DAQ Workshop Summary, Version 2

Kurt Biery

CPAD TDAQ Session (with formatting fixes and summary slides)

23 October 2017

One-Day Workshop on Future DAQ and Trigger

Held on Wednesday of this week

~35 people, representing 6 national labs & 7 universities

Afternoon working groups

Charge

Full charge here: https://indico.fnal.gov/event/14744

My summary:

- *Identify future DAQ and Trigger needs, both in terms of long-term future experiments and in terms of providing tools to support short- and medium-term detector R&D efforts.*
- *Identify promising technologies and techniques to address the needs.*
- *Identify opportunities for collaboration between organizations in future research.*

Agenda

Full agenda here: https://indico.fnal.gov/event/14744/timetable

Summary:

- Cosmic Frontier DAQ Needs, Paul O'Connor (BNL)
- Energy Frontier Trigger and DAQ Needs, Wesley Smith (U. Wisconsin)
- Future DAQ Needs for the Intensity Frontier, Josh Klein (U. Penn)
- Testbeam and Test Stand DAQ Needs, Mandy Rominsky (FNAL)
- Future FPGA Architectures, Ryan Herbst (SLAC)
- Future Process Architectures, Chris Green (FNAL)
- DAQs for cryogenic detectors in Cosmology, Gustavo Cancelo (FNAL)
- Tech Trends and R&D Ideas for a Post HL-LHC Energy Frontier, Paul Padley (Rice)
- Future Technologies for the Intensity Frontier in Trigger and DAQ, Wes Ketchum (FNAL)
- Technologies for Test Beam and Test Stand DAQ, Ryan Rivera (FNAL)
- Participant roundtable
- Working groups and discussion

Talks

Lots of information still to be gathered from the talks.

Summaries slides next.

Apologies for what I’ve missed here.

Summary

CF DAQ Needs, Paul O'Connor

- Next-generation large-scale cosmic acceleration experiments in planning stages
- Several areas of common R&D interest:
  - Digital RF signal generation and processing
  - Algorithms for on-the-fly calibration and data cleaning
  - Analog data transport into/out of cryostats
  - Test stands and calibration
- For galaxy spectroscopic survey:
  - Viability of MKIDs at 10^5-pixel scale
  - Control and alignment of 10^5 fiber positioners
- For CMB:
  - Choice of detector and multiplexing method
- For 21cm:
  - Leveraging high-speed digital sampling and processing to minimize cost and power

Trade off resolution and field of view
Low-res spectra w/color or fiber spectrographs
Multi-color CCDs or MKIDs

Trigger Needs

Economical track-fitting and clustering techniques and technologies

- Algorithm optimization and mapping onto hardware
- Low latency, modest resources

Need to keep p_T thresholds the same with m
more pileup.
Triggerless architecture

Algorithm processing solutions

- Low Latency, modest resources
- Specialized FPGAs, GPUs, CPUs and large memory solutions
- Processor architectures – parallel, time-multiplexed, etc.

Fast Network Solutions

- Needed for data transmission into, out of and between trigger processors
- Needed for trigger configuration, timing, clock, control
- High speed network connectivity, e.g.100GbE and beyond

High Speed Optical Links

Rad-hard

- Link and connector infrastructure: bandwidth, power, mass, size (wireless?)

Infrastructure

- Power and cooling, general form factor (e.g. ATCA)
- Configuration and control infrastructure: IPMC, embedded Linux (e.g. ZYNQ)
- Latency and timing measurement and monitoring infrastructure

Firmware Development

- Management and build infrastructure
- High Level Synthesis tools

DAQ and HLT Needs

Studies of co-processors and their integration into HLT or event builder or network nodes

- GPUs, FPGAs, etc.

New packaging & interconnect technologies and infrastructure (incl. power & cooling): e.g. ATCA …

Evolution of Network Switches

Development of more sophisticated Clock & Control Networks

HLT on the Cloud

- e.g. share resources between HLT & Tier-0

HLT Software

- More sophisticated algorithms, increased occupancy
- Merging of HLT & offline software development

Future DAQ Needs for the Intensity Frontier

Physics Drivers of Future DAQ

- Truly "rare process" searches
- Reduced statistical uncertainties
- Reduced systematic uncertainties
- Need to operate in difficult environments (cold, wet, surface, radiation…)
- Physics breadth
- Paranoia
- Greed

Most experiments are driven by more than one of these.

