ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-TUBPA05

HIGH THROUGHPUT DATA ACQUISITION WITH EPICS*

K. Vodopivec†, B. Vacaliuc, Oak Ridge National Laboratory, Oak Ridge, TN, USA

Abstract

In addition to its use for control systems and slow device control, EPICS provides a strong infrastructure for developing high throughput applications for continuous data acquisition. Integrating data acquisition into an EPICS environment has many advantages. The EPICS network protocols provide for tight control and monitoring of instrument operations through an extensive set of tools. As part of a facility-wide initiative at the Spallation Neutron Source at Oak Ridge National Laboratory, EPICS-based data acquisition and detector control software has been developed and deployed to a growing number of neutron scattering instruments. The software interacts with in-house detector electronics using fast optical channels for bi-directional communication and data acquisition, and is built around "asynPortDriver," which allows the passing of arbitrary data structures between software plugins. This completely modular design enables versatile configuration of data pre-processing plugins, as dictated by the neutron detector type and various instrument requirements. After 3 years of experience capturing average data rates of 1.5 TB per day, this system shows exemplary results of efficiency and reliability.

INTRODUCTION

The Spallation Neutron Source (SNS) [1] provides the most intense pulsed neutron beam in the world for scientific research and industrial development. Up to 20 distinct neutron scattering instruments utilize this beam for scientific research. Coupled with high-resolution detectors, SNS instruments offer unprecedented performance for the neutron-scattering analysis of a variety of materials, ranging from "liquid crystals to superconducting ceramics, from proteins to plastics, and from metals to micelles to metallic glass magnets." [2] While the high neutron flux enables more efficient scientific research, it also represents a big challenge for developing a reliable data acquisition system. The SNS instruments continuously capture data, on the order of tens of millions of events per second (aggregate bandwidth), with some individual beamlines approaching 5 to 10 million events per second locally, and producing about 1.5 TB of pre-processed raw experiment data overall per day.

Reliability and data integrity are also crucial at the SNS, given the megawatts of power required to generate and maintain the neutron beam, and the fact that users travel from all around the world, for limited time frames, to perform their experiments. The users' proposals are carefully evaluated and tightly scheduled to efficiently use the neutron beam. Any failure or downtime from the data acquisition system wastes the users' time (and the facility's beam power) and detracts from important investigations.

Given these stringent constraints, and the breadth of data acquisition hardware and instruments at SNS, a new data acquisition software system has been developed to target these demanding requirements for high-throughput and reliable data collection. This software is called "neutron Event Distributor" (nED). It enables SNS instruments to quickly and completely acquire valid detector events from the neutrons scattered by the users' unique sample materials. Ensuring efficiency and reliability in capturing and storing these experimental data empowers users to focus on optimizing experimental/sample parameters for maximal scientific (or industrial) impact.

Besides efficiency and reliability, a main requirement for this updated data acquisition software was integration with the EPICS control system. An internal feasibility study conducted in 2014 recommended that EPICS should be used for instrument data acquisition and controls. In early 2015 nED was first put into production at a neutron scattering instrument. This instrument is one of the highest data throughput instruments at SNS. The deployment of nED was fully successful. Even with an un-attenuated beam, nED was able to reliably and completely collect the full data rate of neutron events at the start of commissioning. In mid-2017, 70% of SNS neutron instruments have been converted to use nED and are collecting more data than ever before.

SYSTEM ARCHITECTURE

The data acquisition system at the SNS consists of sensors, data acquisition measurement hardware, communication hardware and a computer with data acquisition software (Fig. 1). Neutron detector hardware is connected to read-out electronics that are responsible for detecting electrical signals from detectors, applying discriminator settings and baseline correction, calculating detector "pixel" positions and measuring the time of flight relative to the beginning of the pulse. Given the pulsed beam, capturing the time when each event occurs is a critical component in reconciling the relationships among these events. Digitized events from multiple read-out electronics are sent to a concentrator that groups events into structured data packets. In addition to collecting neutron events, the concentrator also connects to fast meta-data sources to acquire non-neutron events, like pulse magnets used to apply high magnetic fields to the sample, load frames to apply compression on the sample and other sample environment devices. Each group of neutron and meta-data events is packetized by the concentrator, completed with corresponding beam pulse information and sent over the fast optical channel for processing.

