TDAQ Snowmass Community Workshop 8/6/20

Indico agenda: https://indico.fnal.gov/event/44741/

Introduction

Instrumentation Frontier focus on detector technologies and R&D needs for future experiments across the physics frontiers

Milestones along the way:
Letters of Interest due 31 Aug 2020
Virtual Community Planning Meeting: NEW DATE 5-9 October 2020

Aim right now: Letters of Interest to help shape planning

Contributed papers will be part of the Snowmass proceedings. Want to see LOIs help shape up to contributed papers (but LOI not a prerequisite)

Virtual community planning meeting: https://snowmass21.org/2020_oct_cpm
Aiming for four 5-hour days with ...
Plenary sessions on large-scale vision and planning/wrap-up
Breakout sessions, including a focus on inter-frontier communication
Your feedback welcome!

Contributions:

Realtime CRES DAQ for Project 8, Noah Oblath

For neutrino mass measurement, Cyclotron Radiation Emission Spectroscopy
2.4 exabytes/year after digitizer, but trigger and reduce 75 GBs - > 1GB/s (still ~100 PB/yr raw, but window can reduce more)
3-5 year development

From Colliders to Short baseline neutrino, Bill Badgett

Long history with TDAQ topics from SDC on up several generations!

Discussion session: Data links, smart sensors, and fast timing (very front end)

- How do we handle boundaries with Electronics/ASICs on things like neuromorphic sensors? Are we focusing mostly on algorithms, or actually talking the sensors too? Maybe good way to think about it is aiming at reducing the dataflow.
- Wireless an important point here (have an LOI on that already)
- Distributed DAQ systems: how do smart sensors communicate across miles of physical space? Related to microcontrollers for data reduction. Calibration and timing synchronization critical here: related to fast timing for large-scale distributed systems.
    - E.g. have a beam monitor reading in data and want to do some processing on the front-end. How do you do that data reduction/what information do you need to do that data reduction (and how do you get it there)?
- Fast timing talk in ASIC workshop on cosmic array fast timing → expect an LOI, but want to see connections with timing in collider detectors, timing and calibration synchronization to ps level will be a big challenge.
- Very fast timing with tracking to handle large pileup for triggering.

Discussion session: Trigger Algorithms, Streaming DAQ, Architecture (real-time analysis)

How do we see streaming DAQ evolving?
From ATLAS experiment, severe constraints for pixels, outer tracker maybe more possible
Combine system-on-a-chip with FPGAs to do more low-level processing, which may allow a more streaming DAQ concept to be applied to even pixels? Coherent design for the full processing chain.
Physics-enabling: long-lived particles for example, which currently rely on MEt triggers, but using pixel information may open up more sensitivity/phase space; would be good to flesh out these and other ideas integrated highly with architecture (maybe an ideal push for CPM meeting?). Bring more of the high-level trigger algorithms to the low-level.
Disappearing track signatures? How to trigger?

Work closely with ASIC designers for future algorithms and options: need to bring in TDAQ expertise at a very early stage to accommodate long lead-times (move the answer of can that be done from “no” to “maybe”!) → R&D to see what can be done needs to be years in advance: need to focus on the types of problems we want to solve, and not wait for a specific problem to solve too close to when we need it.

Trend in industry to move processing closer and closer → in 10 years, commercial sensors and processing elements may be viable. Should keep a close eye on the industry trends. Can we invite industry experts to come and present their view and feedback?

If a lot of the reconstruction is done in low-level hardware, HLT can build on existing algorithms/L1 trigger more directly. Related to 'data-scouting'/'trigger-level analysis'/'turbo-streaming': do processing and save trigger objects rather than the full raw data.

Asynchronous "L1" designs? Have a more asynchronous model for shipping data out and still be able to assemble and make use of data in higher-level triggering. Allows a trigger with more "burst" capability? Related to 'event-driven electronics' (clock-independent electronics). Drive the electronics when it has something to send and not forced on a clock.

Visionary: blur the lines between what we think of traditionally in what happens in FPGAs vs. computing. Think about how we can re-architect the overall TDAQ system with the technologies we now have available.

Can we lower the barrier for the broader physics community to contribute triggering algorithms? E.g. can we enable a way to have a lot of low-bandwidth triggers? Are there tools that could enable this?

Discussion session: Trigger hardware, heterogeneous computing

What types of beyond-GPU hardware would be most interesting for physics algorithms? Think about this a lot for machine learning, but could be for non-machine learning algorithms as well (e.g. what's the ideal computing hardware for Kalman filters? Sparse ML architectures may demand different types of hardware too).

Adjacent to much of this is networking technologies, and how to move data around more efficiently. Deterministic networking may play a part in stability for getting data where it needs to go.

Keep an eye on the quantum computing world? What connections do we have in TDAQ to future of quantum computing? Should invite someone from that community and explore that interface. (quantum sensing readout needs here: https://indico.fnal.gov/event/44698/)

Discussion session: Autonomous monitoring, control, and calibration

Amount of information in trigger system is very large and more than we can readout: can we do real-time analysis/"learning" of that data to tell us more about what's going on in the detector. E.g. general anomaly detection related to triggering on new physics?

Calibration/alignment be done in-situ, with ability to feed-back into the trigger system. Systems-level needs are important to understand here. Can we update weights for ML algs that are running in TDAQ hardware.

Should think about designing systems that could make use of that information, even if we don't know exactly what we would use it for. Monitoring/calibration information does take up not insignificant bandwidth: more intelligent monitoring can help reduce this/give more bandwidth for physics.

Automating how we digest this data is going to be necessary as we move to larger detectors.

Related to control/real-time systems for accelerators → an interface to explore with Accelerator Frontier and others. (in general, would be good to discuss overlap with accelerator needs for reading out and controlling a large/distributed system)