2014

The NO upsilon A Far Detector Data Acquisition System

Jaroslav Zalesak

Kurt Biery

Gerald Guglielmo

Mark Mathis
College of William and Mary

Recommended Citation

The NOνA Far Detector Data Acquisition System

Jaroslav Zálešák[1,2]**, Kurt Biery**[2]**, Gerald Guglielmo**[2]**, Alec Habig**[4]**,
Robert Illingworth**[2]**, Susan Kasahara**[5]**, Rick Kwarciany**[2]**, Qiming Lu**[2]**,
Gennadiy Lukhanin**[2]**, Stephen Magill**[3]**, Mark Mathis**[8]**, Holger
Meyer**[7]**, Adam Moren**[4]**, Leon Mualem**[6]**, Mathew Muether**[2]**, Andrew
Norman**[2]**, Jonathan Paley**[3]**, Denis Perevalov**[2]**, Luciano Piccoli**[2]**,
Ronald Rechenmacher**[2]**, Peter Shanahan**[2]**, Louise Suter**[3]** and Abigail
Waldron**[9]

[1] Institute of Physics, Academy of Science of the Czech Republic, Na Slovance 2, 182 21
Prague, Czech Republic
[2] Fermi National Accelerator Laboratory, P.O. Box 500 Batavia IL 60510, USA
[3] Argonne National Laboratory, 9700 S. Cass Avenue, Lemont, IL 60439, USA
[4] University of Minnesota Duluth, 1049 University Drive, Duluth, MN 55812 , USA
[5] University of Minnesota, 116 Church Street S.E., Minneapolis, MN 55455, USA
[6] California Institute of Technology, 1200 E California Blvd Pasadena, CA 91125, USA
[7] Wichita State University, 1845 Fairmount St Wichita, KS 67260, USA
[8] College of William and Mary, 102 Richmond Rd Williamsburg, VA 23185, USA
[9] University of Sussex, Sussex House, Brighton, East Sussex BN1 9RH, United Kingdom

E-mail: `zalesak@fzu.cz`

Abstract. The NOνA experiment is a long-baseline neutrino experiment designed to make measurements to determine the neutrino mass hierarchy, neutrino mixing parameters and CP violation in the neutrino sector. In order to make these measurements the NOνA collaboration has designed a highly distributed, synchronized, continuous digitization and readout system that is able to acquire and correlate data from the Fermilab accelerator complex (NuMI), the NOνA near detector at the Fermilab site and the NOνA far detector which is located 810 km away at Ash River, MN. This system has unique properties that let it fully exploit the physics capabilities of the NOνA detector. The design of the NOνA DAQ system and its capabilities are discussed in this paper.

1. Introduction

The NuMI Off-axis ν_e Appearance (NOνA) experiment is a two-detector, long-baseline, oscillation experiment designed to address a broad range of open questions in the neutrino sector. NOνA measures ν_e ($\bar{\nu}_e$) appearance probability and ν_μ ($\bar{\nu}_\mu$) disappearance probability with neutrino and anti-neutrino beams. The ν_e ($\bar{\nu}_e$) appearance experiment investigates: the neutrino masses hierarchy (the mass ordering of ν_3 and the other two); the CP violation phase in the neutrino sector; PMNS matrix mixing angle θ_{13}; and θ_{23} octant (whether $\theta_{23} >$ or $< 45°$). Whereas for the ν_μ ($\bar{\nu}_\mu$) disappearance experiment NOνA will perform precise measurement on the atmospheric oscillation parameters $|\Delta m^2_{32}|$ and θ_{23}. Because NOνA has large detectors and powerful, dead-time-less continuous data readout, it can also be used to study other physics such as neutrino cross sections, neutrino magnetic monopoles, supernova and sterile neutrinos, and other non-standard, exotic neutrino interactions.

Figure 1. The NOνA Far detector is located at a site in Ash River, MN which is 810 km from Fermilab at an angle of 14 mrad of the primary NuMI beam axis. This puts the NOνA detector at an L/E corresponding to the first oscillation maxima. The 14-kton Far Detector (FD) is currently under construction and The 0.3-kton Near Detector (ND) will be located on the Fermilab site in a new cavern, excavated near the existing MINOS Near Detector Hall. Also, the construction of Near detector has already started.