Every SNS instrument provides unique scientific capabilities and they require different neutron detector setups.

* ORNL is managed by UT-Battelle, LLC, under contract DE-AC05-00OR22725 for the U. S. Department of Energy
† vodopiveck@ornl.gov

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-TUBPA05

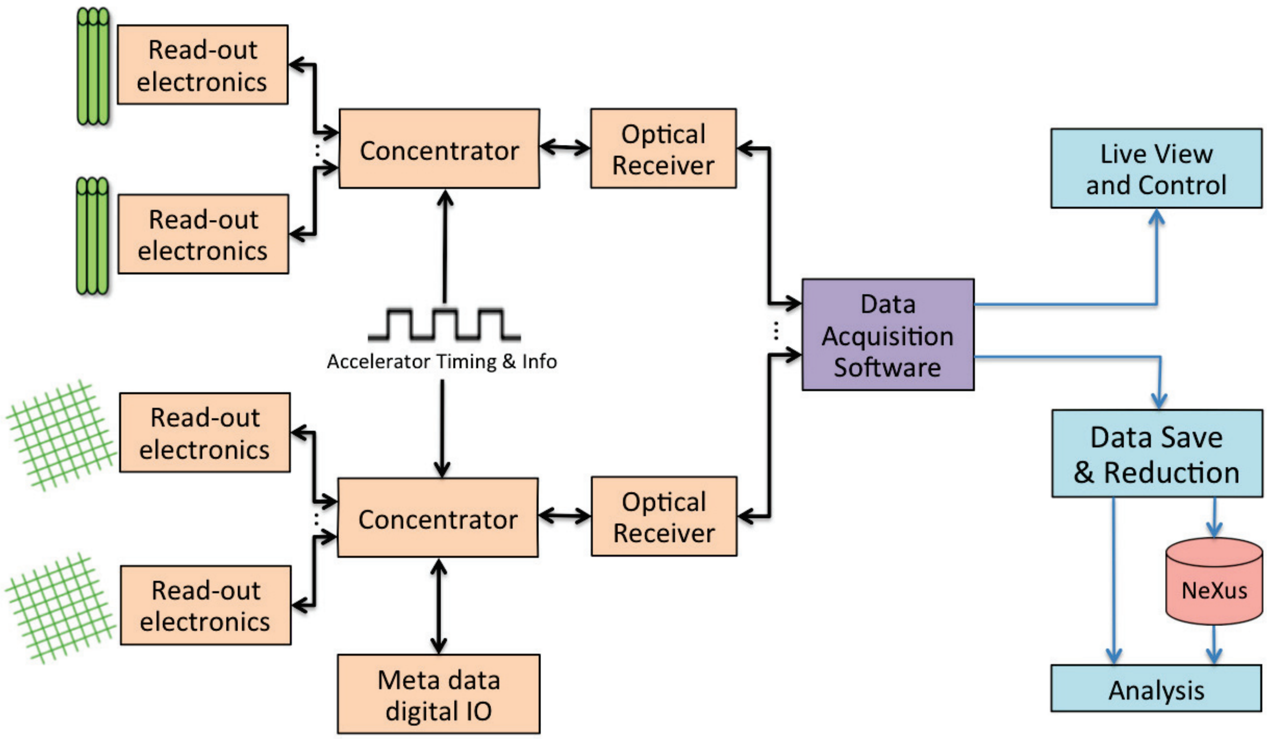

Figure 1: Data Acquisition System components at SNS

Neutron detector configurations at SNS include large detection areas in a variety of three-dimensional shapes. The effective geometry between the sample and the detector are dependent on the neutron beam characteristics (wavelength) resulting in a varying "time of flight" for each event. Supported detector types include:

- Helium-3 Linear Position Sensitive Detector [3]
- Neutron Anger Camera [4]
- Cross Fiber Detector [5]
- Helium-3 2D Detector [6]

The maximum throughput of a single optical channel is 200 MB/s. Due to the varying amount of data required to describe a single neutron event, depending on the detector type, this optical channel throughput can be up to 25,000,000 events per second from a single concentrator. While not many instruments require splitting the optical channels for maximum throughput, some of the instruments' detectors are scattered across a large area making it more convenient to have multiple concentrators to simplify cabling connections.

On the other side of the optical channel sits a PCIe board that puts all incoming optical packets into a Direct Memory Access (DMA) buffer. One PCIe board is needed for each optical channel. With no limitation on the DMA buffer size, nED can multiplex data directly from multiple DMA locations and then export it without making a copy (zero-copy approach). However, most instruments require some form of pre-processing to standardize the event formats from various detector types before they are sent to a live view, and to reduction and analysis. Event pre-processing depends on the detector type and may involve calculating event pixel positions, correcting detector geometrical distortion, applying a time offset and others. The modular design of the nED data acquisition software enables various pre-processing configurations.

Yet the raw data alone is not very useful unless properly presented to the user and reliably saved on permanent storage for subsequent analysis. nED publishes event data through two distinct network channels:

- EPICS7 [7] "PV Access" exchange protocol, with a custom data structure that embeds neutron and meta-data events from a single beam pulse. With a multiple subscriber mechanism, this channel is used simultaneously to serve a live data display tool, for experimental control and visualization, [8] and detector tools for diagnostics, calibration and troubleshooting purposes.
- ADARA [9] uses a custom TCP/IP protocol to connect to the Stream Management Service (SMS), for accumulating data from a variety of data sources and reliably "translating" or organizing and storing all this experimental data in a hierarchical NeXus [10] HDF5 [11] file format, on permanent storage as well as for sending data along to analysis and reduction workflows.

EPICS INTEGRATION

The EPICS control system environment has been used at SNS for accelerator slow controls since the very beginning of operations in 2006. A recent facility-wide initiative is expanding the use of EPICS also for instrument slow control, experiment planning and data acquisition for neutron scattering instruments. Tight EPICS integration across the

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-TUBPA05

data acquisition system has many benefits over the former "legacy" stand-alone solution.

The nED system utilizes EPICS Process Variables (PVs) for real time monitoring and configuration of data acquisition hardware, electronics and software. Configuration changes such as modifying detector sampling parameters, adjusting timing settings according to the selected neutron wavelength, fine-tuning pre-processor parameters and enabling/disabling functionality are performed on demand, usually between experiments. Monitoring for health status, however, is a periodic task that is triggered internally by nED and touches every single component in the data acquisition system to ensure up-to-date system status. Using the PV alarming mechanism, error conditions are reported promptly to the user, in order to take corrective actions or restart the experiment with corrected settings. Precisely timed periodic processing provided by the PV records is used for calculating accurate statistics, and also drives the automatic system checks. Combined with a group of ready-to-use client tools, the EPICS-based nED project is designed into the integrated environment so as to be easy to use for the end user, yet a powerful apparatus for engineers. Due to the large number of detectors at some instruments, the total number of PVs served by a single nED instance can reach over 200,000. Yet there is no noticeable performance impact to the data acquisition side of nED.

