

**FY17 Acquisition Evaluation Committee Report
for the
Lattice QCD Computing Project Extension II
(LQCD-ext II)**

Unique Project (Investment) Identifier: 019-20-01-21-02-1032-00

Operated at
Brookhaven National Laboratory
Fermi National Accelerator Laboratory
Thomas Jefferson National Accelerator Facility

for the
U.S. Department of Energy
Office of Science
Offices of High Energy and Nuclear Physics

Final Version

November 15, 2017

PREPARED BY:

The FY17 Acquisition Evaluation Committee

Rob Kennedy¹, Gerard Bernabeu Altayo², Carleton DeTar, Steve Gottlieb, Don Holmgren, Chulwoo Jung, James Osborn, Amitoj Singh, Chip Watson, Frank Winter, Alexandr Zaytsev

1 Deceased

2 Resigned

FY15 LQCD-ext II User Survey Report Change Log

Version	Description	Date
0.5	Initial draft for review by Committee and Project Manager	5/8/2017
0.9	Revised draft ready for public consumption	5/12/2017
0.95	Revised draft for review by Committee and Project Manager	10/23/2017
1.0	Final report	<i>Planned: 11/13/2017</i>

Table of Contents

- 1 Introduction and Background²
- 2 Deliverable: Recommendation to the Project Manager²
- 3 Deliverable: USQCD-Specific Software Benchmarks³
- 4 Deliverable: Addressing Charge Elements⁴
- 5 Suggestion for Future Acquisition Evaluations⁸
- 6 Appendix: Charge to the FY17 Acquisition Evaluation Committee⁹

1 Introduction and Background

The purpose of the FY17/FY18 Acquisition Evaluation Committee is to review and consider existing and planned BNL institutional cluster computer resources and provide input to the LQCD-ext II Project Manager regarding the alignment of the proposed allocation of computer time and procurement of new computer hardware with the anticipated computing needs of the USQCD scientific program for FY18. The intent is to help ensure that the project is making the most effective use of project resources to further the USQCD scientific program..

This evaluation has taken place in two stages. The evaluation began in Spring 2017, but was paused because of budgetary and policy uncertainties. At that time, it was assumed that BNL would be acquiring separate hardware dedicated to USQCD use. Evaluation resumed in Fall 2017 with the assumption that the acquired hardware would instead become a part of the BNL Institutional Cluster (IC). LQCD participation in the BNL IC gives USQCD access to a wide range of hardware offerings but also imposes some constraints on hardware selection. Since the selection of the most cost-effective hardware depends on negotiations between the LQCD project manager and BNL, the purpose of this evaluation is to determine which options could meet USQCD computing needs and to set forth factors that should be considered in making the hardware selection.

2 Deliverable: Recommendation to the Project Manager

The following recommendation is made by the FY17 Acquisition Evaluation Committee:

The committee is evaluating three BNL cluster architectures for possible USQD use, namely, a proposed, expanded Skylake cluster with an EDR network, the existing KNL cluster, possibly expanded, and the existing GPU cluster, currently being expanded.

We find that all three would be of interest for meeting USQCD computing needs. Currently USQCD has access only to the GPU cluster. To meet present USQCD demand, we would need to keep our GPU usage at approximately the same level, but we would need significant access to either the Skylake expansion or the existing or expanded KNL cluster, or a combination of these two.

With the BNL institutional cluster model, there is flexibility in selecting a mix of all three for USQCD use. However, the schedule for BNL institutional cluster hardware acquisitions depends in part on the emphasis USQCD places on the different architectures. Thus, for this evaluation we consider two alternatives discussed in the accompanying Alternatives Analysis document.

1. Request the Skylake expansion and split our FY18 usage among Skylake, the existing KNL cluster, and the GPU cluster.
2. Request no Skylake time, but split our FY18 usage between a (possibly expanded) KNL cluster and the GPU cluster.

We prefer alternative 1, because it offers the earliest availability of additional computing capacity for the project, and it provides USQCD with the greatest flexibility, in allowing users to choose the architecture on which their codes run best. The enhanced KNL alternative 2 might provide higher performance on the inverters used heavily by the project, but a KNL expansion would come online as much as four months later than a Skylake expansion, and it would not accommodate codes that run better on more conventional cluster CPUs.

Although at present we do not see USQCD user demand for greater GPU usage, this could change in FY19 as FNAL and JLab retire their aging GPU clusters and as more users port their codes to GPUs.