For all of the oscillation measurements, NOvA takes advantage of a two-detector configuration to mitigate uncertainties in neutrino flux, neutrino cross sections, and event selection efficiencies. The NOνA project formally consists of an upgrade to the Fermilab accelerator complex to double the power of NuMI (Neutrinos at the Main Injector) beam line to 700 kW at beam momentum of 120 GeV/c, and the building of two large liquid scintillator based neutrino detectors. The far detector for the NOνA experiment is a 14 kton totally active surface detector that is tuned for the detection of neutrino interactions with an energy of 2 GeV. The near detector shared an identical design as the far detector but is built to have a 1/4 scale cross section of the far detector. A map for NOνA beam and detectors is shown in Figure 1.

The NOνA detectors are fine grained and highly active tracking calorimeters. They consist of plastic (PVC) extrusions filled with liquid-scintillator, with wave-length shifting fibers (WLS) connected to avalanche photodiodes (APDs). APDs have two substantial advantages over other photodetectors: high quantum efficiency and uniform spectral quantum efficiency. The high APD quantum efficiency enables the use of very long scintillator modules, thus significantly reducing the electronics channel count.

The cross-sectional size of detector cells are about $6cm \times 4cm$. Each cell extends the full width or height of the detector, 15.6 m in the FD and 4.1 m in the ND. Extrusions are assembled in alternating layers of vertical and horizontal extrusions, so 3-D hit information is available for tracking. The 14-kton Far Detector has 344,064 cells and the 0.3-kton Near Detector has 18,000 cells. Each plane (cell width) of the detectors is just 0.15 radiation lengths (X_0). This level of granularity helps greatly to separate electrons from π^0 backgrounds.

More information about the NOνA experiment goals and the NOνA detector design can be found in the Technical Design Report [1].

The construction of NOνA is in good progress. The site at Ash River for the Far Detector was completed in 2012 and the Far Detector's assembly and commissioning is very good ongoing as shown in Figure 2: by Oct 7, 2013, 22 of 28 PVC blocks have been installed, 18 blocks of the detector have been filled with liquid scintillator and more than 4 blocks have been outfitted with electronics. The completion of the Far Detector is expected by May, 2014. The construction of the Near Detector is also initiated: cavern excavation for the Near Detector is complete and first PVC blocks are in. The first half of the detector will be installed by the end of 2014 and the second half is expected by summer 2014.
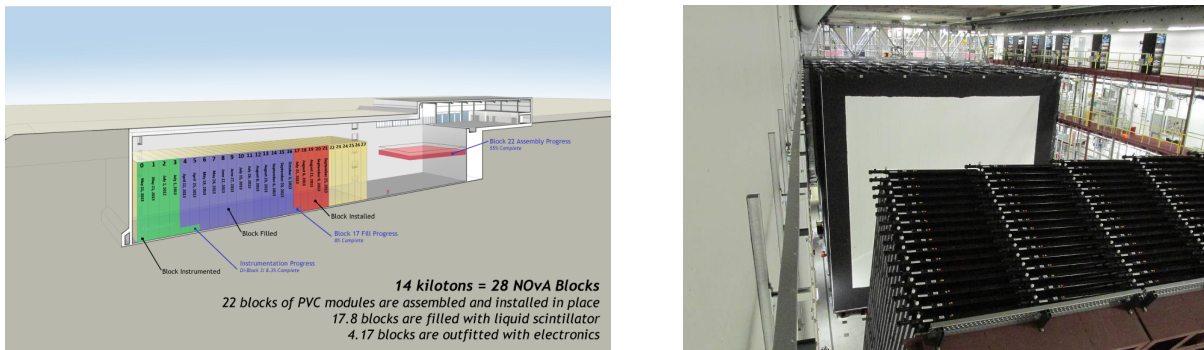
Figure 2. Far detector assembly status (left) and a photo of the one completed di-block installation (right).