The nED code base relies heavily on the EPICS development environment, which provides portability, enhances code quality and results in a faster development cycle. The EPICS build environment supports many platforms and CPU architectures, and while the SNS currently employs primarily the Linux Intel x86 platform for all neutron instruments, such cross-platform compatibility provides the option that other platforms could be used in the future. The EPICS programming interfaces are control-system oriented and have been capitalized on to a great extent. With many interfaces targeted for data processing, such as the efficient threading model and lock-free FIFO buffers, nED has benefited greatly from this design choice. Using the new PV Access data transport protocol, the nED software is able to saturate a Gigabit Ethernet link. The built-in error detection has proven to be a very reliable transport mechanism, with no observation of errors that could be attributed to protocol in the 3 years of operations. Through PV Access, nED supports simultaneous local or remote clients, with throughput limited only by network bandwidth.

MODULAR DESIGN

Due to the versatile detector installations on many instruments at SNS, special care has been taken with the design of nED. Real-time data acquisition of neutron event data and other meta-data is challenging not only due to the high throughput and having to wrangle multiple input sources, but also due to the compute-intensive pre-processing required by various detector types. This must be performed on every event, in order to provide a common event format with time-of-flight and pixel position. In addition, each detector is continuously being monitored and its status is updated through PVs. A scalable and modular design is the answer to this challenge (Fig. 2), as is described next.

The nED data acquisition software is built on the asynPortDriver [12] platform with a software architecture similar to AreaDetector [13]. Specifically, the functional modules are implemented as plug-in blocks that can be seamlessly connected to one or many other blocks. The plug-in connection interface is a bi-directional communication channel that supports an exchange of arbitrary message structures for neutron event packets, with message-oriented communication for read-out electronics and inter-plug-in parameter exchanges. The number of connections is not limited, however testing shows noticeable efficiency degradation at around 100 plug-in connections. Once subscribed, the plug-in will receive all data from its predecessor. The plug-in can filter or modify any data that is sent to its successors. This flexible mechanism allows for arbitrary configurations to address the variety of neutron detector setups.

The detector electronics parameters for sampling and timing, as well as status, are published through a register-like interface. Registers are accessed using a set of read/write commands sent over the optical channel down to read-out electronics, and their responses come back using the same pathway. A convenient interface to map registers to PVs is built into nED, to hide the complexity of packing/unpacking communication packets. This abstraction enables tools for detector calibration and diagnostics to mitigate communication protocol complexities and instead focus on the correlation between the parameters and the produced data. Each register can be broken down into single bit flags that are each represented by a PV. The number of PVs varies between 100 and 500 depending on read-out type and version. The largest neutron scattering instruments at SNS employ up to 400 read-out electronics, rendering an outstanding 200,000 PVs per EPICS Input/Output Controller (IOC).

Event pre-processing is a mechanism for applying detector-specific corrections to neutron events, and converting raw event data into a pixel format common to all detectors. The basic correction applied on almost all instruments is a pixel remapping, where neutron events from a given *physical* detector are positioned into a 2D *logical* space according to the effective "geometry" from a scientific perspective. A more computationally intensive pre-processing operation is done for geometrical distortion correction, as required by scintillators and some Helium-3 detectors, where the calculated event position moves away slightly from its true position due to the effect of neighboring wires. For some detector types, position calculation methods are still being evaluated and tested, and as such they are implemented in software until finalized and ready to be migrated into firmware. These and other pre-processing techniques all have in common an independence or non-correlation between events, which makes them ideal candidates for parallel computing. Due to the great flexibility of the EPICS thread model, the nED pre-processing plug-ins can utilize both system threads and