3 Deliverable: USQCD-Specific Software Benchmarks

A set of USQCD-specific software benchmarks has been developed, and performance data has been collected to assist in evaluating the BNL hardware options. These benchmarks are summarized by Bob Mawhinney in an updated Alternatives Analysis document.

- The following codes and run-time configurations have been used to provide a somewhat “portfolio” view of the performance. In all cases the most heavily used sparse matrix solver provides the quantitative measure of performance. We take this performance measure only as a rough indication of hardware effectiveness, since some calculations may emphasize other algorithms or they may depend heavily on system infrastructure, such as I/O bandwidth.
 - DWF using Grid, 24⁴ local volume, run on 1 and 16 nodes

This highly optimized code represents a significant fraction of USQCD computation.

- MILC code with optimizations, 32^4 local volume, run on 1 and 16 nodes. Optimized single-node code is representative of a significant fraction of USQCD computation.
- MILC code with generic C code, 32^4 local volume, run on 1 and 16 nodes. The code without optimizations is considered to be representative of all non-optimized USQCD code.
- Benchmarks were discussed in Bob Mawhinney's talk at the April 2017 USQCD All-Hands Meeting:
 - <http://www.usqcd.org/meetings/allHands2017/slides/Mawhinney.pdf>
- Benchmark results are also discussed in the updated Alternatives Analysis dated 11/06/2017.

4 Deliverable: Addressing Charge Elements

Element 1: The near- and long-term demand (at a high level) for each hardware architecture in the existing portfolio and how the proposed procurement of additional compute cycles at BNL will augment or complement the existing hardware portfolio;

- Near-term Demand:
 - Anna Hasenfratz and Aida El-Khadra, of the USQCD SPC, state that in the 2017 allocation year proposals, the architecture demand was as follows:
 - CPUs are over-requested by a factor of 2.49
 - GPUs are over-requested by a factor of 0.98
 - KNLs are over-requested by a factor of 2.19
 - We note that the procedure for summarizing the proposal demands does not take into account that some proposals can be served by CPU/MIC or GPU, so while the demand for GPUs relative to their supply is less than other architectures, they may not be as under-requested as the reported values indicate.
- Long-term Demand: In the long-term we see:
 - Continued demand for CPU (whether conventional or MIC) technologies
 - A significant fraction of conventional CPU demand metamorphosing into MIC and GPU demand, when higher performance is available
 - Continued demand for GPU technology
 - Technologies that appear likely to be the most cost-effective for LQCD going forward are:
 - MIC (example today is Intel KNL)
 - GPU (example today is NVIDIA Pascal)
 - CPU (example today is Intel Skylake)
- How will proposed acquisition augment or complement existing hardware portfolio:
 - It is desirable to maintain flexibility in the USQCD portfolio, since scientific goals, algorithms, and software are constantly evolving. The available hardware acquisition choices with the BNL institutional cluster are to add to one of three existing clusters there, namely, NVIDIA/GPU, Intel/KNL, and Intel/Skylake. Whichever choice is made, USQCD would have access to other clusters, but only

if the cycles were available. They are available on half of the KNL cluster (72 nodes) without further additions. USQCD already has access to the GPU cluster (40 nodes), so this capacity is presumed to remain available. However, the current Skylake cluster is fully subscribed to RHIC and Atlas. So, without further addition, Skylake would not be available to USQCD,

- Currently, CPU systems are available only at FNAL and, to some extent, JLab. The need for CPU systems in the near term will become more critical as the FNAL Bc cluster will be retired by June 2018, leaving only the four-year-old FNAL Pi0 cluster as the available CPU-based system. The smaller Jlab system will also be retired and the end of June.
- KNL systems could provide a substitute for CPU systems if performance is adequate on code without the need for too much optimization. Our sparse solver benchmarks show that for optimized software, this option provides higher performance for single-node jobs but node-for-node, roughly comparable performance for multi-node jobs (which still makes KNL more cost effective as the nodes are cheaper).
- At the moment, the demand for GPU systems appears to have been met; they are less over-requested than CPUs and KNLs. GPUs are sufficiently cost-performant to meet LQCD computing project goals provided enough of the software portfolio can run on GPUs. That said, the 12k cluster at Jlab also hits end of life at the end of the current allocation year; 12k is 30% of the current USQCD GPU capacity
- There are questions currently about our ability to support very large memory jobs: on GPU systems due to bandwidth, on KNL systems due to operational robustness at large memory footprint scale (which appears to be improving), and on CPUs systems due to cost-performance. Skylake as an architecture could address this, but only if enough nodes were connected to a fast network fabric. The large-memory cluster at FNAL, the Pi0 cluster goes out of warranty in September 2019. It has 314 nodes of 128 GB memory each (39 TB total), with most jobs running no larger than 128 nodes. To match that would require 100 Skylake nodes at 384 GB per node, more than the project can afford with its residual FY17 funds. The KNL clusters have been operated at 64 nodes (12 TB total memory), and could presumably reach 128 nodes. 256 might be a stretch (needs to be demonstrated), and so for planning purposes it is assumed that larger jobs (above 128*192 GB) large will not be hosted on USQCD resources.

