DOE FY2013 Annual Progress Review

Response to Technical Questions

For the LQCD-ext Project Team

and USQCD Collaboration

May 10, 2013

DOE FY2013 Annual Progress Review
Thomas Jefferson National Accelerator Facility
May 9-10, 2013

8) With regard to the User Survey, is there any correlation between respondent job classifications and satisfaction levels?

In principal, the data exists in the survey results to determine whether such correlations exist. However, mining and analyzing the data from this year's survey at this level would be unethical in that it would break the condition of anonymity that we declared when we announced the survey. Given the small population, it is possible that such an analysis would divulge responses provided by a specific individual.

We agree that such an analysis might provide useful insight and will consider ways of extracting such information from future surveys while also protecting respondent identities.

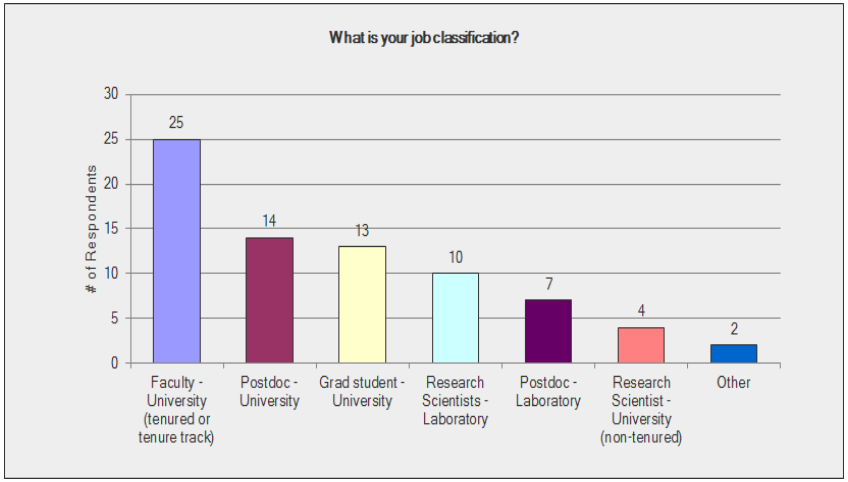

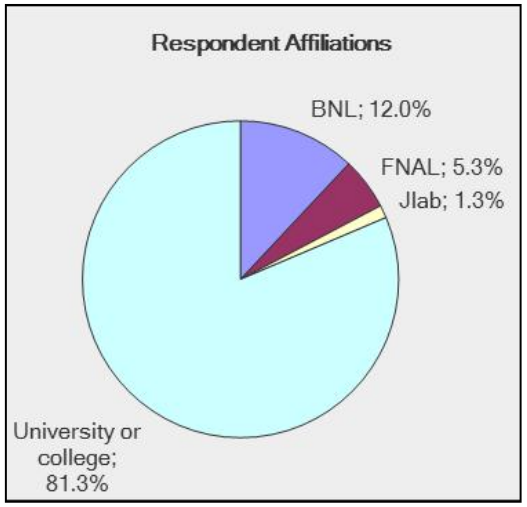

9) Were there any survey questions specific to GPUs?

No, we did not ask any questions specific to GPUs. We did include a question related to which compute facility was used for the majority of a respondent's work, but we did not break down the request further to specify hardware type.

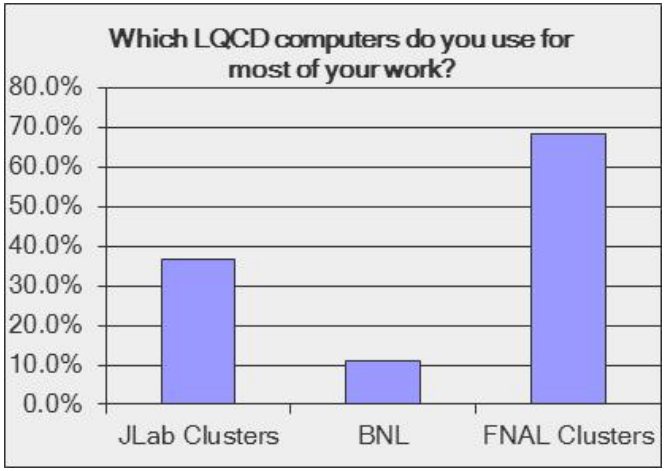

Again, we agree that the additional granularity might provide useful insight in identifying areas of strength and areas requiring improvement. In preparing future surveys, we will consider restructuring the survey in a hierarchical manner that will allow us to correlate ratings and feedback to specific hardware types.