2. DAQ Design

In order to make the oscillation measurements the experiment needs to be able to send a narrow ($\sim 10~\mu s$) pulsed neutrino beam to both the near detector and far detector. In each of these detectors the beam will interact and leave characteristic topologies corresponding to the different neutrino interaction modes (neutral and charged currents). The experiment must have an ability to correlate which interactions in the detectors corresponded with the actual beam pulse and be able to compare the rates of interaction in the near and far detectors along with their energy spectra to determine the rates at which the different neutrino flavor states are (dis-)appearing.

Although designing a data acquisition system (the DAQ) to perform such measurements looks simple there are many challenging difficulties described below.

Because the detector is 810 km away from the primary production target, there is no way to provide a hard timing signal to the far detector site that can arrive fast enough to trigger the readout of the far detector. For this reason the readout needs to happen independent of any knowledge of the beam conditions and the data needs to be buffered long enough for any signals that are going to be sent to propagate to the far detector site. Due to the operations schedule of the beam line there is also no way to clearly predict when a beam spill will occur and thus correlate with the detectors. While the NuMI spills are slotted into each accelerator super cycle at well defined points, the actual time line structure changes based on which experiments and beam line are actually running. It is also difficult to verify that a spill has actually happened until it has occurred and the beam line monitors have measured the flux of protons that were extracted onto the target. The other difficulty in performing the NOνA triggering/beam correlation is that as a surface detector there is alway a significant amount of activity in the detector coming from cosmic rays. Not to be lost in the sea of cosmics the fine timing information for the beam spill window is needed.

As a result of these difficulties, the NOνA DAQ system is based on a design to accommodate a traditional trigger to initiate the readout of the front end electronics, while at the same time the design is able to tolerate extremely long latencies in the propagation of information about the beam spills to the Ash River site.

The solution is based on a continuous readout system in conjunction with an absolute time synchronization of all the readout electronics. The schema consists the following [2]:

- Every channel of readout is instrumented with a TDC in the form of a high resolution time stamp counter.
- Every channel in the detector is synchronized to every other channel in the detector such that their time stamp counters increment synchronously.
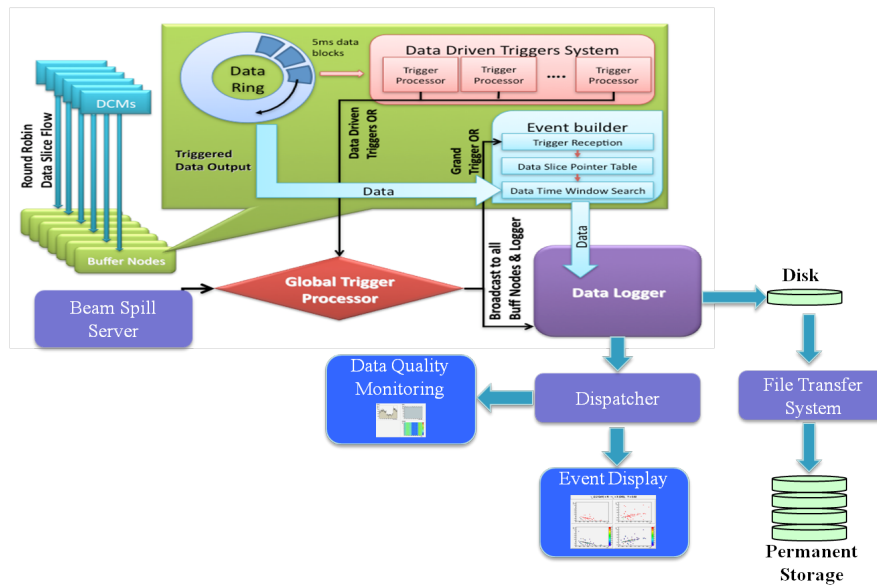
Figure 3. The NOνA DAQ data systems scheme.

- Every time stamp counter is additionally synchronized to an external wall clock such that the value held in the time stamp counter corresponds to the actual wall time.

