FY14 Alternatives Analysis
For the
Lattice QCD Computing Project Extension
(LQCD-ext)

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

Version 1.1

Revision Date
August 20, 2013

PREPARED BY:
Don Holmgren, FNAL

CONCURRENCE:

______________________________ __August 20, 2013____
William N. Boroski Date
LQCD-ext Contract Project Manager

Lattice QCD Computing Project Extension (LQCD-ext)
Change Log: Alternatives Analysis for FY14 Procurement

Revision No.	Description	Effective Date
0.1	Document created from FY13 document.	July 25, 2013
0.2	FY14 scenarios detailed and analysis added	July 26, 2013
0.3	Improvements (scaling discussions)	July 30, 2013
0.4	Added Executive Summary.	Aug 6, 2013
1.0	Added signature approval	Aug 6, 2013
1.1	Minor corrections (typos and changing “asqtad” to “HISQ”	Aug 20, 2013

Executive Summary

This document presents the FY14 analysis of alternatives for obtaining the computational capacity needed for the US Lattice QCD effort within High Energy Physics (HEP) and Nuclear Physics (NP) by the SC Lattice QCD Computing Extension Project (LQCD-ext). The alternatives are constrained to approximately fit within the current budget guidance of the project, ~$3.5M / year for the five years of the project (FY10-FY14), and in particular ~$1.97M for capital procurements in FY2014. This constraint provides adequate funding to meet the basic requirements of the field for enhanced computational capacity, under the assumption of access via the INCITE program to resources at the Argonne Leadership Computing Facility and the Oak Ridge Leadership Computing Facility, and under the assumption that a reasonable fraction of those resources are ultimately allocated to Lattice QCD.

All alternatives assume the continued operation of the existing resources from the FY09-FY13 LQCD Facilities. At present these resources constitute an aggregate conventional resource including the FY13 BlueGene/Q half-rack at BNL of approximately 88 teraflop/s sustained on LQCD benchmarks plus 812 GPUs deployed in 5 clusters with an effective capacity of 118 teraflop/s sustained.

In FY14, viable hardware options are conventional Infiniband clusters, accelerated clusters, and the IBM BG/Q. Specifically, the following options are considered: (1) A half-rack expansion of the BlueGene/Q hardware at Brookhaven National Lab (BNL) plus a conventional cluster at Fermilab (FNAL), (2) A half-rack expansion of the BlueGene/Q hardware at BNL plus a GPU-accelerated cluster at FNAL, (3) A conventional cluster at Fermilab, (4) A GPU-accelerated cluster at FNAL, and (5) A mixture of a conventional and a GPU-accelerated cluster at FNAL, each consuming half of the budget.

The performance per core delivered by an LQCD computing system on a fixed size problem will decrease on computing hardware as the number of cores or GPUs used is increased; the overall decrease in time-to-solution as the number of cores is increased is known as "strong scaling." Of the systems under consideration for FY14, the BlueGene/Q hardware under consideration has the best strong scaling performance. If strong scaling effects are not taken into account, the best alternative is the 50:50 budget split between a conventional cluster and a GPU-accelerated cluster. Although the scenario with only a GPU-accelerated cluster has the best overall price/performance of the alternatives, it would result in a portfolio of resources that would be unbalanced, based on the allocation requests received for the 2013-2014 allocation year in which the most oversubscribed resources were the conventional clusters. The 50:50 budget distribution could be modified to best match the anticipated demands, taking into account the hardware options available for purchase at the beginning of FY14.

If strong scaling effects are taken into account, the overall effective computing capacity of the portfolio of USQCD dedicated hardware including the FY14 purchase would be optimized through the purchase of the BlueGene/Q expansion only if the overall fraction of large jobs exceeds about 27%. Based on actual running on the existing hardware from July 2012 through February 2013, however, demand for large jobs on the clusters is quite small. Of the 27 class-A

proposals for the 2013-2014 allocation period, only 3 requested and were allocated time on the FY13 BlueGene/Q half-rack. We believe that for at least the next year, the existing BG/Q half-rack will meet the community needs for large jobs.

We believe that the 50:50 conventional:accelerated mix will best optimize the USQCD ensemble of dedicated hardware.

Table of Contents

1 Introduction

This document presents the FY14 analysis of alternatives for obtaining the computational capacity needed for the US Lattice QCD effort within High Energy Physics (HEP) and Nuclear Physics (NP) by the SC Lattice QCD Computing Extension Project (LQCD-ext). This analysis is updated at least annually to capture decisions taken during the life of the project, and to examine options for the next year. The technical managers of the project are also continuously tracking market developments through interactions with computer and chip vendors, through trade journals and online resources, and through computing conferences. This tracking allows unexpected changes to be incorporated into the project execution in a timely fashion.

Alternatives herein are constrained to approximately fit within the current budget guidance of the project, ~$3.5M / year for the five years of the project (FY10-FY14), and in particular ~$1.97M for capital procurements in FY2014. This constraint provides adequate funding to meet the basic requirements of the field for enhanced computational capacity, under the assumption of expanding resources at ANL and ORNL already planned by the Office of Science (SC), and under the assumption that a reasonable fraction of those resources are ultimately allocated to Lattice QCD.

All alternatives assume the continued operation of the existing resources from the FY09-FY13 LQCD Facilities Projects until those resources reach end of life, i.e., until each resource is no longer cost effective to operate, or about 4 years for clusters. At present these resources constitute an aggregate conventional resource including the FY13 BlueGene/Q half-rack at BNL of about 88 teraflop/s sustained on LQCD benchmarks plus 812 GPUs deployed in 5 clusters with an effective capacity of 118 teraflop/s sustained. The aggregate project cost of operating these existing systems in FY2014 is approximately $1.32M (for the three sites combined). Replacing and running the flexible computational capacity represented by these existing resources cannot be done for less than its current operating cost.

In FY14, viable hardware options are conventional Infiniband clusters, accelerated clusters, and the IBM BG/Q. Conventional clusters can run codes for all actions of interest to USQCD. Optimized codes for the BG/Q are currently available for the DWF, Wilson, twisted mass and HISQ actions, with code for the clover action under development. Optimized multi-GPU codes are available for the HISQ, Wilson, clover, and twisted mass actions, with code for the DWF action under development.