However, we will need to balance this against the need to keep the survey simple and easy to complete. Otherwise, we will revert back to a more complicated and lengthy survey which may reduce the survey response rate.

13) Should LQCD-ext put aside funds in the LQCD-ext project budget for new hardware ideas?

Yes. An appropriate level of prototyping helps minimize risk for the later large procurements of the project. On the lower end, this might involve purchasing one or two compute nodes of some new type, for example, with a new model processor, with a different brand or speed Infiniband interface, or with a new accelerator. On occasion, a larger quantity of nodes with networking might be necessary to understand behavior of parallel codes; generally 8 to 16 systems are sufficient.

The project could set aside about $50K in each year for such prototyping. In years with little or no prototyping, the balance would be used in the main system purchase.

16) There is a concern that IBM's commitment to the BG/Q family is not certain.

a) Will IBM continue to pursue the Blue Gene line of supercomputers?

At present IBM is seriously developing a Blue Gene/Q follow-on. Of course, the future is difficult to predict.

b) Would the current investment in Blue Gene/Q hardware and software be a waste if IBM fails to continue the Blue Gene line?

The Blue Gene/Q is an IBM product sold in large quantities to many customers throughout the world. It should have the same 4-5 year life cycle as the earlier Blue Gene machines.

The hardware and software investment in Blue Gene/Q has been made to enable important physics calculations on this platform and is already a large success with many hundreds of millions of core hours presently being consumed by USQCD for critical projects. This is not a waste but instead is allowing us to exploit a tremendous opportunity.

Whatever small risk is associated with investment in the product of a single company, IBM, NVIDIA, AMD or Intel, is mitigated by the diversity of platforms targeted by USQCD.

18) What fraction of the groups use 100% of their allocations?

At Fermilab in 2011-2012 (conventional):

177.4M core-hrs were allocated

194.6M core-hrs were used by projects (109.7%)

14 Class A&B projects met or exceeded 100% of their allocations

8 Class A&B projects did not use their entire allocation (36%)

Total unused by these projects: 9.39M -> 5.3% of allocated, 4.8% of used

At JLab in 2011-2012 (conventional):

77.56M core-hrs were allocated

78.22M core-hrs were used by projects (100.9%)

7 Class A&B projects met or exceeded 100% of their allocations

3 Class A&B projects did not use their entire allocations (30%)

Total unused by these projects: 10.02M -> 12.9% of allocated, 12.8% of used

At JLab in 2011-2012 (GPU):

3.476M GPU-hrs allocated

3.972M GPU-hrs used by projects (114.3%)

5 Class A&B projects met or exceeded 100% of their allocations

5 Class A&B projects did not use their entire allocations (50%)

Total unused by these projects: 0.715M -> 20.6% of allocated, 18.0% of used

19) Please present both the BG/Q and Ds cluster weak scaling plots as speedup plots.

