Scientist Interview with Richard Jones

Dr. Richard Jones, in the Physics Department at UConn, has been involved in data-intensive research for the past 10 years. He has been deeply involved in the GlueX collaboration, which consists of about 100 physicists from 30 institutions and 10 different countries. One activity among others that Richard leads in GlueX concerns using Open Science Grid (OSG) to run large-scale Monte Carlo simulations of the GlueX experiment. Prior to the start of physics data collection, the collaboration has organized a series of warm-up exercises called *data challenges*, aimed at developing the capacity to handle the large amounts of data (that the experiment will generate when it starts up at Thomas Jefferson National Accelerator Facility in October 2015). Richard is a Co-PI of the CC-NIE project. He is excited about the benefits that the Science DMZ will bring to his research. Dr. Bing Wang interviewed Richard on Feb. 25, 2014. The following are excerpts of the interview.

1. Could you briefly describe your research on GlueX and why it is data-intensive?

The goal of the GlueX collaboration is to understand the force that holds together protons and neutrons. Both neutrons and protons are made of elementary particles called quarks, which are bound tightly together by a force that nuclear physicists call the “gluonic field”. A beam of extremely high energy gamma rays can vibrate the quarks inside protons and neutrons and oscillate the gluonic bonds between them. GlueX attempts to experimentally observe these oscillations and measure their frequency spectrum. The experiment will be done at the Thomas Jefferson National Accelerator Facility in Newport News Virginia, also known at Jefferson Lab. Scientists came up with the GlueX idea in 1998, and began pitching it to the broader physics community as an opportunity for major scientific discovery. In 2004, the United States Dept. of Energy (DOE) announced the decision to support the project for $300M over the next 20 years. The construction of the detector and accelerator at Jefferson Lab is now nearing completion. It is expected to start commissioning in 2014, with physics data collection starting in 2015.

The volume of data produced by the GlueX experiment is expected to be very large, requiring massive storage and computing resources to handle it. The detector itself generates 4 TB/s. These data are distributed to many parallel processors which compress and filter the data, reducing the volume to the level of 100 MB/s. Experiments at Jefferson Lab typically run for 2500 hours per year, which yields a total volume of filtered data of 1 PB per year. To interpret the data collected from the detector, a similar amount of data must be generated from Monte Carlo simulation that enables the experimenters to connect the hit patterns seen in the detectors with the excitations of mesons that produced them.

2. You have been contributing to the Open Science Grid for a long time. Could you describe your main activities there? Could you describe an example workflow on the Open Science Grid?

My activity with OSG was primarily motivated by the need of the GlueX project. To be able to generate and analyze the required amount of data from Monte Carlo simulation, one or two local computing clusters are not enough. We see grid computing as an opportunity. In 2009, we (UConn, CMU and Indiana University) successfully obtained a grant from the National Science Foundation *Physics at the*

Information Frontier to create a new Virtual Organization (VO) service for GlueX at UConn and begin learning how to use grid tools for this application. Coupling local cluster resources to this service, we have been able to contribute unused cycles on UConn servers to other users on the OSG, and leverage that investment later by gaining access to the entire grid for high-throughput GlueX simulations.