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-TUBPA05

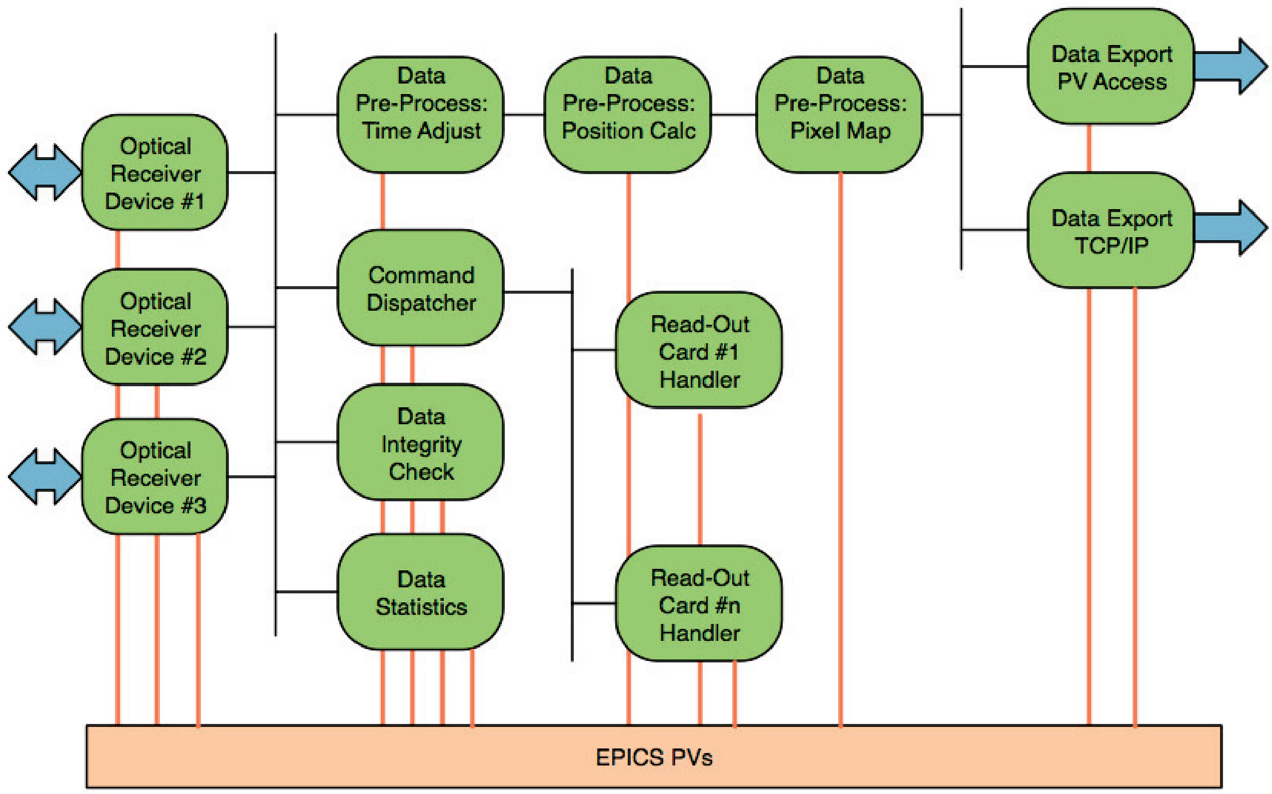

Figure 2: The nED modular design allows arbitrary plug-in connections, from any plug-in to any other plug-in, for numerous configuration possibilities

OpenMP [14]. This ability to use EPICS threads interchangeably with OpenMP provides a new potential for fast data acquisition and event pre-processing.

Many other plug-ins are implemented in support of the data acquisition process. These include: data inspection and verification plug-ins to ensure that the communication channels are error free and no data is being lost or corrupted; beam information extraction plug-ins provide PVs with the absolute time and delivered proton charge for experimental counting purposes; plug-ins calculating inbound and processed data statistics; and read-out electronics support plug-ins to ease adding, replacing or removing detectors.

PERFORMANCE

During the initial design phase, several tests were conducted to verify the suitability of the tight integration between EPICS and asynPortDriver with the data acquisition system. The test metrics included raw data throughput, and measured EPICS' ability for concurrent processing, and handling bi-directional communication with read-out electronics while data acquisition was in progress. There was no significant performance impact detected in any of these categories, and the initial synthetic tests have been proven valid, given the subsequent 3 years of reliable operation, at more than 10 neutron scattering instruments, with others following soon.

