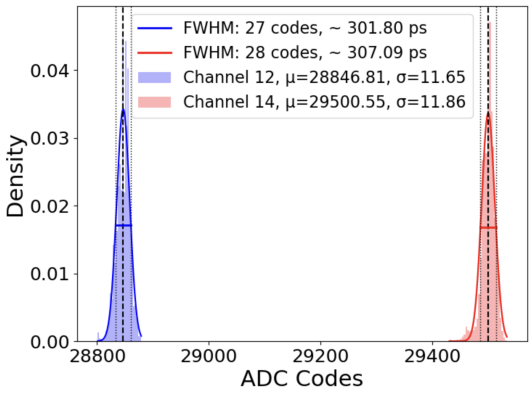


Figure 8: Digitized TVC histograms for two channels with a fixed delay between the prompt (ch. 12) and delayed (ch. 14) pulses.

127 devices per controller [18, p. 13], endpoint and hub overheads quickly reduce this in practice.

To address these challenges, CompactDAQ adopts a hierarchical readout architecture that retains USB as the per-card transport link but introduces a dedicated *data aggregation* subsystem. Each aggregator node is built around a Raspberry Pi 5 Compute Module, which connects to up to 16 CompactDAQ cards via USB and forwards a combined data stream to the main host computer. The aggregator integrates directly onto the 3U Eurocard motherboard that houses the 16 CompactDAQ cards, as illustrated in Figure 7. This design enables CompactDAQ to scale seamlessly from single-card setups to rack-mounted configurations supporting 256 channels per unit. Figure 2(b) shows a 3U rack with a 128-channel configuration and expansion slots for the remaining 128 channels.

We have fabricated 2 prototype 4-card data aggregator motherboards to validate multi-card operation. These 64-channel systems are currently undergoing performance evaluation, with a complete 256-channel motherboard under final design review.

4 Evaluation

4.1 Timing

The timing performance was evaluated in two stages. The first measured the self-timing resolution, which captures the jitter of a single TVC channel. The second quantified the relative timing between channels.

4.1.1 Self-timing. For the self-timing test, the `common_stop` signal was derived from a delayed version of the CFD chip's fast `OR` output. Since each TVC_i starts on the trailing edge of its CFD_i trigger, enabling only one channel ensures that both the TVC start and stop share the same timing reference.

After calibration, the intrinsic timing resolution is obtained from the Gaussian width. The full width at half maximum (FWHM) of the fitted peak is approximately 188 ps ($1\sigma \approx 80$ ps), representing the effective timing precision of the TVC in this configuration. This performance is within the simulated expected timing performance of the CFD and PSD ASICs.

4.1.2 Multi-channel timing. The setup for the multi-channel timing measurement was similar, with the `common_stop` signal again generated as a delayed version of the CFD `OR` output. The input signal was split into two, with one branch delayed by a known amount. For the prompt pulse, both the TVC start and stop were referenced to the same trigger source. For the delayed pulse, however, the start was generated by its own CFD trigger, while the stop remained tied to the prompt pulse's OR signal. Figure 8 shows the resulting histogram, which exhibits two distinct peaks corresponding to the prompt and delayed paths. The applied inter-channel delay was 7 ns, yielding a FWHM of approximately 302 ps ($1\sigma \approx 128$ ps) for the prompt pulse and a slightly larger value of about 307 ps ($1\sigma \approx 130$ ps) for the delayed pulse. The modestly wider distribution observed for both the pulses compared to self-timing is attributed to the slower signal source used in this measurement. The approximate 300 ps inter-channel timing resolution is no worse than that achieved by the current PSD-ASIC based data acquisition system.

4.2 Delay Lines

This section summarizes the tests performed to evaluate the off-label use of the EL9115 active delay line integrated circuits. Since no manufacturer-provided development or evaluation boards were available, these tests were conducted using a custom-designed PCB.

Figure 9 presents the results from this evaluation. The EL9115 was tested using input pulses with rise times of approximately 3.7 ns, significantly faster than the input signals expected under normal operating conditions of the CompactDAQ board, where the pre-amplification stage limits rise times to around 4.5 ns. The purpose of this test was to assess the delay line's performance under worst-case conditions and verify that it remained stable and predictable even when driven by unusually fast transients.

Subplot (a) of Figure 9 shows the overlaid waveforms of the input and delayed output pulses at different delay settings. These measurements demonstrate that the delayed pulses remain well aligned and exhibit only minimal distortion, even at high delay settings. Subplot (b) shows the ratio of the peak-aligned delayed output to the input pulse, computed using a 10 ns sliding-window average to suppress high-frequency artifacts.