The detector runs in a continuous readout mode. In this mode all the hits that appear in the detector above a noise threshold receive a time stamp, are read out, sorted into continuous time windows and stored in a deep hit buffer. When a beam spill occurs at Fermilab it is time stamped with and synchronized to an external wall clock. The spill information with the high resolution time stamp is then transmitted to the far detector site. When the beam spill information is received the hit buffers are searched for any hit data that overlaps with the beam spill and the results are stored through a data logger system to permanent storage. The NOνA DAQ systems scheme is depicted in Figure 3.

The advantage of this system is that it provides a true free running system where the electronics are always live and always digitizing. It also provides a readout that is completely dead-time-less. This means that the whole system becomes sensitive to interactions both from the beam but also from external sources such as cosmic rays, cosmic neutrinos or theorized exotic particles such as magnetic monopoles.

The draw back of this system is that it produces a big amount of data that needs to be continuously buffered. In the case of the NOνA detector the aggregated data rate produced by the front end boards can reach a sustained rate of 4.3 Gb/s which needs to be transferred over a modern switch fabric from the front end systems to a large farm of computers which serve as the buffers for the data while it awaits triggering. This requires that the network and computing resources that are used for the transfers and buffering must be capable of handling these rates.

3. Readout Design

The NOνA readout scheme is arranged in hierarchical topology where at each level the data being read out is merged and sorted into discrete time slices. There are 334,064 detector readout cells at the lowest level.

The detector cell is a PVC tube filled with mineral oil containing scintillator, in each tube there is a wave length shifting (WLS) fiber looped down the length of the cells with both ends of the fiber being read out by a single pixel of an avalanche photo diode (APD). When a charged particle crosses the cell, the light from the scintillator is collected by a WLS fiber and transfered onto an APD producing an amplified electrical signal.
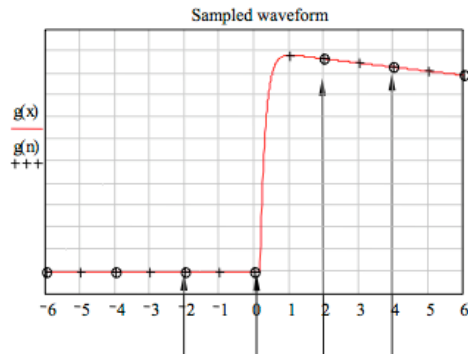
Figure 4. The signal from the APD is amplified and shaped by a custom ASIC on the front end board to produce a well defined waveform that is then sampled with a 2 MHz digitization clock. A dual correlated sampling algorithm is used to establish a threshold for zero suppression of the data at some level above the baseline. There is a time division of 500 ns on the x axis and the waveform $g(x)$ sampled into point-like function $g(n)$ on the y axis, respectively.

The signal from the APD is amplified and shaped by a custom ASIC to produce a waveform shown in Figure 4 which is designed to have a 380 ns rise time and 7 ms fall time. These choices of rise and fall times are intended to minimize the overall noise induced by the electronics. They were determined through measurements of the leakage current and shot noise of the APDs and front end boards. This optimization of the readout was required in order to achieve a signal to noise ratio of 10:1 for signals coming from the far end of the detector cells. The waveforms from the detector cells are sampled by a high speed ADC at a 2 MHz sampling frequency (500 ns time resolution) to obtain multiple sample points along the baseline, rising edge and falling tail of the waveforms. A dual correlated sampling algorithm is then used to establish a rising edge triggered threshold under which the sampling points are zero suppressed. This threshold is set independently for each channel of the detector.
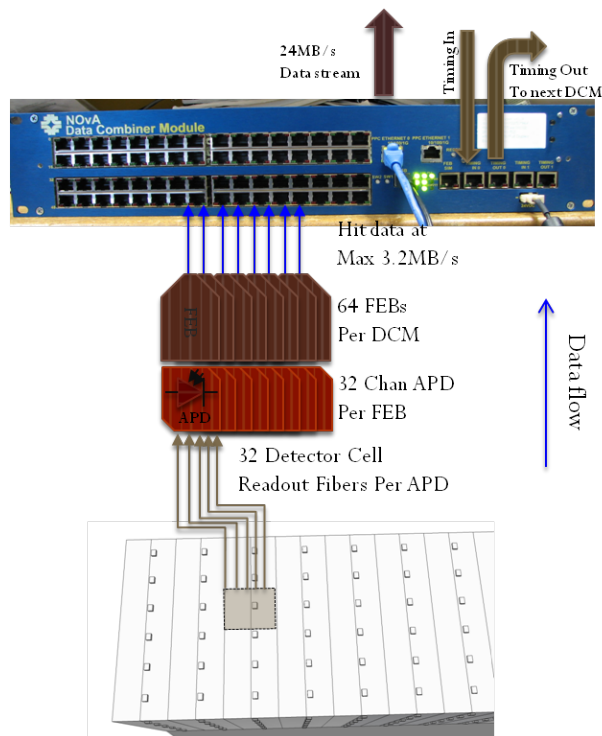