J. Klein, U.Penn

Future DAQ Needs for the Intensity Frontier

Concluding Thoughts

- Breadth of the IF program means no single critical need
- Work on adapting commercial systems may save costs and design efforts
- (But needs to be considered in context of ownership and training)
- Lossless compression specific to IF requirements needed for future plans
- But R&D on new technologies (e.g., trigger-level learning) could yield new expts
- Portable DAQ software and hardware could also be a benefit

GPUs vs. FPG As is a tension between development ease (e. g. CUDA vs. Vivado) and ultimate data bandwidth. (p17)

From the discussion session:

- 'More information, less data' (J. Fast)

Common tools in other areas (simulation); existing multi-experiment HW and SW; advantages and disadvantages.

J. Klein, U.Penn

Test Beam Needs

- Facility instruments integrated with users' DAQ
- Facility info with beamline experts (maybe only a FNAL problem?)
- Beam vs non beam users "Facility teststands"
- Support of whatever product there is
- Training experts
- documentation

Test Stand needs

- Slow control, environmental needs
- Portability
- Flexibility
- Experimenter controlled vs facility offering

Summary

- Possible Inter-Lab Collaborations
  - Shared libraries and open source development
  - Common hardware implementation of generic DAQ platforms
  - Adapt server farm scheduling applications for FPGA resources

- Consider deploying machine learning processing in the mid level DAQ
  - Even in the front end
  - Avoid event building before applying machine learning

- Take full advantage of partial reconfiguration features to quickly deploy dynamic data processing
  - Scripting in FPGAs

Hybrid memory cubes – sharing data between FPGAs
Connecting FPGAs to networks
Higher level languages for firmware development
Coprocessing
Development challenges

Future Processors, Memory, Storage (1)

- Improvements to provide new avenues for event-building, triggering, filtering, transport.
- Commodity now is already different from commodity 10y or even 5y ago (e.g. co-processors, SSD, 56Gb/s InfiniBand). Keep up! What will be commodity in 10y?
- Major drivers: ASCR, AI, ML, “big data.”
- IBM Power9+ could be disruptive, but roadmaps for Intel/AMD x86-64 are evolutionary (no surprises), schedules to stretch as era of Moore’s Law ends. AMD Zen absorption of chipset interesting trend? SOC?
- Co-processors: Tesla, Xeon Phi, AMD Instinct. Definitely AI/ML influence, mostly evolutionary improvements over next 10y.
- Big push toward integrated systems: ”traditional CPU” + co-processors e.g. AMD project 47 “PFLOP on a rack.” Commodity supercomputing?
- Possibly-disruptive new technologies: TPU, neuromorphic, automata processing, quantum (D-Wave annealing, Google / IBM qubits).
- Moore’s Law demise is being addressed both with innovative new uses of transistors (SIMD, SIMT, TPU, AP, neuromorphic, etc.) and with entirely new ideas like qubits.

Future Processors, Memory, Storage (2)

- HDD: density / throughput improvements 10-40 over current best in 10y.
- Tape: density / throughput improvements 300-400 over LTO-7 in 10y.
- New bottleneck to be (already, in server configurations) the bus -> bypass with e.g. NVLink, OminPath, SOC. X4 improvement over next few years, then nothing? Blue-sky: surface plasmons?
- Tb/s networking (Ethernet, IB, Omnipath) about 10y out.
- NAND SSD certain to be eclipsed by 3D-Xpoint, NRAM, ReRAM very soon (<5y).
- DRAM improvements to be relatively minor, could be utterly eclipsed by NRAM and/or ReRAM (lower latency?) over the next 10y: revolution! Major implications for whole system if so. Board architectures, SOC?
- Will need to think very carefully about programming / algorithms to take advantage of all this innovation: Intel TBB, C++ standard, HPX, MPI, OpenACC, OpenMP, Numpy, Julia.

