SBS GEM DAQ Electronics Review

Date: August 2014

Reviewer: David Abbott
Jefferson Lab DAQ Group

Charge:

Guidance and a recommendation on which readout system (INFN/MPD or CERN/SRS) to use for the Rear GEM Tracker using the following guidelines:

- Cost comparison
- Production timelines
- Integration with CODA DAQ system
- Meeting technical requirements (5kHz at 50% occupancy)
- Outside support
- Ability to upgrade

As a starting point I describe the detector systems as I understand them from the documentation presented. There are two main detectors, a Forward GEM tracker consisting of six planes and 41K channels (324 APV cards), and a Rear GEM tracker with 10 planes and a total of 61K channels (480 APV cards). The Forward Tracker has already been established as using the MPD/VME readout being developed by INFN. That is expected to continue.

If we consider the technical requirements of the Rear Tracker alone, I calculate for the specified 3-sample readout mode of the APV25 (1200 bytes/card):

Data rate = 1200 Bytes * 5000 Hz * 50% * 480 cards = 1440 MB/s

This is the total data rate coming off the detector. If we include the Forward GEM as well this generates an additional 972 MB/s. Based on this information my first question is not can either system meet that requirement (my opinion is yes they both should). My expectation is that one does not really need to actually keep all this data. A statement was made in one of the documents that one should actually be able to reduce the data volume by an order of magnitude in hardware. This seems to be a critical piece to consider.

How far through the DAQ system that quantity of data has to be dealt with makes a large impact on the size and scope of the entire system. If the data can be sparsified in the early stages (either in the MPD of the INFN system or in the ATCA/FEC of the CERN system) then the viable DAQ architectures are much more flexible. However, if there is the expectation that one needs to bring all the data together before some filtering can be done, or if one wishes to actually keep the entire set for analysis, then I see some problems.

Any VME/VXS solution for the front-end, where a CODA ROC is running on a VME single board computer, will have an inherent limitation from the VME backplane performance, which is 200MB/s for boards with the JLab VME interface design. This would also require the VME SBC to have an add-on 10G Ethernet interface card. If one uses the standard 1G interface then the Ethernet would become the limitation (~120MB/s). Therefore in order to pull 2.5GB/s off the both GEM detectors using any VME solution will take a minimum of 20-25 crates. This is nowhere in the scope for the project.

If however the MPD cards can execute filtering/suppression at the 90% level independently, then there are a couple VME solutions that can be considered. The 3 VME crate solution that was proposed will work. With 10 or 11 modules per crate the suppressed output data rate from each VME crate would be around 45-50MB/s. This is easily manageable. There is still the option of further aggregation of the data before filtering by implementing the alternate proposal of using the MPDs in a standalone mode in non-VME crates and using the optical outputs to aggregate the data in JLAB SSP modules in VME crates. One SSP board can handle up to 32 MPD output streams at full rate. If firmware could be developed in the SSP to reduce the data volume even more, one could envision readout of one SSP in one VXS crate would be all that was needed to handle the entire Rear GEM tracker.

In the CERN/SRS solution the CODA ROC must run on a PC. This would preferably be a rack-mounted server. It must have at least two PCIx slots. One slot must contain a JLAB PCIx Trigger Interface card. This is the endpoint for the JLAB trigger/clock/sync distribution system (there is a VXS version of this board as well). This interface insures that all front-end detector systems in the experiment are synchronized and that they get all accepted triggers. Some work may be necessary to get clock and trigger signals to the SRS SRU unit) system from this card. This may require development in coordination with the JLAB Electronics group.

The other PCIx slot presumably would hold a dual-port 10G Ethernet (or preferably Infiniband) network interface card. It could use this interface to accept the full bandwidth output from an SRU (even unsurpressed), and could also output the full bandwidth if necessary to some NAS RAID systems over a high-performance network (up to 56Gbit with Infiniband).

The one advantage of the CERN/SRS system proposed is that it provides the option of handing the full bandwidth of the detector with no suppression if necessary, assuming the necessary network and filesystem infrastructure is available (this would be a substantial increase in total costs). If however, data can be filtered at the ATCA/FEC level then the output could be handled by a single SRU and PC (ie one CODA ROC).

Cost Comparisons:
The hardware costs for both readout systems seem comparable. If there is a desire to keep support for full data bandwidth output for the detector then only the CERN/SRS is viable, but costs would rise significantly. Since the Front GEM tracker is already scheduled to use the INFN/MPD system, surely there would be significant cost and time saving with having a common system for both detectors in parts, labor, and overall efficiency. One note with the commercial ATCA system, there is no mention I can see of a spare ATCA crate or modules in the cost estimates. While there may be a number of possible spare VXS crates at JLAB, there will not likely be any ATCA equipment.

Production Timelines:
With the INFN/MPD system it seems the production timelines are within some control by the collaboration. My concern with the CERN/SRS system (in particular the commercial option), is that some timelines for hardware availability and for support for that hardware in case there are problems could be variable. I did not see a timeline for when experiment and detector installation was to begin in the reports, but both systems seem to be in similar stages of development. It is not clear one or the other would be ready sooner.

Integration with CODA:
I discussed some of how both systems would integrate with CODA above. The VME/VXS hardware exists and is tested, however, the PCIx version of the TI board is only in prototype. More work and library development on this card would be needed by the DAQ group. Firmware development in boards such as the SSP if it will be used would also have to be coordinated with the Electronics group.

Meeting Technical Requirements:
The issues with technical requirements were also discussed in more detail above. The CERN/SRS system definitely comes closer to meeting those specs in the most general flexible way. It should operate in both the worst and best case situations.

Outside Support:
The CERN/SRS system looks to depend more on outside support, both for R&D and regular maintenance, if one defines “outside” to include individuals or companies not part of the collaboration or resources like the DAQ and Electronics groups at JLAB. This could lead to unanticipated costs or delays if there are equipment failures or firmware/software bugs.

Ability to upgrade:
It is not clear where the biggest issues will lie for potential upgrades. Robust front-end electronics in radiation areas as well as noise levels for longer or shorter cables may ultimately play the biggest role to overall performance of the detector systems. The ability for one system or the other to easily adapt for these issues should be considered. This is not an area of expertise for this reviewer however.

Recommendation:
If one considers primarily the ability to meet the technical requirements for the detector then in my opinion the CERN/SRS system would be the most flexible system to consider. However, my primary concern is that the collaboration should put real effort into implementing some real, front-end, FPGA-based data sparsification. If this can be done to the level of 90% or better as suggested, then these technical requirements are not nearly as restrictive.

With this requirement eased, I believe the advantages of the INFN/MPD system are greater. Having a uniform readout system for both the Front and Back GEM detectors should result in a more stable and easily maintainable architecture. There is less reliance on outside support, and ultimately I believe it will cost less to implement and operate in the long term.