Currently manifesting R&D proposals

G.Barr 3.Apr.2018

Architecture view:

- Architecture rather clear from TP: Must test
  - CPU and PCIe throughput – how many machines do we need, do they all fit in the CUC?
  - Not just SP-TPC, we have four 'sub-detectors' [SP-TPC, DP-TPC, SP-Photon-detectors, DP-Photon-detectors]
  - Tradeoff between FPGA and CPU for the key features in our architecture: trigger hit finding, buffering, SSDs, network, link specs, memory management.
  - Run control, partitioning, subscribing to trigger decisions, subscribing to dataflow.

UK organization view:

- Ready ourselves for the documents we must provide the funding agency
  - Demonstrate UK organization and capability for the construction, procurement, QA/QC, etc.

Largely the things on the previous page are covered by the R&D plans I know are being talked about, more ideas are welcome.

- [From Kurt] https://docs.google.com/document/d/1L0-ylQ0nsAuc3ukKHqJLfykAyZG2Rk46OSKLZAjHUJs/edit?usp=sharing
  - [A PCIe specific one, subset of Kurt's proposal] https://docs.google.com/document/d/1LMQyjN4cwxU719mPMMeXy7R8fNDbMHc5__X86hXTdk4/edit?usp=sharing
- [Hardware – hopefully superseded by HW group] https://docs.google.com/document/d/1LMQyjN4cwxU719mPMMeXy7R8fNDbMHc5__X86hXTdk4/edit?usp=sharing
- [More specific about ASIC links] https://docs.google.com/document/d/1sFWnmVW5noHjVuzT-Vt7KrBS9HtNXd66AeTOvR6dGCA/edit?usp=sharing
- [SP photon detectors – quite good plan already] https://docs.google.com/document/d/1sFWnmVW5noHjVuzT-Vt7KrBS9HtNXd66AeTOvR6dGCA/edit?usp=sharing

Hardware possible discussion

General goals

1. Layout and board manufacture of standardised high speed links, such as 10Gb/s, 40Gb/s (which is 4 x 10Gb/s) or 100Gb/s (4 x 25Gb/s), or layout of the high speed lanes compatible with current and future PCIe, and/or DRAM. DPM Ultrascale+ upgrade
2. Layout, demonstration and board manufacture of customised high speed links that are interfaced directly with ASICs, as in the cold-ASICs in the DUNE SP TPC.
3. Design and demonstration of a processing chain inside FPGAs. SLAC have such a processing chain (the RCE for ProtoDUNE). Modify to allow continuous triggering. An alternative demonstrator could be a bump-on-cable using a DEV board, (e.g. using DP data)
4. Implementation to allow demonstration of the SLAC two-stage buffering scheme involving SSDs for SNB-triggers.
5. Demonstration of PCIe performance on commercial motherboards (could be SW-group). Do we want a PCIe evaluation board?

Possible proposed R&D step

'jig' board for testing purposes. Like a mini-COB. Has a socket for DPM (the new prototype, or the current generation ones used in ProtoDUNE), SFP+ for the high speed links, FMC connector for the other links from the DPM. Route to goal#1 from above.

A separate project could be to build the FMC to investigate #2 from above; i.e. to make a first stab at an ASIC-interface demonstrator. If we standardise the connections on the FMC, this could also be used on an PCIe evaluation card.

Backup slides

List of potential DUNE DAQ studies in run up to TDR.

DRAFT 0.11

View from: Architecture group

G.Barr 19.Mar.2018

Summary

(1) Continuing development of upgraded DPM using Ultrascale+ and large buffer memory

(2) DUNE specific RTM

(3) COB for DUNE

(4) Development of hardware demonstrator for prototype demonstration: Possible goals

(5) Enumeration of PCIe capability requirements (a.k.a. 'how many PCs do we need')

(6) Studies of trigger hit finding algorithms in FPGAs, GPUs, CPUs

(7) Extend studies of trigger hit finding capabilities by adding varying levels of noise and by using algorithms invented in (6) above, that are easy to translate to FPGA/GPU/CPU implementations

(8) In depth plan for interfacing dual phase TPC

(9) In depth plan for single and dual phase photo detectors

(10) Timing system design

(11) Run control design

(12) Front-end enhancements to software

- TPC-node software design
- Photon-detection-node software design

(13) Trigger farm enhancements to software (Module level and global level)

(14) Back-end enhancements to software

(1) Continuing development of upgraded DPM using Ultrascale+ and large buffer memory

- Currently a SLAC/Oxford, activity, beneficial to expand to other institutes, especially UK
- Status: PCB layout and 10Gb/s signal integrity studies well advanced. Should start prototype building in O(2month)
- Near-term future activities: Working out how to test ((i) on COB, or (ii) custom-test board, see (4) below), prepare firmware.