DAQ synergies

- CMB S4, Optical surveys and quantum computing:
- CCDs and CMOS detectors for large scale DM and Neutrino experiments.
- Silicon photon detectors for neutrino, DM, Intensity interferometry, and HEP.

CCDs for Dark Matter and neutrinos
DAQs for Dark Energy
Silicon Photomultiplier DAQ

Where to put effort and money?

- Large scale, ultra low noise electronics for DM and Neutrinos.
- Low Noise cryogenic amplifiers for DE, CMB and QC.
- Multi channel RF electronics embedded in analog/digital systems.
- Highly complex multi GHz analog and digital design.
- Complex FPGA design.
- High speed and optical links of over 30Gb/s.
- EMC design.
- Fast timing for Cosmology and HEP.
- DAQ software interfaces.

HEP R&D

Tech Trends and R&D Ideas for a Post HL-LHC Energy Frontier, Paul Padley

- To take advantage of the technologies being developed for the commercial world we need to be working with them
- Cant be driven by project's it needs to be somewhat undirected R&D
- Only by "playing with" and becoming familiar with these technologies will we gain the insight to understand how best to apply them to particle physics.

High radiation environment
Power needs, probabilistic computing
Data needs – optical and wireless
Coprocessing with GPUs and FPGAs
No EVB; novel architectures

CMS DAQ R&D for HL-LHC, Remi Mommsen

Summary

Sergio Cittolin © 2009-2018 CERN
(License: CC-BY-4.0)

Future CMS DAQ systems evolve from current architecture

- Run 3 DAQ is a modest upgrade
- Run 4 (phase 2) needs to be scaled up by a factor ~25

Main challenges for the future

- DAQ and TCDS hub with up to 1.2 Tb/s throughput, precision clock and trigger distribution
- Handle 200 Gb/s input & output per computer with high-speed switching network with ~50 Tb/s effective throughput
- Provide enough HLT CPU capacity, possibly assisted by coprocessors

R&D proposals with technical challenges, but also potential game changers for physics performance

- Integrate event building into interface between custom hardware and commercial components
- Develop a versatile software framework straddling the boundary between FPGA, coprocessors and general purpose h/w
- Explore the possibility for large key-value stores
- Partial data acquisition at 40 MHz

Conclusions (or, "OK, what about the future?")

- Josh hit on this in his talk, but some clear general trends for IF experiments
- Overall data rates are high, either from high rates or large event sizes → data reduction common theme via full-event filtering, smarter triggering, and/or compression
  - Experiments addressing this with mixes of firmware and software
- Push towards "commodity" or "commoditizable" components
  - Some: just get data into servers, apply software compression/filtering (maybe using accelerators)
    - "Software is the most flexible"
  - Others: push towards powerful general-purpose readout solutions that are flexible
    - "Software won't be enough"
- "Figuring things out" via prototypes/experience → informing future experiments

GPUs used in Muon g-2 DAQ
Commercial components in Mu2e
Tools used in neutrino experiments
Trigger, Timing, Software

Test Beam and Test Stand DAQ Technologies

- This is a brief **survey** of DAQ technologies for **test beams** and **test stands** currently available at U.S. institutions.
  - Based on input from SLAC, BNL, and FNAL.
- The goal is to raise awareness of existing solutions and to spark discussion toward future collaborations and new developments.
- SLAC (Contact: Carsten Hast, hast@slac.stanford.edu)
  - No standard test beam facility DAQ. Users generally provide their own DAQ. Available tracking telescope, ADCs, flash ADCs, scalars.
- BNL (Contact: Martin Purschke, purschke@bnl.gov)
  - RCDAQ was developed as a general purpose DAQ – used by experiments, test beams, and test stands. Light weight and configured via Linux script.
- Fermilab (Contact: Mandy Rominsky, rominsky@fnal.gov)
  - Test beam facility DAQ based on *otsdaq/artdaq* (general purpose DAQ with web-based configuration and control) commissioning. Available tracking telescopes, wire chambers, Cerenkov, scalars, PREP equipment, CAPTAN FPGA-based test stand hardware.