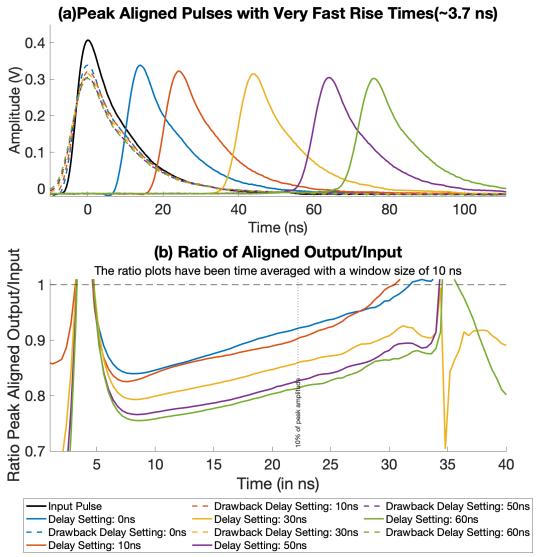


Figure 9: Plots of performance evaluation of the EL9115 delay IC with fast input signals. Note that the x-axes of (a) and (b) have different scales. The ratio plot (b) has been zoomed in to emphasize the shape distortion due to the delay chip.

Ideally, this ratio would appear as a flat line along the $y = 1$ axis, indicating identical shapes and amplitudes between the input and delayed signals. In practice, a small uniform attenuation is expected, resulting in a slight downward offset. The key observation is that this ratio remains linear and monotonic across all tested delays, indicating that the EL9115 introduces negligible nonlinearity or shape distortion. The ratio plot is shown at an expanded scale to emphasize these small deviations and does not share its x-axis with subplot (a). The ratio briefly goes above 1 before stabilizing because the delayed signals have a slightly slower rise time compared to the input pulse. The ratio plot also has artifacts present well after the pulse has decayed to 10% of its peak amplitude. These arise from signal reflections caused by impedance mismatches between our test pulse generator and custom evaluation PCB.

For input pulses with rise times closer to those expected in the CompactDAQ system, the deviation from the ideal flat response becomes even smaller, confirming that the EL9115 meets the required timing and amplitude fidelity specifications for this application.

4.3 Overall Operation

We next demonstrate the operational performance of the CompactDAQ with real-world detectors. The data for this and the following sections were collected using an EJ301 liquid scintillator optically coupled to a photomultiplier tube (PMT). We used a custom-made scintillator cell–PMT assembly consisting of a cylindrical scintillator cell with a 4 inch diameter and 3 inch thickness, and a 3 inch Photonis XP53X2 photomultiplier tube. The entire assembly was packaged inside an aluminum casing. The PMT output signals were fed directly into a single channel of the CompactDAQ PCB-Rev3 for digitization and analysis. The data for different sources were

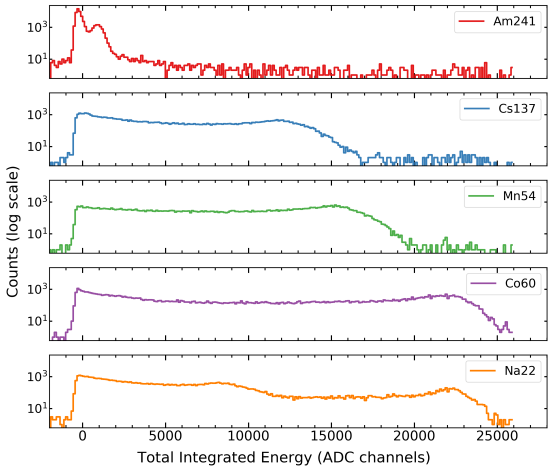


Figure 10: Energy spectra captured by CompactDAQ.

collected for approximately 15 minutes each and all detected events were digitized.

Figure 10 presents the resulting energy spectra obtained with CompactDAQ configured with PSD ASIC gain setting of 1,000 Ω and an integration gate width long enough to capture the full pulse energy. The five subplots show clear, characteristic differences in the energy spectra obtained using known radioactive calibration sources: ^{241}Am, ^{137}Cs, ^{54}Mn, ^{60}Co, and ^{22}Na.

Finally, the same setup was used with a ^{252}Cf neutron source to obtain confirmation of the pulse-shape discrimination capabilities of the CompactDAQ. Figure 11 presents the resulting discrimination plots, generated using two of the integration gates.

In this configuration, Gate C integrates the full pulse, while Gates A and B were configured to capture portions of the pulse tail with slightly different delay and width parameters, with only Gate A plotted. The PSD plots in Figure 11 show the PSD ratio, which is

$$\text{PSD} = \frac{\int Q_{tail}}{\int Q_{total}} = \frac{A}{C},$$

where C represents the total integrated charge and A corresponds to the tail integration from the A integrator. The ratio is plotted against the total pulse energy (ADC_C) on the x-axis for the two subplots, with the ^{137}Cs calibration source in subplot (a) and ^{252}Cf neutron source in subplot (b).