2 FY14 Goals

The project baseline called for deployment in FY14 of conventional capacity totaling 57 TF. In FY13, after adjustments to the plan to increase the amount of storage to be purchased, to adjust labor based on FY10-FY12 experience, to absorb the JLab ARRA project clusters, and to take advantage of the performance of GPU-acceleration, the FY14 goal was revised to increasing conventional capacity by between 19 and 28 TF, and accelerated capacity to between 40 and 60 effective teraflops. Here, the effective GPU TFlops are based on benchmarks developed in FY13 to assess the performance of the NVIDIA GPUs used on the various project clusters on HISQ, clover, and DWF applications. The ranges reflect that the project will choose the relative ratios of conventional and accelerated resources so that the resulting total portfolio of hardware best matches USQCD needs. In FY14, the project will also decommission 8.4 teraflops of conventional cluster capacity (the JPsi cluster at Fermilab); note that the project had planned in

FY13 to decommission this cluster in part to free computer room space at Fermilab for the planned FY13 hardware, but elected to operate this system one additional year as the space was not needed due to the decision to deploy a half-rack of BG/Q hardware at BNL.

Sustained performance on conventional clusters is defined as the average of single precision DWF and improved staggered ("HISQ") inverter performance on jobs utilizing 128 MPI ranks. "Linpack" or "peak" performance metrics are not considered, as lattice QCD codes uniquely stress computer systems, and their performance does not uniformly track either Linpack or peak performance metrics across different architectures. Note that GPU clusters or other accelerated architectures are evaluated in such a way as to take into account the Amdahl's Law effect of not accelerating the full application, or of accelerating the non-inverter portion of the code by a smaller factor than the inverter, to yield an "effective" sustained teraflops, or an equivalent cluster sustained performance.

The goal for FY14 is to install these new resources by May 31, 2014 and release them to full production by July 1, 2014.

3 Hardware Options

> *Each year the project will optimize the next procurement to yield an ensemble of hardware resources that achieves the highest performance for the portfolio of projects that USQCD intends to execute. This may include procuring two different types of computer systems in a single year.*

The following types of hardware are considered in this analysis:

1. Expansion of the current half-rack (512 nodes) IBM BG/Q system, deployed at Brookhaven National Laboratory, to a full rack (1024 nodes).
2. A conventional cluster, based on Intel or AMD processors with Infiniband communications, deployed at Fermilab.
3. An accelerated cluster, based on Intel or AMD processors with Infiniband communications and with NVIDIA GPU accelerators, deployed at Fermilab.

BG/Q

Benchmarking data for DWF and HISQ inverters on BG/Q hardware have been provided to the LQCD-ext project. For DWF, the average of the performance of Ls=8 and Ls=16 Möbius and Shamir implementations running using 4^4 per core local volumes (each node of a BG/Q has 16 cores) is 59.5 GF/node, single precision. Single precision HISQ performance is 26 GF/node. Note that these figures correspond to 29% and 12.7% of peak, respectively. The average of DWF and HISQ performance, used in the rest of this document, is 42.75 GF/node or 2.67 GF/core.

Based on a cost estimate from BNL, the price of the half-rack expansion (512 additional nodes) of BG/Q to the project will be $1.276M. This includes G&A and the first year maintenance cost.

Conventional Clusters

USQCD has tracked price/performance on LQCD Infiniband-based conventional clusters deployed at Fermilab or JLab since 2005. The plot below shows these cost trends, along with exponential fits to two subsets of the data. Also included are data and an extrapolation line for GPU-accelerated clusters.