Participant Roundtable

Lots of good information, apologies for glossing over it here.

Link to slides:

- https://indico.fnal.gov/event/14744/session/4/contribution/13/material/slides/0.pdf

Fermilab

Working Groups

Not enough time to discuss all of the needs that were presented.

Topics that were selected:

- Firmware development
- Test beam and test stands
- Digital Radio

Notes from the groups on the next few slides.

Firmware development discussion notes (1)

Communication between groups

- Sharing of knowledge, experiences/lessons learned, and developments
- Forum for requesting support
- Virtual workshops

Firmware repositories and support for tools

- Expand and unify existing work in this area
- Sharing of experience with select dev tools (training, too)

Future R&D pursuits

- Intelligent switch technology; configuration tools; tools and approaches to common modules; signal integrity tools; forward error correction and data compression; optimization of systems based on blocks

Firmware development discussion notes (2)

In the area of shared modules, we need specific guidance and help from DOE on the following:

1. Appropriate license for developed code. This license would have to be consistent with the license of the host labs and universities. This may be the biggest hurdle.

2. Guidance on determining which modules are considered supportive with no commercial value This will allow labs to save time sharing supportive features.

Suggest DOE call for proposals to provide the universities and labs the resources to develop the skill sets necessary for development in the following areas:

1. Co-processing using FPGAs in a mixed software/hardware environment.

2. Neural network applications for FPGAs.

Firmware development collaboration

Copy of Angelo Dragone's slide

HEPIC.ORG

& HEPFW.ORG?

HEPIC 2017

Key scopes

Remove Legal Barriers:

- Enable sharing of design blocks
  - Common ground for foundry NDAs (Eg. TSMC, ARM)
- Establish clear and consistent guidelines on design blocks sharing among partners
  - Acknowledgment of credits
  - Understanding of critical needs related to multi-purpose institutions

Create more opportunities for collaboration and increase communication:

- Organize yearly workshops
- Catalog of designs / Data management
- Promote Cross-institution MPWs
- Lists information on groups and projects
- Maintains a Twiki and FAQs repository
- Provide Education and Training (course material to share)
- Advertise Jobs / Fellowship opportunities
- Lists ideas for student projects

by increasing communication and lowering or removing barriers for collaboration any costs incurred by the organization will be offset by its clear benefits

Similar issues with FW Dev as ASIC design

Notes from Testbeam/Teststand working group (1)

- Need to find DAQ system balance between being generic/flexible and becoming too large and impossible to maintain
  - (User) community should converge on ~3 different, light weight systems
  - FW & SW need to be free
  - Support and maintenance needs to be guaranteed (at a lab?, funding?)
  - Examples include: MIDAS, OTSDAQ, artdaq, DRS4
- Groups often forced into certain products by projects, cannot always freely choose their architecture
- Commercial products not usually an option: too expensive, not flexible enough for our FW & SW needs
- Need for a user facility/teststand, where different HW platforms can be tried out before deciding on a specific architecture for an experiment
- Need for sharing of HW resources, e.g. PREP at Fermilab

Notes from Testbeam/Teststand working group (2)

- Need to form user+developer groups per DAQ system
  - Meet periodically
  - Share and contribute code and libraries
  - E.g. HEPSoftware Foundation (common repository): http://hepsoftwarefoundation.org/
- Testbeam specific needs:
  - Find common solutions across facilities for beam & environmental information
  - Ideally provide the 2-3 most popular DAQ systems for users
  - CAVEAT: Fermilab testbeam will soon be the only one available for our community while DESY, CERN and SLAC beams are being upgraded! Increased resources needed for additional DAQ R&D and user support needed (~2 FTE)

Additional idea from Angela, Mandy

- Possibility of dedicated ECA-like funding opportunity for TB/TS infrastructure
- Develop, maintain, and support a small number of DAQ frameworks across the labs, build support structures for the use of these at universities.
- Early career award proposals typically have a physics impact. What if one slot per N years were reserved for TB/TS infrastructure development/support…