Element 2: Alternate computing architectures that may better meet USQCD needs, considering compatibility with the existing hardware portfolio and infrastructure, as well as the existing software portfolio;

- Other options include IBM OpenPower, ARM/ARM64, and, possibly, FPGAs. However, since BNL already offers three viable computing choices, there is no compelling need to consider these alternative options further.

Element 3: The availability of production software for use by the USQCD collaboration to effectively utilize the capabilities of the proposed procurement of additional compute cycles;

- CPUs: essentially all USQCD software can run on CPUs, provided it is not optimized to some other architecture, though it may not necessarily run at peak performance.
- KNL: most USQCD software can run on KNL, provided it is not optimized to some other architecture. This may require running in a backward-compatibility memory mode, but that can be arranged with the local site administrators for specific jobs.
- GPU: some USQCD software can run on GPUs provided it is optimized for that architecture. Large-memory algorithms however are not as effective on GPUs due to limited total memory per host in the packaging that we could afford.

Element 4: The ability of the proposed acquisition, along with the existing hardware portfolio, to meet the established time-based performance goals for the computing project;

- Because more USQCD code already runs on CPU systems, predicting portfolio performance on this architecture is easy. Acquisition of Intel Skylake CPUs can meet time-based performance goals subject to available LQCD funding. As mentioned above, the current BNL Skylake RHIC/ATLAS cluster has been installed, but with a network too slow to be of much use for most USQCD computing. The vendor has provided a short-term option for the purchase of additional Skylake nodes, which would make it feasible to start operations in the second quarter of FY18. BNL would connect them with a low latency robust network, which would then make it useful for USQCD calculations of size up to 64 nodes (12 TB).
- KNLs can meet time-based performance goals subject to available LQCD funding. Performance depends on the degree to which users are able to optimize their code, as with the Grid and QPhiX software strategies. Network bottlenecks have limited the performance of even these highly optimized software strategies for multi-node applications. Although Intel is working on solutions, a prudent cost-performance calculation should be based on present-day performance. The enhanced KNL option would require a lengthier procurement process compared with joining the existing Skylake procurement, with SkyLake taking of order 2 months and KNL of order 6 months.
- GPUs can meet the time-based performance goals, subject to available LQCD funding and provided enough of the software portfolio can run on GPUs. The GPU option would also require a lengthier procurement process compared with Skylake.
- A large fraction of USQCD computing requires high bandwidth, low latency networking, such as Infiniband or Omnipath. Some applications require high I/O bandwidth, and some high memory capacity. Thus, the above options are of interest, provided the acquisition includes networking and memory suitable to lattice QCD calculations.

Element 5: The capability of the project team to effectively support the computing hardware in the proposed acquisition, in terms of 1) meeting system uptime target goals; and 2) supporting the user community in the use of the newly acquired hardware;

- BNL has enough experience with the three hardware options that uptime and maintenance considerations are already translated into a cost per unit of hardware and a cost per core-hour. BNL purchases additional unallocated nodes as spares to cover downtime. Thus, systems with lower reliability will require a larger number of spare nodes which translates to a higher hardware cost. Thus, operational robustness is less of a concern to this committee. We offer a few comments.
- CPU systems: These are an established architecture. We consider them to have the lowest risk of not meeting uptime goals and the most supportable hardware for the user community.
- KNL systems: These systems have known issues that are being addressed: for example, Fast Memory fragmentation in cache mode, hangs when rebooting to change memory modes. This represents only about a 5-10% performance hit at this time, but it requires additional administrative effort to address these issues. We appear to be making progress resolving some of the operational issues seen on KNL systems. For example, Christoph Lehner has successfully run 64-node jobs on JLab's 16p cluster. An upcoming round of BIOS upgrades and chip screening/replacement are expected to improve reliability.
- GPU systems: These have some added risk compared to CPUs, based on experience with the BNL-IC system (NVIDIA K80s) deployment and Pi0g (NVIDIA K40s) mid-life, due to packaging and heat management.