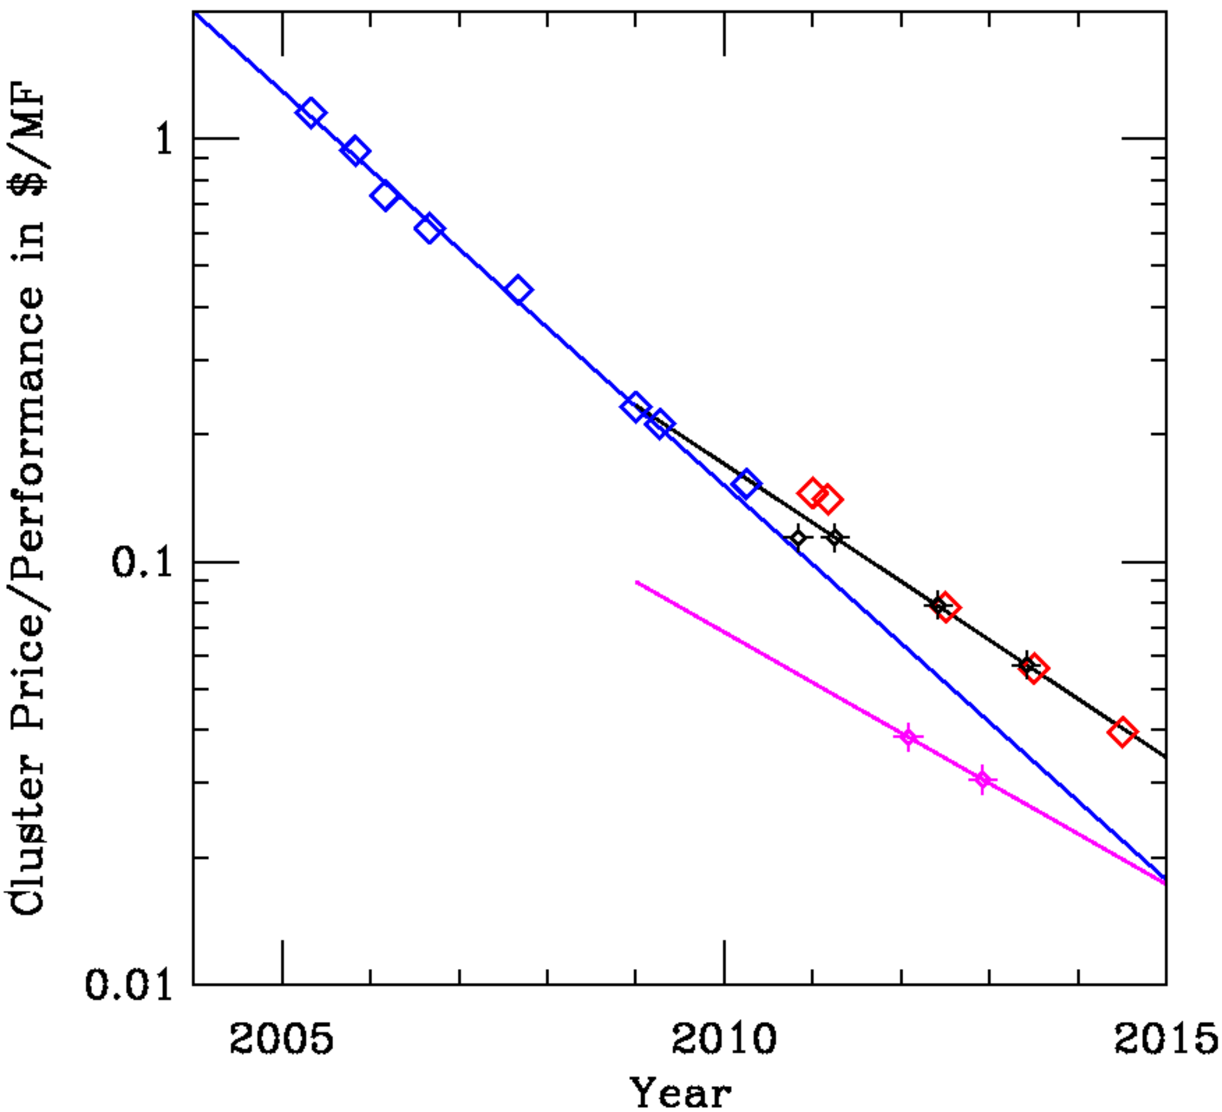

Here, the blue line is the least-squares fit to the clusters purchased between 2005 and 2011, shown as blue diamond symbols. The red diamond symbols are baseline goals used in the project plan. The black line is the fit to the points from 2009 through the FY13 cluster, Bc. The magenta line connects the points corresponding to the latest GPU clusters, Dsg and 12k; this line will be used for accelerated-cluster projections later in this document.

Using the two fits, in mid-2014 the price/performance values corresponding to the blue and black lines are $0.0234/MF and $0.0417/MF. An FY14 conventional cluster will likely have price/performance between these values, but to be conservative we will project performance using the poorer (larger) figure. For an FY14 purchase, new Intel "Ivy Bridge" dual- and quad-socket systems will be available. Because these will be available at a higher memory clock speed (1866 MHz) than was previously available (1600 MHz) with Sandy Bridge processors, we expect a boost in performance on lattice QCD applications compared to prior Intel systems, such as those used on 12s. The introduction of a new processor family will also put downwards price pressure on Sandy Bridge systems. Both developments argue for the projected drop in price/performance.

GPU Accelerated Clusters

For those calculations for which optimized software is available, GPU-accelerated clusters offer a substantial improvement in price/performance compared with conventional clusters. The latest clusters, Dsg at Fermilab (delivered January 2012) and 12k at JLab (delivered November 2012) have price/performance values of $0.0385 and $0.0306, respectively, based on a suite of application benchmarks that measure throughput with HISQ, clover, and DWF actions. By mid-FY14, using an extrapolation through these two data points, a GPU-accelerated cluster is estimated to have $0.020/MF price/performance. By the time of purchase, NVIDIA will have introduced a re-spin of the "Kepler" GPU that was used on 12k. Based on a similar re-spin on the prior generation "Fermi" chips, the new chips should have a performance increase for similar cost, and the older accelerators will decrease in price.

Scaling

The BG/Q exhibits flat weak scaling within a rack (or half-rack), so sustained performance, as a function of job size (expressed in core count), does not vary up to the 8192 cores (2.67 GF/core) in a half-rack. A weak scaling measurement is shown in the plot below. Note that flat scaling on the BG/Q extends to calculations utilizing many racks.

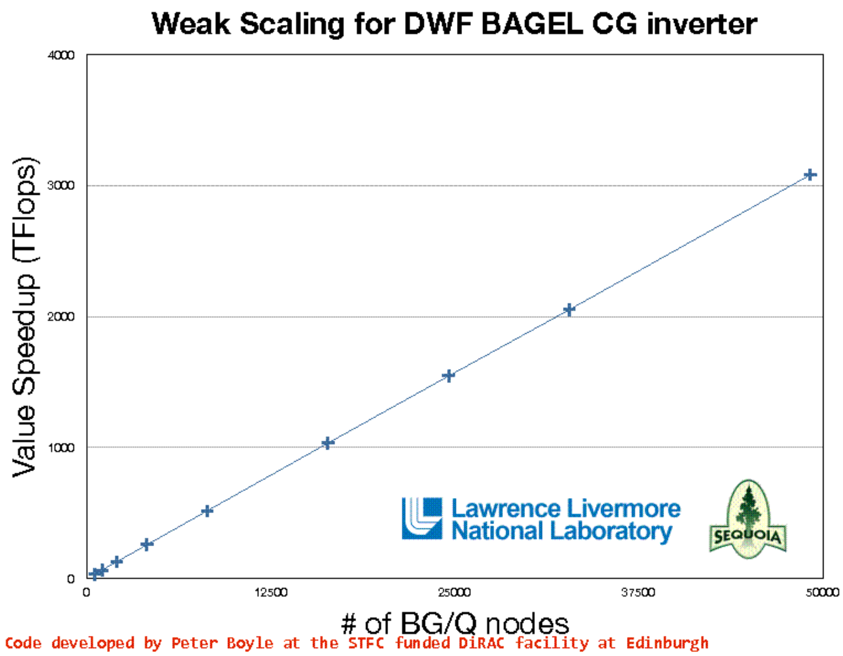

Because of longer network latencies, a conventional Infiniband cluster does not have flat weak scaling, but rather performance per core drops as job core counts increase (in weak scaling, the work accomplished per core is fixed as the number of cores are increased). The plot below shows weak scaling on the improved staggered inverter on the Ds cluster as a function of local (per core) lattice sizes, as the problem size is varied corresponding to jobs sizes of between 32 and 4096 cores. The Ds cluster, purchased in 2010, has 2.0 GHz AMD Magny-Cours nodes, each with 32 cores. The cores each sustain 1.59 GF on parallel (128 process) jobs.