Recent measurements in controlled laboratory environment were tailored to check the maximum throughput of the nED software. The setup consisted of two optical input channels and two network output channels - PV Access and the custom ADARA network protocol. There were clients connected to both output channels making sure no data was lost. Neutron event pre-processing was excluded from this test configuration, as it was not within the scope or intent of these specific tests. Pre-processing is known to consume a lot of CPU resources, depending on the given algorithm implementation. For analyzing and tuning this aspect of nED system performance, there are a number of previously documented optimization mechanisms that can be applied, or else algorithms can be ported to the firmware of the read-out electronics.

While the maximum required neutron event throughput of any SNS instrument is about 10,000,000 events per second (which translates to 80 MB/s), the maximum nED capacity is actually quite ample, with extra bandwidth to spare by at least a factor of 5. Given the simplified testing setup, the back-to-back measured throughput was 50,000,000 events per second (or 400 MB/s), which is the maximum capacity of two full optical channels. The CPU usage during the test was merely 3-5% on a dual Intel(R) Xeon(R) CPU E5-2690 system. The test results show that the throughput of the nED software is limited only by the communication channels rather than by any CPU resources, and that the system

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-TUBPA05

is likely to handle even higher data rates given enhanced operating environments.

SUMMARY

A brief overview has been given of the design and EPICS integration of an advanced data acquisition software, nED, suitable for efficiently and reliably supporting all neutron detector types and detector installations at SNS instruments. The nED system has been in use for almost 3 years and is currently serving more than 10 instruments. We find the measured sustained throughput of 50,000,000 events per second more than adequate for current instrument needs.

ACKNOWLEDGMENT

This work was supported by the U.S. Department of Energy under contract DE-AC0500OR22725.

REFERENCES

[1] Spallation Neutron Source (SNS), http://neutrons.ornl.gov/sns

[2] Neutron Science at Oak Ridge National Laboratory (ORNL), https://neutrons.ornl.gov/about

[3] R. A. Riedel *et al.*, "Design and performance of vacuum capable detector electronics for linear position sensitive neutron detectors," *Nuclear Instruments and Methods in Physics Research* A 664 (2012), 366–369.

[4] R.A. Riedel, C. Donahue, T. Visscher, C. Montcalm, "Design and performance of a large area neutron sensitive anger camera," *Nuclear Instruments and Methods in Physics Research* A 794 (2015), 224–233.

[5] C. L. Wang *et al.*, "Wavelength-shifting-fiber scintillation detectors for thermal neutron imaging at SNS," *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 2011 IEEE, http://ieeexplore.ieee.org/document/6152489/

[6] J.Fried *et al.*, N.A.Schaknowski, G.C.Smith, B.Yu, "A large, high performance, curved 2D position-sensitive neutron detector," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, Volume 478, Issues 1–2, 1 February 2002, Pages 415-419.

[7] L. Dalesio *et al.*, "EPICS 7 Provides Major Enhancements to the EPICS Toolkit," presented at ICALEPCS 2017, Barcelona, Spain, this conference.

[8] EPICS areaDetector driver for V4 neutron event data, https://github.com/areaDetector/ADnED

[9] G.M. Shipman *et al.* "Accelerating Data Acquisition, Reduction, and Analysis at the Spallation Neutron Source," *2014 IEEE 10th International Conference on eScience* http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6972268

[10] The NeXus Data Format for Neutron, X-Ray and Muon Science, http://www.nexusformat.org

[11] The HDF5 Data Model, https://support.hdfgroup.org/HDF5/

[12] Asynchronous Driver Support, http://www.aps.anl.gov/epics/modules/soft/asyn/

[13] EPICS Software for Area Detectors, http://cars9.uchicago.edu/software/epics/areaDetector.html

[14] OpenMP, http://www.openmp.org