An example workflow is

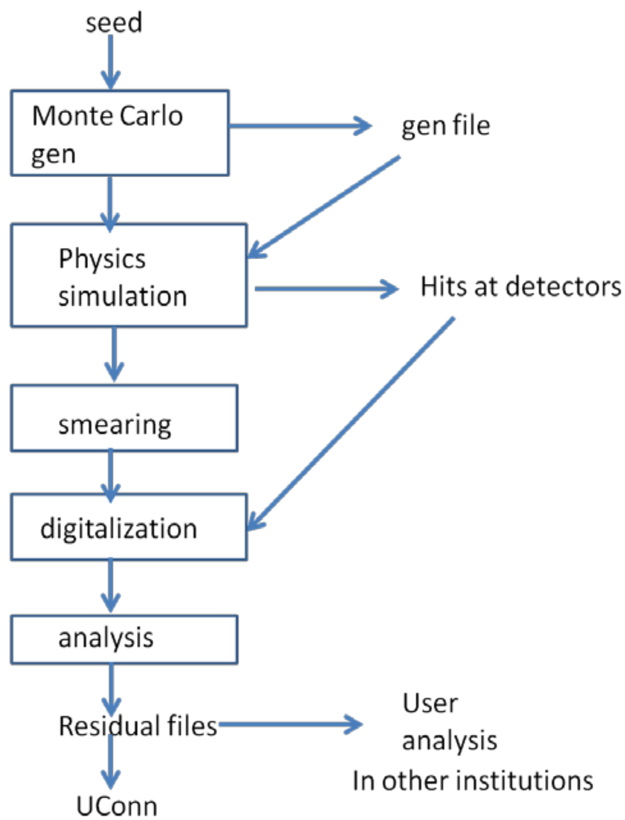

3. Have you encountered any difficulties (in terms of networking) while using the Open Science Grid in the past?

We have encountered many issues with network bottlenecks and firewalls. It is difficult to debug issues related to firewalls when submitting jobs at a large scale (50,000 jobs). Other issues involved actual bugs in software that only showed up in large scale production. For instance, rare circumstances can cause a few jobs to go into infinite loops, causing the jobs to be cancelled by the grid job monitor and returned to the run queue. After many hours of production, the run queues eventually fill up with these non-terminating jobs, leading to efficiency close to zero.

4. You mentioned that the first data challenge was in December 2012. What institutions were involved in the challenge? How much data did you generate? How much was transferred to/from UConn? What are the challenges involved? How do they inform the second challenge that is going to happen soon?

The first data challenge was around Christmas time in 2012. It generated 10TB over 10 days (1 Petabyte raw data, reduced to 10TB data after analysis). Data were generated using several scientific computing facilities, including the OSG, the Jefferson Lab computer center, and a computer cluster at CMU. It turned out that the majority (90%) of the data were generated on the OSG. At the peak, we were able to run about 10,000 jobs simultaneously on the grid. Using the OSG was a great experience. The amount of work we were able to finish makes a convincing argument that the OSG is very helpful for GlueX project.

The amount of data that were transferred to UConn and stored online for consumption by the GlueX collaboration was around 10TB. These data include those returned from OSG and those pulled from Jefferson Lab and CMU. Since that time, the data at UConn have been downloaded by several other GlueX institutions for their research. Due to various issues that limited our production during the first data challenge, the running efficiency on the OSG was only around 50% (i.e., the number of useful CPU hours consumed divided by the total wall clock occupied by jobs on a cpu core). We are aiming to increase the efficiency to 90% in later challenges.

The first data challenge provided many important insights. One by-product of this exercise was a very large sample of simulated physics data, which has been very useful to debug the software. Analysis of these data also revealed a gap in the ability of the experiment to distinguish between a few similar-looking physics reactions. Based on this, GlueX collaborators at MIT have developed a new proposal for an upgrade to the GlueX apparatus that will help to resolve the ambiguities and close the gap. This proposal was recently approved by the DOE.

6. Could you describe a bit more about the second data challenge? How much data are expected to be generated? How much data are expected to be transferred to/from UConn?

The second data challenge is expected to take place later this year. The goal will be to generate 50 TB of data. The reason for larger amount of data is because the simulation has become much more realistic (e.g., adding background noise), leading to larger data files.

7. What are your expectations in the science DMZ project to facilitate your data-intensive research?

Data produced from simulation will grow as the simulation becomes more and more realistic. Larger files mean longer-lasting network connections for each transfer, and also more active connections at once. And also, as we push up the scale, the production will take up a larger part of the firewall connection tables and may reveal new limiting factors. The Science DMZ comes at a very convenient time to help with resolving these potential issues.

8. For the peer institutions that you are working with in GlueX, do you transfer data back and forth?

We have pulled data from Jefferson Lab and CMU (after the first data challenge). Other institutions have pulled data from UConn.

9. What are the timelines for other near-term data challenges?

The closest one will be in March 2014. The next one will be in December 2014, after getting early data from the first GlueX commissioning run. Future data challenges will depend on whether or not new production bottlenecks emerge, as we ramp up to full simulation and analysis rates.