Figure 5. Data Flow: a signal from 32 single cells goes to 1 APD and a group of 64 APDs is connected to one data concentrator module, the DCM. Thus the DCMs aggregate data from 2048 channels of a localized readout region of the detector and produce a 24 MB/s output data stream that is transmitted to a large computing farm over standard gigabit Ethernet.

The data streams from the front end boards (FEB) are sent to a data concentrator module (DCM) which aggregates and sorts the data. Each APD is connected to one FEB which contains one of ASIC, ADC and FPGA for the digital signal processing. The DCM is a custom built

single board computer based on a PowerPC platform. The DCM takes input data streams from 64 FEBs which are grouped in a localized geographic region of the detector, as shown in Figure 5 and time orders them into windows corresponding to 50 μs intervals, called microslices. The DCM organizes this data inside of a large Xilinx FPGA. When the data has been successfully organized into the microslice an interrupt is generated on the PowerPC. This causes a custom Linux kernel module to read the data out of the FPGA into the normal system memory of the PowerPC platform. The data is then further organized by event building software that runs on the embedded linux system.

The DCM event builder constructs larger 5 ms long time units (called millislices) which are optimized for ethernet network transmission. A limitation of handling the data rate of 24 MB/s comes primarily from the speed of the memory bus.

The DCMs are synchronized through the use of a sophisticated timing system [3]. This system provides a stable master clock line as well as command and SYNC lines, that permit the time stamp counters that are present on the front end boards, DCMs and the timing system to be loaded and synchronized with a universal time based on a link to the global positioning system (GPS). The timing system is capable of operating at 16/32/64 MHz and can achieve unit to unit synchronization that is accurate to within one clock cycle. This system is also used to time stamp the beam spill information coming from the NuMI accelerator complex at Fermilab.

Once the data streams are built by the DCMs, all the data from a specific time window is transmitted to a buffer farm of O(100) 16-core commodity Linux nodes that is housed at the far detector site. Each DCM has a list of available buffer nodes within the computing cluster to which it can send data. The DCMs send the data from each 5 ms millislice to the list of buffer nodes in a specific round-robin pattern. All of the DCMs use the same pattern so that all the data from across all of the geographic regions of the detector from a given 5 ms time interval sent to the same buffer, which gives each buffer node a complete snap shot of the detector.

When the round pattern is complete the DCMs cycle back to the first buffer node in the sequence. For the nominal far detector cluster size of about 200 buffer node computers, this cycle repeats once every second. When the buffer node receives a new set of millislices the older ones are not immediately discarded, rather a circular buffer with a fixed event depth is maintained and events in it are not discarded until the buffer is full. In the original baseline designs the depth of the circular buffers for the nodes were set to allow for a system keeping of 20 s worth of data. The foreseen capability of keeping data in the shared RAM opens up the possibility to perform extremely high latency trigger decisions, such as would be required if a supernova were detected by an external observatory and there was a need to extract the corresponding time interval from the NOνA experiment.

4. DAQ Control and Monitoring systems

An inseparable part of the data acquisition are the control and monitoring systems, software applications allow us to keep the whole DAQ system running healthy. The scheme of major parts of such system is displayed in Figure 6.