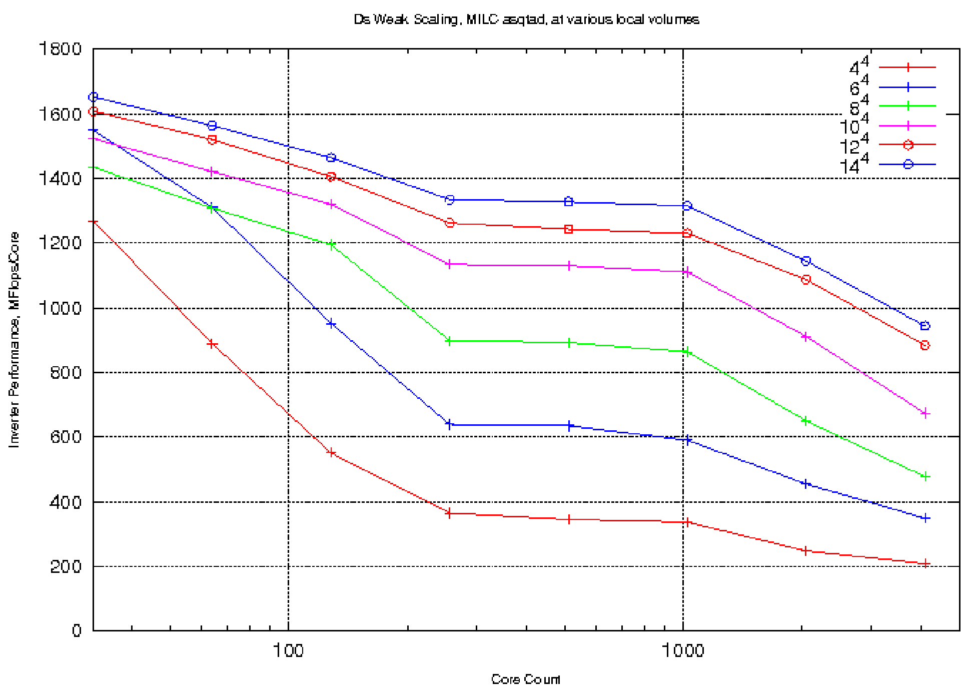

The USQCD cluster ratings use job sizes of 128 processes. For local volumes of 14^4, performance per core of the HISQ inverter on jobs from 128 through 4096 cores on Ds is as follows:

- 128 cores: 1463 MF
- 256 cores: 1334 MF (91.1% of 128-core performance)
- 512 cores: 1327 MF (90.7% of 128-core performance)
- 1024 cores: 1314 MF (89.9% of 128-core performance)
- 2048 cores: 1144 MF (78.2% of 128-core performance)
- 4096 cores: 943 MF (64.4% of 128-core performance)

For jobs requiring large core counts, either because of memory or time-to-solution considerations, the Ds cluster drops to 64.4% of its 128-core job efficiency at 4096 cores in the case of weak scaling. In strong scaling, where a fixed size lattice is solved on an increasing number of cores so that volume per core drops, the drop in per core performance relative to 128-core performance will be even greater; this will be true for all hardware under consideration, but the amount of decrease will be greater for architectures that show larger rates of decrease with increasing core count in weak scaling data. From the weak scaling Ds plot, we can estimate the drop in strong scaling performance for given lattice sizes. For example, a 14^4/core problem on 128 cores spread out to 4096 cores would have a factor of 32 reduction in local volumes, to 5.89^4/core. Using the 14^4 and 6^4 weak scaling lines, the 14^4/core problem achieves 1463 MF/core using 128 total cores, and would achieve approximately 350 MF/core using 4096 cores. For this example, aggregate performance increases by a factor of 7.66 from 0.187 TF to 1.43 TF when core counts are increased by a factor of 32 from 128 to 4096; the performance per core drops by 76%. For 2048 core jobs, the estimated drop per core is 62%, and for 1024 core jobs, the estimated drop is 41%.

Since the Ds cluster was purchased, the LQCD-ext project has procured and installed the 12s cluster (FY12) at JLab, and the Bc cluster (FY13) at Fermilab. Bc is similar to Ds, with 32 AMD cores per node. Each core runs at 2.8 GHz, and faster memory is used, but because these processors are two generations newer ("Abu Dhabi", model 6320, instead of the "Magny Cours", model 6128 processors on Ds) and cores in the new models share a single faster SIMD floating point unit, rather than having a floating point unit per core, sustained performance per core on Bc is only slightly better than Ds (1.77 GF/core versus 1.59 GF/core) and weak scaling behavior is similar. The Intel E5-2650 2.0 GHz processors used on the 12s cluster, by contrast, have better memory bandwidth per core and each core contains a floating-point unit. Per core, on LQCD codes the 12s cluster sustains 2.98 GF, nearly twice as much as the Ds cluster. Therefore for similarly sized jobs in terms of aggregate sustained TFlops, the process count on 12s is half that used on Ds. Because of the higher performance per socket on Intel clusters, weak scaling performance at similar aggregate job throughput is better than on Ds-like clusters. However, Intel systems are more costly; on the Bc cluster purchase, the bid cost per core for Intel-based solutions was more than twice that for AMD-based solutions, and so the award was made to a vendor who supplied an AMD-based cluster. For the purposes of the scaling discussions in the rest of this document, we will assume an AMD-based solution as a worst case.

For GPU-accelerated systems, weak scaling plots are not readily available, but the plots below show, respectively, strong scaling of the HISQ inverter on a dual-GPU per node cluster similar to Dsg, and strong scaling of the Clover inverter using the Chroma code on Cray supercomputers at

Oak Ridge National Laboratory. On the HISQ plot, performance/GPU on 128 GPUs is 81.5% of performance/GPU on 64 GPUs, and performance/GPU on 256 GPUs is 64.0% of performance/GPU on 64 GPUs. An important consideration to note, however, is that for the analysis computations to be performed on any GPU-accelerated cluster under consideration in this document, job size in terms of GPU count will nearly always be determined by the minimum number of GPUs necessary for the required memory footprint. In the discussion of alternatives below, we use a GPU count of 256, which is larger by a factor of 4 than the largest GPU analysis jobs run in the past year on the GPU-accelerated LQCD-ext clusters. Scaling analysis using the larger count is therefore a worst-case.