Element 6: The alignment of the computing hardware in the existing portfolio with vendor technology roadmaps; and with the technology roadmaps of leadership-class facilities at which USQCD collaboration members run scientific software codes;

- Conventional CPU architectures are not on the LCF roadmaps, but clearly will continue to be evolved by the vendors. The new Intel Skylake architecture itself is of interest because it supports the AVX-512 instruction set. The new AMD Ryzen CPUs appear on paper to be a worthy competitor to the Intel line.
- KNL is currently a major element of Intel's roadmap. We do expect to see some feature mixing and convergence between conventional Xeon and Xeon Phi line over time.
 - KNL is used by NERSC's CORI Phase II, Argonne LCF Theta.
 - KNL has been in the field long enough that we have significant production experience with it now.
 - The Argonne LCF Aurora has been postponed, along with a decision on its architecture, so at present it isn't known whether KNL is on the path to future LCF hardware.
- NVIDIA Pascal GPUs are also in vendor and LCF roadmaps, and are available now.

- BNL IC allocation could maintain the total GPU performance level for now as older GPU clusters at FNAL and JLab are retired. Cost of doing so is unknown. BNL has begun to acquire NVIDIA Pascal GPUs for its IC.
- The next generation in the Tesla series, the NVIDIA Volta, has been announced. It is advertised to deliver 5X the performance of the Pascal, although relative LQCD performance remains to be measured. NVIDIA says it will be available in early 2018.

Element 7: Which USQCD-specific software benchmarks should be used in making the best-value assessment during the cluster evaluation process.

- Please see Section 3 “Deliverable” above.

5 Suggestion for Future Acquisition Evaluations

Suggestion: For future hardware selection and negotiated allocation agreements, it is essential to document the software portfolio by creating a list of USQCD production jobs with their performance characteristics and resource requirements. This will require effort from both users and the site managers to accomplish.

6 Appendix: LQCD-ext II FY17 Acquisition Evaluation Committee Charge

February 1, 2017
Revised February 20, 2017
Revision 2
October 5, 2017

Purpose and Background

On an annual basis, the LQCD-ext II Computing Project has typically executed one or more large purchases of computing hardware to augment the existing hardware portfolio operated by the project. The hardware portfolio is used by USQCD in support of its scientific program.

In fiscal year 2017, the project planned to execute two distinct hardware acquisitions.

1. The first acquisition occurred in October 2016: \$410K in project hardware funds were used to purchase an extension of the Knights Landing cluster implemented at JLab earlier in 2016, as outlined and in accordance with the FY16 Hardware Acquisition Plan. The purchase order was executed in October and the new hardware has been received, installed, and placed into production use.
2. The second acquisition was scheduled to occur in calendar year 2017, wherein \$750K in project hardware funds were to be used to purchase and deploy new computing hardware at Brookhaven National Laboratory (BNL). Following standard project practices, Bob Mawhinney prepared an initial draft of the FY17 Acquisition Plan with input from the rest of the project team. Bob is one of the LQCD-ext II co-Site Architects for BNL.

An LQCD Acquisition Review Committee was formed in January 2017 and a charge letter was issued on February 1, 2017. The committee, under the chairmanship of Rob Kennedy, began working in response to the charge.

On June 9, 2017, work by the initial committee was suspended due to the impact of the FY18 President's Budget Request (PBR) on future funding for the LQCD Computing Project. By this time, the committee had obtained a set of USQCD-specific benchmarks, analyzed potential hardware options, and begun preparing a summary of their work along with a recommendation on how to proceed with the second of the FY17 acquisitions.

Over the course of summer 2017, the DOE Office of High Energy Physics (OHEP) strongly encouraged the LQCD Project to transition from a dedicated cluster model to a new operating model under which project funds would be used to purchase computing cycles from institutional clusters (ICs) operating at BNL and FNAL.

On September 13, 2017, the LQCD Federal Project Director instructed the LQCD Contractor Project Manager to reconvene the FY17 acquisition process, with specific guidance to use the \$750K budget to purchase compute time on the BNL Institutional Cluster.

Due to various circumstances, committee membership has changed slightly and the committee is now being chaired by Carleton DeTar.

To reflect the shift from purchasing dedicated hardware to evaluating available computing options, the committee name has been changed to the LQCD Acquisition Evaluation Committee.