When the data in the buffer node farm is cross-checked by the global trigger processor as a good event, it is sent to the Data Logger (nodes with 22 Terabyte of RAID array). The Data Logger enables to read its shared memory segments by a Memory Viewer application and send information about triggered events into an Event Dispatcher. The Memory viewer displays raw data with highlighting to indicate the different data block types. The Event Dispatcher sends data to monitoring applications: Online Monitor for producing and viewing histograms of chosen event characteristics, and to the Event Display showing an on-line 3D-reconstruction of an event.

To control the whole DAQ system we have adopted several software applications running independently running on different nodes:
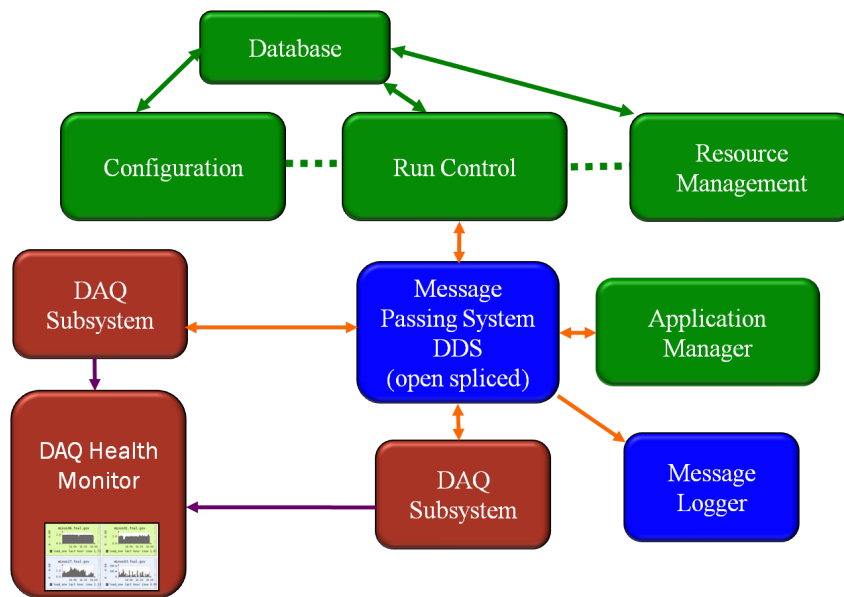
Figure 6. The NOνA DAQ control and monitoring systems scheme.

- Resource Manager - keeps track of which resources are in use, DCMs and the other nodes. It allows to run the whole DAQ system in several, independent partitions, for example one partition can be used for taking physics data on an established part of the detector, while another is used for checkout and commissioning.
- Run Control - is used to select resources for a given partition and to execute configuration steps, and start and stop run.
- Application Manager - starts, stops and monitors Applications and Message Service Daemons directly on selected nodes.
- Message Facility - servers as a destination of all communication among various systems and helps to analyze these massages to be input on the Automatic Error Recovery application.
- DAQ Monitor - generates warnings and alarms when metrics vary outside configured limits, using web based, open-source Ganglia cluster monitoring package to display these metrics.

All of these application communicate via a dedicated massage passing system. We have been using the open source software application OpenSplice, implementation of the Data Distribution Service (DDS) for Real-Time Systems (provided by Prismtech company) [4].

5. Conclusion

The NOνA experiment has designed and implemented a set of DAQ and readout systems for the Far Detector which are capable of meeting the requirements imposed by the core nova neutrino measurements. The system is projected to solve the problem of performing site to site correlation of beam and detector data, as well as designing a method for time synchronizing over 300 thousands channels of readout to high precision. The challenges being faced by the NOνA DAQ are now the production and deployment of the system on the 14 kton far detector and its ability to scale with both the channel counts and readout rates that will be encounter there.

References

[1] D.S. Ayers et al., The NOνA Technical Design Report. 2007.
[2] A Norman, Journal of Physics: Conference Series **396** (2012) 012035.
[3] E Niner, "Synchronization of the 14 kTon NOνA neutrino detector with the Fermilab NuMI beam", to be published in Journal of Physics: Conference Series, 2013.
[4] WWW page: http://www.prismtech.com/opensplice