

Heavy Photon Experiment engineering run DAQ/offline readiness review

Committee:

Graham Heyes (chair), Sandy Philpott, Chris Cuevas

Observers:

Veronique Ziegler, Mark Ito, David Lawrence, Gagik Gavalyan

July 3, 2014

Initial comments

The HPS experiment is anticipating its first run late in 2014 and an experiment readiness review in July 2014. This internal review was called to give the HPS collaboration feedback on the state of readiness of the data acquisition hardware and software, calibration software, offline analysis, and ensure that the requirements for computing and storage are in line with what can be delivered by JLab.

The review took the form of several presentations over a two hour period during the HPS collaboration meeting. Although the committee members were already somewhat familiar with HPS no additional material was presented in advance of the presentations.

The charge to the committee is included in the appendix. In the remainder of this document will the focus topics listed in the charge are covered.

Data acquisition hardware

Charge:

Is hardware in place for high rate data acquisition?

If not is there an acceptable completion and commissioning schedule?

Observation:

Not all of the required hardware is in place. A timeline was presented that showed an expected completion date in August. This would give adequate time for testing since parts of the hardware are already tested.

One of the presentations focussed on the new SVT readout hardware. This hardware is quite different from any existing DAQ system at JLab with the ROC code running processors embedded in the custom electronics rather than a commercial VME board. Running the ROC on the embedded ROCs may present unforeseen issues. It is critical to get the new SVT readout hardware on site to give sufficient time for testing and integration with the rest of the DAQ. There are a lot of complex custom hardware components in the system, running the whole system together should be done soon. It was not made clear to the committee if the SVT detector itself will be ready for the initial runs, if it is not then the DAQ requirements for the initial runs are much simpler. However, having the SVT readout electronics in place before the run will be useful for integration tests.

Recommendations:

The committee recommend that a detailed plan for having the SVT readout in place and tested should be produced with a full detector test, including SVT readout electronics (even if there is no SVT detector) taking place at least a month before beam.

Data acquisition system

Charge

Has the DAQ system demonstrated readiness for high rate data acquisition?

If not is there an acceptable completion and commissioning schedule?

Observation:

Readiness of the DAQ was not completely demonstrated to the review committee. The data acquisition system is based on the 6 GeV era CODA 2.6 backend with updated ROC code. The use of the older CODA version has the benefit of being a stable and well understood software base. The target event rate of 50 kHz at the expected event size has not been demonstrated using the old CODA backend. The concern here is that, unlike CODA 3.0, the event builder in the older version of CODA has a single processing thread performing all of the event building. Similarly the event recorder has a single thread writing the data to disk. Essentially, at two points in the system the entire data stream passes through a single thread. It is not clear if this will be a limitation using the updated computers available to HPS. In recent years processors have increased in compute capacity by adding cores while the speed of an individual core has remained relatively constant.

Recommendation:

The committee recommends a rate test running at 50 kHz with the correct number of ROCs to prove that the requirements can be met.

JLab and IT resources

Charge

Are available JLAB resources adequate to transport and store the data?

Do the required resources match with the available JLAB computing resources?

If not what is needed?

Observations:

As far as reconstruction is concerned, yes. The software seems to be in very good shape with data volumes and processing loads which, while non trivial, are manageable and within the scope of what can be delivered by JLab.

In one of the presentations we were told that the plan was to output the reconstructed data in two formats DST and LCIO. One of the reviewers questioned this choice and the response was that the DST is a summary and a much more portable format for use offsite while the LCIO is a richer, lossless, format that can be referred to without having to go back to the raw data and repeat the reconstruction. This seemed reasonable response.

The situation regarding the resources required to perform simulation was less clear from the presentations.

An estimate was presented that 500 job slots would produce in half a year a simulated dataset equivalent to one week of raw data. With three weeks of running in 2015 this implies 270,000 cpu-days to simulate the complete dataset. This was concerning since this simulation requirement would be larger than the resources that could reasonably be allocated to HPS.

Some arguments were made that the final requirement will be smaller but nothing concrete was presented during the review. In the response to the review closeout a more detailed calculation was provided which showed requirements that are much more reasonable with a total of 30,000 cpu-days for simulation.

Recommendation:

The HPS collaboration should firm up the simulation and reconstruction requirements and present a detailed plan of what resources will be needed and when. The tape and disk requirement should be included in that plan.

Data calibration, alignment and quality checks Charge

Is the software for data calibration, alignment and quality checks ready?

Is the reconstruction software ready?

If not is there an acceptable completion and commissioning schedule?

Observation

The initial schedule that was shown had overrun issues and a revised schedule was presented that simply re-baselined by slipping it forward in time. The presenter commented that this was clearly unacceptable to the HPS group and the task list has since pared down to remove the niceties and leave the “should” and “must” tasks that would lead to an October 1 ready date. The revised and pared down schedule was not presented.

Recommendation

The HPS collaboration should clearly show the revised task list, what was discarded, why and the impact if any. The critical path to an October 1 readiness should be clearly identified.

Data analysis workflow and tools

Charge:

Are adequate plans in place for rapid data analysis?

Observations:

The committee suggest that HPS look into code profiling if they are not already doing it or document their process if they are, this was not clear from the presentations.

The area of software validation was not talked about in the presentations. Knowing that modifications to software do not have an undesirable effect on data quality is important. It would be useful to see a plan of how software quality checks will be applied during data analysis.

It wasn't clear to the committee what the expected momentum resolution is and what is the current resolution of the reconstruction. In the response to the closeout it was stated that at 2.2 GeV the resolution seen is ~4.5% which is in line with the value used to calculate the physics reach in the proposal. They state that the resolution of the reconstruction is expected to improve.

Conclusion

The committee heard many good things and was impressed overall by the presentations and the current state of readiness. The HPS group clearly understands the issues and many of the concerns of the committee were due to the limited time for presentations and the lack of supporting material before the review.

A lot of the software is in very good shape and the hardware is close to being ready, testing schedules are important. The monitoring software was very nice, having all the histograms on the same GUI was a good feature. One of the reviewers asked if there be automated histogram generation? If not then it would be a useful feature.

Will HPS be ready? Probably, but there are open questions. Manpower estimates were not clearly communicated. This is part of producing the overall plan between now and beam.

The state of the code and supporting documentation is good. One of the review observers did check out some of the code and managed to get it to compile.

Overall we were very impressed by the progress made so far.

Appendix 1

Charge to the review committee

1. Is hardware in place for high rate data acquisition? If not is there an acceptable completion and commissioning schedule?
2. Has the DAQ system demonstrated readiness for high rate data acquisition? If not is there an acceptable completion and commissioning schedule?
3. Are available JLAB resources adequate to transport and store the data?
4. Is the software for data calibration, alignment and quality checks ready? Is the reconstruction software ready? If not is there an acceptable completion and commissioning schedule
5. Do the required resources match with the available JLAB computing resources? If not what is needed?
6. Are adequate plans in place for rapid data analysis?