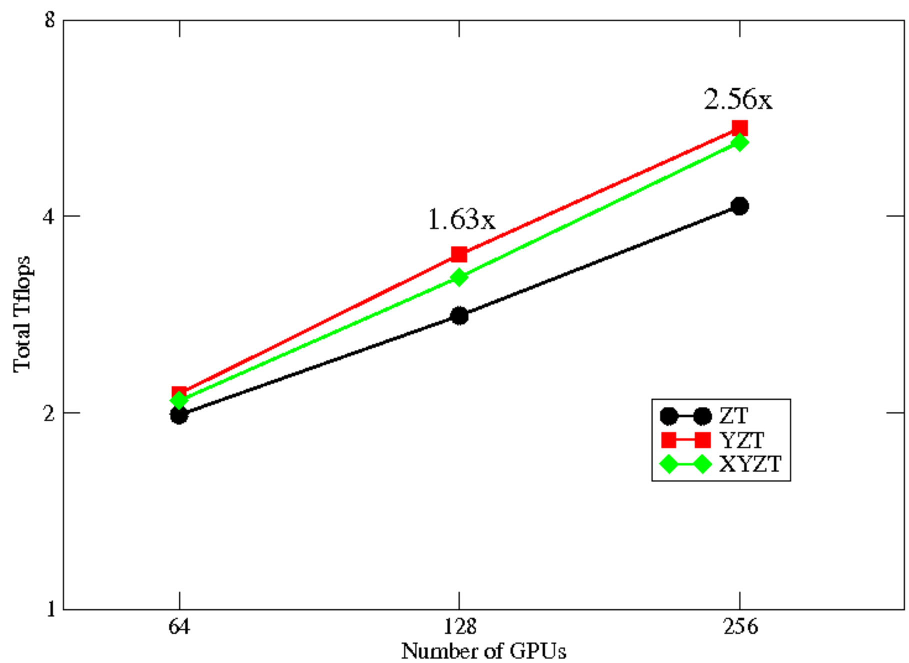

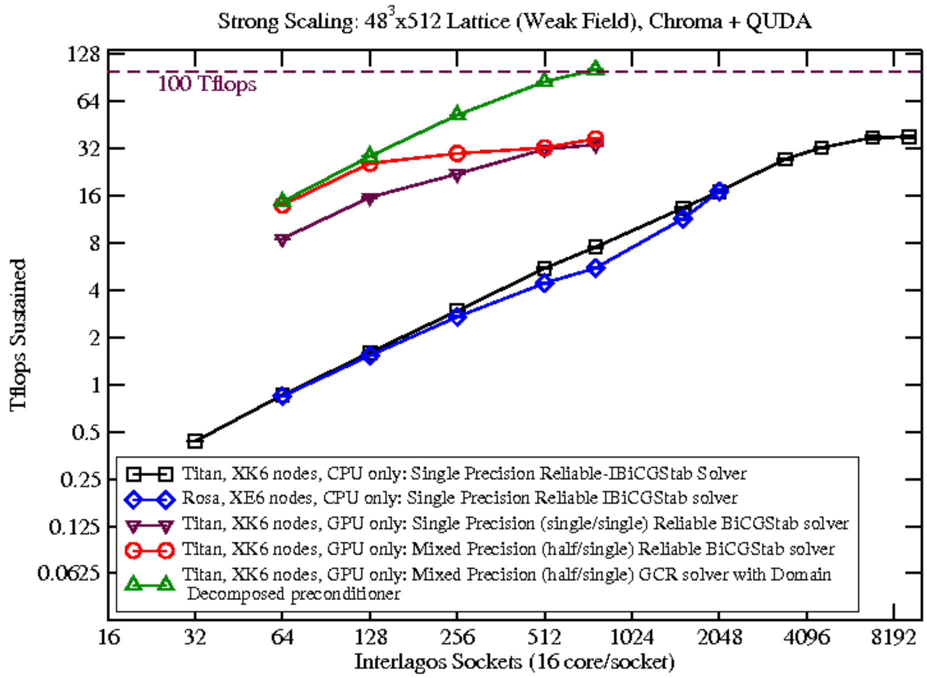

4 Alternatives

The following sections summarize the alternative technologies considered to achieve some or all of the stated performance goals of this investment for FY14.

4.1 *Alternative 1: A half-rack of BG/Q deployed in Q1 2014, and a conventional cluster released to production by July 1, 2014.*

Expand the existing half-rack of BG/Q to a full rack in the first calendar quarter of 2014, with the expansion hardware capable of sustaining at least 21.9 teraflop/s for a total M&S cost of $1.276M, and deploy a conventional cluster by July 1, 2014 capable of sustaining at least 16.6 teraflop/s for a total M&S cost of $0.694M, for a total of 38.5 teraflop/s incremental capacity.

The incremental three-year lifecycle cost of this alternative is estimated as follows:

- Procure and install a 21.9 TF BG/Q ($1.276M)
- Operations of the BG/Q at $70K/year for a total of $0.210M
- 2nd and 3rd year hardware maintenance for the BG/Q estimated at $114K per year.
- Procure and install a 16.6 TF conventional cluster ($0.694M)
- Incremental procurement and operations costs for the conventional cluster for three years ($117K + $75K + $75K = $0.267M)
- Three-Year Lifecycle cost: $2.675M ($1.714M BG/Q and $0.961 cluster)

Analysis: The hardware costs for this alternative are within the FY14 project budget. Because limited funds would be available after the BG/Q purchase, only a conventional cluster is purchased rather than a mixture of conventional and accelerated clusters. This alternative therefore meets only the conventional hardware deployment goal of 19 to 28 TF, and does not address the GPU-accelerated deployment goal of a new machine with effective throughput of between 40 and 60 TF.

For 128-core jobs, the price/performance of the BG/Q portion of the lifecycle cost is $0.078/MF, and that of the conventional cluster is $0.058/MF, with the overall price/performance $0.069/MF. For 2048-core (1.59 GF/core) jobs, the price/performance of the BG/Q portion is unchanged, and that of the conventional cluster increases because of strong scaling losses (38% performance/core compared to 128-core jobs) to an estimated $0.155/MF, with an overall price/performance of $0.092/MF.

4.2 Alternative 2: A half-rack of BG/Q deployed in Q1 2014, and an accelerated cluster released to production by July 1, 2014.

Expand the existing half-rack of BG/Q to a full rack in the first calendar quarter of 2014, with the expansion hardware capable of sustaining at least 21.9 teraflop/s for a total M&S cost of $1.276M, and deploy a GPU-accelerated cluster by the end of FY14 of capacity to 32.6 effective teraflop/s for a total M&S cost of $0.658M.

The incremental lifecycle cost of this alternative is estimated as follows:

- Procure and install a 21.9 TF BG/Q ($1.276M)
- Operations of the BG/Q at $70K/year for a total of $0.210M
- 2nd and 3rd year hardware maintenance for the BG/Q estimated at $114K per year.
- Procure and install a 34.7 TF (effective) accelerated cluster ($0.694M)
- Incremental procurement and operations costs for the accelerated cluster for three years ($117K + $75K + $75K = $0.267M)
- Three-Year Lifecycle cost: $2.675M

Analysis: The hardware costs for this alternative are within the FY14 project budget. Because limited funds would be available after the BG/Q purchase, only an accelerated cluster is purchased rather than a mixture of conventional and accelerated clusters. This alternative does not meet the GPU-accelerated deployment goal or the conventional cluster deployment goal.

The price/performance of the BG/Q portion of the lifecycle cost is $0.078/MF, and that of the accelerated cluster is $0.028/effective MF, with the overall price/performance $0.047/MF. For large GPU-count jobs, based on strong scaling measurements the performance per GPU drops by 36% from 64-GPU-count to 256-GPU-count problems. There is at least as large a fractional performance drop from small GPU-counts (4 to 8) to 64-GPU-count jobs. Assuming an overall performance drop due to strong scaling of 60%, for large GPU count jobs the price/performance of the accelerated cluster increases to $0.069/MF, with an overall price/performance including the BG/Q of $0.075/MF.