Digital Radio

- Interesting application across fields/agencies
  - OHEP: CMB, optical survey going into near infrared, quantum sensors
  - BES: X-Ray Spectroscopy
  - NASA, NIST
- Cold electronics options MKIDS, TES/Squids
  - Sensor might be different, but (warm) RF electronics control/DAQ is very similar
- Warm RF electronics needs R&D to scale up to larger arrays and higher performance
  - Is (relatively) in infant stage
  - Need higher multiplexing ratios to scale up to larger focal planes
  - Technically challenging
    - RF electronics circuits
    - Algorithms to scale up non-linearily with number of channels to be processed
      - Firmware
- University and lab collaboration
  - Opportunity to work together
  - Strong intellectually contribution from University with lab KA25 technical support and leadership to architect and implement systems, testing, making it work with university leadership
- 21-cm telescope
  - Again scaling issue to develop software defined radio architecture and implementations which can be cost and power efficient at a 10e6 antenna level
- KA25 support is essential to demonstrate the feasibility of the next round of experiments

Other discussion notes

Possibility of university participation in KA25 projects

Use more deep learning, already in trigger/DAQ, on different HW platforms

Enable faster, better data processing closer to front-end

- “more information, less data”
- Subject to the paranoia and greed mentioned by Josh

Summary (1)

Many needs were presented, and readers are encouraged to refer to the original presentations to see the full breadth.

In the discussions at the workshop, the following areas were identified for future investigation. These include both evolutionary and revolutionary research. This is not a prioritized list; it simply follows the approximate order in which the topics were introduced.

1. Common warm RF electronics control/DAQ for future CF experiments
   - Need for higher multiplexing ratios to scale up to larger focal planes
   - Technical challenges include the development of the RF circuits, algorithms for data cleaning and calibration, and firmware development
2. Rad-hard, high-speed optical components
   - Very high radiation environments at future colliders
   - Improvements needed in bandwidth, power, mass, and size
3. Wireless communication for data transfers
   - Would reduce material in detector, fewer cables and connectors
   - Tb/s possible by 2020 over short distances

Summary (2)

4. Optimum locations in the data stream for compression & filtering
   - FPGAs or GPUs or software close to the detector?
   - Came up in IF discussions, but concepts may be useful elsewhere
   - Develop trustworthy tools and algorithms that increase the fraction of valuable information in the data
     - “more information, less data”; avoid paranoia, greed, and saving everything
5. Coprocessor investigation and development
   - CPUs with GPUs or FPGAs
   - Likely needed for higher level triggering; closer to the detector may be valuable, too
   - GPUs used before event building in Muon g-2; CMS exploring track seeding in GPUs for HLT
   - Industry trends making coprocessors common
   - High level synthesis development tools for programming

Summary (3)

6. Firmware development collaboration and future development
   - Communication between groups
   - Firmware repository and support for tools
   - Suggestions for future R&D pursuits
   - Licensing issues (are there ways to share vendor IP between labs? what open source license should we use for blocks that we develop?)
   - Developing expertise in co-processing and machine learning uses for FPGAs
7. Test beam and test stand common tools, support, and communication
   - Identify and support a small number of DAQ systems
   - User test facility where different HW platforms can be evaluated
   - Communication between user groups
   - Fewer test beams during some time periods (e.g. CERN LS2 2019/2020)
     - Should this translate to increased support at Fermilab test beam?
   - Funding for testbeam-specific infrastructure development

Thanks to all participants

Very relevant and useful talks. Good discussions. Notes from participants appreciated.

Special thanks to the organizing committee:

Hucheng Chen (Brookhaven National Laboratory)
Eric Church (Pacific Northwest National Laboratory)
Angela Fava (Fermi National Accelerator Laboratory)
Gunther Haller (SLAC National Accelerator Laboratory)
Klaus Honscheid (The Ohio State University)
Petra Merkel (Fermi National Accelerator Laboratory)
Alan Prosser (Fermi National Accelerator Laboratory)
Robert Tschirhart (Fermi National Accelerator Laboratory)
Robert Wagner (Argonne National Laboratory)
Peter Wittich (Cornell University)
Sergio Zimmermann (Lawrence Berkeley National Laboratory)