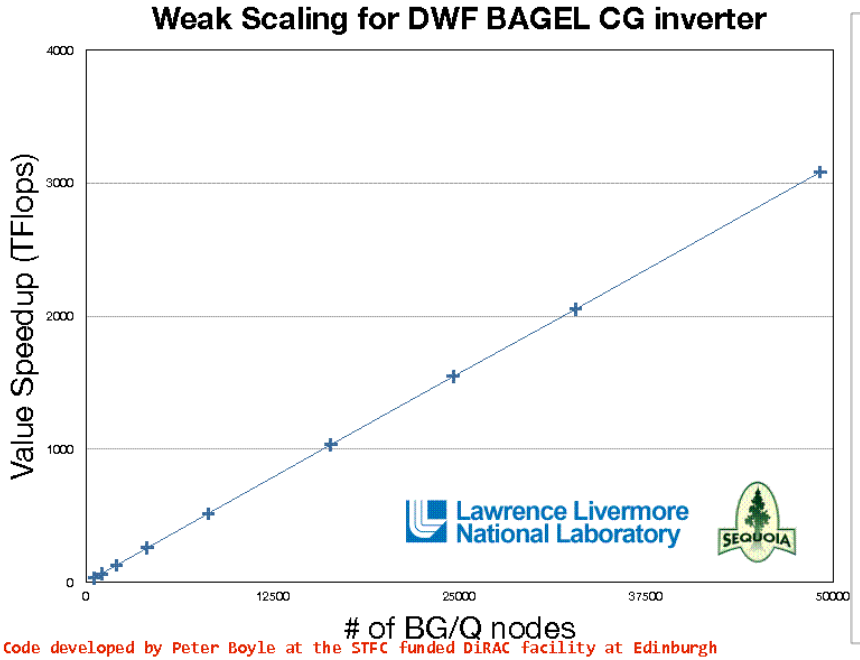

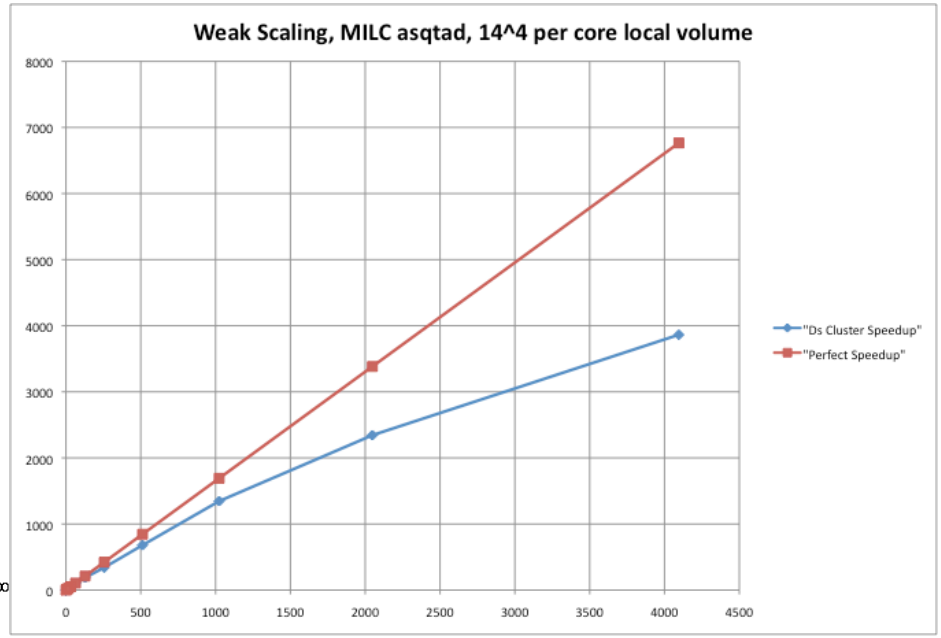

20) Intel's MIC architecture will likely be a significant hardware option going forward. Is it integrated into the project's plans?

We agree, based upon Intel's obvious commitment, the apparent ease (relative to GPUs) of performing initial code ports, and the very encouraging performance results seen so far on LQCD code by the JLab group.

Since software development is out of scope for the LQCD-ext project, we would only deploy large MIC-based resources if software were available for a large enough fraction of the USQCD community's work. Software porting and optimization for the MIC architecture is included, however, in the LQCD SciDAC-3 program. Assuming the availability of ported code, we would deploy a Phi-accelerated cluster if it satisfied our criterion of best optimizing the scientific production of our portfolio of hardware in comparison with the alternatives.