4.3 Alternative 3: A conventional cluster released to production by July 1, 2014.

Deploy a conventional cluster by the end of June 2014 capable of sustaining at least 47.2 teraflop/s with an M&S cost of $1.97M.

The incremental three-year lifecycle cost of this alternative is estimated as follows:

- Procure 47.2 TF in FY14 ($1.97M)
- Incremental procurement and operations costs for three years ($234K + $151K +$151K = $0.536M total)
- Three-Year Lifecycle cost: $1.970M + $0.536M = $2.506M

Analysis: The hardware costs for this alternative are within the FY14 project budget. The overall price/performance for 128-core jobs is $0.053/MF. For 2048-core (1.59 GF/core) jobs, the price/performance of this alternative increases to $0.140/MF, assuming a 62% drop in performance per core relative to 128-core jobs.

This alternative exceeds the FY14 conventional cluster goal, but it does not address the GPU-accelerated cluster goal.

4.4 Alternative 4: A GPU-accelerated cluster released to production by July 1, 2014.

Deploy a GPU-accelerated cluster by the end of June 2014 sustaining 98.5 effective TF, with an M&S cost of $1.970M.

The incremental lifecycle cost of this alternative is estimated as follows:

- Procure a 98.5 effective TF GPU cluster in FY14 ($1.970M)
- Incremental procurement and operations costs for three years ($234K + $151K +$151K = $0.536M total)
- Three-Year Lifecycle cost: $1.970M + $0.536M = $2.506M

Analysis: The hardware costs for this alternative are within the FY14 project budget. The overall price/performance is $0.025/effective MF. Assuming a 60% drop in per GPU performance for large GPU-count jobs, the overall price/performance is $0.064/effective MF for 256-core problems.

This alternative exceeds the FY14 GPU-accelerated cluster goal, but does not address the conventional cluster goal.

4.5 Alternative 5: 50:50 (by budget) mixture of Conventional and GPU-Accelerated Clusters

Deploy a conventional and a GPU-accelerated cluster by the end of June 2014 capable of delivering respectively at least 23.6 TF and 49.2 effective TF, at an M&S cost of $1.970M.

The incremental lifecycle cost of this alternative is estimated as follows:

- Procure a 23.6 TF conventional cluster in FY14 ($0.985M)
- Procure a 49.2 effective TF GPU-accelerated cluster ($0.985M)

- Incremental procurement and operations costs for three years ($234K + $151K +$151K = $0.536M total)
- Three-Year Lifecycle cost: $1.970M + $0.536M = $2.506M

Analysis: The hardware costs for this alternative are within the FY14 project budget. The 128-core job price/performance of the conventional cluster portion is $0.053/MF, and that of the GPU-accelerated cluster $0.025/effective MF, for an overall price/performance of $0.034/MF. For 2048-core (1.59 GF/core) jobs, the conventional cluster price/performance increases to $0.140/MF, and assuming a 60% drop in per GPU performance for large GPU-count jobs, the price/performance for the GPU-accelerated cluster increases to $0.064/effective MF for 256-GPU problems. The overall price/performance for such large jobs is $0.087/MF.

This alternative would meet both the FY14 conventional and GPU-accelerated cluster deployment goals.

4.6 Alternative 6: Status Quo (no additional deployment in FY14)

Continue to operate the existing project clusters deployed at FNAL and JLab.

The cost of this alternative is $1.28M in FY2014 to operate the existing facilities. The incremental cost of this alternative (new investment) is $0.

Analysis: This alternative is included only for completeness and would not be capable of providing the necessary computational capacity to achieve the scientific goals of this project. Specifically, it would not leave USQCD with sufficient capacity to exploit the configuration generation capability of the supercomputers that DOE and NSF will or have released to production during FY13.

4.6 Other Alternatives

Other alternatives may be relevant for consideration in future years. These were not considered for detailed analysis at this time, as their current state of maturity was not deemed sufficient. Each of these alternatives functions as a co-processor with limited memory size, and so could most likely only be used to accelerate floating-point intensive kernels. The alternatives include:

- Intel Xeon Phi processor based systems: these will be similar to GPUs initially (PCI based accelerators). Software maturity is currently not sufficient to consider large-scale deployment for LQCD production.
- Hybrid processors (CPU cores + accelerated cores): while such systems are beginning to emerge at the low end for low power devices (tablets), these are still future products for the High Performance Computing space.

5 Discussion

The goal of this alternatives analysis is to select the purchase scenario which best optimizes the portfolio of USQCD dedicated resources. All of the scenarios considered place hardware either at Fermilab (conventional and GPU-accelerated clusters), or at both Fermilab (clusters) and BNL (a half-rack expansion of the existing BG/Q).

Since a half-rack expansion of BG/Q has a fixed price, all of the alternatives either direct the entire \$1.970M FY14 hardware budget to Fermilab, or direct \$1.276M to BNL for the BG/Q and \$0.694M to Fermilab. As covered at the FY13 project annual review, and endorsed by the review panel, the determination of the optimal split in Fermilab funds between conventional and accelerated clusters can be delayed into FY14. This flexibility allows the project to take into account information that is not currently available, such as the release schedule, performance, and pricing of hardware from the various vendors of interest (AMD, Intel, NVIDIA, and systems integrators utilizing components from these manufacturers).

Of the hardware alternatives, GPU-accelerated clusters have the narrowest applicability to LQCD calculations. This results from software availability and the suitability of LQCD algorithms to heterogeneous computing platforms; future software development from outside of the LQCD-ext hardware project will likely increase the fraction of calculations that can take advantage of GPU acceleration, but at the present much of the “low hanging fruit” has been harvested. For those calculations that can take full advantage of GPU acceleration, the gain in cost efficiency is very high. At present, based on the ratio of allocation requests to available resources, conventional cluster resources are more heavily oversubscribed than are GPU-accelerated or BG/Q resources. For the 2013-2014 allocation year, there were 21 cluster allocation requests with an aggregate resource total of 155% of the available resources. For GPU-accelerated resources, 8 requests had an aggregate resource total that was 124% of the available resources. For the BlueGene/Q half-rack, 3 requests had an aggregate resource total that was 126% of the available resources. The purchase of only GPU-accelerated hardware in FY14 would significantly over-provision this type of hardware, and the resulting USQCD hardware portfolio would not be balanced against demand.

