

DOE FY2016 Annual Progress Review

Response to Technical Questions

Bill Boroski, Rob Kennedy, Amitoj Singh, and Chip Watson

for the

LQCD-ext II Project Team
and USQCD Collaboration

June 29, 2016

DOE FY16 Annual Progress Review

Jefferson Lab

June 28-29, 2016

7) Multi-architecture: How do we get useful information from users on the choice of architecture (Chip)

New architectures are typically explored by one or more members of the collaboration. For example, already in 2008, 2 LQCD postdocs at BU were exploring GPUs, following the success of a group in Germany in using GPUs for LQCD. The success of these postdocs in software development laid the groundwork for a large uptake of GPUs at Jefferson Lab in 2009 as part of an ARRA funded computing activity at JLab, complementary to the LQCD-ext project.

For Xeon Phi, the story is similar: Jefferson Lab attracted Intel's attention by our large uptake of GPUs, and Intel wanted to understand why GPUs worked for us, and how they could better compete. That led to initial studies by Balint Joo and Intel researchers in running LQCD on Knights Corner, and then later on Knights Landing.

In an earlier example, LQCD's engagement with IBM came through Columbia University's work on QCDOC (QCD on a chip), which helped inform IBM's development of the initial BlueGene chip, with Columbia University's participation on part of that chip's design. Once the BlueGene supercomputers began to be deployed at large scale, the larger USQCD community moved to use them.

7) Multi-architecture: How do we get useful information from users on the choice of architecture.....cont'd

In each case, the same story unfolds:

1. an LQCD researcher takes an early exploratory look at a chip
2. key kernels are ported to test the chip's capability
3. a larger group of people become involved, especially once the chip reaches the point of being announced to be coming as a Supercomputer at large scale; more code is developed, and the LQCD community is almost always ready with high performance code day 1
4. the LQCD Computing project buys that architecture once there is both enough interest and enough software to justify its deployment (GPUs were the exception in that they were perceived as valuable for capacity running independent of deployment at large scale).

We do not have a formal process per se but Site Architects at each site work closely with the users and vendors to explore new hardware architectures. Often it is computer experts who make the first two steps even before the collaboration at large indicates an interest. Interest grows later because of anticipated large scale deployments.

Consequently, the collaboration doesn't really lead the selection of new architectures, rather it confirms the viability and usefulness of a new architecture as it emerges within the collaboration. This interaction takes place at annual "All Hands" meetings, as well as in workshops (2009 GPU workshop at JLab) and virtual committee meetings (2016 committee to review readiness of USQCD to exploit KNL).

9) Describe the FY17 Acquisition Process

We have a documented procurement strategy and a well-defined process for determining the optimum acquisition and deployment plan for every acquisition we undertake. This process is executed on an annual basis and has been exercised successfully many times.

As a reference, the following excerpt is taken from the document “LQCD-ext II Acquisition Strategy, v 1.3”, which is posted on the review website.

Procurement Strategy

LQCD-Ext II will procure as many as four separate lattice QCD computing systems, one in each of the final four years of the project. We consider a mixed conventional and GPU-accelerated cluster purchase to be a single procurement, as these would take place at a single host laboratory typically using a single purchase contract. If appropriate, the hardware budget from two years might be used to procure a single larger system. **The guiding principle of all of these procurements is that the most cost effective hardware will be deployed, where effectiveness is judged by the quantity of science (and of course, quality of science in terms of the reliability of the numerical results) that will be produced during the lifetime of the individual lattice QCD system.** In addition to commodity hardware and GPU-accelerated clusters, similar to those deployed during LQCD-ext, we will evaluate alternatives such as the IBM BlueGene family of computers, traditional supercomputers such as the Cray series, and other hardware suitable for lattice QCD calculations that may emerge.

9) Describe the FY17 Acquisition Process (2)

Assuming CR16-01 is approved, we will execute two procurement actions in FY17

1. Procuring the second portion of the JLab acquisition (using ~1/3 of the FY17 hardware budget).
2. Procuring and deploying the first portion of a new compute cluster at BNL (using ~2/3 of the FY17 hardware budget).

JLab Acquisition

We will assess the performance of the FY16 KNL machine to confirm that the system meets or exceeds performance requirements.

1. If the system performs as expected, we will exercise the options on the FY16 contract according to the timeline presented in the FY16 Acquisition Plan.

Sep 2016	Evaluation and Recommendation
Oct 2016	Expansion award
Dec 2016	Delivery & Acceptance Testing
Jan 2017	Release to Operations

2. If the system fails to perform as expected, we will piggyback onto our standard acquisition process to consider alternate options.

9) Describe the FY17 Acquisition Process (3)

BNL Acquisition (cont'd)

The BNL acquisition process will follow the procedure documented in our Acquisition Strategy document and refined in our Annual Acquisition Plans

The Project Manager is responsible for developing the FY17 Acquisition Planning Process. [Acquisition Planning Process](#) (example)

The BNL Site Architects are responsible for developing the FY17 Alternatives Analysis and Acquisition Plan.

The alternatives analysis document estimates the cost and benefits of multiple known options as well as a “do nothing” option for completeness. This feeds into the RFI process in the fall. Additionally, as part of building up a procurement RFP, details are refined to constrain the range of responses that vendors will submit, with details on network bandwidth, memory speed, etc. to help guide the vendors in their optimizations.

[FY16 Alternatives Analysis](#) (example)

The Acquisition Strategy, Acquisition Planning Process, and Alternatives Analysis documents will be updated for FY17 as appropriate based on the outcome of the CR16-01 change request process.

9) Describe the FY17 Acquisition Process (4)

BNL Acquisition (cont'd)

Excerpt from the [Acquisition Strategy](#) document (posted on FY16 review website)

All procurements will be performed by the host laboratory chosen for the particular hardware deployment. Such purchases will utilize firm, fixed-price contracts. The typical sequence for new deployments will be:

1. In consultation with the USQCD community (through the Executive Committee and Scientific Program Committee), determine anticipated usage profiles for new deployments (*e.g.*, distribution of job types and sizes, file I/O requirements)
2. Complete preliminary design
3. Issue a Request for Information (RFI) to likely vendors
4. Evaluate the RFI responses and complete a final design
5. Obtain host laboratory purchase approvals via the local requisition process
6. Issue a Request for Proposal (RFP) to likely vendors
7. Evaluate RFP responses and award purchase contract
8. Approve sample node and/or sample scalable unit (rack)
9. Test and approve vendor-integrated final system (acceptance test)
10. Operate final system in “friendly user” mode and tune the configuration
11. Release the final system to users

9) Describe the FY17 Acquisition Process (5)

Summary

We have used this acquisition process successfully since 2009, making adjustments annually to improve process effectiveness and efficiency.

Example: Last year we initiated the Acquisition Review Committee to provide an independent review of the proposed acquisition alternatives to confirm that USCQD needs would be met and that the new machine would be usable by a suitable fraction of the community. This year, we will incorporate the functions of this committee into our standard acquisition process and establish the committee earlier in the process.

Our acquisition process has worked well for many years to procure and deploy hardware systems that have been highly utilized by the USQCD community.