Subplot (a) shows a singular gamma band as expected for the ^{137}Cs source which does not have any neutron decays. The sharp drop off of events in the gamma band indicates the ^{137}Cs Compton edge which is at ≃477 keVee.

The ^{252}Cf PSD plot in subplot (b) exhibits excellent PSD performance, showing the expected separation between neutron and gamma events. Here, the upper band corresponds to neutron interactions, while the lower band represents gamma events detected from the ^{252}Cf source. The red vertical line at total energy of 14000 ADC_C channels on both subplots marks the ^{137}Cs Compton edge location. Interpolating from that, subplot (b) demonstrates γ/n pulse shape discrimination at least down to about 40 keVee.

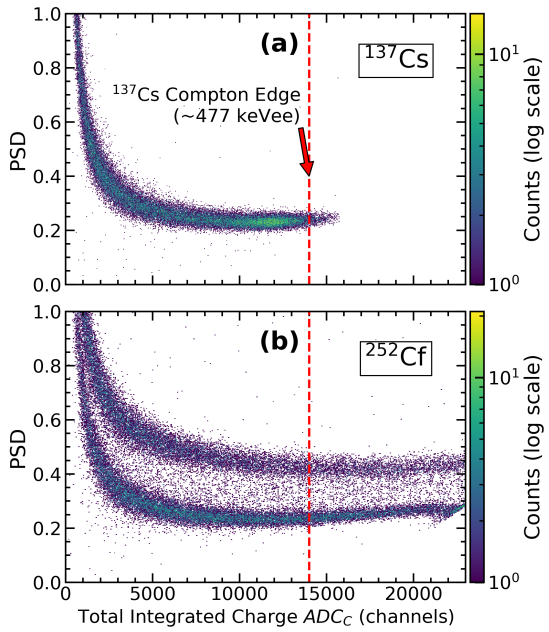


Figure 11: PSD plots: subplot (a) using ^{137}Cs calibration source, and subplot (b) using ^{252}Cf neutron source.

5 Conclusions and Future Work

In this paper, we have presented CompactDAQ, a compact, modular, and extensible analog ASIC–based data acquisition system for pulse-shape discrimination. CompactDAQ combines the proven PSD and timing capabilities of the PSD ASICs with their extremely terse data rates into a standalone, desktop-scale 16-channel device that can scale to support detector arrays with thousands of channels. We achieve this by integrating the new custom CFD ASIC to generate precise triggers and incorporating on-board high bandwidth analog delays by the off label use of video delay lines to eliminate the bulk and wiring complexity associated with existing analog ASIC based systems. Evaluation of the system demonstrates stable and minimally distorted operation of the EL9115 delay lines, expected timing performance of the PSD TVCs, accurate energy spectra across multiple calibration sources, and clear n/γ separation using a ^{252}Cf source. Collectively, these results show that CompactDAQ preserves the fidelity of established ASIC-based PSD systems while substantially reducing their physical and operational footprint.

A clear next step for this work is the deployment of CompactDAQ in a full 16-card configuration. We have recently received the volume production run of CompactDAQ PCB Rev4, and the corresponding 256-channel aggregation motherboard is in final design review. Together, these components will enable synchronized, rack-scale operation with unified power, timing distribution, and USB aggregation.

CompactDAQ functions as a versatile research testbed. An immediate opportunity is applying machine-learning techniques to correct timing and shape errors introduced by nonlinear or reduced-bandwidth lower power delay elements. Another important direction is to leverage the extensible software framework to close the loop between configuration and data acquisition. Because the system provides full access to digitized data and exposes all analog configuration parameters, many of the labor-intensive calibration procedures that traditionally consume days of experimental setup—such as DAC offset alignment, gain tuning—can be automated. Classical closed-loop optimization approaches can be used to iteratively adjust PSD offsets or integration parameters to converge on optimal operating points without manual intervention.

Acknowledgments

This work is supported in part by the NNSA through the Center for Excellence in Nuclear Training and University Based Research (CENTAUR) under grant number DE-NA-0004150. It is also supported by the LLNL-LDRD Program under Project No. 25-LW-034 (prepared by LLNL under Contract DE-AC52-07NA27344).

References

[1] R. Aryaeinejad, J.K. Hartwell, and D.F. Spencer. 2005. Comparison Between Digital and Analog Pulse Shape Discrimination Techniques for Neutron and Gamma Ray Separation. In *Nuclear Science Symposium Conference Record*. IEEE, 500–504. doi:10.1109/NSSMIC.2005.1596302

[2] J.B. Birks. 1964. *The Theory and Practice of Scintillation Counting: International Series of Monographs in Electronics and Instrumentation*. Vol. 27. Elsevier.