Conventional clusters have the widest applicability to LQCD calculations of the three hardware alternatives. In contrast to GPU-accelerated clusters, optimized applications are available and in production for all fermion actions of interest. Unlike BG/Q hardware, very small (*e.g.* single computer) jobs can be run easily.

The table below shows price/performance for the various hardware options. For all hardware types, both a small job and a large job estimate are shown. The small job estimate assumes that all conventional resource (cluster and BG/Q) jobs use 1024 or fewer cores. The large job estimate assumes that all jobs use 2048 cores (1.60 GF/core); as discussed above, for an Intel-based cluster a 1024-core job would have higher throughput than a 2048-core job on an AMD-based cluster like Ds, so this is a worst-case analysis. The large job estimates use degradation assumptions of 62% for clusters, 60% for GPUs, and no degradation for BG/Q. Smaller scaling penalties would apply for an Intel-based cluster or for more typical 64-count GPU jobs. In this table, the BG/Q and conventional cluster performance estimates use the standard HISQ:DWF average. The BG/Q estimates use measured performance values (see Section 3 above), and the conventional cluster estimates use the upper price/performance trend line (the fit to USQCD clusters from 2009 to 2013).

Considering small jobs only, the best overall price/performance figures are those of the GPU-accelerated cluster and the 50:50 conventional/GPU-accelerated cluster mix. Considering large jobs only, the best overall price/performance figures are those of the GPU-accelerated cluster and the additional half-rack BG/Q plus accelerated cluster. As noted above, the purchase of only a

GPU-accelerated cluster would leave the overall portfolio of USQCD unbalanced against allocation requests.

Scenario	HW Cost	Ops Cost	Total Cost	Perf Small	P/P Small Jobs	Perf Large	P/P Large Jobs
Half Rack BG/Q	$1.276	$0.438	$1.714	21.9	$0.078	21.9	$0.078
Conventional Cluster	$0.694	$0.267	$0.961	16.6	$0.058	6.3	$0.152
Overall	$1.970	$0.705	$2.675	38.5	$0.069	28.2	$0.095
Conventional Cluster	$1.970	$0.536	$2.506	47.2	$0.053	17.9	$0.140
Half Rack BG/Q	$1.276	$0.438	$1.714	21.9	$0.078	21.9	$0.078
Accelerated Cluster	$0.694	$0.267	$0.961	34.7	$0.028	13.9	$0.069
Overall	$1.970	$0.705	$2.675	56.6	$0.047	35.8	$0.075
GPU Cluster	$1.970	$0.536	$2.506	98.5	$0.025	39.4	$0.064
50% Conventional	$0.99	$0.27	$1.253	23.6	$0.053	9.0	$0.140
50% GPU	$0.99	$0.27	$1.253	49.2	$0.025	19.7	$0.064
Overall	$1.970	$0.536	$2.506	72.8	$0.034	28.6	$0.087

In communications in July 2013 from the USQCD Executive Committee and from the chair of the Scientific Program Committee, we learned that they anticipated in FY13 and later years that a significant fraction of the jobs to be run on dedicated hardware operated by this project will be large, requiring 2K and higher core counts. To study the consequences of a mixture of large and small jobs types, we examine the portfolio of dedicated hardware under two scenarios: the purchase of the half-rack expansion of BG/Q plus a conventional cluster, and a conventional cluster. The capacity of the hardware in these two resources will join with the existing Bc, Ds, 9q, 10q, and 12s clusters (considering conventional resources only) and the half-rack of BG/Q installed at BNL in FY13. We posit that only large jobs would run on the BG/Q, and that the per-node performance of the BG/Q is the same for large and small jobs. We also suppose that the effective capacity of the cluster resources is independent of job size for small jobs, and for large jobs the capacity of the cluster resources is reduced by the same 62% factor observed on the Ds cluster when comparing 2048-core HISQ-action jobs with 128-core jobs, and by the same 76% factor when comparing 4096-core HISQ-action jobs. The plots below show the effective total capacity, summing the existing cluster resources and the new resources in the two scenarios, as a function of the fraction of large jobs. In this model, all large jobs are assigned to the BG/Q hardware unless the fraction of large jobs exceeds the ratio of BG/Q to aggregate conventional cluster capacity; in this latter case, as many large jobs as possible are assigned to the BG/Q, and the rest are run on clusters.