Charge

The Acquisition Evaluation Committee is asked to review the work completed by the original committee, consider changes that have occurred in the hardware landscape, and provide input into the FY17 computing hardware planning process to ensure strong alignment of the LQCD hardware portfolio with the anticipated computing needs of the USQCD scientific program. The intent is to help ensure that the project continues to make the most effective use of computing hardware funds to support and advance the scientific program.

Each committee member is asked to review supporting materials, provide input, actively participate in committee discussions, and where possible, provide USQCD-specific code samples that will help the project benchmark the performance of candidate hardware against USQCD needs. The committee is asked to consider:

- The near- and long-term demand (at a high level) for each hardware architecture in the existing portfolio and how the proposed procurement of additional compute cycles at BNL will augment or complement the existing hardware portfolio;
- Alternate computing architectures that may better meet USQCD needs, considering compatibility with the existing hardware portfolio and infrastructure, as well as the existing software portfolio;
- The availability of production software for use by the USQCD collaboration to effectively utilize the capabilities of the proposed procurement of additional compute cycles;
- The ability of the proposed acquisition, along with the existing hardware portfolio, to meet the established time-based performance goals for the computing project;
- The capability of the project team to effectively support the computing hardware in the proposed acquisition, in terms of 1) meeting system uptime target goals; and 2) supporting the user community in the use of the newly acquired hardware;
- The alignment of the computing hardware in the existing portfolio with vendor technology roadmaps; and with the technology roadmaps of leadership-class facilities at which USQCD collaboration members run scientific software codes;
- Which USQCD-specific software benchmarks should be used in making the best-value assessment during the cluster evaluation process.

Deliverables

- A set of USQCD-specific software benchmarks that can be used to evaluate the performance of candidate computing architectures.
- A brief, written report summarizing the evaluation committee's analysis of potential hardware architectures and an assessment of how effectively each potential architecture will meet the computing needs of the scientific program and augment the overall hardware portfolio.
- Recommendation(s) to the Project Manager on how best to proceed with the hardware acquisition.

Timeline

- The evaluation committee should assemble a set of appropriate benchmarks that can be used to effectively evaluate candidate hardware architectures.
- The evaluation committee should complete its full analysis and provide a final written report with recommendations to Bill Boroski, LQCD-ext II Project Manager, no later than November 2, 2017.

Membership

The evaluation committee comprises members of the LQCD-ext II project and USQCD Collaboration with an appropriate mix of relevant technical and scientific expertise to effectively evaluate the merits of the proposed acquisition plan. In addition, the committee includes a member of the USQCD Scientific Program Committee, so the interests, needs and input of that committee are represented and factored into the process.

The Chair of the committee is Carleton DeTar. The membership of the FY17 evaluation committee is as follows:

NAME	PROJECT ROLE	AFFILIATION	EMAIL
Carleton DeTar, Chair	Collaboration Representative	U. of Utah	detar@physics.utah.edu
Steve Gottlieb	Collaboration Representative	Indiana U.	sg@indiana.edu
Don Holmgren	HPC Consultant	FNAL	djholm@fnal.gov
Chulwoo Jung	Collaboration Representative	BNL	chulwoo@bnl.gov
James Osborn	Collaboration Representative	ANL	osborn@alcf.anl.gov
Amitoj Singh	Site Architect	FNAL	amitoj@fnal.gov
Chip Watson	Site Architect	JLab	watson@jlab.org
Frank Winter	Collaboration Representative	JLab	fwinter@jlab.org
Alex Zaytsev	Site Architect	BNL	alezayt@bnl.gov

Supporting Documentation

The following documentation will be provided to the evaluation committee as documents and information becomes available.

- LQCD-ext II Acquisition Strategy
- LQCD-ext II FY17 Acquisition Plan
- Performance Goals and Milestones for the LQCD-ext II Computing Project
- Anticipated Computing Needs of the Scientific Program (2017-2021)
- Performance data on USQCD applications running on the actual FY16 production hardware to compare actual performance against early benchmarks.

Requests for additional information should be made to the chairperson of the evaluation committee.

Revision History

Revision #	Description of Change	Date	Author
0	Original version	02/01/17	W, Boroski
1	Updated membership list. Frank Winter replaced Balint Joo. Alex Zaytsev replaced Shigeki Masawa.	02/20/17	W. Boroski
2	Updated membership table to reflect new chairperson and committee membership. Updated document to reflect shift in strategy from purchasing dedicated systems to purchasing time on institutional clusters.	10/05/17	W. Boroski