(2) DUNE specific RTM

- Design and review current technology for proposed 1.28Gb/s interface
- Near-term future activities: study and investigate more possibilities

(3) COB for DUNE

- Are there enhancements to the COB design?
- Institutes needed to take ownership of COB (production and elements of design)
- Component obsolescence up to 2024 (also ATCA crate components?)
- [This isn't a very good list, I feel a new institute needs preparation]

(4) Development of hardware demonstrator for prototype demonstration: Possible goals

- Simple tester for DPM (see (1) above)
- Some form of 'bump on cable' trigger primitive finder for e.g. dual phase to prototype FPGA-based trigger hit algorithms (see (6)) and run at ProtoDUNE
- In parallel with electronics consortium, develop the interface to cold electronics
- Cheaper version of a test stand for institutes who want to develop firmware? Part of CE QA/QC? Part of our QA/QC?
- This may be one board, or possibly more optimal to have two separate R&D paths; e.g. do CE-interface with dev board with an FMC, do bump on cable with DEV board
- Advantageous to demonstrate the interface with ASIC-based electronics early. Facilitates demonstration of capability to funding agencies to build stuff.

- Playing data into a test system... Important.

(5) Enumeration of PCIe capability requirements (a.k.a. 'how many PCs do we need')

- How many PCs do we need?
- Is a specialized motherboard necessary to maximize PCIe card usage; does this push us into niche vendor territory ?
- Start with experience and measurements from FELIX developers
- Test loading a PC with multiple data sources and sinks (GPUs, NICs, FELIX or other FPGA boards, etc.) and run them together
- Identify where the bottlenecks are and whether specialized motherboards will help.

(6) Studies of trigger hit finding algorithms in FPGAs, GPUs, CPUs

- Pick a fairly simple algorithm and test in each of the above platforms
- Pick something a bit simpler and something a bit more complex and test those as well
- Compare efficacy of algorithms chosen with studies in (5) below
- Since Altera and Intel merged, in future there will be Combi FPGA+Xeon. Does this help, e.g. to speed up decompression/unpacking?

(7) Extend studies of trigger hit finding capabilities by adding varying levels of noise and by using algorithms invented in (6) above, that are easy to translate to FPGA/GPU/CPU implementations

- E.g. make a plot of expected efficiency/purity as a function of noise added. At what point does the noise kill the experiment. How big a gap is there between the noise tolerable by a simple algorithm and the noise tolerable by a complex algorithm.
- Overlay ProtoDUNE noise
- Study how much the coherent noise contribution plays a role in the performance

(8) In depth plan for interfacing dual phase TPC

- Challenge is that data is proposed to be delivered compressed and without trigger hits. Expensive to un-compress in CPUs to find hits.
- Consider 'bump on cable' FPGA with 10Gb in 10Gb out which adds the missing info.
- Later, the bump on cable functions could be incorporated into either an FPGA on the dual-phase ADC card or an FPGA on a PCIe card

(9)In depth plan for single and dual phase photo detectors

- Challenge is that trigger requirements and data volume are not well defined, but there are some good baseline numbers for us to make a start with
- Plan how to store data with similar indexing system to TPC data waiting for trigger decisions from TPC.
- Do the maths to calculate a plausible coincidence requirement to generate triggers from photo-dets (time window, adjacency, n-hits threshold achievable, fake rate)
- Calculate whether random coincidences from single 'dark' hits or correlated patterns from radioactivity accumulating are the bigger background.

(10)Timing system design

- [Add some things, continuing from discussion for TP text]

(11) Run control design

(12)Front-end enhancements to software

- TPC-node software design
  - FELIX software
  - Integration into artDAQ
  - Oxford ‘Escalator’ software
  - Subscription based connection model (avoid run-start for any change)
  - Error detection
  - Inclusion of GPU processing option
  - Memory management (freeing memory in time for new data)
  - Trigger decision delivery
  - Does the ‘’bump-on-cable’ trick described in (2) and (6) ‘fix’ the compatibility between SP and DP TPC data?
  - [Critical] Assess how efficient the stuff runs to determine how many APAs can be processed in one node
- Photon-detection-node software design
  - Track latest photon detector readout system proposals; is the buffering or hit finding in the software node, or in hardware?
  - All the bullet points from the TPC node software above

(13)Trigger farm enhancements to software (Module level and global level)

- Integration with artDAQ
- Expand this section - Real time techniques are likely to be needed
- Investigate whether escalator protocol (Oxford idea) could help
- What is the expected latency. Can we achieve e.g. 1s latency?
- [Critical] Speed and fan-in architectural considerations

(14)Back-end enhancements to software

- Enhancements to artDAQ
- HLT architecture and throughput
- Disk writing technologies; estimation of disk storage capacity needs at FD
- Archival software choice
- Handshake with Fermilab