[3] B. S. Budden, L. C. Stonehill, A. Warniment, J. Michel, S. Storms, N. Dallmann, D. D. S. Coupland, P. Stein, S. Weller, L. Borges, M. Proicou, G. Duran, and J. Kamto. 2015. Handheld readout electronics to fully exploit the particle discrimination capabilities of elpasolite scintillators. *Nucl. Instrum. Methods Phys. Res. A.* 795 (2015), 213–218. doi:10.1016/j.nima.2015.06.004

[4] S. Cheshire and M. Baker. 2002. Consistent overhead byte stuffing. *IEEE/ACM Transactions on Networking* 7, 2 (2002), 159–172. doi:10.1109/90.769765

[5] G.L. Engel, M.J. Hall, J.M. Proctor, J.M. Elson, L.G. Sobotka, R. Shane, and R.J. Charity. 2009. Design and performance of a multi-channel, multi-sampling, PSD-enabling integrated circuit. *Nucl. Instrum. Methods Phys. Res. A.* 612, 1 (2009), 161–170. doi:10.1016/j.nima.2009.10.058

[6] F. Q. L. Friesen and C. R. Howell. 2020. A functional form for liquid scintillator pulse shapes. *Nucl. Instrum. Methods Phys. Res. A.* 955 (2020), 163302. doi:10.1016/j.nima.2019.163302

[7] C. Fu, A. Di Fulvio, S. D. Clarke, D. Wentzloff, S. A. Pozzi, and H. S. Kim. 2018. Artificial neural network algorithms for pulse shape discrimination and recovery of piled-up pulses in organic scintillators. *Annals of Nuclear Energy* 120 (2018), 410–421. doi:10.1016/j.anucene.2018.05.054

[8] R.L. Garnett and S.H. Byun. 2024. NeutralNet: Development and testing of a machine learning solution for pulse shape discrimination. *Applied Radiation and Isotopes* 211 (2024), 111384. doi:10.1016/j.apradiso.2024.111384

[9] G.F. Knoll. 2010. *Radiation Detection and Measurement*. John Wiley & Sons.

[10] A. Lintereur, J. Ely, J. Stave, and B. McDonald. 2012. *Neutron and Gamma Ray Pulse Shape Discrimination with Polyvinyltoluene*. Technical Report PNNL–21609, 1059208, NN2001000. PNNL. doi:10.2172/1059208

[11] L. F. Miller, J. Preston, S. Pozzi, M. Flaska, and J. Neal. 2007. Digital pulse shape discrimination. *Radiation Protection Dosimetry* 126, 1-4 (Aug. 2007), 253–255. doi:10.1093/rpd/ncm052

[12] B. Orabutt. 2018. *Design and Analysis of a Multi-Channel Discriminator Integrated Circuit for Use in Nuclear Physics Experiments*. Master's thesis. Southern Illinois University Edwardsville, Edwardsville, Illinois.

[13] B. Orabutt, R.D. Chamberlain, J. Elson, G. Engel, F. Delaunay, and L.G. Sobotka. 2021. Design of Mixed-mode Systems for Pulse-shape Discrimination. In *Proc. of 64th Int'l Midwest Symp. on Circuits and Systems*. IEEE, 990–994. doi:10.1109/MWSCAS47672.2021.9531905

[14] S. V. Paulauskas. 2022. A Guide to List-Mode Data Streams in Pixie-16 Hardware. https://xia.com/09/a-guide-to-list-mode-data-streams-in-pixie-16-hardware/

[15] D.P. Scriven. 2023. TexAT-TPC and a Neutron Detector Array, TexNeut. https://cyclotron.tamu.edu/tpc2023/slides/5_Thu_AM/Scriven_TPC_Workshop_2023.pdf

[16] D.P. Scriven, G.V. Rogachev, J. Bishop, G. Christian, L. Sobotka, and A. Alafa. 2022. *Progress in Research, Cyclotron Institute, Texas A&M University*. Technical Report 2021-2022. Cyclotron Institute, Texas A&M University. IV–40–42 pages.

[17] D. P. Scriven, G. Christian, G. V. Rogachev, C. E. Parker, L. G. Sobotka, S. Ahn, G. Chubarian, S. Ota, E. Aboud, J. Bishop, E. Koshchiy, and A. G. Thomas. 2021. A position and pulse shape discriminant *p*-terphenyl detector module. *Nucl. Instrum. Methods Phys. Res. A.* 1010 (Sept. 2021), 165492. doi:10.1016/j.nima.2021.165492

[18] USB-IF. 2000. *Universal Serial Bus Specification, Revision 2.0*. http://archive.org/details/USB-2.0