GLUEX L3 readiness review

On July 22, 2016, a mini-review was held to assess progress in the implementation, testing and deployment of the Level 3 trigger system required for high luminosity ($5x10^7$ photons/sec) running in Hall D with the GLUEX detector in the Fall of 2018. The committee was comprised of Amber Boehnlein, Graham Heyes and Stephen Wood. David Lawrence presented the status and goals for the implementation, testing and deployment of the L3 trigger system required for high luminosity running in the Fall of 2018. The presentation itself was thorough and demonstrated understanding of the goals and challenges. However, it was clear from the presentation that the level to which GLUEX understand the detector, trigger and DAQ is in a state of flux as they study the data and operational experience from the Spring 2016 run. While the committee recommends that progress in a number of areas be carefully monitored by management, and challenges still remain, they believe that high luminosity running of E12-13-003 in 2018 is achievable.

Observation

The proposal for a L3 trigger for Hall D has passed a number of technical tests including parsing the packaged events, running JANA in a filtering mode, and providing for a framework to do detailed studies on 'real' data and simulation data. Preliminary algorithms have been developed by experts who have previously worked on LHC-b.

Analysis of Spring 2016 data

The goal for the low luminosity Spring 2016 running was to operate at a luminosity of 10^7 γ/s with a DAQ design goal of 300 MB/s at that luminosity. The achieved luminosity was $0.8x10^7$ γ/s but the data rate was 800 MB/s. Study of the event size from Spring 2016 data showed that a contribution from accidentals causes a linear increase in event size with beam current. This adds a quadratic term to the data rate as a function of current. The goal for Fall 2018 is a luminosity of $5x10^7$ γ/s. An extrapolation from the Spring 2016 data gives a 9 GB/s data rate from the detector at this luminosity. This is a factor of three larger than the design goal based on previous calculations.

fADC125

The fADC125 has performance issues when operated at rates higher than originally specified. This will need to be fixed to handle the rates anticipated for Fall 2018.

L1 trigger

The L1 trigger was not discussed in the presentation but there was some discussion with members of the GLUEX group that L1 is not yet fully optimized and that a reduction in rate by an unknown factor may yet be possible. Since the event and data rates strongly drive the design and

cost of the DAQ, storage (both disk and tape) and offline analysis, understanding L1 should be given priority. This should be possible using the low luminosity Spring 2016 data.

Crate level readout

If L1 does not significantly improve, the Fall 2018 rates will lead to data rates from individual crates that will exceed the available bandwidth of the 1 GB/s network link from the single board VME processor. There was speculation that this link could be upgraded to 10 GB/s. This is technically feasible but untested. It is not clear that the VME based processor has the capacity to manage both the readout of the crate and the TCP stack on a 10 GB/s link. If this is to be considered it must be tested.

Even with 10 GB/s link, the VME backplane itself has a bandwidth limitation that may be reached at these rates. Options for mitigation (for example distributing front end modules over a greater number of crates to reduce per crate rates) should be evaluated on the basis of technical feasibility and cost.

System level readout

The GLUEX DAQ design goal was 3 GB/s from the detector and 300 MB/s to storage after a factor of ten L3 reduction for high luminosity running. If the extrapolation from low luminosity running is correct the new goal will be 9 BG/s from the detector with only a factor of four L3 reduction and 2.25 GB/s to storage. So a 7.5 times larger data rate to storage than the initial design goal. The modular nature of CODA gives flexibility in the DAQ architecture. The optimum architecture to handle these rates must be re-evaluated, implemented and tested.

A study should be made to evaluate the impact on offline analysis resources requested by GLUEX. Once this evaluation is complete the results must be discussed with SciComp in IT so that they can plan and provision for data storage and analysis.

L3 strategies

The presentation contained several slides describing L3 strategies. It seemed from the presentation that the relative timings of the different algorithms, and hence the rates per compute node that could be processed by a L3 farm, are reasonably understood for data acquired so far. However, it is not clear if the timings are predictable for the higher multiplicity events expected at higher luminosities. Also the original factor of ten event and data rate reduction between the L3 input and output has been reduced to a factor of maybe four.

Are L3 algorithms that reduce event sizes in addition to filter events being considered?

The Spring 2016 data along with reasonable extrapolation to high luminosity events sizes should be used to more accurately determine both the rates that can be handled per node and the reduction factor that can reasonably be expected.

L3 system design

The L3 farm system architecture as described in the presentation was designed around a farm of modest size, 50 nodes, and input rate of 3 GB/s. Current funding only covers a system of this scale. Does this architecture scale to 9 GB/s input rate? What is the true number of nodes that will be required to handle this rate? Careful attention should be given to whether the architecture as presented is cost-effective or even affordable. Are there alternatives that may make more sense in light of the unanticipated increase in event size?

The appropriate location of the L3 farm should be considered (Hall D vs. Computer Center). Given the cost of the L3 farm, provisions for use of this resource between running periods should be made.

Conclusion

GLUEX still faces many challenges before they can accept the rates that they believe the high luminosity run in the Fall of 2018 will present. Fortunately, they have a competent team and several opportunities to test strategies before that time. That said, the anticipated rates may require resources that are currently unfunded. The group should urgently study the factors that influence the DAQ, L3 trigger, data storage and analysis.