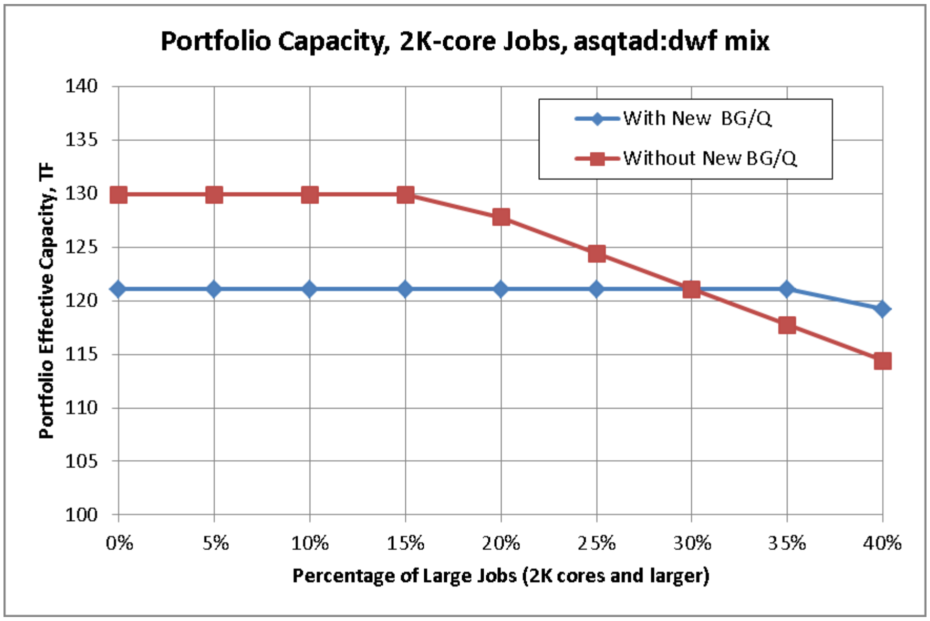

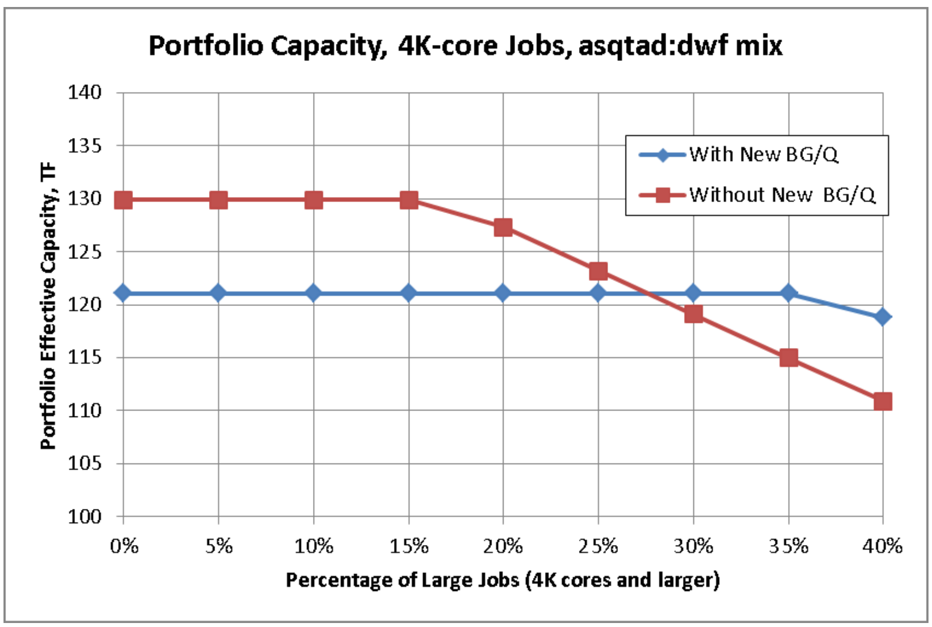

In both cases in the plots above, for the conventional-only hardware ("Without New BG/Q") as the large-job fraction increases, the effective capacity of the USQCD dedicated hardware portfolio begins to decrease once fraction of large jobs exceeds the capacity of the existing half-rack of BG/Q. In the "With New BG/Q" half-rack curves, the effective capacity is independent of the fraction of large jobs until that fraction exceeds the ratio of BG/Q capacity to non-BG/Q capacity

(35%); beyond that fraction, increasing numbers of large jobs are run on the conventional resources. From this model, if the percentage of large jobs exceeds between 27% and 30%, then the effective capacity of the ensemble of USQCD dedicated hardware is larger for the alternative with the BG/Q half-rack addition.

The two histograms below show the statistics for jobs run on the conventional clusters at Fermilab and JLab from July 2012 through February 2013. Although there is some weight at job sizes using 1024 cores and larger at Fermilab, it is clear that large a large fraction of production on these clusters does not involve the large (2048+) core-count jobs analyzed in the model discussed above. This may indicate that the transition to large jobs has occurred more slowly than estimated in the summer of 2013. Further, an Intel-based conventional cluster purchased in FY14 would tend to push the weights of the bins in the Fermilab histogram to the left, as for jobs with similar TFlops throughput, assuming that each host has sufficient memory, half the core-count would be needed on the FY14 cluster.

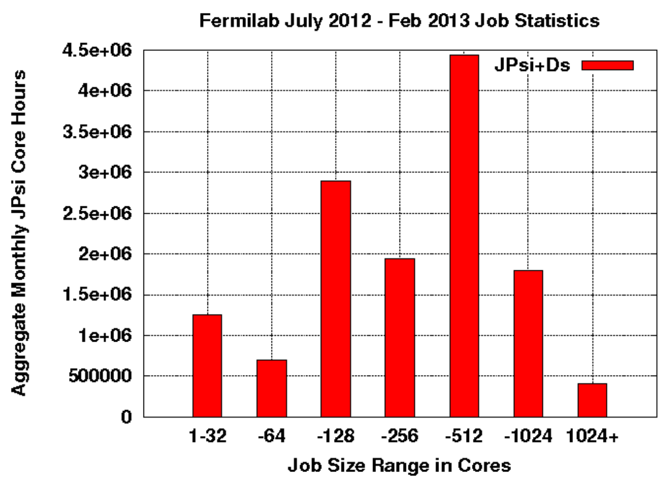

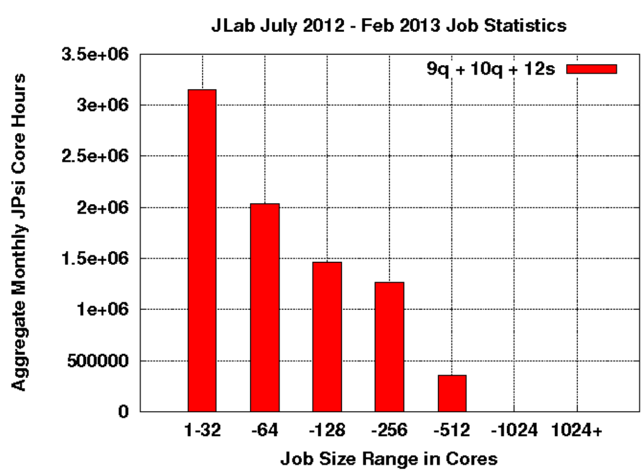

6 Conclusion

If strong scaling effects are not taken into account, the best alternative is given in section 4.5 above, a 50:50 budget split between a conventional cluster and a GPU-accelerated cluster. Although the scenario with only a GPU-accelerated cluster has the best overall price/performance of the alternatives, it would result in a portfolio of resources that would be unbalanced, based on the allocation requests received for the 2013-2014 allocation year. The oversubscription rates for the different types of hardware (request divided by allocation) were 155%, 126%, and 124%, respectively, for conventional, BG/Q, and GPU-accelerated resources. The planned retirement in FY14 of the 5-year old Fermilab JPsi conventional cluster will exacerbate this imbalance. Compared to the GPU-only option, the 50:50 mix has the next best price/performance. The 50:50 budget distribution could be modified to best match the anticipated demands, taking into account the hardware options available for purchase at the beginning of FY14.

If strong scaling effects were taken into account, the overall effective computing capacity of the portfolio of USQCD dedicated hardware would be optimized through the purchase of the additional BG/Q half-rack only if the overall fraction of large jobs exceeded about 27%. Based on actual running on the existing hardware from July 2012 through February 2013, however, apparent demand for large jobs on the clusters is quite small. The FY13 BG/Q half-rack was released to production quite recently, on July 1, 2013. Of the 27 class-A proposals for the 2013-2014 allocation period, only 3 requested and were allocated time on this half-rack. We believe that for at least the next year, the existing BG/Q half-rack will meet the community needs for large jobs.

We believe that the 50:50 conventional:accelerated mix will best optimize the USQCD ensemble of dedicated hardware.